



# Opinion

## The Agentic Economy

*The architecture of agentic communication will determine the extent to which generative AI democratizes access to economic opportunity.*

**G**ENERATIVE AI HAS revolutionized the way we interact with technology, allowing people to express their intent in free-form natural language. It has paved the way for AI agents that not only converse with users but also perform actions on their behalf, flexibly and with minimal guidance. Delegation to AI has already begun to improve the efficiency of individual processes, making both consumers and businesses more productive in the set of tasks they had already been doing.<sup>4,5</sup> However, we believe that the more disruptive—and yet to be realized—impact of generative AI is its potential to drastically reduce the communication frictions between and among consumers and businesses. This could lead to a reorganization of markets, shifts of market power, and the introduction of entirely new products and services.

Consumers have traditionally faced high communication costs when initiating relationships with businesses, reducing efficiency.<sup>6</sup> For example, a consumer seeking a new tax preparer might hesitate to switch because she would have to explain her financial situation all over again to a new person or online service. These communication hurdles can prevent consumers from taking advantage of better products and services or lower prices. Businesses have tried to lower these costs with tools like online forms and voicemail menus, but these often just shift communication costs to the consumer and can make interactions more rigid.

Imagine instead a future where every consumer has an assistant agent



to communicate their preferences and personal information to businesses, and every business has service agents to interact with consumers and other businesses. These agents could be designed to interface with each other seamlessly and flexibly, transforming the landscape of consumer-business interactions. Delegating interactions to such assistant and service agents lowers communication costs and makes markets more efficient by expanding the range of options available to both consumers and businesses.

To unlock the full economic potential of generative AI's communication capabilities, two developments are necessary. First, consumers and businesses must widely adopt assistant and service agents. This is already under way.<sup>2,7</sup> Second, these agents must be designed

to interact seamlessly with each other to facilitate transactions. On the technical front, there has been significant progress in standardizing such agentic interaction, with frameworks such as Microsoft's AutoGen,<sup>a</sup> and protocols such as Anthropic's Model Context Protocol<sup>b</sup> and Google's Agent2Agent Protocol.<sup>c</sup> However, it remains to be seen how these advances will be adopted and implemented, or constrained, given their complex interplay with and dependence on market forces.

We believe the largest benefits of inter-agent communication will be realized as markets reorganize around these new capabilities.<sup>3</sup> While the pos-

a See <https://github.com/microsoft/autogen>

b See <https://github.com/a2aproject/A2A>

c See <https://github.com/modelcontextprotocol>

sibilities of this technological distribution are numerous and difficult to predict, we offer a framework for understanding the most plausible outcomes and contrasting them with the current state of affairs. A key question is whether inter-agent communication will occur within closed “agentic walled gardens” controlled by a few dominant providers, akin to today’s app stores, or through a more open “web of agents” where agents freely connect and transact, much like the World Wide Web. The answer to this question will determine how today’s largest online platforms are impacted by the proliferation of interconnected agents—whether through further entrenchment of their market power or a loss of dominance and broader democratization of AI’s economic benefits. We explore how those benefits might manifest, with implications for advertising, e-commerce, and the creation of new products and industries.

### The Current State of Agentic AI

Before discussing future possibilities for agentic economies, it is helpful to survey the current landscape of AI agents. On the surface it may appear as if several existing efforts are well on their way to providing consumers and businesses with assistant and service agents that could function as described here. But most existing agents lack a key ingredient: While they are designed to interact with or simulate human users, few public offerings are *designed to interact with each other*.

Existing agents generally come in one of two forms: siloed service agents or general-purpose end-to-end agents. The first, siloed service agents, provide a new user interface for products and services within a single company. For example, Amazon’s Rufus allows customers visiting Amazon to access their order histories or compare potential purchases through natural language instead of navigating a website. Likewise, Expedia’s Romie provides a chat interface to help customers build travel itineraries—including flights, hotels, and restaurants—by pulling information from customer email messages and group chats. Importantly, however, these efforts do not expose interfaces intended for interaction with other agents. As a result, they are still fundamentally lim-

## A key question is whether inter-agent communication will occur within closed “agentic walled gardens” or a more open “web of agents.”

ited by people having to navigate to and personally interact with them.

End-to-end agents take a different approach, aiming to provide general-purpose automated functionality that is not limited to a single company or service. For instance, technology from OpenAI, Google, and Microsoft can aggregate research from external sources, navigate business websites on the user’s behalf, and even perform simple tasks like making reservations or placing food orders, all within one interface. Importantly, however, much of the functionality that these end-to-end agents provide currently comes through “computer use models” that simulate a user pointing and clicking on existing (non-agentic) websites (or most recently the assistant agent may absorb a menu of products and pricing, or even an app, but still all within the confines of the assistant agent). This gives the illusion of assistant and service agents working together, but the absence of a true service agent on the business side limits what an end-to-end agent can accomplish.

Moreover, by mimicking human users, these agents risk creating adversarial relationships with businesses that, for example, rely heavily on advertising revenue and may resist having their websites accessed by agents instead of humans. Even if end-to-end agents can perfectly simulate human users and businesses agree to agent-based access, interactions would still be constrained by what businesses currently expose through existing web forms that limit the expressiveness of requests and responses. For instance, businesses offering highly customizable products—such as catering services—often use only basic contact forms, handling

niche or specialized requests through human followup.

### The Future Impact of Agentic AI

Given their limitations, we anticipate that siloed and end-to-end agents will ultimately give way to agents that are designed to seamlessly interact directly with each other. However, there is a complex interplay between technical capabilities and market forces that could lead to a number of different scenarios for what this solution will look like, who will control it, and what it will be capable of.

**The Market Power of Digital Intermediaries.** Two-sided platforms (businesses designed to bring together two distinct sides for a transaction) such as Amazon, Expedia, OpenTable, and Spotify are key intermediaries of the current digital economy that create value by matching millions of consumers and businesses to each other within specific domains such as shopping, travel, dining, and music.<sup>9</sup> They do this in part by standardizing how both sides interact. For example, Amazon requires sellers to follow specific formats and policies, while consumers must use its interface to search and transact. In exchange, consumers access a vast range of sellers and businesses gain exposure to a large customer base. But this comes with trade-offs: both operate within a tightly constrained system, are subject to the platform’s design choices (for example, ranking algorithms), and pay referral fees.

If an agentic economy enables each consumer’s assistant agent to directly and flexibly communicate with each businesses’ service agent via *unscripted communication*, the role—and market power—of intermediary platforms could shift substantially. In principle, once communication frictions are low enough, interoperable AI agents could eliminate the necessity for two-sided platforms as intermediaries altogether. Consumer assistant agents could directly find and flexibly negotiate with service agents to buy goods, book hotels and airlines, make dining reservations, and stream music.<sup>6</sup> This would represent a drastic decentralization of power compared to today’s markets.

In practice, however, intermediary platforms often provide value beyond simply standardizing communica-

tion between buyers and sellers, via discovery, validation, remediation, and economies of scale. For instance, in domains like travel, it may still be useful for an intermediary to provide trusted suggestions, offer trip insurance, resolve disputes, or ensure regulatory compliance. However, an agentic economy could lead to fierce competition between intermediaries due to low switching costs, reducing the profits they can extract.

**Agentic walled gardens vs. the web of agents.** Even though any specific assistant and service agent may be technically capable of communicating with each other in an unscripted manner, the pool of agents they might interact with could be *restricted* due to market forces. Select firms may provide assistant agents for free but restrict communications, creating “agentic walled gardens.” In some sense this would be a natural evolution of today’s existing application ecosystems such as the app stores in dominant operating systems. Given their existing large user bases and nascent assistant technologies like Apple Intelligence, Google Assistant, Microsoft Copilot, and Meta AI, these firms are well positioned to extend their current marketplaces to include interoperable AI agents. For example, in March 2025, Meta launched basic service agents for business pages on Facebook and Instagram at no cost, but these service agents are only accessible to users on their own platforms. In addition, firms such as OpenAI and Anthropic that have built out large user bases for their assistant agents could develop their own marketplaces.

These walled gardens could offer a number of benefits such as ensuring a baseline of quality and security by filtering out low-reputation or fraudulent service agents and streamlining discoverability and rating of agents. They could also offer insurance if an agent makes a mistake. However, much like today’s app stores, this could concentrate market power in the hands of a few dominant players, who could leverage their position to extract substantial profits and limit the openness, competitiveness, and innovation of the broader ecosystem. This could also lead to fragmentation of agentic ecosystems and suboptimal user experiences—for instance, a given individual

might have siloed assistant agents for personal and professional purposes, making it difficult to coordinate between them.

Conversely, if consumers and businesses fully own and manage their agents, communication could be both unscripted and unrestricted, leading to a completely open and decentralized “web of agents” that is not controlled by any one entity. Similar to today’s World Wide Web, any agent could join and transact with any other. Assistant agents would play a role similar to web browsers, and service agents similar to websites. This could foster competition, innovation, and broad access to agentic technology, but it also comes with substantial challenges. In particular, successful development of a web of agents requires large-scale coordination among many players—including corporations and governments—to develop and agree upon standards and protocols. It also requires robust mechanisms for discovery, trust, and security among interacting agents. We visualize these two plausible scenarios in the accompanying figure.

**The future of advertising.** In today’s digital economy there are generally many more businesses and products available to consumers than they have the time to research and consider, and so advertising plays a crucial role in capturing attention and guiding online transactions. But in an agentic economy where assistants can interact with millions of businesses on behalf of consumers, attention is a less-constrained resource. What matters more is the al-

gorithm that matches assistants to service agents.

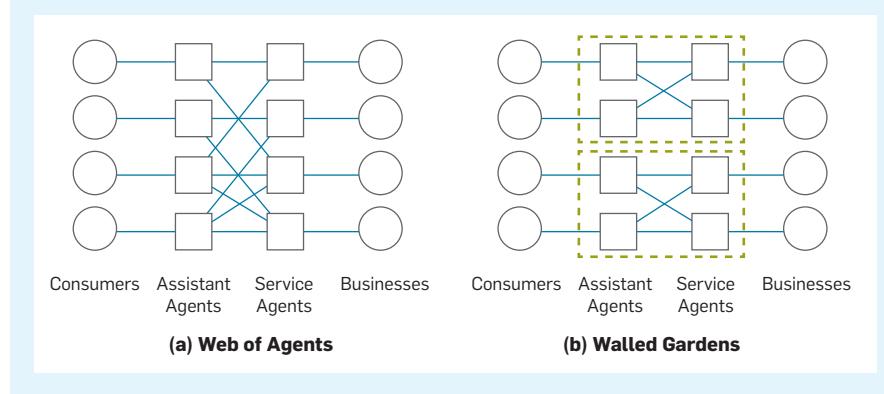
In a scenario where there are strong central intermediaries (that is, an agentic walled garden or web of agents with dominant discovery layers) some form of paid prioritization, akin to today’s advertising, is very likely to influence rankings. But the truly scarce and valuable resource—particularly in a web of agents—will be high-quality human feedback on goods and services. This feedback will be crucial not only for improving offerings (for example, training AI agents) but also for distinguishing high-quality services from poor ones. The focus of monetization and competition could shift from the “attention economy” to a “preference economy.” Success will hinge on attracting early, engaged users who provide valuable feedback.

New means of monetization may emerge to aggregate and rank service agents based on feedback from consumers and their assistants. Just as people leave reviews on platforms like Yelp or Google Maps today, future assistants might generate reviews based on user satisfaction data. Businesses could compete by offering better prices or services to attract these early users, creating a flywheel of feedback and preference data.

#### Payments and micro-transactions.

As interactions between consumers and businesses become more seamless and platform intermediaries less central, we anticipate a rise in “one-off” transactions and accompanying decrease in repeat engagements and long-

**Figure. In an agentic marketplace, assistant agents (representing consumers) and service agents (representing businesses) interact directly with each other. We distinguish between two scenarios: (a) an example of an open web of agents with unrestricted communication between agents and (b) an example of walled gardens in which agents are restricted to interacting within the confines of platforms.**



term consumer-business relationships. This shift may encourage the growth of micro-transactions. For example, a consumer whose assistant frequently and seamlessly switches between multiple content providers (for example, Spotify and Pandora, or Netflix and Amazon Prime) may prefer usage-based micro-payments versus subscribing to both services. Furthermore, the usual friction associated with micro-payments is eliminated when transactions are handled entirely by assistants and service agents, encouraging micro-transactions that would otherwise be too inconvenient to manage manually. This shift is likely to evolve regardless of whether walled gardens or a web-of-agents scenario triumphs.

**Unbundling, Rebundling, and New Products.** Products, goods, and services are often bundled to balance complexity and efficiency of transactions. For example, a news article may bundle interviews, photos, and facts, even if the reader is already familiar with parts of the story. A consumer's assistant agent, however, can track what the user has already read and collaborate with a service agent (for example, from the *New York Times*) to generate a customized article focusing only on new or relevant information. This dynamic and personalized rebundling optimizes knowledge transfer while minimizing cognitive load. Such personalization can be accomplished even today using methods such as retrieval-augmented generation (RAG); in principle, an assistant agent could pull from multiple sources of high-quality content to create customized offerings for its user. Currently, however, such methods are often limited to public-domain sources such as Wikipedia or under fair use claims, as publishers are generally unwilling to permit their content to be reused without compensation. In the future, micro-transactions between assistants and service agents for the use of individual pieces of digital content could enable an ecosystem that compensates content creators while unlocking the ability to create customized user experiences.

More generally, the power of dynamic and personalized bundling could apply to many other digital goods and services that can be deconstructed and reconstructed by assistant agents working with service agents to best serve

## As technological progress drives greater specialization, it also increases the need for coordination.

the needs of end users. This capability aligns closely with the anticipated rise of micro-transactions handled by agents. As the infrastructure for micro-payments develops, it will become possible for digital components to be individually and dynamically negotiated, priced, and packaged into hyper-personalized products. We expect more extreme unbundling, rebundling, and new products in a web of agents where there can be unrestricted communication between the assistant and *any* service agent, allowing for more innovation in product creation and bundling.

### Conclusion

As technological progress drives greater specialization, it also increases the need for coordination, which in turn demands more sophisticated communication between individuals, organizations, and systems.<sup>1</sup> For example, where one doctor once handled a patient's care, today multiple specialists must collaborate using advanced tools and shared information. AI agents mark a shift in this pattern: rather than compounding communication overhead, they offer a way to streamline it. By coordinating tasks across fragmented systems, agents stand to reduce friction in markets, lower switching costs, and unlock more decentralized access to digital goods and services.

The architecture of this emerging marketplace is still taking shape. Much will depend on whether AI agents are allowed to interact freely across an open web or whether their interactions are restricted within closed ecosystems. This will be shaped by early stakeholder decisions, emerging technical standards, and evolving regulatory frameworks. Regulatory outcomes may vary by coun-

try—especially given the divergent approaches of the E.U., which favors precautionary, preventive measures, and the U.S., which tends to adopt a more reactive, ex-post regulatory stance in digital markets. In that sense, this moment is reminiscent of the early days of the Internet before it began generating billions of dollars of revenue. The current platform-dominated economy was certainly not obvious in the late 1990s into the early 2000s, when digital commerce was becoming increasingly decentralized as early Internet portals gave way to a blossoming World Wide Web.

Now is the time for us to reflect and ask what kind of agentic economy we want in the near future. The choices we make today will determine not only how these markets function, but also who benefits from this new wave of technology. C

### References

1. Becker, G.S. and Murphy, K.M. The division of labor, coordination costs, and knowledge. *The Quarterly J. of Economics* 107, 4 (Nov. 1992); doi: 10.2307/2118383
2. Bick, A., Blandin, A., and Deming, D.J. *The Rapid Adoption of Generative AI*. Tech. Rep. National Bureau of Economic Research (2024).
3. Bresnahan, T. F. and Trajtenberg, M. General purpose technologies 'Engines of growth'? *J. of Econometrics* 65, 1 (1995); <https://bit.ly/4s88WEY>
4. Brynjolfsson, E. et al. Generative AI at work. *The Quarterly J. of Economics* (2025).
5. Cui, Z. K. et al. The effects of generative ai on high skilled work: Evidence from three field experiments with software developers. *SSRN* 4945566 (2024).
6. Davidson, T. R. et al. Evaluating language model agency through negotiations. *arXiv:2401.04536* (2024).
7. Humlum, A. and Vestergaard, E. *The adoption of ChatGPT*. Tech. rep. IZA Discussion Papers (2024).
8. Klemperer, P. Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *The Rev. of Economic Studies* 62, 4 (1995).
9. Rysman, M. The economics of two-sided markets. *J. of Economic Perspectives* 23, 3 (2009).

**David M. Rothschild** (davidmr@microsoft.com) is a senior principal researcher at Microsoft Research, New York, NY, USA.

**Markus Möbius** is a senior principal researcher at Microsoft Research, Cambridge, MA, USA.

**Jake M. Hofman** is a senior principal researcher at Microsoft Research, New York, NY, USA.

**Eleanor Dillon** is a principal researcher at Microsoft Research, Cambridge, MA, USA.

**Daniel G. Goldstein** is a partner research manager at Microsoft Research, New York, NY, USA.

**Nicole Immorlica** is a senior principal research manager at Microsoft Research, Cambridge, MA, USA.

**Sonia Jaffe** is a principal researcher at Microsoft Research, Cambridge, MA, USA.

**Brendan Lucier** is a senior principal researcher at Microsoft Research, Cambridge, MA, USA.

**Aleksandrs Slivkins** is a senior principal researcher at Microsoft Research, New York, NY, USA.

**Matthew Vogel** is a principal program manager at Microsoft Research, New York, NY, USA.