4 SEPTEMBER 2026

Commentaries

New Series

WHO IS THE *FAST SECOND* IN THE ARTIFICIAL INTELLIGENCE INDUSTRY? AND WHY DOES IT MATTER?

Simone Vannuccini*

COMMENTARY
NO.051 NS/2026

*Chair of Economics of Artificial Intelligence and Innovation - Université Côte d'Azur/GREDEG CNRS

Senior non-resident Fellow at Fondazione CSF

If you are an attentive follower of the daily developments around artificial intelligence (AI), you will almost certainly have experienced FOMO (fear of missing out), along with a *short blanket* dilemma: where should you direct your attention when everything seems to be happening everywhere all at once? And which lens can help make sense of a phenomenon in which short- and long-term business decisions are so strongly intertwined with political ideologies and priorities?

Consider the news coming from the United States (US) alone. In the span of just a few months, under an Administration whose decision-making is consistently nationalist and protectionist, but often erratic, and sometimes even willingly self-harming, we have witnessed the launch of a new AI strategy, the US *AI Action Plan*. This prominently features the goal of establishing an *American AI Export Program*, explicitly aimed at undermining partners' technological sovereignty. Then, next to the export controls – mostly on semiconductors – imposed against China and the attempt at influencing the partners' consensus on supply chain security and dependencies through the Pax Silica initiative, the Government also targeted one of the major American AI labs, Anthropic, briefly and *creatively* designated a "supply chain risk" and thus forced to temporarily halt the release of its Fable model. After that came accusations of *adversarial distillation* of AI models, again directed mostly at China, and a proposal to restrict access to open-weight models as a way of protecting the domestic AI-lab oligopoly – an idea that generated a backlash from a broad coalition of companies led by Nvidia (the company is so invested in the open paradigm that it just acquired for 13 billion dollars HuggingFace, the major open model repository). In parallel, a tidal wave of opposition to data-centre construction gathers momentum and may well become a decisive issue in several important contests in the upcoming midterm elections.

At the same time, alongside several releases of frontier models, China has spelled out its vision for AI in a widely dissected keynote address by Xi Jinping at the World AI Conference (WAIC) in Shanghai. It outlines a very different perspective, with China presenting itself as a champion of global AI cooperation in contrast to the aggressively assertive and divisive US behaviour. Meanwhile, the alarmist narrative promoted by AI companies, lobbying organisations and ideological interest groups continues to dangerously weave its way into the European intellectual and policy debate, most recently through the Europe 2031 initiative.

The events involving Anthropic are illustrative of the manifold forces that have made AI a key political battleground. On the one hand, the justification for the ban concerned the cybersecurity risks associated with the capabilities of the company's new model. These risks have some basis in reality, but they were first inflated and then downplayed, following a

script that has more to do with marketing than with security and that has characterised the industry since its early days. If you need a reminder, something similar happened with most major model releases going back to GPT-2, following a pattern and strategy that Gary Marcus dubbed "scare, hype, release. And repeat."

On the other hand, part of the story was a toddler-level petty reaction by the Administration, retaliating against Anthropic's refusal to allow unrestricted use of its frontier models for national-security purposes by the very demurely renamed *Department of War*. The reaction turned Anthropic and its CEO into something close to heroic figures overnight, despite there being very little heroic about either. In other words, the technology is prone to being misunderstood and hyped, while political actors are prone to misuse it to justify objectives that have relatively little to do with AI itself.

All these events matter, all are interdependent, and all will shape the next phases of AI development. Advances in the AI industry are increasingly essential to geopolitical goals in a more conflictual and distrustful global environment; and the geopolitical importance attached to AI is becoming essential to sustaining an industry whose underlying economics remain precarious. AI labs and hyper-scalers alike are chain-linked in a loop of cross-subsidies that burns cash in abundance while generating losses rather than profits. The result is an increasingly risky bubble that has not yet burst partly because investors are either already too deeply involved in the industry or unwilling to recognise that AI can simultaneously be a transformative technology and a bad financial investment.

However, the more we look at AI through the lens of international relations and treat it as a geopolitical problem, the easier it becomes to forget that it is also, fundamentally, an industrial problem. If we momentarily set aside the concerning fact that *America is now one big bet on AI* – and the extent to which this helps explain many of the events unfolding around the technology – what remains is a classic race among a handful of companies for market dominance. The remainder of this commentary takes a closer look at this particular piece of the much larger AI puzzle. However, even when focusing on market dynamics, we will be brought back to AI politics at the very end of the piece.

An interesting question is whether the AI market we see today will look anything like the AI market of the near future. This is important, as answering it may help us identify key forces and tensions in the field and understand how they interact; in turn, this could help guide policymaking already today. Part of the answer lies in understanding whether its protagonists will remain the same. After all, market structures are shaped by the actors that succeed – or simply survive – and by the strategies they pursue. Think of Microsoft in the computer industry, becoming dominant by shifting value capture from hardware production to software licensing, or TSMC emerging as a quasi-monopoly in chip fabrication by gatekeeping a fundamental layer of the supply chain and helping transform a vertically integrated industry into a specialised production network.

The bet I make here is that, in AI, the leading actor of today – the company OpenAI – will not be the leading actor of tomorrow. To support this bet, I draw on a little-cited strategic framework that is more than twenty years old, developed by Constantinos Markides and Paul Geroski in their book *Fast Second*. More recently, journalist Ian Bogost has touched on a similar argument in *The Atlantic* when discussing the future of the AI industry.

The thesis of the book is that "the early pioneers of radically new markets are almost never the ones that scale up and conquer those markets". This is

not because "the pioneers are small or insignificant players with no resources or bad management. And it's not because their products are inferior to the products that latecomers introduce". Rather, first and second movers possess different capabilities. First movers have what it takes to *colonise*, in the authors' terminology, a market: they leverage a radical innovation to open up a new market niche. Fast seconds, by contrast, are usually better at *consolidating* that market.

The theory helps reconcile seemingly contradictory evidence about innovation and market leadership. On the one hand, we have the so-called *innovator's dilemma*, according to which incumbents – usually large firms – are displaced by smaller disruptors through a kind of self-reinforcing mechanism: established companies continue to do precisely what made them successful in the first place and, in doing so, miss radical innovations lying off the beaten path. This is not necessarily a strategic mistake, but rather a feature of the dynamics of innovation and corporate growth.

Yet relatively few incumbents seem to be completely disrupted, and what we often observe – especially in digital markets – is instead the remarkable uncontestability of dominant actors. Conversely, small, dynamic and often venture-backed start-ups disproportionately explore and colonise new paths, but face very low survival rates over the medium and long run. The key tension is that new-market creation is primarily a supply-side story: often, it is a matter of operationalising an idea through technology. Consolidation is instead a demand-side endeavour, requiring the capabilities and strategies needed to achieve large-scale user acquisition and retention and, crucially, to turning that into a stable and economically viable mass market. Very few firms are good at both.

Can the *Fast Second* theory help us predict what is going to happen in the AI industry? The context fits rather well. AI labs rushed to compete for, rather than in, the market, making consolidation a priority unusually early in the industry's life cycle. Yet AI presents some distinctive challenges. Platform dominance in digital markets has traditionally been built on network effects (or network externalities): the value of a platform to each user increases with the size – or quality – of its user base. Greater use can, in turn, accelerate efficiency gains and steepen learning curves. By contrast, as Cory Doctorow put it, AI has terrible unit economics: network effects do not work quite as intended and may even work against AI companies' profitability. Given how current AI systems work – notably the high compute costs of inference – and the business models prevailing in the industry, every additional user can add more in costs than they bring in revenues. Scaling up and consolidating the market can become a deadly business.

Against this backdrop, we can try to assess the future of OpenAI. The company ticks many of the boxes on the market-pioneer checklist. It did not invent AI as a technology, nor was it responsible for introducing the Transformer architecture that underpins today's Large Language Models (LLMs). But it was the first actor to successfully colonise the new mass market opened up by LLMs. In other words, OpenAI got the supply side right – even if doing so required considerable trial and error and the backing of strategic partnerships, most importantly with Microsoft.

OpenAI has gone to great lengths to consolidate its foothold after colonising the market. In terms of consumer adoption, it has succeeded remarkably: ChatGPT is now used by hundreds of millions of people. Yet its performance in acquiring commercial users – a crucial source of revenues from AI-model access – has been less impressive. At the same time, OpenAI has been trying relentlessly to drive down costs. Seen from this

perspective, its repeated calls for pauses in AI development can sound less like moral initiatives to govern existential risk and harness AI capabilities for humanity than pleas for breathing room in which to recoup enormous expenditures.

The pricing wars that have made AI inference progressively cheaper can likewise be understood as a classic attempt to attract demand after a new market has been created on the supply side, even as users' overall AI bills can rise when reasoning models engage in increasingly long loops of token processing. Yet, as some major venture-capital investors (that certainly aren't radical AI critics) have pointed out, AI remains an unfinished *proto-market*. Moreover, as Bogost notes, as AI becomes commoditised, the industry's core product is likely to become increasingly anonymous and interchangeable. The share of value attached to any specific company may therefore decline, further undermining OpenAI's attempts at consolidation. Finally, OpenAI lacks many of the complementary assets that could insulate its products from competition.

OpenAI's attempt to act simultaneously as coloniser and consolidator has exposed the company to an unprecedented level of financial commitments. Without presenting the back-of-the-envelope calculations here – I recommend Ed Zitron's analysis for that – the basic problem is: OpenAI would need to generate an extraordinary amount of future profits simply to make its existing commitments sustainable.

So, if we subscribe to Markides and Geroski's theory, who might be the fast second that ultimately reaps most of the benefits from AI? There are several candidates for the role. The obvious place to start is the constellation of Chinese AI labs releasing increasingly capable, state-of-the-art open-weight models. But this immediately takes us back to the geopolitical arena that we deliberately set aside for this market-focused analysis – even if geopolitics evidently remains crucial to the story.

If we stay within the US, what are the possibilities? Zitron has argued that an OpenAI bankruptcy could produce several outcomes, including the company being absorbed or acquired by one of the hyperscalers – Microsoft being the most obvious candidate – or even merged with Anthropic, particularly if the latter succeeds in its initial public offering, an outcome that is itself far from guaranteed. Such an outcome would carry a certain irony: Anthropic was itself born as a *disagreement spin-off* from OpenAI.

Anthropic is indeed a plausible candidate for the title of fast second. In the Fast Second framework, consolidators succeed when they move at the right time or find the right angle from which to capitalise on a market pioneered by others. Anthropic was a second mover in LLMs generally, but arguably a pioneer in LLMs specialised for coding, one of the relatively few AI use cases to have demonstrated value. Yet Anthropic shares many of OpenAI's vulnerabilities, including enormous financial exposure.

A second possibility is to look at the major incumbent that repeatedly struggled to catch up with the model-performance frontier despite pouring massive investments into the field and that, in many respects, encapsulates the innovator's dilemma: Google. Consolidation is precisely where incumbent firms can get back into the game, and Google fits the bill of the textbook fast second remarkably well: it has the deep pockets needed to subsidise its own survival while demand grows and competition plays out. More importantly, it possesses the complementary assets that pioneers lack: vast amounts of data generated across its pervasive digital services and a prominent presence across the technology stack, from consumer

applications to infrastructure and compute. These advantages arguably put it ahead of other incumbents, including Meta. Google has also easily avoided the main danger facing fast-second firms: not being fast enough. It has been present since the beginning of the frontier LLM *wars* – to echo the *browser wars* of an earlier Internet era – even if it was not the company that successfully opened the market niche.

Google may well emerge as the fast-second winner in the competition for the AI market. But the debate around a potential US ban on open-weight models illustrates why identifying the fast second may not be enough to guarantee either continued innovation or a competitive expansion of the market as a whole if competition policy fails. Lobbying the Government for protection suggests that the major actors in the American AI-model market may already be favouring strategies aimed at protecting their positions through non-market means rather than competing through innovation and product quality. The danger was already identified clearly by Lina Khan, former Chair of the Federal Trade Commission under the Biden Administration, in a forceful essay urging policymakers to "stop worshipping the American tech giants" after the release of Chinese model DeepSeek.

Finally, a third possibility is to consider the US government itself as a fast second. This would represent an unconventional extension of Markides and Geroski's theory. But at a time of renewed state intervention in the economy and a broader revival of industrial policy, the idea is not as far-fetched as it might initially sound. Indeed, experts such as Sarah Myers West and Amba Kak at the AI Now Institute have been warning for some time about the possibility that the current Administration could end up bailing out the AI industry if the bubble bursts. There has even been discussion of the US government taking equity stakes in AI companies. Such intervention could be justified on the grounds that AI companies have become "too big to fail" – although in this case *too big* would refer less to systemic financial risk, as it did with banks during the 2007–08 global financial crisis, than to the strategic and anticipated future value attached to AI which, however, is itself very likely to be overestimated (thus undermining the very justification for bailouts).

This possibility casts a more ambiguous light on the future of AI. A publicly owned AI sector could, in principle, be conceived as a distributed, democratic and open utility-like asset oriented towards the public interest – the kind of alternative that some have advocated for Europe, and that I have argued for elsewhere as a distinctive approach for the European Union's AI industrial policy. But it could equally become something much darker: a tool of a government building an infrastructure of societal control and a state-sponsored slop-production machine acting to undermine the very social contract on the rubble of a market that failed to consolidate itself. Given the current state of US and global politics, my fears tilt towards the latter.