


Could Gen AI End Incumbent Firms' Competitive Advantage?

by Toby E. Stuart

November 21, 2024



Bloomberg Creative/Getty Images

Summary. In business strategy and entrepreneurship, it is necessary to have a protective “moat” to sustain high performance. As gen AI systems advance in capability, they promise to transform and disrupt industries far beyond traditional software, to... [more](#)

In 2011, Marc Andreessen, the eponymous founding partner of venture capital firm Andreessen Horowitz, wrote [an editorial](#) famously proclaiming that “[software is eating the world.](#)” His argument: four decades after the invention of the microprocessor and two decades since the creation of the internet, the transformational software technologies that had the potential to reshape the industrial landscape were finally ready for full-scale deployment. Boy, history has proved him right.

Fast forward 13 years, gen AI is supercharging Andreessen's prediction — and harkening a transformative shift in nearly every aspect of human endeavor. As gen AI systems advance in capability, they promise to transform and disrupt industries far beyond traditional software. This includes automating many — if not most — complex cognitive tasks and radically reshaping knowledge work, accelerating R&D and problem-solving across fields, and enabling hyper-personalization of products and services. If we continue on our current course, near future generations of AI will become agentic, capable of acting in a goal-directed manner, interacting with humans and other automated agents to make and execute decisions in business, finance, and consumption contexts. Finally, AI promises to unearth entirely new technological capabilities we can't yet envision.

There are many definitions of “artificial general intelligence” (AGI) right now. An encompassing one is that AGI has arrived in full force when the majority of tasks we now perform on a computer become automated. Whatever label we put on this, what's clear is that no major social or economic system is likely to remain unchanged when that occurs. The software-driven disruption Andreessen foresaw will pale in comparison to the changes now on the horizon.

Predicting the future is a challenge even in steady days, to say nothing of historically unprecedented moments. It's relatively easy to forecast near-term developments of the technology. It's advancing at an extraordinary pace, it's already achieved brehtaking multimodal capabilities, it's becoming agentic, it's likely to hallucinate less. Compute will increase and algorithms will continue to improve. But what does all of this mean for organizations and their leaders? In a time of unprecedented foment, it is ironic that the most logical way to chart a course

ahead may be to look to the theories that have explained our past. Filtering expected technical changes through the cornerstone frameworks we use to understand our current world may be the best approach to anticipating how a General Purpose Technology (GPT) with unprecedented capabilities will impact the choices leaders must begin to make today.

One thing is certain though. AGI will uproot the market positions, unit economics, competitive dynamics, and business strategies of *many* incumbent organizations. In so doing, AGI also will radically alter the allocation of profits across firms and industries.

Transforming Moats

In business strategy and entrepreneurship, it is necessary to have a protective “moat” to sustain high performance. In a competitive economy in which entrepreneurial firms are free to enter new markets, moats are the reason incumbents maintain profitability. Companies like Apple, Alphabet, Amazon, Meta, Uber, Eli Lilly, NVIDIA, JP Morgan, TSMC, Oracle, PepsiCo, and McKinsey have thrived due to their strong brands, economies of scale, strong network effects, IP portfolios, product design capabilities, high switching costs, proprietary data, and human capital advantages. In a word, their moats.

If the capabilities of AI systems continuing to evolve, these could soon be challenged. And given what we know now, I believe that is nearly certain to happen.

There are three pillars that have driven AI forward: compute power, model development, and training data. Right now, they are progressing at an astonishing rate. The amount of computational power devoted to training models has doubled approximately every *6 months* in the last decade. That’s four times the rate of advancement in underlying semiconductor technology described by Moore’s Law. In fact, compare a seminal moment in the development of deep learning models, which

occurred in 2012 when researchers trained the image recognition model AlexNet on two NVIDIA graphics cards for about a week. AlexNet was a remarkably successful GPU implementation of a 60 million parameter convolutional neural network to perform computer vision. On a present day, state of the art NVIDIA Blackwell GPU, one could train this model in 5 minutes or so. To put this into current context, Meta just announced it is training the next-gen open source foundation model Llama 4 on a cluster of more than 100,000 state of the art GPUs. At this blistering pace of advancement, we should expect fast progress in the next few years.

The computational power is just part of the equation, though. There is also rapid progress in algorithmic efficiency and capability, in training data, and in synthetic (machine-generated) data. AI algorithmic development is likely to become recursive. As the systems become more capable, they increasingly can be used to optimize themselves. This creates a potential feedback loop of accelerating improvement. With the recent release of OpenAI's GPT-o1 Preview, we have a glimpse of the progress that will be made with chain of reasoning, in which AI systems follow a logical "chain of thought" to reach a conclusion (versus next-word prediction). Simplistically speaking, chain of thought refers to a technique in which AI models generate responses by explicitly reasoning through intermediate steps before arriving at a final answer. Instead of trying to solve a problem directly in one step, the model simulates a multi-step reasoning process, similar to how humans might think through a problem.

Although these models require greater compute resources, they promise to vastly improve model capabilities and reliability.

The significance of this progress is hard to grasp. Tracing the arc of progress forward, gen AI systems will continue to improve in nearly all forms of cognitive work including knowledge synthesis, number crunching, comparing documents or podcasts, writing, idea generation, simulation, multimedia (and many other forms

of) content creation, decision making, medical diagnosis, task execution, and on. Many recent papers show that challenging cognitive work already is performed as a collaboration between humans and AI, which will reduce human capital-based barriers to entry in many markets. Many VC investors already have begun to imagine a future in which startups will no longer need the world's top engineering talent to create novel software. Think of the implications of that: In the near future, the limiting factor for high potential, VC-fueled startups may not be software developers. Likewise, strategy consultancies won't require legions of MBAs from elite business schools for data collection, insight, and analysis to populate consulting teams. Financial institutions such as investment banks, hedge funds, and insurance companies, will find themselves in the same position. Much of the traditional marketing industry players can expect to confront a disrupted landscape.

Looking ahead, *some* — but not all — long-standing moats will topple as we advance toward AGI.

Declining Moats

By lowering moats in many industries, AI will promote new entry. It also will compel incumbents to think strategically and act swiftly to protect their positions. Many companies will face strong competition, often from new entrants that will find it far easier to gain footholds in traditionally difficult-to-enter markets. In fact, many of the new entrants to these spaces will arrive with products and services that are superior to those offered by incumbents.

Service as Software.

We are about to witness a true fusion of the \$5+ trillion white collar labor market and the \$250 billion enterprise software market. AI agents already have begun to join the workforce: They've entered as customer service and leasing agents, among others, but the gen AI consultants, software developers, nurses, tutors, corporate trainers, radiologists, pathologists, investment

bankers, paralegals, designers, medical billers, and so on also are coming fast. Without a major overhaul in business model, many professional service firms face a difficult road in the intermediate and long term. In the short term, however, gen AI may create significant, additional work for the professional services sector, as we seek expert opinions to collectively sort out the legal, safety, security, bias, and compliance questions surrounding gen AI, and as organizations “sense make” about when, how, and in what products, services, and internal functions to adopt gen AI.

The overarching story for the professional services sector will be the transition from “software as a service” to “service as a software.” SaaS revolutionized software delivery and completely transformed the enterprise IT stack. Rather than buying and installing bulky programs, software buyers now pay subscription fees to access cloud-based applications. This made software more accessible, scalable, and constantly updateable.

While SaaS created novel software tools, AI adds an unprecedented layer of intelligence on top of these systems. Pre-trained LLMs with access to the enterprise’s data will fine-tune to analyze, synthesize, and interpret them, as well as generate novel content and present and execute findings. Gen AI’s analytic, simulation, content creation, and decision-making capabilities will encroach on the customer value proposition of many professional service firms. Expect to see companies stand up AI “workers” that are fine-tuned on a company’s proprietary data, which will offer an impressive level of customization of expert knowledge catered to client- or company-specific needs. This process is already underway, with a number of companies such as Moody’s, Accenture, and Zoom blazing the trail.

Take the major management and strategy consulting firms. Their moats include stellar reputations, talent pipelines that source the best and brightest university graduates, a depth of internal knowledge accrued from prior consulting projects, and well-positioned alumni networks that reach into many executive

suites. However, these resources come at a very high cost to clients. My prediction is that these firms' human capital advantage and internal knowledge base will shrink in value as the "service as software" world arrives, and their brands and relationships are unlikely to be enough to protect industry incumbents from competitive pressure. The same applies to leading companies that work in sectors like tax preparation, which have brand and retail distribution advantages over their smaller competitors. These organizations provide wonderful services, but for tax prep, financial advice, and accounting services, upstart tax-focused "service as software" entrants will provide faster, less expensive, high quality and more convenient service.

Content was King.

In 1996, another tech luminary, Bill Gates, wrote an essay titled "Content is King" to convey his view that content would be a fundamental driver of value on the internet. Many at the time argued that the internet, by lowering the moat to distributing content, would shift value from distributors of creators of content.

In the gen AI world, I anticipate that content-based businesses, such as book publishers, media companies, modeling agencies, film studios, and distributors will experience declining moats. The historical moats for these businesses have been economies of scale in content creation, strong relationships with content producers, expert curatorial judgement, existing content libraries, existing brands, and in many cases, in-house distribution capabilities.

We are on the verge of the "mass customization" of content. The next generation of multi-modal AI systems promise to create high quality, on-demand books, articles, games, and films that are nearly instantaneously produced and finely customized to individual consumer preference. Synthesia, HeyGen, OpenAI's Sora, Luma AI, and many companies are showing us the path

from text to high-quality video. In a recent paper that hints at what is to come, Google presented GameNGen, a stable diffusion-based generative AI that can actually recreate an immersive video game after watching 100s of millions of frames of people playing it.

The surreal translation of this: it's possible to develop a sophisticated computer game without source code. Gen AI can generate story ideas, possible plot structures, and character development with a few sentences of prompting. Extending this capability forward suggests its potential to generate book-length manuscripts on the topics of our choosing, of the length of our choosing, written in styles or read to us or viewed in languages, voices, or faces we will also choose. This will cause the cost of content production to precipitously drop, and the line items of expense to generate content may begin to shift from the creator economy to compute and data center costs. Likewise, ancillary workflows, such as fact checking and copy-editing manuscripts, will roll into automated content production. With these developments, the moats that protect content producers and distributors will shrink. (Note: this will not be true for live entertainment, such as sports, which is likely to create even more consumer value in the future.)

AI agents trained in consumers' individual preferences also will take over much of the curatorial function currently performed by many brands. We have already seen the power of this shift in social media. TikTok and Instagram's explosive growth and addictive properties are now partly based on algorithms that helps to match inconceivably vast, ever-expanding video libraries to individual consumer tastes. This level of AI-driven matching of consumer preferences to goods and services will happen at scale, across many product and media content categories.

Cracks in the foundation of the Ivory Tower.

As an industry segment, the upper echelon of higher education has been incredibly stable for decades and barriers to entry for new competitors are prodigious. The prominent universities have 100+ year old brands, outstanding student bodies, large, internationally renowned research and teaching faculties, expensive PPE and facilities, and incredibly far-reaching and impressive alumni networks. These institutions long have been the admission ticket to lucrative careers.

But even here, I expect change. The instructional capabilities of the LLMs already are remarkable. Online learning platforms have real potential to provide high-quality education at scale, and AI tutors will offer personalized learning experiences. Scholarly research already has begun to escape from the Ivory Tower, as a great deal of the progress in gen AI itself is emerging from highly capitalized corporate labs. Finetuned AI models have begun to draft scientific articles in styles preferred by scientific journals, and generative AI systems appear to have the potential to lower the moat to cutting edge academic research.

Employers too are likely to continue to shift to hiring for skills rather than recruiting based on university logos on CVs, as AI systems will also enable greater visibility into pools of prospective employees and their job-related experiences and skills. As these changes set in, we may see a decline in the demand for on-campus university educations and the emergence of new providers and new models for post-secondary education.

Opportunities

Many companies will face declining moats, but some firms and industries will become much more profitable as their protective barriers stand up to the AI-drive changes afoot. What moats might continue to hold?

Brands: In an AI-first world, brands will remain important for their coordinating effect rather than for their traditional role of

signaling product quality or on grabbing consumers' scarce attention. People will choose brands based on shared consumption values. The luxury goods producers across most product categories likely will continue to hold their moat. I don't expect Hermes's Birkin bags to fall in price. Likewise, in the gen AI world in which high-quality fakes populate the digital (and physical) world, customer trust will become even more important. Brands that have meaning to consumers and can assure the provenance of their goods and services will prosper.

Data: Companies with large, hard-to-simulate data sets for fine-tuning models will have a competitive edge. Expect to hear a lot about data-based moats as AI capabilities expand. The volume and quality of proprietary data in fields like healthcare and finance will become indispensable to customize agentic AI models to industry-specific needs. As privacy and data use regulations intensify, companies with legally obtained, high-quality datasets will gain an advantage.

Intellectual Property: IP will continue to forestall competition in certain contexts, but the value of copyrighted content is likely to decline, while biopharmaceutical and deep tech IP may remain a valuable barrier to entry. This said, expect modifications to the IP system as generative AI raises core questions about what kinds of IP can be patented and copyrighted. And expect the faster development of copycat technologies that circumvent IP protection.

Operational Effectiveness: Firms that adopt and deploy new AI systems ahead of competitors will thrive. We are entering a period in which innovations will arrive at an unprecedented pace. Leaders who craft agile operations that are able to rapidly experiment with and implement AI technologies will capture new market opportunities faster. It'll finally be time to stop talking about operating effectiveness as table stakes for keeping up with competition, and instead start thinking about it as a competitive moat.

Lobbying and Government Relations: Companies, industries, and professions that successfully leverage the political and legal systems to preserve their economic interests will maintain advantages longer. Expect to see these efforts in law, medicine, the ACE professions, and in heavily unionized occupations. As AI threatens to disrupt labor markets, particularly in regulated industries like law and healthcare, organizations that maintain strong government ties may be able to influence legislation in their favor, slowing down or redirecting automation trends.

New Business Models: Major technological changes always create new business models. We are already seeing this with the foundation models. The rise of AGI will accelerate the emergence of new subscription models for AI services, where businesses will charge for usage in novel ways. Perhaps we will see more outcome-based pricing. Entirely new ecosystems, like large marketplaces for data and for access to AI agents, will emerge and thrive. To facilitate automated transactions, blockchain-based systems may finally find product-market.

Personal Relationships: In the long run machines may indeed take over decision-making in many contexts, but in the near to intermediate term humans will have the final word on purchase decisions. A network of strong client relationships will help to protect competitive positions. Even in industries where AI can outperform humans in technical tasks, trust-based relationships between businesses will maintain value. Emotional connections among humans will continue to play a role in building long-term partnerships.

With total wages in the U.S. alone amounting to approximately \$12.5 trillion per year, companies that maintain their competitive moats while adopting AI systems to augment their high-wage workers with generative AI productivity multipliers will experience unprecedented profitability. By understanding these dynamics, business leaders can better navigate the seismic shift

brought about by AGI, positioning their companies for success in an AI-first future.



Toby E. Stuart is the Helzel Chair in Entrepreneurship, Strategy and Innovation; Faculty Director of the Berkeley Haas Entrepreneurship Program; the Associate Dean for External Affairs; and Faculty Director of the Institute for Business Innovation at the Haas School of Business, UC Berkeley.



Read more on **AI and machine learning** or related topics **Generative AI**, **Strategy**, **Technology and analytics** and **Competitive strategy**