Tomorrow's Affairs
Analysis of today
Assessment of tomorrow

By: The Editorial Board

Washington enters the AI race with no brakes

When the White House postponed the signing of an executive order on 21 May concerning the oversight of the most advanced AI models, it became clear that Washington does not yet have an answer to how to accelerate the development of artificial intelligence while maintaining at least minimal state control over systems whose security capabilities are advancing faster than current regulations.

The proposed regulation was not particularly radical: the Pentagon would be given a deadline to assess its own vulnerabilities, while companies developing the most advanced models would be required to submit them for a security check before public release.

It was neither a ban on development nor an attempt to tie the American technology industry to bureaucracy.

It was the bare minimum of government oversight into systems whose capacities are already approaching a point where an error is no longer just a business or reputational issue, but a matter of national security.

The Trump administration did not want to make even that minimum an obligation.

The reason is not difficult to discern. The prevailing view in Washington is that any significant restrictions could slow US AI companies at a time when China is rapidly developing its own models, chips, and computing infrastructure.

That assessment has political logic but carries a serious security risk. America is trying to maintain its technological advantage by delaying the creation of clear rules for technology that will soon become part of critical infrastructure, military planning, cyber operations, financial systems, and public administration.

That is the essence of the problem: Washington no longer regulates AI solely to protect citizens, markets, or the electoral process. It regulates, or avoids regulating, because of China.

China has changed the American calculation

A few years ago, discussion about artificial intelligence in the West focused mainly on ethics, disinformation, copyright, privacy, and the social consequences of automation.

These questions have not disappeared but have been overshadowed by a more rigid strategic logic.

In the American administration, everything is now measured by whether a regulation helps or hinders the preservation of an advantage over Beijing.

That is why AI chips, cloud systems, computing power, and the most advanced models are no longer ordinary products of the technology market. They have become instruments of state power.

American bans on the export of advanced chips to China, restrictions on Chinese companies' access to key technological layers, and increased control over cooperation with the Chinese sector show that Washington no longer views AI as just another digital industry, but as the foundation of future economic and military advantage.

> If AI is tightly regulated, American companies risk losing some of their speed

In this context, companies such as NVIDIA, AMD, Microsoft, Google, OpenAI, Anthropic, and xAI are no longer merely private actors.

Their decisions on model development, access to computing power, security testing, and new product launches directly affect America's position in technological competition with China.

The state therefore treats them not only as entities to be monitored, but also as a resource that must not be slowed down.

This creates a contradiction that the Trump administration cannot conceal. If AI is tightly regulated, American companies risk losing some of their speed.

If it is poorly regulated, the state is left without clear insight into systems that could be used for attacks against itself.

A country that relies on the goodwill of companies for the security of its most advanced AI systems is effectively admitting that it lacks a sufficiently robust legal framework for a technology it considers strategic

Voluntary oversight is not regulation

Washington is attempting to circumvent this contradiction through voluntary agreements with industry.

Google DeepMind, Microsoft, xAI, OpenAI, and Anthropic have agreed to some form of cooperation with the US Center for AI Standards and Innovation, allowing government institutions to preview the most advanced models before their public release.

On paper, this appears to be a reasonable compromise between innovation and security. In practice, it is a weak mechanism for a technology of such power.

A voluntary agreement is not a law. It does not impose a clear legal obligation, provide reliable sanctions, or resolve the issue of liability if a model that has passed assessment is misused or if the company supplies incomplete data.

The company may cooperate as it chooses, narrow the scope of cooperation, delay providing information, or effectively negotiate every aspect of oversight.

In other sectors, US self-regulation has followed the same pattern: industry accepts voluntary rules until they seriously interfere with the business model.

With AI, this problem is more dangerous because the same capability can be used for both defence and offence.

A model that identifies vulnerabilities in older software systems can help protect critical infrastructure, but it can also assist someone seeking to attack that infrastructure. The difference lies not in the model itself, but in the user, approach, and intent.

Therefore, a company's decision to limit the public release of a risky model may be responsible, but it cannot serve as the basis for national policy.

The next company may not do the same, especially if it believes the market, investors, or competition will penalise caution.

A country that relies on the goodwill of companies for the security of its most advanced AI systems is effectively admitting that it lacks a sufficiently robust legal framework for a technology it considers strategic.

Cyberspace does not wait for legislators

In mid-May, Congress addressed the practical side of the problem. A bipartisan group of representatives asked the National Cyber Director Sean Cairncross for a coordinated federal response to AI-generated cyber threats, clearer protocols for detecting vulnerabilities, and a stronger defensive approach to tools that accelerate the identification of weaknesses in information systems.

This was not an ideological intervention, but a response to a technical reality that American institutions find increasingly difficult to ignore.

The most advanced models can already accelerate the discovery of software vulnerabilities.

> Between the identification of a weakness and the development, verification, distribution, and installation of a patch, there is often enough time for the vulnerability to become an operational attack opportunity

However, the process of removing them cannot keep pace. Between the identification of a weakness and the development, verification, distribution, and installation of a patch, there is often enough time for the vulnerability to become an operational attack opportunity.

AI does not necessarily reduce this gap; it can also widen it, as it enables attackers to search a large number of systems more quickly and find mistakes that previously required more time, expertise, and human effort.

This is why the regulatory gap is not abstract. It is evident in questions such as who knows what the models can do before they are widely deployed, who has the authority to halt a launch if the risk becomes too great, who bears responsibility if the model is misused, and what happens when a company claims to have acted responsibly, and the state has no independent mechanism to verify this.

Washington does not yet have a convincing answer to these questions. Delaying regulation only makes them more apparent.

Europe is late in another way

The European Union faces the opposite problem. Brussels has adopted the most ambitious regulatory framework for AI, but its implementation is lagging, slowing, and colliding with the capacity of the regulators.

The AI Act remains the most detailed attempt yet to pre-define rules for high-risk systems, biometric surveillance, transparency, and accountability.

However, there is a significant gap between the text of the law and its actual implementation, as companies continue to develop the market faster than institutions can build oversight.

Brussels may write the standards but applies them in a market heavily dominated by American companies. It is regulatory ambition without a proper technological base

Delays in implementing key provisions for high-risk systems show that Europe has direction but lacks operational speed.

The reasons are partly technical: the standards are not finalised, national regulators are not equally prepared, and companies are requesting more time to adapt.

There is also a political dimension. American tech companies are lobbying strongly against rigid European regulation, and European governments know that an overly strict framework could further weaken Europe's already fragile position in the global AI race.

That is why the European model is not simply "strict regulation", as critics often claim.

It is a slow attempt to introduce rules in a sector where Europe does not have a leading industrial position.

Brussels may write the standards but applies them in a market heavily dominated by American companies. It is regulatory ambition without a proper technological base.

America fears that the rules will slow its companies. Europe fears that without rules it will have no control over technology it did not create. Both fears are rational. No system provides a complete answer.

Japan chooses a softer path, China waits for no one

Japan is seeking to avoid both American regulatory improvisation and the heavy European legislative apparatus.

Its approach relies on principles, recommendations, reputational pressure, and cooperation between the state and the economy, without strong punitive mechanisms. In the short term, this can be advantageous.

Companies invest more easily, implementation is not delayed by procedures, and the state retains flexibility.

However, such an approach has the same fundamental weakness as US voluntary arrangements: without a binding mechanism, the system relies on companies.

That may work in a stable business environment, but AI does not thrive in such conditions. It evolves under the pressure of competition, capital, geopolitical rivalry, and increasingly shorter launch cycles.

> A state that treats AI as a strategic asset can rapidly link regulation, industrial policy, security surveillance, and ideological control

China is a special case. Its model is not transferable to democratic systems, but it demonstrates an important fact: a state that treats AI as a strategic asset can rapidly link regulation, industrial policy, security surveillance, and ideological control.

That model carries its own risks and limitations, but it does not suffer from the same institutional indecisiveness that now affects the US, the EU, and Japan.

The result is a world without a common regulatory minimum. There is the American push for speed, the European tendency to apply rules late, the Japanese reliance on reputational discipline, and Chinese state control.

Companies developing the most advanced models will learn to move between these systems faster than countries can agree on how to harmonise them.

The real advantage is shifting to companies

The main consequence of this fragmentation is not legal disorder, but a shift in power relations. OpenAI, Anthropic, Google DeepMind, Microsoft, xAI and other companies possess capital, engineers, infrastructure, models and development continuity.

States have outdated laws, regulators who often lack technical expertise, and political processes that move more slowly than markets.

This allows companies to engage in regulatory arbitrage. They can tailor launches for different markets, negotiate data access, exploit differences between jurisdictions, and influence standards before they become mandatory.

When rules are delayed, market leaders do not wait; they establish practices that often later become standards.

The European Union will continue to promote the AI Act as a global standard, but its real strength will only be tested when the delayed provisions come into force

Therefore, the issue of AI regulation is not merely a dilemma between innovation and control; that is too simplistic.

The real question is whether democracies can establish even a minimal common framework for technology use in cyber defence, military systems, elections, finance, health, education, and administration.

If they cannot, the rules will be shaped by the business decisions of the strongest companies and by states' reactions after the first major incident.

The most likely short-term outcome in the US is a relaxed executive order that will formalise part of the existing voluntary agreements, but without full obligations or clear responsibilities for all actors.

This will allow the administration to show it has responded, while the industry retains room for manoeuvre.

The European Union will continue to promote the AI Act as a global standard, but its real strength will only be tested when the delayed provisions come into force.

Japan will remain attractive to companies seeking a more flexible framework. China will continue to build its own more closed system.

A more significant change is unlikely to result from a regulation, summit, or industry pledge, but rather from an incident large enough to reveal the cost of the regulatory vacuum.

This could be an AI-assisted cyber-attack on critical infrastructure, misuse of models in the financial system, serious manipulation of the electoral process, or harm in a sector where automated decision-making has direct consequences for human lives.

Then, the question postponed today will return in a much more difficult form: who knew what the model could do, who had the authority to stop it, and who bears responsibility when the state can no longer claim it was not warned.

Trump's delay of AI regulation is therefore not a minor episode in American technology policy.

It is an early indication that the world's most powerful technological nation has not resolved the relationship between market speed, national security, and public accountability.

As long as that relationship remains undefined, those who develop the models will write the rules, while governments will attempt to catch up only when the consequences become politically unavoidable.