

Against the Logic of “Pacing” AI

Dario Amodei's recent argument in [We Must Pace the Frontier](#) begins with a problem that is relatively straightforward: artificial intelligence is advancing extremely quickly, while our ability to make increasingly powerful systems safe, controllable, and predictable is not advancing at the same pace. His proposed solution is not to stop AI development, but to "pace the frontier." That is, to allow AI companies to continue developing increasingly powerful systems while requiring them to meet increasingly demanding safety standards before moving to the next stage. The basic idea is that technological capabilities should not advance substantially faster than our ability to control them.

The problem becomes more complicated in Amodei's estimation, however, because the United States is not developing AI in isolation. China is pursuing advanced AI as well, and Amodei worries that the United States could put itself at a disadvantage if it imposes constraints on its own development while China continues advancing without comparable restrictions. His concern is therefore not simply that AI could become dangerous, but that the United States could fall behind a geopolitical competitor that is less willing to slow its development in response to safety concerns. In this framework, the United States has to find a balance: it needs enough restraint to develop the safety mechanisms necessary for increasingly powerful AI, but not so much restraint that it loses its technological lead.

I believe this argument rests on an assumption that treats constraints on technological development primarily as a source of friction. If one country imposes more constraints on its technological development than another, the assumption goes, the first country should, all else being equal, advance more slowly. The country with fewer restrictions can move faster, develop more quickly, and potentially gain a competitive advantage. Amodei's concern about China therefore depends partly on the possibility that China will benefit from being less constrained while the United States risks suffering from constraining itself.

But there is another possibility: *the ability to operate under constraints can itself become a technological capability*. A guardrail does not necessarily have to function like a speed limit. It can become part of the infrastructure that allows a system to operate at greater speed and complexity. Once the mechanisms for managing risk become integrated into a technology, they may cease to be merely restrictions imposed upon it. They can become part of what makes further development possible.

The history of other complex technologies provides some reason to take this possibility seriously.

Aviation, for example, did not become more capable simply because engineers removed restrictions from aircraft. It became more capable as engineers developed better ways of controlling complex systems, including redundancy, instrumentation, simulation, testing, monitoring, and procedures for dealing with failure. These mechanisms impose requirements on aircraft development, but they also make possible a scale and level of complexity that would otherwise be impossible. A modern aircraft is not capable because it operates without constraints. It is capable because an extraordinary number of constraints have been incorporated into a system that can operate reliably within them. Safety, at a certain point, becomes part of the technology's productive capacity.

The same principle could eventually apply to artificial intelligence. The important breakthrough may not simply be a model capable of performing a new task. It may be a system that can become progressively more capable while also becoming more observable, testable, controllable, and resistant to catastrophic failure. Such a system would not necessarily be slower than one developed without those properties. It might ultimately be capable of doing things that a less controllable system could not safely be permitted to do.

This distinction matters when we consider the question of China. Amodei's analysis treats China as a potential "unpaced" competitor that continues advancing while the United States deliberately restrains itself. But there is no reason to assume that China's response to the risks of advanced AI will simply be to remove constraints. China has its own reasons to develop mechanisms for controlling increasingly powerful AI. As AI becomes embedded in energy systems, telecommunications, manufacturing, transportation, finance, government services, and military systems, a major AI failure would not simply harm an AI company. It could threaten the infrastructure and economic systems on which the country itself depends. The more deeply AI becomes integrated into society, the greater the incentive to make increasingly powerful systems reliable enough to operate without destabilizing the systems around them.

China's development of clean energy provides a useful example of how this dynamic can work. China has pursued an extraordinarily rapid expansion of solar power, wind power, batteries, and electric vehicles while simultaneously confronting the infrastructure problems created by that expansion. China's clean energy investment reached more than \$625 billion in 2024, and the country reached its 2030 target for installed wind and solar capacity six years early. At the same time, the rapid deployment of renewable energy has created challenges involving grid integration, transmission, storage, and electricity system management, leading to additional investment in those areas.

The significance of this example is not that China's energy policy has been uniformly successful or sustainable. China continues to invest heavily in coal, and its enormous expansion of manufacturing capacity has produced problems such as overcapacity and intense competition among producers. The relevant point is the relationship between technological expansion and the problems that expansion creates. Renewable energy creates new demands on the electricity grid. Those demands create incentives to build transmission, storage, and other forms of infrastructure capable of accommodating the new energy system.

In other words, technological development can generate the conditions for its own next stage. A new technology creates problems that existing systems were not designed to handle, and solving those problems requires new technological and infrastructural capabilities. The solution is not necessarily to reduce the original technology. It can be to build the capacity needed to support it.

China's response to foreign technological restrictions provides a different version of the same phenomenon. The United States and its allies have increasingly attempted to restrict China's access to advanced semiconductor technologies. These restrictions have clearly created obstacles for Chinese companies, and it would be inaccurate to claim that China can effortlessly reproduce every technology it is denied. *Nevertheless, the restrictions have also created powerful incentives for China to develop domestic alternatives.* [Recent analysis](#) from the Center for Strategic and International Studies argues that semiconductor export controls have given additional impetus to China's drive for technological localization, including efforts to develop domestic capabilities in semiconductor manufacturing equipment, chips, materials, and other parts of the supply chain.

The Huawei experience illustrates this dynamic particularly well. After the United States imposed restrictions on Huawei's access to American technology, Huawei and its Chinese partners intensified efforts to reduce their dependence on foreign components and technologies. By 2024, Huawei had introduced products incorporating increasingly advanced Chinese manufactured semiconductors, while Chinese firms were working to reproduce capabilities that had previously depended on foreign suppliers. The important point is not that technological restrictions are ineffective. They can cause substantial short term disruption and make technological development more difficult. The point is that a restriction can also reveal a dependency and create a powerful incentive to eliminate it.

This complicates the conventional understanding of technological competition. A constraint can impose an immediate cost while simultaneously creating pressure to develop capabilities that did not previously exist. In the energy example, rapid deployment creates pressure for better grids, storage, and transmission. In the

semiconductor example, restrictions create pressure for domestic manufacturing and supply chains. In both cases, the response to a technological constraint is not necessarily retreat. It can be adaptation.

That possibility changes how we might think about Amodei's argument about AI. Suppose the United States decides that frontier AI development must slow because it needs time to develop better methods of testing, monitoring, interpretability, cybersecurity, containment, and alignment. If China responds by pursuing increasingly capable AI while also developing increasingly sophisticated mechanisms for controlling those systems, then the competition is not simply between a "paced" United States and an "unpaced" China. It becomes a competition over who can solve the problems created by greater capability.

That is a fundamentally different conception of technological competition. It does not ask simply which country can remove obstacles and move fastest. It asks which country can develop the technological and institutional capabilities necessary to *operate at the edge of what is possible without being overwhelmed by the problems that edge creates*.

This is where the concept of pacing becomes somewhat misleading. The central question is not whether technological development should proceed quickly or slowly. Nor is it whether safety requirements should be imposed or removed. The deeper question is whether safety and control can *become part of technological capability itself*. If they can, then the country that develops effective guardrails may not be the country that falls behind. It may be the country that learns how to build more powerful systems precisely because it has learned how to operate them safely. The ultimate advantage, in other words, may very well end up belonging not to the country that has figured out how to ignore constraints, but to *the country that has figured out how to turn constraints into capabilities*.