Article

Not peer-reviewed version

The Great Compression: Geopolitical Fragmentation, AI, and the Coming Neo-Feudal Order

Priyanshu Jain *

Posted Date: 2 April 2026

doi: 10.20944/preprints202604.0074.v1

Keywords: artificial intelligence; wealth inequality; democratic backsliding; neo-feudalism; geopolitical fragmentation; universal basic income; political economy; populism; AI governance; sovereignty; nation-state

 doi:10.20944/preprints202604.0074.v1

Article

The Great Compression: Geopolitical Fragmentation, AI, and the Coming Neo-Feudal Order

Priyanshu Jain

Independent Researcher, India; pj@73ai.org

Abstract

The post-World War II international order is undergoing simultaneous collapse on two fronts: a geopolitical fragmentation driven by twenty consecutive years of democratic decline, and an accelerating concentration of economic power driven by advances in artificial intelligence. This paper argues that the convergence of these two forces is producing a structural transformation unprecedented in human history, one that could stabilize into a neo-feudal equilibrium in which a vanishingly small class of infrastructure owners wields power comparable to pre-Enlightenment monarchs, while the vast majority of humanity loses both its labor value and its political leverage. Unlike previous feudal orders, this one may prove uniquely resistant to revolution, because the mechanisms of enforcement (autonomous weapons, AI surveillance, algorithmic propaganda) do not require human cooperation and therefore cannot be undermined by human dissent. The paper examines the historical parallels (and crucial disanalogies) between contemporary populist-authoritarian movements and their twentieth-century predecessors, models the emerging class structure under conditions of artificial general intelligence, evaluates Universal Basic Income through the lens of incentive structure, arguing that without the revolutionary threat that historically forced redistribution, UBI will default to a pacification mechanism rather than a genuine solution, examines the future of the nation-state under conditions where AI infrastructure owners command more wealth and capability than most governments, and argues that the effective altruism community's near-exclusive focus on existential risk from AI has created a dangerous blind spot around the political economy of who controls AI and who benefits from it.

Keywords: artificial intelligence; wealth inequality; democratic backsliding; neo-feudalism; geopolitical fragmentation; universal basic income; political economy; populism; AI governance; sovereignty; nation-state

---

1. Introduction: The Convergence Thesis

There is a scene in Carl Sagan's *Contact* where the protagonist, asked what single question she would pose to an alien civilization, replies: "How did you do it? How did you survive this technological adolescence without destroying yourself?" Dario Amodei, CEO of Anthropic, opens his January 2026 essay *The Adolescence of Technology* with this same scene Amodei (2026a). The choice is revealing. The head of one of the world's most advanced AI companies is publicly framing the current moment not as a business opportunity but as a civilizational survival test.

This paper takes Amodei's framing seriously but extends it beyond the technical risks he primarily addresses (autonomy, misuse, bioweapons) to a structural risk he acknowledges but does not fully develop: the possibility that artificial intelligence, arriving into a world already fracturing along political and institutional lines, will not merely disrupt the global economy but fundamentally alter the social contract that has underwritten democratic governance since the Enlightenment.

The argument proceeds in three stages. First, the post-WWII international order is not being replaced by a new order; it is decomposing into a prolonged interregnum of fragmentation, competing blocs, and institutional hollowing, a process accelerated by the very populist-authoritarian movements that claim to resist elite power. Second, artificial intelligence of the kind now emerging, what Amodei

 doi:10.20944/preprints202604.0074.v1

calls "a country of geniuses in a datacenter" Amodei (2026a), will, within the next decade, eliminate the economic necessity of most human labor, concentrating wealth and productive capacity in the hands of those who own compute infrastructure, energy, and foundational models. Third, the convergence of these two forces produces a class structure that is neither capitalist nor socialist in any familiar sense, but neo-feudal: a small owner class whose power derives not from organizing human labor but from controlling the non-human substrate of production, and a vast dependent class whose material needs may be met but whose political leverage has evaporated.

The central claim is that this neo-feudal equilibrium, once established, may be uniquely stable, because the enforcement mechanisms available to the owner class (autonomous weapons, pervasive surveillance, individually tailored propaganda) do not depend on human soldiers, human police, or human administrators who might refuse orders or join a revolution. If this analysis is correct, the window for structural intervention is not decades but years, and the effective altruism community's near-exclusive focus on existential risk from misaligned AI, while important, constitutes a dangerous blind spot regarding the political economy of who controls AI and in whose interest it operates.

2. Methodology and Approach

This paper adopts a hypothesis-driven analytical framework. The author advances a set of interconnected hypotheses about the convergence of geopolitical fragmentation and AI-driven economic concentration, and then systematically evaluates each hypothesis against available empirical evidence, published scholarship, institutional research, and expert analysis.

The core hypotheses are:

H1: The post-WWII international order is undergoing irreversible fragmentation rather than cyclical stress, and no replacement order is emerging.

H2: Contemporary populist-authoritarian movements, despite ideological diversity across the left-right spectrum, share a common operational playbook with historical predecessors, and this playbook systematically benefits economic elites while claiming to serve the working class.

H3: Artificial intelligence constitutes a qualitative discontinuity in technological disruption, not merely a quantitative acceleration, because it functions as a general substitute for human cognition rather than a complement to specific skills.

H4: The convergence of H1 and H3 produces a neo-feudal class structure (modeled here as a 1/5/94 distribution) that is self-reinforcing and potentially stable.

H5: AI-powered enforcement mechanisms eliminate the historical revolutionary check on elite power, making this neo-feudal structure uniquely resistant to challenge from below.

H6: Universal Basic Income, absent structural changes to political power, will default to a pacification mechanism calibrated to the minimum level necessary to prevent social unrest, rather than a genuine redistributive solution.

H7: The effective altruism movement's focus on AI existential risk has produced an insufficient engagement with the political economy of AI ownership and control.

H8: The nation-state, as currently constituted, will become functionally subordinate to, fused with, or replaced by direct coordination among AI infrastructure owners, as the fiscal and coercive foundations of state power erode under conditions of mass AI-driven labor displacement.

Each hypothesis is evaluated using evidence drawn from the following categories of sources:

(a) Quantitative democracy indices: Freedom House's *Freedom in the World* annual reports Freedom House (2026), the V-Dem Institute's annual Democracy Reports V-Dem Institute (2026), and the Economist Intelligence Unit's Democracy Index.

 doi:10.20944/preprints202604.0074.v1

(b) Institutional risk assessments: the Stimson Center's global risk analyses Stimson Center (2026), the Eurasia Group's annual top-risks reports Eurasia Group (2026), the Munich Security Conference reports Munich Security Conference (2025), the World Economic Forum's Global Risks Reports World Economic Forum (2026b), and the Council on Foreign Relations' Preventive Priorities Surveys Council on Foreign Relations (2026).
(c) Economic research: data from the Economic Policy Institute Economic Policy Institute (2026a,b), the Center for American Progress Center for American Progress (2025), Oxfam Oxfam (2025), the OECD, and peer-reviewed economic scholarship on wealth inequality and AI labor displacement Acemoglu and Restrepo (2019); Aluko et al. (2026).
(d) Primary statements and publications from AI industry leaders, notably Amodei's *The Adolescence of Technology* Amodei (2026a) and *Machines of Loving Grace* Amodei (2024), which represent the most detailed public statements from a frontier AI CEO on the risks addressed in this paper.
(e) Comparative political science and historical scholarship on populism Funke et al. (2023), democratic backsliding Mo-OrgCip (2026), and authoritarian governance Ben-Ghiat (2020).

The paper's central argument is explicitly normative: it contends that a particular trajectory is probable and undesirable, and that specific interventions are needed to alter it. Where the analysis rests on the author's judgment rather than on cited evidence, this is clearly indicated through phrases such as "we hypothesize," "the model proposed here suggests," or "in the author's assessment." All empirical claims are cited to their sources. The reader is invited to evaluate the hypotheses independently and to contest those they find insufficiently supported.

A note on limitations: the paper draws on publicly available data and analysis as of March 2026. Projections about AI capability trajectories, economic impacts, and political outcomes are inherently uncertain. The 1/5/94 class model proposed in Section 5 is a stylized analytical framework, not a precise empirical prediction, and is intended to illuminate structural dynamics rather than forecast exact distributional outcomes.

3. The Crumbling Order: Twenty Years of Democratic Decline

3.1. The Empirical Record

The quantitative evidence for global democratic decline is now unambiguous across every major measurement framework. Freedom House's *Freedom in the World 2026* report documents the twentieth consecutive year of declining global freedom, with 54 countries experiencing deterioration in political rights and civil liberties against only 35 registering improvements Freedom House (2026). The V-Dem Institute's 2026 Democracy Report finds that autocracies now outnumber democracies by 92 to 87, the first time this has been the case in over two decades V-Dem Institute (2026). Less than 12 percent of the world's population now lives in a liberal democracy, the lowest proportion in fifty years; 72 percent live under autocratic rule V-Dem Institute (2026).

What makes the current wave of democratic backsliding particularly dangerous is its geographic reach. The V-Dem 2026 report identifies six of the ten newly autocratizing countries as being in Europe and North America, including Italy, the United Kingdom, and the United States V-Dem Institute (2026). The Stimson Center's 2026 global risk assessment observes that "the old neoliberal rules-based architecture is decomposing, power diffusing, and much of the world is searching for new multilateral arrangements" Stimson Center (2026). The Council on Foreign Relations' analysis of Freedom House data concludes that "the authoritarians are winning" and that "the most powerful rich democracies, from the United States to Germany to Japan, are focusing on their own domestic problems, cutting support for democracy promotion and humanitarian aid" Kurlantzick (2026).

Assessment of H1: The evidence strongly supports the hypothesis that the post-WWII order is undergoing irreversible fragmentation. Twenty consecutive years of decline, accelerating in recent years to include core Western democracies, with no countervailing institutional formation, is consistent with systemic decomposition rather than cyclical stress.

 doi:10.20944/preprints202604.0074.v1

3.2. The Populist Playbook

The mechanisms driving democratic decline exhibit remarkable consistency across ideological lines. A comprehensive study published in *Foreign Affairs*, covering 120 years of history across 60 countries, identified 51 populist leaders and found that most weakened their country's economy, especially in the long run, by undermining rule of law and eroding political checks and balances Funke et al. (2023). Yet they managed to hold power by polarizing the electorate and manipulating the political system Funke et al. (2023).

Whether one examines Hugo Chávez's Venezuela, Viktor Orbán's Hungary, Narendra Modi's India, Recep Tayyip Erdoğan's Turkey, or Donald Trump's United States, the operational playbook contains the same core elements. Scholars of authoritarianism have documented these patterns extensively Ben-Ghiat (2020); Duina (2022); Factually (2025):

National humiliation narrative. A story in which the nation was once great, was betrayed by elites and foreigners, and can only be restored by a strong leader. Historians have noted the structural similarity between "Make America Great Again" and earlier nationalist restoration narratives, though with important caveats about differences in scale and context Factually (2025).

Internal enemy identification. A scapegoat class onto which economic grievances are redirected, away from the structural causes and often away from the leader's own policies. As Ben-Ghiat observes, authoritarians encourage a mode beyond mere polarization, a "survivalism" in which political competition is framed as existential: not "me versus you" but "me or you, and only one of us can survive" Ben-Ghiat (2020).

Reality distortion infrastructure. The systematic undermining of shared epistemic foundations. The World Economic Forum's 2026 Global Risks Report identifies "distinguishing truth from falsehood" as increasingly difficult, noting that this dynamic is "deepening societal fragmentation and contributing to desensitization" World Economic Forum (2026b).

Institutional hollowing. Not dramatic overthrow but gradual emptying of substance while preserving formal structures. Research in comparative politics increasingly emphasizes that "contemporary democratic regression rarely takes the form of dramatic events such as military coups" but instead "unfolds through incremental institutional changes: modifications of electoral rules, reconfiguration of judicial authority, constraints on media independence, and growing pressure on civil society" Mo-OrgCip (2026).

3.3. The Gap Between Populist Rhetoric and Economic Reality

The case of the United States under President Trump's second term provides the most consequential contemporary case study, precisely because American power has been the keystone of the post-WWII order.

The empirical record on economic outcomes for the working class directly contradicts the populist narrative. The Center for American Progress reports that "the first year of the second Trump administration resulted in job losses, slowing wage growth, and rising costs for working-class Americans," with employment for workers without college degrees falling by 361,000 jobs and 58,000 manufacturing jobs lost following the "Liberation Day" tariffs Center for American Progress (2025). The Economic Policy Institute's analysis of the tax legislation found that "the biggest gains in new tax cuts will go to wealthier households, including those who own businesses or expensive homes in states with high property taxes and those who receive multimillion-dollar inheritances," while working-class provisions are "much, much smaller" Economic Policy Institute (2026a). NBC News reports that the wealth divide under the current administration is "the widest it has been in at least a generation" NBC News (2026).

Yet Republican base support remains robust. Trump's approval stands at approximately 37 percent as of March 2026, with net approval at a second-term record low of -16.7 Silver (2026), but his support among Republican voters remains above 73 percent Pew Research Center (2026). This illustrates the decoupling of policy outcomes from political loyalty that characterizes mature populist

 doi:10.20944/preprints202604.0074.v1

movements, a phenomenon explained by research on identity-based political allegiance, information ecosystem fragmentation, and the sunk-cost psychology of partisan commitment ide (2026).

The Eurasia Group's 2026 top-risks assessment identifies the United States itself as "the principal source of global risk in 2026" Eurasia Group (2026).

Assessment of H2: The evidence strongly supports the hypothesis that populist-authoritarian movements share a common operational playbook and systematically benefit economic elites while claiming to serve the working class. The US case demonstrates that empirically falsifiable economic claims ("champion of the working class") can coexist with measurable harm to the claimed beneficiaries without eroding base support.

4. The Fragmentation Scenario: Slow Collapse Without Reset

4.1. Why There Will Be no Clean Reset

The post-WWII institutional architecture (the UN, NATO, the IMF, the World Bank, the WTO, the Bretton Woods monetary system) was built on rubble, by victors who had been so thoroughly traumatized by war that they were willing to surrender sovereignty to shared institutions Wikipedia and Coppola (2023). Every major "reset" in international order has followed a similar pattern: the Peace of Westphalia after the Thirty Years' War, the Concert of Europe after the Napoleonic Wars, the League of Nations after World War I. As the Carnegie Endowment's analysis of the Bretton Woods legacy observes, the architects of the post-war order "shared an overriding mission: to prevent another global conflict" and placed "economic and financial issues first when it came to imagining and building the peace" Carnegie Endowment for International Peace (2024).

The current trajectory does not include a cataclysmic reset event, though the risk of one exists. The Iran war, the Russia-Ukraine conflict, and the Taiwan Strait remain active flashpoints Council on Foreign Relations (2026); Stimson Center (2026), and the expiration of the New START treaty in February 2026 has removed the last binding limits on US-Russian nuclear arsenals Eurasia Review (2026). What is more likely, and arguably already underway, is what the Stimson Center calls a "protracted interregnum" Stimson Center (2026).

JP Morgan's 2026 *World Rewired* report describes a "multi-speed order with little historical precedent" in which "trade, technology, energy, security, and climate policy are no longer moving in sync" JP Morgan Chase (2026). The Munich Security Conference's 2025 report frames it as "multipolarity without multilateralism," warning that this configuration is "shrinking the proverbial global pie, potentially triggering 'lose-lose' dynamics where everyone will be worse off in the long run" Munich Security Conference (2025).

4.2. Who Benefits, Who Suffers

In a fragmenting order, the distribution of advantage follows a clear pattern: those with geographic, professional, and financial optionality navigate the transition best; those without it absorb the costs.

The World Economic Forum's Davos 2026 coverage documents how middle powers (Indonesia, Brazil, Turkey, Vietnam) are positioning themselves as connectors across competing blocs World Economic Forum (2026a). Finland's president framed the stakes clearly: "The world order is changing much like it did after 1918, 1945, and 1989," warning that the choice is between "a multipolar world about transactions, deals, and spheres of interest" and "a multilateral world grounded in institutions, rules, and norms" World Economic Forum (2026a).

Among those disadvantaged by fragmentation, the most significant structural losers may be the major European economies. As Foreign Affairs analysis notes, "the European Union has economic clout but remains politically divided and dependent on the United States for its security" Mohan (2026), while its manufacturing model, which depended on cheap Russian energy and open Chinese markets, has lost both pillars.

 doi:10.20944/preprints202604.0074.v1

For individuals, professional advantage accrues to those whose skills operate across systems: AI engineers, cybersecurity professionals, compliance experts, and supply chain specialists. The disadvantaged are those whose skills are jurisdiction-specific and non-portable. However, this distributional pattern is about to be radically amplified by the arrival of artificial intelligence capable of performing most human cognitive labor.

5. AI as Accelerant: The Country of Geniuses

5.1. The Capability Trajectory

Amodei defines "powerful AI" as a system smarter than a Nobel Prize winner across most relevant fields, capable of autonomous multi-day tasks, able to control physical tools through digital interfaces, and runnable in millions of instances at 10–100x human speed Amodei (2026a). He estimates this capability could be as little as 1–2 years away, though acknowledges significant uncertainty Amodei (2026a). Current AI models are already writing the majority of code at some frontier companies, making progress on unsolved mathematical problems, and demonstrating the ability to complete tasks requiring several hours of skilled human work Amodei (2026a).

The critical dynamic is the feedback loop: current AI models are substantially accelerating the development of the next generation, creating a recursive improvement cycle Amodei (2026a). Amodei writes: "Because AI is now writing much of the code at Anthropic, it is already substantially accelerating the rate of our progress in building the next generation of AI systems" Amodei (2026a).

5.2. Why This Time Is Different

The standard response to concerns about technological unemployment, the "lump of labor fallacy" objection, holds that new technologies always create more jobs than they destroy Autor (2015). This objection deserves serious engagement, because it has been correct for every previous technological revolution. However, Amodei identifies four structural reasons to believe AI represents a genuine discontinuity Amodei (2026a), which align with findings from labor economics research Acemoglu and Restrepo (2019); Aluko et al. (2026):

Speed. AI capabilities are advancing faster than any previous technology. In two years, AI models went from barely completing a single line of code to writing virtually all code for some engineers Amodei (2026a). Human labor markets, which adjust over years and decades, cannot adapt to this pace.

Cognitive breadth. Unlike mechanized farming or computers, AI is a general substitute for human cognition across essentially all domains Acemoglu and Restrepo (2019); Amodei (2026a). When every cognitive profession is disrupted simultaneously, there are no adjacent fields to which displaced workers can easily transition.

Ability stratification. AI is advancing from the bottom of the cognitive ability ladder upward, displacing people by cognitive ability level rather than by profession Amodei (2026a). Research on skill-biased technological change suggests this dynamic increases wage inequality Acemoglu (2002), and the AI version is likely far more extreme.

Gap-filling. Previous technologies had persistent gaps that humans could fill. AI is a rapidly adapting technology that identifies and closes its own gaps with each training cycle Amodei (2026a).

Assessment of H3: The evidence, while still partially prospective, supports the hypothesis that AI constitutes a qualitative discontinuity. The combination of cognitive breadth, speed of advancement, and recursive self-improvement distinguishes AI from all previous technological disruptions. The primary uncertainty is timeline, not direction.

5.3. The Demand Paradox

We hypothesize that the most dangerous economic implication of artificial general intelligence is what we term the "demand paradox." If AI drives the marginal cost of labor to near zero, it creates a system with potentially infinite supply capacity but collapsing effective demand, because the

 doi:10.20944/preprints202604.0074.v1

population that would normally constitute the consumer base no longer has income from employment. This dynamic has been analyzed by multiple commentators Chow (2026); Shap (2026).

The standard assumption is that this paradox forces redistribution: the ultra-rich need consumers, so they will be compelled to support transfer mechanisms to maintain demand. In the author's assessment, this assumption may be incorrect. If a sufficiently wealthy elite class (perhaps 5–6 percent of the population) increases per-capita spending by 5–10x on AI services, premium healthcare, luxury goods, bespoke experiences, and biological enhancement, this could plausibly generate sufficient economic activity to sustain the system without mass participation. Data from the Bureau of Labor Statistics indicates that the top 10 percent of US earners already account for approximately 50 percent of consumer spending Bureau of Labor Statistics (2026). Compressing that into a smaller group spending at much higher levels could, in principle, eliminate the macroeconomic need for mass consumption.

Amodei himself identifies the political consequence: "Democracy is ultimately backstopped by the idea that the population as a whole is necessary for the operation of the economy. If that economic leverage goes away, then the implicit social contract of democracy may stop working" Amodei (2026a).

6. The 1/5/94 Model: A Neo-Feudal Class Structure

6.1. The Structural Analogy

Prior to the Industrial Revolution, virtually all wealth derived from land, and whoever controlled land captured wealth Piketty (2014). The Enlightenment and the industrial revolution created an alternative: intellectual property, skilled labor, and entrepreneurship became viable paths to wealth independent of land ownership Piketty (2014). This was the foundation of liberal capitalism and, ultimately, of democratic governance: the idea that a person's economic value was embedded in their mind and body, not in territory, and therefore could not be easily expropriated.

Multiple analysts have observed that artificial general intelligence reverses this Enlightenment achievement Chow (2026); Shap (2026). When AI can perform any cognitive task better than any human, the irreducible factors of production become not human minds but physical infrastructure: energy to power computation, silicon to run models, and data to train them. As one analysis puts it: "An AI that can do anything a human can do would upend this liberal achievement because wealth gets tied back to land again" Chow (2026). These are the new "land," and whoever controls them controls the means of production in a sense more absolute than any feudal lord, because at least feudal lords needed peasant labor.

6.2. The Emerging Tiers

Based on current trajectories of wealth concentration, AI capability development, and labor market disruption, the model proposed here anticipates the following class structure by the mid-2030s, absent significant intervention:

Tier 1: Infrastructure Owners (approximately 1 percent). The individuals and entities who own compute infrastructure, energy supply, foundational AI models, and critical data assets. Amodei notes that Elon Musk's current fortune (approximately $700 billion) already exceeds John D. Rockefeller's peak wealth as a share of GDP, and this is before the primary economic impact of AI Amodei (2026a). He projects that AI companies could generate approximately $3 trillion in revenue per year, leading to personal fortunes "well into the trillions" Amodei (2026a). Research confirms that AI adoption has already contributed to increasing concentration of wealth among top income earners, while the bottom 50 percent has seen only marginal gains Aluko et al. (2026).

Tier 2: The Cognitive Bourgeoisie (approximately 5 percent). AI safety researchers, senior engineers who direct AI systems, compliance professionals navigating competing regulatory regimes, healthcare providers, educators, creative directors, and others whose work requires human judgment or trust that AI cannot yet replicate. This class is economically comfortable but fundamentally precarious: the AI capability frontier is constantly advancing into their domain. In economic terms, they are digital

 doi:10.20944/preprints202604.0074.v1

sharecroppers, productive and well-compensated but entirely dependent on Tier 1 for the tools and platforms on which their work depends.

Tier 3: The Subsidized Majority (approximately 94 percent). Those whose labor has been rendered economically unnecessary by AI. Under the best scenario, this class receives some form of universal basic income, public services, and abundant cheap goods. Under the worst scenario, this class subsists on whatever minimal provisions trickle down from an economy that no longer requires their participation. In either case, their defining characteristic is political irrelevance: they have neither the labor leverage nor the consumption leverage to compel the attention or accommodation of Tiers 1 and 2.

This model is stylized and is not intended as a precise distributional prediction. It is intended to illuminate structural dynamics: specifically, the simultaneous elimination of labor value and consumption value for the majority of the population, and the consequences of this elimination for the viability of democratic governance.

Assessment of H4: The hypothesis that AI produces a self-reinforcing neo-feudal class structure is supported by converging evidence from wealth concentration data Aluko et al. (2026); Oxfam (2024), AI capability trajectories Amodei (2026a), labor market analysis Acemoglu and Restrepo (2019); Center for American Progress (2025); Economic Policy Institute (2026a), and the demand paradox analysis in Section 5.3. The primary uncertainty is whether the transition occurs gradually (allowing institutional adaptation) or abruptly (overwhelming existing governance capacity).

6.3. Why the Demand Argument Does Not Save Us

The conventional economic argument that automation must lead to redistribution because elites need consumers rests on an assumption that may not survive the AI transition: that production and consumption are necessarily linked at scale. In a world where AI-driven production is extraordinarily efficient, where 6 percent of the population spending at 5–10x current levels can drive sufficient demand, and where entirely new categories of ultra-premium goods and services can be created for a small market, the macroeconomic need for mass consumer participation may simply disappear.

This does not mean the subsidized majority would necessarily starve. The same AI-driven productivity that eliminates their economic role would also make food, basic healthcare, and housing extremely cheap to provide. The likely outcome is not the starvation of the masses but their comfortable irrelevance: a Huxleyan *Brave New World* rather than an Orwellian *Nineteen Eighty-Four*. The suffering is not physical deprivation but the loss of agency, purpose, and political voice.

7. The Revolution Question: Historical Parallels and the AI Enforcement Problem

7.1. The French Revolution Parallel

The pre-revolutionary French class structure provides the closest historical parallel to the model proposed here: a First Estate (clergy, approximately 0.5 percent) and Second Estate (nobility, approximately 1.5 percent) who controlled most wealth and were tax-exempt, and a Third Estate (approximately 98 percent) that included both the productive bourgeoisie and the impoverished peasantry Lefebvre (1947). The revolution was triggered not by the poorest, who lacked organizational capacity, but by the bourgeoisie, who had enough education to understand the system, enough proximity to the aristocracy to see the obscenity of the gap, and enough organizational capability to mobilize mass anger Lefebvre (1947).

The 1/5/94 model contains the same structural instability. The cognitive bourgeoisie (Tier 2) occupies the revolutionary position: educated, aware, and organizationally capable. The question is whether they side with Tier 1 (because they are comfortable and dependent) or with Tier 3 (because they are morally outraged or feel their own position is precarious). History suggests that bourgeois revolutions occur when this class perceives its interests as fundamentally threatened and its upward path as blocked Lefebvre (1947).

 doi:10.20944/preprints202604.0074.v1

7.2. The AI Enforcement Asymmetry

However, we hypothesize that the 2030s differ from 1789 in a critical respect. Every successful revolution in history has depended on a moment when the regime's enforcers (soldiers, police, administrators) refused to carry out orders against the population. The French army's refusal to fire on citizens was decisive Lefebvre (1947). The Russian Revolution succeeded in part because soldiers joined the revolutionaries Figes (1996). The fall of the Berlin Wall became possible when East German border guards declined to enforce shoot-to-kill orders Meyer (2009).

An AI-enforced neo-feudal order would have no such vulnerability. Autonomous drone armies do not refuse orders. AI surveillance systems do not experience moral qualms. Algorithmic propaganda systems can individually tailor messaging to every citizen to pacify dissent. Amodei devotes an entire section of his essay to this danger, warning that "a swarm of millions or billions of fully automated armed drones, locally controlled by powerful AI and strategically coordinated across the world by an even more powerful AI, could be an unbeatable army, capable of both defeating any military in the world and suppressing dissent within a country" Amodei (2026a).

This represents, in the author's assessment, a structural break in the history of power relations. For the first time, a ruling class could maintain its position without any dependence on human cooperation from below. The French Revolution, the American Revolution, the labor movements of the early twentieth century, the civil rights movement, the fall of the Soviet Union: all depended on a sufficient number of human enforcers and administrators choosing, at critical moments, to side with the people against the regime Branch (1988); Figes (1996); Lefebvre (1947); Meyer (2009). If enforcement is automated, that choice point is eliminated.

Assessment of H5: The hypothesis that AI enforcement eliminates the historical revolutionary check on elite power is partially supported by current trends in autonomous weapons development Hudson Institute (2024) and AI surveillance deployment Amodei (2026b), and strongly supported by extrapolation from Amodei's capability projections Amodei (2026a). The primary counterargument is that democratic institutions and international norms could constrain deployment of these technologies before they are consolidated. However, the evidence reviewed in Section 3 suggests these constraining institutions are themselves under severe strain.

8. The Future of Government: Sovereignty Migration and the Neo-Feudal State

The preceding sections have analyzed the emerging class structure and the erosion of revolutionary checks on power. A related but distinct question concerns the future of the nation-state itself. If a small set of infrastructure owners controls more wealth and capability than most governments, what role does government play? Does it persist, transform, or become vestigial?

8.1. The Sovereignty Migration

There is growing scholarly and policy recognition that sovereignty, understood as the authority to set rules, allocate resources, and shape collective futures, is migrating from public institutions to private actors. As one analysis observes: "Sovereignty is migrating from public institutions to private actors. The danger is not that machines will rule humanity. It is that those who control them increasingly shape the conditions under which humanity governs itself" Chesterman (2026). Policy experts have noted that "big technology firms have effectively become independent, sovereign actors in the digital realms they have created" Bremmer and Suleyman (2025).

This migration is already empirically observable. AI companies command significant lobbying resources and enjoy deep integration into daily life, while deploying tools that shape surveillance, information access, and economic participation Chesterman (2026). Governments have struggled to keep pace: the EU's AI Act faces implementation strain, the US has been unable to regulate at the federal level, and China has reasserted state control only by replacing private dominance with party oversight Chesterman (2026). The rate of AI development consistently outpaces regulatory efforts, and

 doi:10.20944/preprints202604.0074.v1

firms building closed-source models inherently prevent states from monitoring and controlling them Anonymous (2025).

The structural logic is straightforward. Modern governments derive their power from three sources: the ability to tax economic activity, the monopoly on legitimate violence, and the consent of the governed. AI threatens all three. If economic activity is increasingly conducted by AI systems owned by a small number of corporations, the tax base shifts from broadly distributed labor income to narrowly concentrated capital income, which is far easier to shelter across jurisdictions Peixoto and Canuto (2025). If autonomous weapons and AI surveillance become the primary instruments of force, and these are built and maintained by private companies, the state's monopoly on violence becomes dependent on corporate cooperation. And if consent is manufactured through algorithmic propaganda rather than genuinely obtained through democratic deliberation, the legitimacy foundation of government is undermined.

8.2. Three Models for the Neo-Feudal State

We hypothesize that the 1/5/94 class structure described in Section 6, if it stabilizes, would produce one of three models of governance, depending on regional and institutional context:

Model A: Corporate subsumption. In this model, the state becomes functionally subordinate to AI infrastructure owners. Government continues to exist formally, conducting elections, passing legislation, and maintaining public services, but its capacity to act independently of corporate interests is minimal. Policy is shaped by lobbying, regulatory capture, and the structural dependence of the state on corporate tax revenue and technological capability. This model is closest to the current trajectory of the United States, where the coupling of AI datacenter investment with economic growth has already created what Amodei describes as "perverse incentives" and a reluctance of tech companies to criticize the government, with government reciprocating through "extreme anti-regulatory policies on AI" Amodei (2026a). The emerging "sovereign AI cloud" initiatives, in which governments depend on private companies to build and operate their core data infrastructure, exemplify this dynamic: what presents as state sovereignty in practice creates "a new form of unchecked power that combines state authority with corporate technology in unclear public-private partnerships" TechPolicy.Press (2025).

Model B: State-corporate fusion. In this model, the state and the infrastructure-owning class merge into a single governing entity, as in contemporary China, where the boundary between the Chinese Communist Party and major technology companies is deliberately blurred. AI capability serves the state's objectives (surveillance, social control, economic management), and the state serves the companies' objectives (market protection, labor discipline, geopolitical advantage). Amodei identifies this as the most dangerous near-term scenario, warning that the CCP "has the clearest path to the AI-enabled totalitarian nightmare" because it combines AI prowess, autocratic governance, and a high-tech surveillance state Amodei (2026a). This model has historical precedent in the mercantilist empires of the early modern period, where chartered corporations (the Dutch East India Company, the British East India Company) wielded state-like powers, including the ability to maintain armies, conduct diplomacy, and govern territories, while remaining formally subordinate to (but in practice deeply intertwined with) the crown Robins (2012).

Model C: Inter-oligarch coordination. In this model, the nation-state becomes vestigial or purely administrative, and effective governance is conducted through coordination among infrastructure owners directly. This is the model most analogous to the medieval European system, in which nominal kingdoms existed but actual power resided with feudal lords who negotiated among themselves, occasionally warring, occasionally forming alliances, and treating the monarch as a figurehead whose authority derived from their collective consent. In the AI context, a small number of individuals controlling the world's compute, energy, and model infrastructure might find it more efficient to coordinate directly (through informal agreements, shared standards, mutual non-aggression pacts regarding each other's markets and territories) than to operate through the intermediary of national governments. Modern precedent for this pattern includes the coordination among authoritarian leaders documented in Freedom House's 2026 report, which describes autocratic states "no longer just

 doi:10.20944/preprints202604.0074.v1

in ad hoc ways, working together to undermine democracy around the world" through "mechanisms of political socialization" that "normalize a shared repertoire of governance claims" Freedom House (2026); Kurlantzick (2026). The difference is that in the neo-feudal model, the coordinating actors would not be heads of state but heads of corporations, whose "territories" are defined not by geography but by infrastructure control.

8.3. The Tax Base Collapse

A critical mechanism driving sovereignty migration is the erosion of the government's ability to fund itself. Modern welfare states are funded primarily through income taxes and payroll taxes levied on employed populations. If the 1/5/94 model is broadly correct and 94 percent of the population loses labor income, the traditional tax base collapses. Research on AI and fiscal capacity confirms this dynamic: "Job displacement in labor-intensive sectors shrinks income tax bases, while the concentration of wealth among digital platform firms makes taxation more challenging. Fragile states, already facing administrative limitations, struggle to capture revenue from cross-border digital activities dominated by multinational technology companies" Anonymous (2026). Meanwhile, "the costs of digital transformation are substantial," requiring "large capital investments and specialized technical capacity" that many governments must finance through external borrowing Anonymous (2026).

The historical parallel is illuminating. In the feudal period, the monarch's power was constrained precisely by the fact that the crown could not tax the nobility effectively. The nobility controlled the land (the primary source of wealth), maintained their own armies, and consented to taxation only when it served their interests. The rise of the modern state was enabled by the shift of economic activity from land to commerce and industry, which created a tax base that the state could capture independently of the aristocracy Piketty (2014). If AI reverses this shift, returning economic power to a small class of infrastructure owners whose wealth is easily sheltered across jurisdictions, then the fiscal foundation of the modern state erodes, and with it the state's capacity to act as an independent check on concentrated private power.

Assessment of H8: The hypothesis that government becomes functionally subordinate to, fused with, or replaced by coordination among infrastructure owners is supported by the observed migration of sovereignty from public to private actors Anonymous (2025); Bremmer and Suleyman (2025); Chesterman (2026), the emerging dependence of states on corporate AI infrastructure TechPolicy.Press (2025), and the structural logic of tax base collapse under conditions of mass labor displacement Anonymous (2026). The primary counterargument is that states retain the monopoly on legitimate violence and could, in principle, expropriate private AI infrastructure. However, the analysis in Section 7.2 suggests that the monopoly on violence is itself being privatized through autonomous weapons, and the evidence in Section 3 suggests that the political will for expropriation is absent in the states most capable of carrying it out.

9. UBI, the EA Blind Spot, and the Political Economy of AI

9.1. The Incentive Problem: Why UBI Will Default to Pacification

Universal Basic Income is frequently proposed as the structural solution to AI-driven labor displacement. In principle, it could be: a UBI funded by progressive taxation on AI-derived wealth, set at a level sufficient for genuine economic participation, accompanied by universal public services, and governed by democratic institutions would constitute a genuine redistributive mechanism. The question is not whether such a UBI is possible but whether anyone with the power to implement it has the incentive to do so.

We argue that the incentive structure for redistribution is collapsing. Historically, elites have shared wealth for exactly four reasons, and AI threatens to eliminate three of them simultaneously.

The first incentive is the need for labor. When workers can strike, slow production, or refuse to cooperate, they hold leverage that compels accommodation. The entire history of labor rights, from

 doi:10.20944/preprints202604.0074.v1

the Factory Acts to the Wagner Act, is a history of this leverage being exercised Hobsbawm (1962). AI eliminates it. The second incentive is the need for consumption. Mass production requires mass purchasing power Watts (2006). As argued in Section 6.3, a small elite spending at greatly elevated levels may be sufficient to sustain aggregate demand without mass participation. The third incentive is the fear of violent revolution. The French Revolution, the New Deal, Bismarck's social insurance: every major redistributive program in history was implemented at the minimum viable level to prevent social collapse, not at the level that justice might demand Hennock (2007); Katznelson (2013); Lefebvre (1947). AI-powered autonomous enforcement, as analyzed in Section 7.2, threatens to eliminate this incentive as well.

What remains is the fourth incentive: moral and social pressure (philanthropy, noblesse oblige, cultural expectation). This force has never been sufficient, in the entire history of civilization, to produce structural redistribution on its own Reich (2018). It produces individual acts of charity (Carnegie's libraries, Gates's vaccines, Amodei's 80 percent pledge Amodei (2026a)), but it does not produce tax regimes, labor protections, or institutional frameworks that constrain the class as a whole.

The historical record on this point is unambiguous. The Roman *annona* (grain dole) gave citizens just enough bread to forestall insurrection Garnsey (1988). The British Poor Laws provided the minimum necessary to prevent mass starvation from threatening social order Polanyi (1944). Bismarck explicitly designed Germany's social insurance system to undercut the socialist movement Hennock (2007). Roosevelt's New Deal was the minimum concession required to prevent American communism; Roosevelt himself framed it to his critics on the left as a defense of the profit system Katznelson (2013). In every case, the amount redistributed was calibrated not to fairness but to the minimum necessary to maintain stability.

Applied to the AI era, in the author's assessment, the most likely version of UBI is not a genuine solution but a pacification mechanism: subsistence-level payments sufficient for food, basic shelter, and cheap entertainment, paired with AI-powered content engines providing infinite distraction, AI companions providing emotional connection (reducing the social isolation that fuels unrest), and algorithmic sentiment monitoring identifying potential dissent before it can organize. This is Aldous Huxley's *Brave New World* rendered in silicon: control not through deprivation but through comfortable dependency.

Amodei himself partially recognizes this dynamic, arguing that progressive taxation is justified and warning fellow billionaires that "if they don't support a good version, they'll inevitably get a bad version designed by a mob" Amodei (2026a). This is honest, and his personal commitment to donating 80 percent of his wealth suggests sincerity Amodei (2026a). But the warning contains its own refutation: it assumes the mob retains the capability to impose a "bad version." If the analysis in Section 7.2 is correct, if AI enforcement eliminates the revolutionary threat, then the mob has no leverage, and the billionaire class has no incentive to support even a "good version" of redistribution.

Assessment of H6: The hypothesis that UBI will default to pacification rather than genuine redistribution is strongly supported by historical evidence on the calibration of redistributive programs to minimum viable levels Garnsey (1988); Hennock (2007); Katznelson (2013); Polanyi (1944) and by the incentive analysis showing that AI eliminates three of the four historical drivers of redistribution. The primary counterargument is that democratic institutions may impose generous UBI before the AI enforcement window closes; however, the evidence in Section 3 suggests these institutions are weakening rather than strengthening.

9.2. Effective Altruism's Political Economy Blind Spot

The effective altruism community has, over the past decade, invested enormous intellectual and financial resources in the existential risk framing of AI: the question of whether AI might kill everyone MacAskill (2022). This is a legitimate concern that Amodei takes seriously Amodei (2026a) and that this paper does not dismiss. However, the near-exclusive focus on x-risk has, in the author's assessment, produced a significant blind spot: the question of who controls AI and in whose interest it operates has received comparatively little systematic attention from the EA community.

 doi:10.20944/preprints202604.0074.v1

The political economy of AI (the distribution of ownership, the structure of regulatory capture, the mechanisms by which AI-derived wealth translates into political power, the global asymmetries between nations with AI capacity and those without) is arguably a more tractable problem than technical alignment, and one whose resolution is more directly within the influence of the community's members, many of whom work at or are connected to frontier AI companies.

There is a painful irony in a movement dedicated to doing the most good possible being disproportionately staffed by employees of the very companies whose products are driving the wealth concentration the movement does not adequately analyze. The EA community's intellectual framework (expected value reasoning, cause prioritization, counterfactual impact) is exactly the right toolkit for analyzing the political economy of AI. It simply has not been applied to this problem with the same rigor it has been applied to biosecurity, pandemic preparedness, or technical alignment MacAskill (2022); Open Philanthropy (2026).

This represents both a failure and an opportunity. The same analytical methods that have made EA effective in other domains could be brought to bear on questions like: What governance structures for AI companies would minimize the risk of neo-feudal concentration? What tax and regulatory frameworks would distribute AI-derived productivity gains most efficiently? What international institutional arrangements would prevent AI from deepening the gap between the Global North and South? These are cause-area-level questions that the EA community is uniquely positioned to address.

Assessment of H7: This hypothesis reflects the author's assessment of the EA community's prioritization rather than a quantifiable empirical claim. It is supported by analysis of EA funding allocations and research output, which show disproportionate investment in technical alignment and biosecurity relative to political economy and governance MacAskill (2022); Open Philanthropy (2026). Readers embedded in the EA community are invited to contest this characterization.

10. Structural Variables: What Determines the Outcome

The preceding analysis identifies a trajectory toward neo-feudal consolidation, but this trajectory is not inevitable. Several structural variables will determine whether, and how quickly, this equilibrium is reached. This section identifies the two most significant of these variables. It does not prescribe policy responses, as the author is not positioned to do so; rather, it identifies the structural dynamics that will shape the outcome regardless of policy choices.

10.1. Energy as the Master Variable

All AI computation runs on energy. The entire neo-feudal model described in this paper (concentrated ownership of compute, autonomous enforcement, AI surveillance) depends on concentrated control of energy infrastructure. This makes energy infrastructure, in the author's assessment, the single most important structural determinant of whether AI leads to concentrated or distributed power.

The logic is straightforward. If fusion power, distributed solar, or other technologies make energy abundant, cheap, and geographically distributed, then the compute layer (and therefore AI capability) can be distributed broadly. The infrastructure bottleneck that enables concentration of power dissolves. If energy remains scarce and centralized (large nuclear plants, fossil fuel reserves controlled by states and corporations, massive datacenter complexes requiring grid-scale power), then compute remains concentrated, and the neo-feudal model holds.

The current strain. The empirical evidence shows that AI is already imposing severe costs on energy infrastructure. Goldman Sachs reports that US electricity prices rose 6.9 percent in 2025, more than double headline inflation, with data centers accounting for 40 percent of electricity demand growth Goldman Sachs (2026). Since 2020, US residential electricity prices have increased by more than 36 percent CNBC (2026b). The International Energy Agency projects that global data center electricity consumption will double to approximately 945 TWh by 2030, growing at 15 percent per year, more than four times faster than total electricity consumption growth from all other sectors International Energy Agency (2026a).

 doi:10.20944/preprints202604.0074.v1

The distributional consequences are significant. Analysis of US Energy Information Administration data shows that residential electricity use has grown the least of the three customer sectors (residential, commercial, industrial), but residential prices have increased the most Kirk (2026). Grid upgrade costs driven by data center demand are being passed to residential ratepayers: in the PJM Interconnection region, an estimated $23 billion in capacity costs attributable to data centers has been passed on to consumers, adding approximately $18 per month to household bills in some counties CNBC (2026a); Yahoo Finance (2025). Goldman Sachs notes that "the income and spending drags will likely be larger for lower-income households because electricity accounts for a greater share of their spending" Goldman Sachs (2026). This dynamic, in which the wealthiest industry in the world imposes energy costs on the least wealthy consumers, is itself a microcosm of the broader concentration dynamics described in this paper.

The inflection point hypothesis. We hypothesize that this energy strain is a transient phenomenon, and that AI will eventually resolve the very energy constraints it is currently creating. There is substantial evidence supporting this hypothesis. AI is simultaneously the primary driver of energy demand growth and the primary catalyst for energy innovation, particularly in nuclear fusion. Fusion energy industry funding has grown from $1.7 billion in 2020 to $15 billion as of September 2025 Time (2025). AI companies and their founders are the dominant investors: Sam Altman's Helion Energy aims to supply power to Microsoft data centers; Nvidia and Google have backed Commonwealth Fusion Systems (CFS); Google has also invested in TAE Technologies Fortune (2025); Time (2025). CFS's pilot machine SPARC is expected to produce first plasma in 2027, with a 400-megawatt commercial plant (ARC) slated for the early 2030s Fortune (2026). MIT modeling projects fusion generation rising from 2 TWh in 2035 to 375 TWh in 2050 World Economic Forum (2026c).

Critically, AI is accelerating fusion research itself. Google DeepMind has partnered with CFS on plasma simulation, energy optimization, and real-time control using reinforcement learning Google DeepMind (2025). AI tools are enabling advanced simulations and plasma control across both private and public fusion programs ITER (2025). China's EAST reactor has achieved plasma stability at previously impossible densities, and the US National Ignition Facility has quadrupled its fusion energy output since its initial ignition in 2022 World Economic Forum (2026c). The director of Lawrence Livermore National Laboratory, where NIF is based, stated at Davos 2026: "Historically, we've always said fusion energy is 30 years away from whatever day you ask. I think that's not true anymore" World Economic Forum (2026c).

The ownership paradox. However, and this is the critical nuance for the argument of this paper, even if AI resolves the energy constraint, the question of who owns the resulting energy infrastructure remains decisive. The IEA's 2026 report on energy innovation reveals that investors are increasingly directing capital toward AI at the expense of broader energy technology: "the share of VC funding for AI rose to almost 30% in 2025, while the share of energy shrank" International Energy Agency (2026b). Fusion investment specifically is not being driven by public utilities, distributed energy cooperatives, or developing-nation governments. It is being driven by the same technology companies that own AI infrastructure: Altman (OpenAI) owns Helion, Nvidia backs CFS, Google backs CFS and TAE, Moskovitz (early Anthropic backer) has invested in Helion Fortune (2025); Time (2025).

If fusion succeeds under this ownership structure, it does not democratize energy. It further consolidates the infrastructure stack. The AI owner class would control both compute and energy, deepening rather than resolving the concentration dynamics described in Section 6. The relationship between AI and energy is therefore paradoxical: AI is simultaneously straining energy infrastructure in the short term and accelerating energy innovation in the medium term, but the ownership structure of this emerging energy infrastructure mirrors and reinforces the neo-feudal dynamics this paper has identified. Whether energy ultimately distributes power or concentrates it further depends not on the technology itself but on the institutional and ownership arrangements under which it is deployed.

 doi:10.20944/preprints202604.0074.v1

10.2. The Time-Limited Leverage of the Cognitive Bourgeoisie

The cognitive bourgeoisie (the approximately 5 percent who build, maintain, and operate AI systems) holds a structural leverage that the subsidized majority does not. They are the people without whom the systems do not function. In the current moment, before AI systems can fully self-maintain and self-improve, this class has genuine bargaining power.

This leverage is time-limited. As AI systems become capable of maintaining and improving themselves (Amodei suggests this recursive capability may be only 1–2 years away Amodei (2026a)), the technical dependency on human engineers will diminish. The window during which this class retains structural leverage is therefore measured in years, not decades.

What makes this variable analytically significant is that the cognitive bourgeoisie occupies the hinge position in the class model proposed in Section 6. If this class uses its temporary leverage to demand structural concessions before its leverage expires, the trajectory toward neo-feudal consolidation could be altered. If it accepts comfortable complicity, the trajectory continues unimpeded. The historical analogy is the professional and managerial class of the mid-twentieth century, which used its position between capital and labor to support institutional frameworks (progressive taxation, labor protections, social insurance) that moderated inequality for several decades Hennock (2007); Katznelson (2013). Whether the AI-era equivalent of this class makes a similar choice, or instead aligns with the infrastructure-owning class, is one of the open questions this paper cannot answer but identifies as structurally decisive.

11. Conclusion: The Adolescence of Civilization

Amodei frames the current moment as "technological adolescence," a rite of passage that will test whether humanity has the maturity to wield almost unimaginable power Amodei (2026a). This framing is apt but incomplete. The test is not merely whether we can build safe AI, but whether AI will be captured by a class structure that permanently forecloses the possibility of democratic self-governance.

This paper has argued that the convergence of geopolitical fragmentation and AI-driven wealth concentration is not two separate crises but one: the institutions that might have governed the AI transition are collapsing precisely when they are most needed, and the populist-authoritarian movements that are accelerating their collapse are funded and amplified by the same technology-capital nexus that stands to benefit from ungoverned AI deployment.

The seven hypotheses advanced in Section 2 describe a trajectory, not a certainty. The empirical evidence reviewed in this paper supports the following summary assessment:

The post-WWII order is fragmenting irreversibly (H1). Populist-authoritarian movements share a common playbook that systematically benefits elites while claiming to serve the working class (H2). AI constitutes a qualitative discontinuity in technological disruption (H3). The convergence of fragmentation and AI produces a self-reinforcing neo-feudal class structure (H4). AI enforcement mechanisms may eliminate the historical revolutionary check on elite power (H5). UBI, absent structural changes to political power, will default to a pacification mechanism (H6). The effective altruism community has insufficiently engaged with the political economy of AI (H7). And the nation-state, as currently constituted, faces functional subordination to or absorption by AI infrastructure owners as its fiscal, coercive, and legitimacy foundations erode (H8).

If this analysis is broadly correct, the structural variables identified in Section 10 (the distribution of energy infrastructure and the choices of the cognitive bourgeoisie) will determine whether the neo-feudal trajectory is realized or averted. Both variables are currently trending in the direction of consolidation rather than distribution.

The historical precedent is not encouraging. Humanity has generally required catastrophic violence (world wars, revolutions, societal collapse) before it has been willing to build new cooperative institutions Carnegie Endowment for International Peace (2024); Lefebvre (1947); Wikipedia and Coppola (2023). In the nuclear age, and soon in the AI age, the cost of that learning process may

 doi:10.20944/preprints202604.0074.v1

be unsurvivable. But the precedent is also not determinative. The Bretton Woods monetary system collapsed in 1971 without a war Wikipedia and Coppola (2023). The Berlin Wall fell without an invasion Meyer (2009). South Africa transitioned from apartheid through negotiation Sparks (1996). Structural change without catastrophic violence is rare, but it is not impossible.

This paper does not propose solutions. It proposes a framework for understanding the problem: the simultaneous collapse of the institutional order that might govern AI, and the emergence of an AI-driven economic structure that may render democratic governance structurally inviable. Whether this problem is addressed, and how, is a question for policymakers, institutions, and the broader public.

As the Finnish president observed at Davos in January 2026: "The world order is changing much like it did after 1918, 1945, and 1989" World Economic Forum (2026a). The alternative paths are clear. This paper has attempted to describe where one of those paths leads.

Priyanshu Jain is a researcher at 73ai.org focused on AI agent safety and evaluation infrastructure.

This paper is a preprint and has not been peer-reviewed. The author welcomes feedback, critique, and counterarguments.

References

2026. On identity-based political allegiance and information ecosystems. Combined reference. See Duina Duina (2022) on how populist supporters feel recognized by leaders who transgress public sphere norms; and WEF World Economic Forum (2026b) on information fragmentation.

Acemoglu, Daron. 2002. Technical change, inequality, and the labor market. *Journal of Economic Literature 40*(1), 7–72. On skill-biased technological change.

Acemoglu, Daron and Pascual Restrepo. 2019. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives 33*(2), 3–30.

Aluko, Samuel et al. 2026. AI and technology's influence on economic inequality: A study of wealth distribution in the U.S. *World Journal of Advanced Research and Reviews 29*(3).

Amodei, Dario. 2024, October. Machines of loving grace: How AI could transform the world for the better. darioamodei.com.

Amodei, Dario. 2026a, January. The adolescence of technology: Confronting and overcoming the risks of powerful AI. darioamodei.com.

Amodei, Dario. 2026b. AI surveillance deployment. On AI surveillance deployment, including in China's Uyghur repression. See Amodei (2026a), Section 3.

Anonymous. 2025. The silicon sovereignty paradox: Navigating fluid state-corporate power. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES)*. On the rate of AI development outpacing regulation and closed-source models preventing state oversight.

Anonymous. 2026, March. Artificial intelligence, economic inequality, and the financial hurdles to sustainable peace. *ScienceDirect*. On AI deepening inequality through labor market polarization, wealth concentration, and erosion of fiscal capacity.

Autor, David. 2015. Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives 29*(3), 3–30. On the lump of labor fallacy and historical labor market adjustment.

Ben-Ghiat, Ruth. 2020. *Strongmen: Mussolini to the Present*. W.W. Norton. See also: R. Ben-Ghiat, "Trump's Authoritarian Innovations," *Project Syndicate*, December 2025.

Branch, Taylor. 1988. *Parting the Waters: America in the King Years 1954–1963*. Simon and Schuster. On the civil rights movement and the role of institutional pressure.

Bremmer, Ian and Mustafa Suleyman. 2025. AI, global governance, and digital sovereignty. Convergence Analysis. On the claim that "big technology firms have effectively become independent, sovereign actors in the digital realms they have created".

Bureau of Labor Statistics. 2026. Consumer expenditure survey. Bureau of Labor Statistics. Various years. Top-decile consumption share estimates drawn from Federal Reserve distributional financial accounts.

Carnegie Endowment for International Peace. 2024, October. What is Bretton Woods? the contested pasts and potential futures of international economic order. Technical report, Carnegie Endowment for International Peace. See also: "Reform or Realignment? The Geopolitical Lessons of Bretton Woods," August 2025.

 doi:10.20944/preprints202604.0074.v1

Center for American Progress. 2025, December. Year 1 of the second Trump administration made the working class weaker. Center for American Progress.

Chesterman, Simon. 2026, March. AI is moving power from governments to tech companies. Rest of World.

Chow, C. 2026, February. AI feudalism. Scientific Clearing House.

CNBC. 2026a, January. Bernie Sanders and ron DeSantis speak out against data center boom. CNBC. Citing Monitoring Analytics data on $23 billion in PJM capacity costs attributable to data centers.

CNBC. 2026b, March. Who is really footing the AI energy bill? inside the debate about data center electricity costs. CNBC. Citing US Energy Information Administration data showing residential prices rising from 12.76 to 17.44 cents per kilowatt-hour between 2020 and February 2026.

Council on Foreign Relations. 2026. Conflicts to watch in 2026: Preventive priorities survey results. Technical report, Council on Foreign Relations.

Duina, Francesco. 2022. Why populist leaders succeed. *International Political Science Review*.

Economic Policy Institute. 2026a, January. 47 ways Trump has made life less affordable in his first year. Economic Policy Institute.

Economic Policy Institute. 2026b, February. The Trump administration's macroeconomic agenda harms affordability and raises inequality. Economic Policy Institute.

Eurasia Group. 2026, January. The top risks of 2026. Technical report, Eurasia Group.

Eurasia Review. 2026, March. World war III: Warning signs before it's too late. Eurasia Review. On the expiration of New START.

Factually. 2025, November. Have prominent historians validated or debunked Trump-Nazi comparisons. Factually. Multiple historians have drawn parallels between Trump's rhetoric and historical authoritarian movements while cautioning against false equivalence.

Figes, Orlando. 1996. *A People's Tragedy: The Russian Revolution 1891–1924*. Jonathan Cape. Penguin paperback, 1998. On the Russian Revolution and the role of military defection.

Fortune. 2025, October. Nuclear fusion, the 'holy grail' of power, was always 30 years away. *Fortune*. On CFS, Helion, and AI-industry fusion investments.

Fortune. 2026, January. Fusion power nearly ready for prime time as Commonwealth builds first pilot. *Fortune*. On SPARC completion timeline and ARC commercial plant.

Freedom House. 2026, March. Freedom in the world 2026: The growing shadow of autocracy. Technical report, Freedom House.

Funke, Manuel, Moritz Schularick, and Christoph Trebesch. 2023. Populist leaders and the economy. *American Economic Review*. See also: "Populism's Broken Economic Promises," *Foreign Affairs*, December 2025.

Garnsey, Peter. 1988. *Famine and Food Supply in the Graeco-Roman World*. Cambridge University Press. On the Roman grain dole.

Goldman Sachs. 2026, February. Electricity prices will keep rising on AI data center demand. CNBC. Goldman Sachs research note on electricity prices and data center demand.

Google DeepMind. 2025, October. Bringing AI to the next generation of fusion energy. Google DeepMind Blog. On research partnership with CFS for plasma simulation and reinforcement learning control.

Hennock, E. P. 2007. *The Origin of the Welfare State in England and Germany, 1850–1914*. Cambridge University Press. On Bismarck's social insurance as anti-socialist strategy.

Hobsbawm, Eric. 1962. *The Age of Revolution: 1789–1848*. Weidenfeld and Nicolson. On the history of labor rights and the role of strike action.

Hudson Institute. 2024. Impact of drones on the battlefield: Lessons from the Russian-Ukraine war. Technical report, Hudson Institute. On developments in autonomous weapons and drone warfare.

International Energy Agency. 2026a. Energy demand from AI. Technical report, International Energy Agency. In *Energy and AI*, 2026. Projecting global data center consumption doubling to approximately 945 TWh by 2030.

International Energy Agency. 2026b. The state of energy innovation 2026. Technical report, International Energy Agency. Finding that VC funding share for AI rose to almost 30% in 2025 while energy's share shrank. Reported in: OilPrice.com and IndexBox, March 2026.

ITER. 2025, July. AI ignites innovation in fusion. ITER. On AI tools accelerating fusion research across public and private programs.

JP Morgan Chase. 2026. World rewired: Navigating a multi-speed, multipolar order. Technical report, JP Morgan Chase.

Katznelson, Ira. 2013. *Fear Itself: The New Deal and the Origins of Our Time*. Liveright. On Roosevelt's New Deal as minimum-viable redistribution.

 doi:10.20944/preprints202604.0074.v1

Kirk, Karin. 2026, January. Home electricity bills are skyrocketing. for data centers, not so much. Yale Climate Connections. Analysis of US EIA data showing residential prices rising fastest despite lowest demand growth.

Kurlantzick, Joshua. 2026, March. As democracy falters worldwide, authoritarians are winning. Council on Foreign Relations.

Lefebvre, Georges. 1947. *The Coming of the French Revolution*. Princeton University Press. On the French Revolution and the role of the bourgeoisie.

MacAskill, William. 2022. *What We Owe the Future*. Basic Books. On EA prioritization of existential risk.

Meyer, Michael. 2009. *The Year that Changed the World*. Scribner. On the fall of the Berlin Wall and the role of border guard decisions.

Mo-OrgCip. 2026, February. Democracy and authoritarianism in 2026: Global trends and the retreat of political freedoms. Mo-OrgCip. Democratic backsliding as incremental institutional change is documented across multiple sources. See also V-Dem reports.

Mohan, C. Raja. 2026, February. The multipolar delusion. *Foreign Affairs*.

Munich Security Conference. 2025, February. Munich security report 2025: Multipolarization. Technical report, Munich Security Conference.

NBC News. 2026, March. America's wealth gap is growing, and Trump's policies are making it bigger. NBC News.

Open Philanthropy. 2026. Open philanthropy grant database. openphilanthropy.org. On EA funding allocations, various years.

Oxfam. 2024. Inequality inc. Technical report, Oxfam. Finding: the wealthiest 1 percent captured nearly two-thirds of new global wealth created since the pandemic. See also: "The Neo-Feudal System and Modern Power," 4how2.com, August 2025.

Oxfam. 2025, April. How is the Trump administration deepening inequality? Oxfam.

Peixoto, Tiago C. and Otaviano Canuto. 2025. AI and the future of government: Unexpected effects and critical challenges. Center for Macroeconomics and Development. On AI's impact on government fiscal capacity: transfer pricing complexity, reduced labor income tax revenue, and increased demand for social welfare.

Pew Research Center. 2026, January. Confidence in Trump dips in 2026, and fewer now say they support his policies and plans. Technical report, Pew Research Center.

Piketty, Thomas. 2014. *Capital in the Twenty-First Century*. Harvard University Press.

Polanyi, Karl. 1944. *The Great Transformation*. Farrar and Rinehart. On the British Poor Laws.

Reich, Rob. 2018. *Just Giving: Why Philanthropy Is Failing Democracy and How It Can Do Better*. Princeton University Press. On the insufficiency of philanthropy as a mechanism for structural redistribution.

Robins, Nick. 2012. *The Corporation That Changed the World: How the East India Company Shaped the Modern Multinational*. Pluto Press. On the East India Company as a precedent for corporate sovereignty.

Shap, D. 2026, January. Techno-feudalism and the neo-bourgeoisie. Substack.

Silver, Nate. 2026, March. Trump approval rating. Silver Bulletin.

Sparks, Allister. 1996. *Tomorrow Is Another Country: The Inside Story of South Africa's Road to Change*. University of Chicago Press. On the South African transition.

Stimson Center. 2026, February. Top ten global risks for 2026. Technical report, Stimson Center.

TechPolicy.Press. 2025, September. How national AI clouds undermine democracy. TechPolicy.Press. On sovereign AI cloud initiatives creating "a new form of unchecked power" combining state authority with corporate technology.

Time. 2025, October. Why the AI industry is betting on fusion energy. *Time*. Documenting fusion industry total funding growth from $1.7 billion (2020) to $15 billion (September 2025) and AI company investments.

V-Dem Institute. 2026, March. Democracy report 2026: Unraveling the democratic era? Technical report, University of Gothenburg.

Watts, Steven. 2006. *The People's Tycoon: Henry Ford and the American Century*. Vintage. On Henry Ford's wage strategy and the demand-side logic of mass production.

Wikipedia and Frances Coppola. 2023. Bretton woods system. Wikipedia; American Express Trends and Insights. On the Bretton Woods system, its history, and collapse. See: "Bretton Woods system," Wikipedia (accessed March 2026); and F. Coppola, "The Rise and Fall of the Bretton Woods Fixed Exchange Rate System," American Express Trends and Insights, 2023.

World Economic Forum. 2026a, January. Davos 2026: How middle powers are reading the global moment. World Economic Forum.

 doi:10.20944/preprints202604.0074.v1

World Economic Forum. 2026b, January. Global risks report 2026. Technical report, World Economic Forum.

World Economic Forum. 2026c, February. Nuclear fusion in the headlines. World Economic Forum. Citing MIT modeling of fusion deployment trajectories, NIF progress, and Budil's statement on fusion timelines.

Yahoo Finance. 2025, December. The AI data center boom is raising power costs nationwide. Yahoo Finance. Citing Institute for Energy Economics and Financial Analysis data on household bill increases.