

7 OCTOBER 2026

Commentaries

New Series

YOU SAY: IT'S AI APOCALYPSE!

I SAY: IT'S "RUINOUS COMPETITION" ONCE AGAIN

Simone Vannuccini*

I begin straight away with my central claim: if you are able to dodge the apocalypse-speak and magical-thinking bullets flying around the current debate on artificial intelligence (AI), what is left looks remarkably like an attempt at cartel-building by an oligopoly of companies. The potential impact of this anticompetitive conduct is amplified by its being explicitly sanctioned by the aggressive protectionist policy posture of the US Administration. If that is indeed the case, the real story to look at bears a striking resemblance to the past: the use of public-interest narratives to justify coordination among industry rivals. In the AI case, however, the public-interest justification for anticompetitive conduct is built around an unprecedented claim: the very survival of humanity.

In order to dismantle this narrative and bring the discussion back to AI market dynamics, let's proceed step by step. In this [long summer of AI Hype](#), as Timnit Gebru and Emily Bender aptly dubbed it, we have witnessed a rapid sequence of statements and manifestos by prominent AI and tech actors, from Dario Amodei to Mark Zuckerberg and Bill Gates. These wordy essays tend to converge on rather generic claims about AI's never-seen-before transformative potential, its inevitability, and the absolute need to govern it – although how this should be done is never clearly outlined. These texts have a very short news cycle: they remain in readers' active memory for just a few hours or days and are quickly forgotten because, in essence, they are remarkably shallow. The underlying tone of the conversation they set, however, lingers much longer and influences the way our societies see AI.

The key argument that sticks, and is continuously relaunched, is that AI poses an existential risk (x-risk) to humanity. This could take several forms, all involving "rogue" AIs, most recently "swarms" of AI agents, pursuing some harmful path through cybersecurity, biological warfare or damage to critical infrastructure. I will not even discuss here the crucial question – well expressed elsewhere (see [here](#) and [here](#)) – of how any of these harms could materialise without human intervention of some kind, either launching AI systems or failing to keep them in check. Suffice it to say that AIs now have impressive capabilities to execute tasks, but no real autonomy; they do not independently place themselves in the world: they must be deployed, given access and tasked by humans, and therefore cannot literally "go rogue." This point is rarely discussed openly, except by a front of critical AI thinkers working overtime to contain the flood of AI hype, because acknowledging it would make clear that liability lies with producers and users, not with the tool.

The reasoning underlying the AI-going-rogue apocalypse-speak is that if the technology entails even a small, tiny probability of human extinction – the infamous $p(\text{doom})$, literally the probability of doom – then we should take the problem seriously now, before it is too late and $p(\text{doom})$ becomes

COMMENTARY
NO.056 NS/2026



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

Senior non-resident
Fellow at Fondazione
CSF

The opinions
expressed here do
not necessarily
represent the
Fondazione CSF

much larger. This reasoning rests on two fundamental assumptions: that $p(\text{doom})$ exists at all (is non-zero), and that it will increase. These assumptions are the stuff of sci-fi dreams and nightmares: they are believed rather than proved or critically analysed. Still, the $p(\text{doom})$ argument has gained traction. Even more importantly, the idea that we should act now to prevent harms in the far future has found support across the political spectrum, reaching discussions at the latest United Nations General Assembly and producing rather unlikely alliances, such as that between far-right US ideologue Steve Bannon and progressive Senator Bernie Sanders. The latter has even translated this concern into regulatory action, proposing a bill named the "[Ban Superintelligence Act](#)".

The problem with the $p(\text{doom})$ story is that it is not just a fictional story. More dangerously, it is the fundamental punchline promoted by a series of ideologically charged, well-financed and well-networked groups that reduce human activity and endeavour to a simplistic calculative logic and a one-dimensional assessment of risk, while carrying with them a baggage of misogyny, racism and other views that have produced real-world, rather than imagined, harms. I recommend reading the material on the [TESCREAL bundle](#) to explore for yourself these darker roots of the $p(\text{doom})$ discussion whose promoters [display all the tell-tales of a cult](#). In a sense, the philosophy underlying $p(\text{doom})$ is a perfect expression of the dominant Silicon Valley ethos and worldview that inspires cohort after cohort of ambitious, scrappy young "tech bros" striving to achieve what they call "situational awareness" and gain an edge in an ultra-competitive environment. The bottom line is rather lame and sad: the belief that the direction of societal preferences can be shaped through a few banal back-of-the-envelope calculations with little more value added than the Economics 101 exercises in which students derive the expected utility of a representative consumer over an infinite time horizon. Substitute the infinite time horizon with the long-term future and the representative consumer with society, and there you have the supposedly "iron-clad" logic of $p(\text{doom})$. But exercises of this kind are used in economics courses to teach abstraction and analytical modelling; they are not blueprints to be copy-pasted into the real world.

My strong claim in this context is that we should question the very probabilistic premise of x -risk. Existential risk can be assigned a probability only once AI-driven extinction is included among the possible states of the world, however unlikely it might be. But once we start constructing hypothetical states of the world and assigning microscopic probabilities to them that compound over extremely long time horizons, almost anything can generate the apocalypse. What matters, then, is who decides which states of the world deserve to be considered, what probabilities and time horizons are attached to them, and which goals these choices ultimately serve, especially when those making them do not necessarily represent the public interest.

Recent cybersecurity incidents involving AI agents might appear, by contrast, to show that AI systems are getting out of control. Yet any software system tasked with executing potentially harmful actions can produce harmful outcomes. In cybersecurity, AI can identify and exploit backdoors, steal information, communicate, or poison code much as other malicious software does in some form, albeit with a much higher degree of flexibility and autonomy. We would not infer from the behaviour of a computer virus that it is intelligent, let alone superintelligent; nor should harmful behaviour by AI agents, by itself, establish that they have gone rogue.

YOU SAY: IT'S AI
APOCALYPSE!
I SAY: IT'S "RUINOUS
COMPETITION" ONCE
AGAIN

The opinions
expressed here do not
necessarily represent
the Fondazione CSF



Fondazione CSF

If we want to apply economic logic somewhere, this is where to do it, rather than in the computing of x-risk. AI companies have strong incentives to underinvest in safeguards when the costs of harmful outcomes are not fully borne by them nor the behaviour is punished by the law. The resulting incidents can then be presented as evidence of the exceptional risks posed by AI, reinforcing the alarmism on which claims of AI exceptionalism are built and, in turn, helping to sustain economic momentum and demands for special treatment of the industry and its actors.

AI exceptionalism, as Lina Khan has pointed out, is the fundamental flywheel enabling much of what we have discussed above. Given the supposedly unique, revolutionary, transformative, never-seen-before nature of AI, the argument goes, exceptional measures are needed. But AI exceptionalism is itself another in the long list of faith-like assumptions surrounding the technology: many have accepted it, while very few have tested it or seriously put it into question. Put this assumption aside for a moment, and the core narrative informing the AI debate starts to look familiar: arguments framed in terms of the public interest are weaponised by a small group of companies to advance their own private interests, including greater control over the market and the economic viability of their still-unproven business models.

What AI industry leaders claim today sounds very similar to the arguments made by Standard Oil at the turn of the twentieth century, when Rockefeller spoke of “ruinous competition”: industrial “combinations” were presented as a way to save the industry by replacing the chaos of competitive turbulence with a rational organisation of the market. The key argument was that, in an industry displaying dysfunctional features, coordination was necessary not only to make business viable for the producers, but ultimately for society’s good. In other words, attempts to build an uncontestable monopoly could be presented along the lines of: “We are not doing it for us – we are doing it for you.” The ruinous-competition defence was not confined to Standard Oil. It also appeared during the early expansion of the railways, where the parallel with AI may be even stronger: proponents of coordination argued that destructive competition threatened the solvency of the industry and, with it, its ability to continue providing essential services to customers. Of course, the arguments made then and now are not identical. But the difference only makes the comparison more compelling: in the past, industrialists equated the public interest with the survival and stability of an industry; today, AI leaders connect it to the survival of humanity itself. That is a quite a radical escalation in the traditional ruinous-competition argument, and possibly the first time that the end of the whole humankind is used as a justification for business strategy!

The similarities between the cartels of the past and today’s AI-wannabe-cartel are therefore hard to miss, even if today’s ruinous-competition argument comes dressed in the language of x-risk. Dario Amodè’s essay on “[pacing the frontier](#),” which called for industry-wide coordination to slow the pace of frontier-AI development, quickly attracted support from other prominent AI actors. The antitrust implications of such coordination were sufficiently apparent that the issue soon reached the courts. In September 2026, a group of AI users filed a class-action antitrust lawsuit, [Buist v. Anthropic](#), against Anthropic, OpenAI, Google and SpaceX/xAI, alleging that the companies had agreed to coordinate the pace of frontier-AI development. The complaint argues that such coordination amounts to

YOU SAY: IT'S AI
APOCALYPSE!
I SAY: IT'S "RUINOUS
COMPETITION" ONCE
AGAIN

The opinions
expressed here do not
necessarily represent
the Fondazione CSF



Fondazione CSF

restricting innovation and output in violation of Section 1 of the Sherman Act. Whether an unlawful agreement exists remains to be established in court; what matters for the argument here is that the tension between AI “pacing” and competition has already moved from speculation to an explicit antitrust dispute.

In sum, the “pacing” argument amounts to a considerable extent to safety-washing: fears of an AI apocalypse provide cover for a classic and very “un-exceptional” move towards industry consolidation. This does not mean necessarily that AI industrialists do not believe in x-risk – one can be ideological and try to steer a market at the very same time. The leadership of Anthropic, for instance, is very much into the apocalyptic ideology, and recently even tried to convince the Pope about it. One difference from past ruinous-competition arguments is that policymakers did not always fall so easily for the story, and sometimes actively opposed it – the US authorities eventually turned antitrust law against Standard Oil and other trusts. This time around, however, raising the stakes from fears of industrial or social disruption to fears about the survival of humanity has found much friendlier ears among politicians in power: even European Commission President Ursula von der Leyen recently echoed the pacing the frontier narrative in her State of the Union Address. More importantly, at this historical juncture, attempts by the AI oligopoly to restrict competition have found powerful allies in nationalist governments, for which protectionism and support for the industry barons of the day can become complementary strategies. With the Trump administration in particular, we are witnessing an emerging alignment between actors seeking greater concentration of market power and those pursuing greater concentration of political power.

This alignment between market and political power is crystal clear in AI. The current US Administration has embraced the apocalyptic and exceptionalist rhetoric around the technology while pursuing a clear policy goal: supporting domestic AI companies in exporting their technology and turning technological dependencies abroad into a source of geopolitical leverage. This because less domestic competition can be considered instrumental to more competition on the international stage, especially against Chinese challengers of American incumbents. Trump himself first engaged in one of his greatest hits – suggesting mockery nicknames – proposing that AI should instead be called “Superior Intelligence,” “Extreme Intelligence,” or “Supreme Intelligence,” before converging on an executive order formally renaming AI “Super Intelligence” across the Administration and on instituting a Super Intelligence Force, tasked to contrast regulatory capture but not corporate capture. Indeed, more consequential than the forced renaming is the “White House Accord on Super Intelligence”, which gives political sanction to industry self-regulation and risks legitimising precisely the kind of coordination among AI labs discussed above.

As mentioned, the Administration has strong incentives to remain permissive towards concentration in the domestic AI industry because the exceptionalism narrative promoted by these companies is also instrumental to its disruptive international agenda. Michael Kratsios, the Science Advisor to the President of the United States, recently summarised this position in a post following his speech at the UN: “A prosperous future will not be secured by a global regulator. It will be secured by sovereign nations that adopt super intelligence, responsible companies that build it, and free people who refuse to be ruled by fear”. The Administration’s priorities are

YOU SAY: IT'S AI
APOCALYPSE!
I SAY: IT'S "RUINOUS
COMPETITION" ONCE
AGAIN

The opinions
expressed here do not
necessarily represent
the Fondazione CSF



Fondazione CSF

condensed there: opposition to global governance initiatives, a nationalist and protectionist posture, and a willingness to entrust the development of the technology to a handful of powerful domestic firms.

In summary, AI exceptionalism feeds a narrative in which industrial coordination over AI development can be justified in the name of safety and the public interest, while at the same time serving the much more familiar purpose of consolidating and concentrating market power. In this sense, today's "pacing" arguments resemble the ruinous-competition arguments advanced by cartels and industrial combinations in the past. What makes the current situation more troubling is the radicality of the justification and the political support it has attracted, while the users of AI remain exposed to much more immediate and tangible harms: bias, misuse and cognitive offloading, just to name a few. In this context, it is not AI models that have gone rogue. Rogue are the companies seeking to make their markets increasingly uncontestable through non-market means and end-of-the-world narratives, and the governments enabling this discourse when this aligns with their own protectionist agendas. The existential risk worth confronting, then, is not the $p(\text{doom})$, but that emerging from the erosion of multilateral governance and shared international rules.

The AI-apocalypse narrative persists partly because it is evocative to many, but also because there are powerful incentives to keep believing in it. To borrow Upton Sinclair's famous maxim, "it is difficult to get a man to understand something, when his salary depends on his not understanding it." Today, many "salaries"—or, more accurately, investments—depend on the bet on a continued expansion of AI. As in a case of cognitive dissonance, it may be easier to compute a speculative $p(\text{doom})$ than to confront the much more mundane economics of an industry still driven by financial frenzy rather than productive use cases.

If behind the AI-apocalypse narrative hides the risk of a government-sanctioned cartel, then the sensible way to bring AI production and deployment back into the public interest is through good old competition policy. This prescription is rather unexceptional, but so is AI (again, AI is remarkable, but not the promised "intelligence explosion"). The tools to pursue anticompetitive conduct already exist; what is lacking is the political will to start using them. The industry will resist this perspective, claiming that AI is a technology and that regulation should target its applications rather than the technology itself – Jensen Huang of Nvidia has recently made [exactly this point](#). But if we want to prevent ruinous-competition arguments from ruining our current and future welfare, it is time to strip current AI of its sci-fi technological mystique and start seeing it for what it also is: a series of products, services, and markets.

YOU SAY: IT'S AI
APOCALYPSE!
I SAY: IT'S "RUINOUS
COMPETITION" ONCE
AGAIN

The opinions
expressed here do not
necessarily represent
the Fondazione CSF



Fondazione CSF