Trevor Willis

Analytical Essay — Primary Threat to Global Security

The primary threat to global security today is not a weapon, a regime, or a conflict zone. It is a race without rules, the accelerating competition between the United States and China to develop and deploy advanced artificial intelligence without a shared framework to govern how that technology should be built, controlled, or constrained. The danger is not AI itself. It is the combination of transformative capability and a near-total governance vacuum.

To understand why this is the primary threat, it helps to understand what makes a threat primary. It must be global in scope, not regional. It must be capable of cascading into other domains, including economic, military, and political. And it must be urgent enough that delay meaningfully increases the risk of catastrophic outcomes. Uncoordinated AI development between the two most powerful states in the world meets all three criteria.

The US-China competition in AI is not a future scenario. It is already underway. Both governments have made frontier AI development an explicit national priority. Both are investing heavily in military applications. Both are exporting AI capabilities to third countries as instruments of geopolitical influence. Critically, neither has agreed with the other or with the broader international community on any meaningful norms governing the development or deployment of the most powerful AI systems. This is the governance vacuum. It is not simply an absence of policy but an absence of the international architecture that would make policy binding, verifiable, or enforceable across state lines.

The risks this creates are concrete, and events from this past week make them clear. On February 27, 2026, the Trump administration designated Anthropic, one of the leading American AI safety companies, a "supply chain risk" and banned federal agencies from using its products, a designation previously reserved for foreign adversaries. The conflict arose because Anthropic refused to allow its AI systems to be used for fully autonomous weapons or mass domestic surveillance of American citizens, insisting on explicit contractual guarantees. Hours after the ban, OpenAI struck a deal with the Pentagon under pressure, and CEO Sam Altman later acknowledged the agreement looked "rushed" and "opportunistic." This situation illustrates the central problem: even within a single country, there is no settled framework for what ethical guardrails on AI in national security contexts should look like, who enforces them, or what happens to companies that insist on them. If the United States cannot resolve these questions domestically, the challenge of establishing them internationally, across adversarial states with fundamentally different values, becomes even more daunting.

The governance vacuum compounds each of these risks. International institutions designed for the post-World War II order, such as the UN, and existing arms control frameworks were not

built to address this problem. They are slow, consensus-dependent, and increasingly contested. No equivalent to the Nuclear Non-Proliferation Treaty exists for AI, and the conditions that made nuclear arms control possible, including limited actors, visible infrastructure, and clear deterrence logic, do not straightforwardly apply to software that can be copied, concealed, and deployed invisibly. The institutions responsible for managing this technology are not keeping pace with it. Addressing this threat requires different approaches on different timelines. In the short term, the priority must be to establish baseline ethical guardrails, starting domestically and building outward. The Anthropic-Pentagon dispute makes clear that even the most basic questions, such as whether AI should be permitted to make lethal decisions without human oversight, remain unresolved. Resolving them and enshrining them in enforceable agreements, rather than informal assurances, is a prerequisite for any meaningful international framework. Back-channel technical dialogues between the US and China should run in parallel to reduce the risk that AI-related miscalculations escalate into conflict. In the long term, the goal must be to build the institutional architecture that does not yet exist, one capable of setting and enforcing global ethical standards for AI development and deployment. This means multilateral forums specifically designed for AI governance, verification mechanisms that build confidence across state lines, and the inclusion of smaller states already receiving AI exports from both the US and China. The companies building this technology have shown, through their own internal debates, that they understand where the red lines should be. The challenge is to construct the international framework that makes those red lines binding, beyond the reach of any single government's political pressure. The primary threat to global security is not that AI will be used against us. It is that the race will outpace the rules, and that by the time the world agrees on the necessary guardrails, the decisions that matter most will already have been made