DOI: 10.51845.39.2.4

The Ethics of Autonomous War: Power, Responsibility, and the Limits of Restraint

by Lipton Matthews

On April 18, 2026, Palantir Technologies posted to X a 22-point manifesto drawn from CEO Alexander Karp and head of corporate affairs Nicholas Zamiska's book *The Technological Republic: Hard Power, Soft Belief, and the Future of the West*.[1] Titled "The Technological Republic, in Brief," the post accumulated over 32 million views within days, igniting fierce debate across political, academic, and technological circles.[2] The manifesto covers terrain ranging from the reinstatement of compulsory national service to the alleged cultural failures of postwar European disarmament. But two of its 22 points cut to the intellectual core of the entire enterprise.

The first: "The question is not whether AI weapons will be built; it is who will build them and for what purpose. Our adversaries will not pause to indulge in theatrical debates about the merits of developing technologies with critical military and national security applications. They will proceed."[3] The second: "If a U.S. Marine asks for a better rifle, we should build it; and the same goes for software. We should as a country be capable of continuing a debate about the appropriateness of military action abroad while remaining unflinching in our commitment to those we have asked to step into harm's way."[4] These propositions, taken together, constitute a coherent and largely correct theory of what democratic states owe both to their strategic position and to the individuals who bear the costs of their military decisions.

The first argument is not, at bottom, a philosophical claim. It is a description of documented reality. Russia and China have not paused to conduct ethical seminars on the appropriate role of machine autonomy in lethal decision-making. In 2017, Vladimir Putin declared that whoever leads in artificial intelligence "will become the ruler of the world."[5] China's 2017 Next Generation AI Development Plan formally designated artificial intelligence a stra-

tegic national priority, a focal point of great-power competition, and a pillar of the People's Liberation Army's drive to become a world-class military with advanced warfare capabilities.[6] These are not rhetorical commitments. China is developing systems that allow dozens of autonomous drones to coordinate attacks without human input. Russia has already incorporated autonomous targeting features into its Lancet loitering munition, which has been deployed in Ukraine.[7]

Ukraine itself offers the most vivid evidence that the era of autonomous weapons is not approaching—it has arrived. Drones now account for roughly 70 to 80 percent of battlefield casualties in that conflict.[8] Ukraine's strategic response has been to accelerate the replacement of human fighters with autonomous crewless systems, driven, as analyst Kateryna Bondar has observed, by the need to "conserve a limited human force and overcome vulnerabilities such as fatigue, stress, and the limited capacity to process and fuse large amounts of data from various sources and sensors."[9] General Valerii Zaluzhnyi, former commander of Ukraine's armed forces, put it plainly: "The revolution in military technology based on unmanned systems and artificial intelligence has completely changed the nature of war and continues to evolve."[10]

Against this backdrop, the suggestion that democratic nations should abstain from developing autonomous weapons rests on a foundational error: the belief that unilateral restraint is contagious. History provides no example in which a major power's voluntary technological disarmament persuaded its rivals to reciprocate. The nuclear age, the ballistic missile revolution, precision-guided munitions—at every inflection point, the question was not whether the technology would proliferate but which powers would possess decisive capability first, and under what governance frameworks it would be used. The AI weapons era follows the same structural logic, compressed into a faster timeline.

What changes when democracies decline to develop autonomous weapons is not the global trajectory of the technology. It is the identity of those who hold the decisive advantage. A world in which China or Russia leads in autonomous military systems, unchallenged by democratic powers with competing capabilities, is not a more ethical world. It is a world in which the norms governing the use of lethal autonomous systems, the accountability structures that surround targeting decisions, and the geopolitical balance that has underwritten whatever remains of the liberal international order are all set by powers with no meaningful commitment to those norms. Karp and Zamiska's formulation—that "hard power in this century will be built on software"[11]—is not a prescription for aggression. It is a statement of competitive reality that serious people cannot afford to ignore.

The affirmative case for autonomous weapons development extends beyond geopolitical competition. Proponents

advance three distinct arguments: that such systems can reduce soldier casualties by substituting machine risk for human risk in high-threat environments; that sufficiently sophisticated autonomous targeting, operating without fatigue, fear, or the psychological distortions of combat, could in principle apply the laws of war more consistently than any human operator under fire. Furthermore, autonomous systems give democratic nations a strategic edge by shifting military power away from a few costly, hard-to-replace platforms toward large numbers of inexpensive, expendable assets.[12]

Because these systems can be mass-produced and treated as "consumable munitions," democracies can absorb losses and quickly regenerate capability without the same financial or political strain. This plays to their strengths in innovation, industrial capacity, and economic scale, while also reducing reliance on human operators and limiting the domestic backlash associated with casualties

On the casualty question, the logic is straightforward. A drone conducting reconnaissance over heavily defended territory does not put a pilot at risk. An autonomous loitering munition searching for armor in a communications-denied environment does not require a soldier to advance into contested ground. To the degree that autonomous systems can substitute machine risk for human risk in defined mission sets, they represent a genuine moral gain—not a retreat from ethical standards but an expression of the oldest ethical principle in the conduct of war: do not send human beings into danger if the objective can be achieved otherwise.

The precision argument carries weight as well. Much of the moral outrage that attaches to contemporary warfare concerns civilian casualties produced by imprecise weapons or rushed decisions under fire. Human soldiers, operating under stress, sleep deprivation, and the fog of battle, make targeting errors with catastrophic consequences. A sufficiently sophisticated autonomous system, processing sensor data without fatigue and without the emotional disruptions that combat induces, could apply rules of engagement with greater consistency. This is precisely why the U.S. Department of Defense's 2020 *Ethical Principles for Artificial Intelligence* explicitly calls for responsible, traceable, and governable AI development—a framework designed not to eliminate human judgment but to augment and constrain autonomous action within a defined ethical envelope.[13]

Cost matters too. Defense contractor Kratos has claimed it can produce several hundred Valkyrie units annually at a per-unit cost of two to five million dollars, less than manned aircraft, traditional drones like the American Reaper, and in some cases less than individual missiles.[14] Asymmetric cost advantages of this magnitude change the strategic calculus in ways that go beyond any single procurement decision. A democratic nation that can field

ten autonomous systems for the cost of one manned platform has options that a nation dependent on expensive, vulnerable, human-crewed systems simply does not.

The ethical case against autonomous weapons is neither trivial nor dishonest. Critics advance several objections that deserve genuine engagement. The most technically serious is the distinction problem: the difficulty that current AI systems face in reliably distinguishing combatants from civilians in complex, contested environments. Robotics scholar Noel Sharkey has argued that machines cannot interpret the subtle human cues necessary to make that distinction reliably, raising serious doubts about whether autonomous weapons can comply with the international humanitarian law principle of distinction.[15] A 2024 United Nations analysis reinforced this concern, noting that the ambiguity surrounding AI-based targeting decisions muddies legal accountability and may leave no individual legally liable for violations.[16] The International Committee of the Red Cross has similarly warned that AI increases the risk that operators will be unable to understand, predict, or control a weapon's effects—potentially rendering it indiscriminate under international humanitarian law.[17]

These are real problems that demand real solutions. But they are not arguments against developing autonomous weapons. They are arguments for developing them responsibly—which is precisely the position that Karp and Zamiska, along with the DoD's own governance framework, advance. The alternative—declining to develop such systems while adversaries proceed without restraint—does not resolve the distinction problem. It ensures only that the autonomous weapons systems operating in future conflicts will be Russian or Chinese, built under governance frameworks far less attentive to civilian protection than those of democratic militaries. The question, as the manifesto insists, is who builds these weapons and for what purpose. Abstention is not an ethical answer to that question.

The accountability objection—that there is no one to hold responsible when an autonomous system commits a war crime—is real but overstated. Chain-of-command responsibility, rules of engagement, procurement decisions, and doctrine development all create accountability structures that operate independently of whether a particular trigger is pulled by a human or an algorithm. Commanders who deploy autonomous systems that foreseeably commit violations bear responsibility for those deployments. The legal framework requires development and adaptation, not prohibition. As the DoD's governance approach recognizes, the goal is ensuring that AI-enabled systems are reliable, auditable, and aligned with military law—adaptive regulation rather than a ban that only self-disarms the nations most likely to enforce it.[18]

The escalation concern—that removing the human cost of initiating hostil-

ities will make leaders more willing to use force, and that autonomous systems cannot perform the nuanced signaling that deterrence requires—is the most speculative of the objections. It deserves serious attention in both doctrine and system design. But the Ukraine conflict, in which autonomous drone warfare has coexisted with sustained and consequential human deliberation about escalation thresholds, suggests that the two domains can be managed in parallel. The risks of autonomous systems in warfare are manageable technical and governance challenges. The risk of ceding autonomous military advantage to authoritarian rivals is a strategic and civilizational danger of a different order entirely.

The second argument from the manifesto is in some respects the more morally demanding of the two, because it asks democratic societies to hold two difficult positions simultaneously. A nation must be capable of debating the justice of its military commitments vigorously and publicly—that capacity is one of the defining features of liberal democracy. But that debate cannot be permitted to degrade the capabilities of the soldiers who are already in the field, already in harm's way, already bearing costs that the debaters are not.

The point is not abstract. In 2018, Google employees organized a protest that forced the company to decline renewal of its Project Maven contract with the Department of Defense—a contract for AI tools that analyzed drone surveillance footage in support of military operations. Palantir subsequently took on that work. The employees exercised a genuine moral prerogative; they believed their labor should not contribute to operations they opposed. But the operations did not stop. The soldiers conducting them continued. The effect of the protest was not to prevent harm. It was to ensure that the people bearing risk on behalf of the country did so with somewhat less capable tools.

This is the practical expression of what Karp and Zamiska are identifying. The Silicon Valley resistance to defense work is documented and substantial.[19] The same companies that built the commercial infrastructure of the modern internet—infrastructure that exists within the security guarantees provided by American military power—have repeatedly declined to apply their most advanced capabilities to defense applications. The manifesto's opening proposition—that Silicon Valley owes a moral debt to the country that made its rise possible, and that its engineering elite has an affirmative obligation to participate in the defense of the nation—is not nationalism dressed up as philosophy.[20] It is an argument about reciprocity, about what it means to benefit from a system of security and then decline to help maintain it.

The deeper moral logic concerns what it means to ask someone to risk their life on behalf of the state. Both conscription and voluntary service represent a form of social contract: the soldier accepts exposure to danger in

exchange for the nation's commitment to support that soldier with every available resource. To withhold technological capability from a deployed force not because it would be operationally harmful but because the engineers who could build it find the policy that sent that force into the field politically objectionable—is a violation of that contract. It conflates two separate questions: whether force was justly authorized, and whether those who bear the cost of that force deserve the best efforts of their country. Those who bear the cost always do.

None of this is an argument for unlimited military spending, unquestioning deference to military judgment, or the suppression of political dissent. The debate about the appropriateness of military action must continue, and must be vigorous. Democratic societies that cannot subject their military commitments to sustained critical scrutiny are headed toward strategic disasters. But the debate about policy and the commitment to those who execute policy are not the same debate. Confusing them imposes costs paid not by the decision-makers but by the people they sent downrange.

Karp and Zamiska resist, ultimately, a form of moral aestheticism—the preference for the appearance of clean hands over the difficult work of maintaining power in a world that does not reward restraint by only one party. Their assertion that "the ability of free and democratic societies to prevail requires something more than moral appeal" and that hard power in this century will be built on software has been contested by idealists of every era and confirmed by the outcomes of every serious conflict between democracies and their adversaries.[21] The nations that survived the twentieth century did so not by appealing to the better natures of their enemies but by outproducing, outfighting, and out-innovating them.

The ethical obligations that attach to this power are genuine. Autonomous weapons must be developed under governance frameworks that take seriously the requirements of distinction, proportionality, and accountability. The question of meaningful human control over lethal decision-making is not theater—it is both a legal obligation and a moral one, and it requires serious institutional investment. None of this, however, requires democratic powers to stand aside while adversaries who share none of these commitments proceed without constraint.

The two propositions from the Palantir manifesto address, in compressed form, the most consequential questions in contemporary democratic theory: what a free society owes to those it deploys, and what it must be willing to build in order to remain free. The answers are uncomfortable precisely because they refuse the consolation of abstention. They insist that power and responsibility are inseparable, and that in the age of algorithmic warfare, the choice to disengage is itself a choice—one whose consequences democratic societies will not be able to avoid bearing.

Lipton Matthews *is a research professional and podcaster whose work explores history, politics, and culture. He holds an MSc in History from the University of Glasgow and is an Emergent Ventures Fellow. His writing has been featured in a range of publications, including* American Greatness, Merion West, The Federalist, The Epoch Times, *and the* Polish Journal of Political Science.

1. Brent D. Griffiths, "Palantir's summary of CEO Alexander Karp's manifesto is generating buzz. Read the 22 bullet points," *The Business Insider*, April 20, 2026.
2. Marco Quiroz-Guiterrez, "Palantir Published a mini-manifesto calling some cultures 'harmful and middling'," *Fortune*, April 22, 2026.
3. Cheyenne Macdonald, "Palantir posted a manifesto that reads like the ramblings of a comic book villain," *Endgadget*," April 19, 2026.
4. Karp and Zamiska, manifesto point 7.
5. Aaron Bateman, "Russia's Quest to Lead the World in AI is doomed," *Defense One*, June 12, 2019.
6. State Council of the People's Republic of China, "In their Own Words: "A Next Generation Artificial Intelligence Development Plan," China Aerospace Studies Institute, July 2017.
7. Bohuslav Romanenko, "Ukraine's Intelligence reports major upgrades to Russian Lancet drone," *The New Voice of Ukraine*, March 26, 2026.
8. David Kirichenko, "Artificial Intelligence's Growing Role in Modern Warfare," *War Room Online Journal*, August 21, 2025.
9. David Kirichenko, "How AI is eroding the norms of warfare," AI Frontiers, May 20, 2025.
10. General Valerii Zaluzhnyi, quoted in Kirichenko.
11. Joe Sommerland, "Surveillance tech giant Palantir posts manifesto on X dealing with national service, post-World War II attitudes and diversity," April 20, 2026.
12. Amitai Etzioni, Oren Etzioni, "Pros and Cons of Autonomous Weapons Systems," *Military Review*, May-June 2017.
13. C. Todd Lopez "DOD Adopts 5 Principles of Artificial Intelligence Ethics," *Defense.gov*, Mar. 11, 2020.
14. Riley Simmons-Edler et al., "AI-Powered Autonomous Weapons Risk Geopolitical Instability and Threaten AI Research," arXiv, 2024, arxiv.org/html/2405.01859v1.
15. Noel Sharkey, quoted in "The Backlash Against Military AI: Public Sentiment, Ethical Tensions, and the Future of Autonomous Warfare," TRENDS Research and Advisory, Sep. 24, 2025, trendsresearch.org/insight/the-backlash-against-military-ai.
16. *Ibid*.
17. International Committee of the Red Cross, "We Cannot Let AI Be Deployed on the Battlefield Without Oversight and Regulation," statement to the UN Security Council, September 26, 2025.
18. Shamsa Al Qubaisi, "Governing Lethal Autonomous Weapons in a New Era of Military AI," Trends Research and Advisory, August 3, 2025.
19. "Defense Tech Is Blowing Up Silicon Valley's Beliefs." *The Economist*, Feb 13, 2025.
20. Karp and Zamiska, manifesto point 1, *Engadget*, Apr. 19, 2026.
21. Karp and Zamiska, manifesto point 4.