Article

Not peer-reviewed version

A Framework for Understanding the Convergence of Geopolitical Fragmentation and AI-Driven Economic Concentration

Priyanshu Jain *

Posted Date: 7 April 2026

doi: 10.20944/preprints202604.0074.v2

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.v2

Article

A Framework for Understanding the Convergence of Geopolitical Fragmentation and AI-Driven Economic Concentration

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, and examines the future of the nation-state under conditions where AI infrastructure owners command more wealth and capability than most governments.

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

The post-World War II international order is fracturing. Twenty consecutive years of democratic decline, the rise of populist-authoritarian movements across the political spectrum, and the systematic hollowing of multilateral institutions have produced not a transition to a new order but a prolonged interregnum of fragmentation and competing blocs. Simultaneously, artificial intelligence is advancing toward a capability threshold that could eliminate the economic necessity of most human labor within a decade, concentrating productive capacity in the hands of those who own compute infrastructure, energy, and foundational models. This paper argues that these are not two separate crises but one.

The convergence of geopolitical fragmentation and AI-driven economic concentration produces a structural risk that extends beyond the technical concerns most commonly associated with AI safety (autonomy, misuse, bioweapons) to a question of political economy that has received insufficient attention: 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

 doi:10.20944/preprints202604.0074.v2

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 Dario Amodei, CEO of Anthropic, describes as "a country of geniuses in a datacenter" [1], 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.

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 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 [2], the V-Dem Institute's annual Democracy Reports [3], and the Economist Intelligence Unit's Democracy Index.

(b) Institutional risk assessments: the Stimson Center's global risk analyses [4], the Eurasia Group's annual top-risks reports [5], the Munich Security Conference reports [6], the World Economic Forum's Global Risks Reports [7], and the Council on Foreign Relations' Preventive Priorities Surveys [8].

(c) Economic research: data from the Economic Policy Institute [9,10], the Center for American Progress [11], Oxfam [12], the OECD, and peer-reviewed economic scholarship on wealth inequality and AI labor displacement [13,14].

 doi:10.20944/preprints202604.0074.v2

(d) Primary statements and publications from AI industry leaders, notably Amodei's *The Adolescence of Technology* [1] and *Machines of Loving Grace* [15], 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 [16], democratic backsliding [17], and authoritarian governance [18].

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 [2]. 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 [3]. 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 [3].

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 [3]. 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" [4]. 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" [19].

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.

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 [16]. Yet they managed to hold power by polarizing the electorate and manipulating the political system [16].

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 [18,20,21]:

 doi:10.20944/preprints202604.0074.v2

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 [20].

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" [18].

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" [7].

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" [17].

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 [11]. 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" [9]. NBC News reports that the wealth divide under the current administration is "the widest it has been in at least a generation" [22].

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 [23], but his support among Republican voters remains above 73 percent [24]. This illustrates the decoupling of policy outcomes from political loyalty that characterizes mature populist movements, a phenomenon explained by research on identity-based political allegiance, information ecosystem fragmentation, and the sunk-cost psychology of partisan commitment [25].

The Eurasia Group's 2026 top-risks assessment identifies the United States itself as "the principal source of global risk in 2026" [5].

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.

 doi:10.20944/preprints202604.0074.v2

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 [26]. 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" [27].

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 [4,8], and the expiration of the New START treaty in February 2026 has removed the last binding limits on US-Russian nuclear arsenals [28]. What is more likely, and arguably already underway, is what the Stimson Center calls a "protracted interregnum" [4].

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" [29]. 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" [6].

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 [30]. 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" [30].

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" [31], while its manufacturing model, which depended on cheap Russian energy and open Chinese markets, has lost both pillars.

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 Cognitive Discontinuity

5.1. The Capability Trajectory

The threshold that matters for the argument of this paper is what might be called "powerful AI": 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 [1]. This capability could be as little as 1–2 years away, though significant uncertainty remains [1]. 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 [1].

 doi:10.20944/preprints202604.0074.v2

The critical dynamic is the feedback loop: current AI models are substantially accelerating the development of the next generation, creating a recursive improvement cycle. As Amodei notes in *The Adolescence of Technology*: "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" [1].

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 [32]. This objection deserves serious engagement, because it has been correct for every previous technological revolution. However, there are four structural reasons to believe AI represents a genuine discontinuity [1,13,14]:

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 [1]. 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 [1,13]. 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 [1]. Research on skill-biased technological change suggests this dynamic increases wage inequality [33], 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 [1].

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 population that would normally constitute the consumer base no longer has income from employment. This dynamic has been analyzed by multiple commentators [34,35].

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 [36]. Compressing that into a smaller group spending at much higher levels could, in principle, eliminate the macroeconomic need for mass consumption.

The political consequence is stark. As Amodei warns: "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" [1].

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 [37]. The Enlightenment and the industrial revolution created an alternative:

 doi:10.20944/preprints202604.0074.v2

intellectual property, skilled labor, and entrepreneurship became viable paths to wealth independent of land ownership [37]. 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 [34,35]. 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" [35]. 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. Elon Musk's current fortune (approximately $839 billion [38]) already exceeds John D. Rockefeller's peak wealth as a share of GDP, and this is before the primary economic impact of AI. Industry projections suggest AI companies could generate approximately $3 trillion in revenue per year, leading to personal fortunes "well into the trillions" [1]. 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 [14].

Tier 2: The Cognitive Bourgeoisie (approximately 5 percent). This tier is defined not by job function but by *agency*: the intrinsic, obsessive drive to originate directions, create, lead, and push boundaries. Research on proactive personality [39], individual performance distributions [40], and scientific productivity [41] converges on a consistent finding: a small minority of individuals (in the low single-digit percentages) accounts for a vastly disproportionate share of innovative and initiative-driven output. O'Boyle and Aguinis [40], analyzing 633,263 individuals across five industries, found that the top 5 percent produced 26 percent of total output, with performance following a Paretian rather than normal distribution. Price's Law [41] holds that the square root of contributors produces half of all contributions. These are the researchers who identify which questions matter, the founders who see what should exist and will it into being, the leaders whose relentless drive others organize around. Their ranks include AI safety researchers, senior engineers, compliance architects, healthcare innovators, educators, and creative directors, but the defining trait is not the role but the agency with which it is inhabited. This class will be economically secure and highly compensated: the infrastructure owners need their capacity for direction-setting and novel problem-solving, much as a king needs a wise council. They are, however, structurally dependent on Tier 1 for the compute, energy, and data infrastructure on which all their work rests. The implications of this dependency for collective action are examined in Section 10.

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.

 doi:10.20944/preprints202604.0074.v2

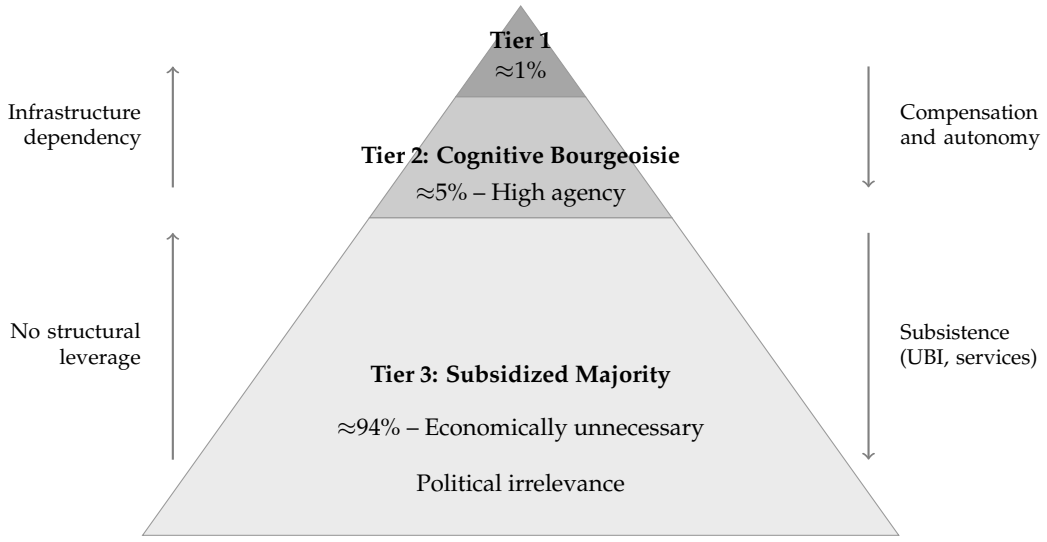

Figure 1. The 1/5/94 class model. Unlike pre-AI class structures, the base does not support the apex: the infrastructure-owning class does not depend on the labor or consumption of the subsidized majority. Arrows indicate the direction of structural dependency (left) and resource flows (right). The cognitive bourgeoisie receives substantial compensation but remains dependent on Tier 1 infrastructure.

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 [14,42], AI capability trajectories [1], labor market analysis [9,11,13], 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 [43]. The revolution was triggered not by the poorest, who lacked organizational capacity, but by the

 doi:10.20944/preprints202604.0074.v2

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 [43].

The 1/5/94 model contains a superficially similar structure. The cognitive bourgeoisie (Tier 2) occupies the position analogous to the revolutionary bourgeoisie: educated, aware, and organizationally capable. However, as Section 10 demonstrates, the analogy breaks down on the question of incentives. The pre-revolutionary French bourgeoisie faced acute status incongruence: economic power without political power or social recognition [44,45]. The cognitive bourgeoisie of the AI era faces no such mismatch. They are economically powerful, socially respected, and professionally courted. History suggests that bourgeois revolutions occur when this class perceives its interests as fundamentally threatened and its upward path as blocked [43]. Neither condition applies here.

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 [43]. The Russian Revolution succeeded in part because soldiers joined the revolutionaries [46]. The fall of the Berlin Wall became possible when East German border guards declined to enforce shoot-to-kill orders [47].

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. The technology for this is not speculative: autonomous drone swarms capable of coordinated strikes with minimal human oversight are already under active development by military contractors and state actors [48]. A swarm of millions of fully automated armed drones, locally controlled by powerful AI and strategically coordinated by an even more powerful system, could constitute an unbeatable enforcement mechanism, capable of both defeating conventional military forces and suppressing domestic dissent without requiring a single human soldier to choose between orders and conscience.

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 [43,46,47,49]. 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 [48] and AI surveillance deployment [50], and strongly supported by extrapolation from current AI capability trajectories [1]. 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

 doi:10.20944/preprints202604.0074.v2

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" [51]. Policy experts have noted that "big technology firms have effectively become independent, sovereign actors in the digital realms they have created" [52].

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 [51]. 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 [51]. The rate of AI development consistently outpaces regulatory efforts, and firms building closed-source models inherently prevent states from monitoring and controlling them [53].

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 [54]. 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 "perverse incentives" and a reluctance of tech companies to criticize the government, with government reciprocating through "extreme anti-regulatory policies on AI" [1]. 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" [55].

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). This may represent the most dangerous near-term scenario. 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 [1]. 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 [56].

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

 doi:10.20944/preprints202604.0074.v2

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 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" [2,19]. 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" [57]. Meanwhile, "the costs of digital transformation are substantial," requiring "large capital investments and specialized technical capacity" that many governments must finance through external borrowing [57].

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 [37]. 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 H7: 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 [51–53], the emerging dependence of states on corporate AI infrastructure [55], and the structural logic of tax base collapse under conditions of mass labor displacement [57]. 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 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

 doi:10.20944/preprints202604.0074.v2

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 the Factory Acts to the Wagner Act, is a history of this leverage being exercised [58]. AI eliminates it. The second incentive is the need for consumption. Mass production requires mass purchasing power [59]. 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 [43,60,61]. 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 [62]. It produces individual acts of charity (Carnegie's libraries, Gates's vaccines, Amodei's 80 percent pledge [1]), 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 [63]. The British Poor Laws provided the minimum necessary to prevent mass starvation from threatening social order [64]. Bismarck explicitly designed Germany's social insurance system to undercut the socialist movement [60]. 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 [61]. 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, in *The Adolescence of Technology*, 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" [1]. This is honest, and his personal commitment to donating 80 percent of his wealth suggests sincerity [1]. 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 [60,61,63,64] 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. Situating the Convergence Thesis in Existing Literature

The questions this paper addresses (how AI reshapes power, inequality, governance, and world order) have generated substantial literatures across multiple disciplines. This subsection maps the principal research streams, identifies their contributions, and clarifies the gap the convergence thesis

 doi:10.20944/preprints202604.0074.v2

occupies. The purpose is not comprehensive review but positioning: explaining what this paper draws from, and where its analysis departs from, existing work.

The economics of AI and inequality. Mainstream economics has produced rigorous models of AI's impact on labor and wealth. Acemoglu and Restrepo demonstrate that recent AI development is biased toward automation rather than new task creation, driving wage stagnation and declining labor share [13]. Acemoglu's "The Simple Macroeconomics of AI" finds that AI's benefits accrue disproportionately to capital owners [65]. Korinek and Stiglitz model AI's implications for income distribution and propose tax-based interventions [66]. Acemoglu and Johnson's *Power and Progress* provides the broadest historical argument: technology has consistently benefited elites unless countervailing institutions redirect innovation [67]. This literature identifies the mechanisms of concentration clearly. Its limitation, for the present analysis, is that it generally assumes institutional capacity to redirect AI remains intact: that governments can tax, regulate, and redistribute if they choose wisely. The convergence thesis challenges this assumption: the institutions that might redirect AI are themselves fragmenting (H1), and the populist movements accelerating that fragmentation systematically benefit the very elites who profit from ungoverned AI deployment (H2).

Techno-feudalism and AI political economy. A growing literature analyzes AI-driven concentration not as a market failure but as a structural transformation of the economic order. Varoufakis argues that digital platforms have replaced market-based profit with feudal-style rent extraction [68]. Durand traces how information networks push the digital economy toward feudal logics of rent and personal domination [69]. Srnicek documents how platforms accumulate power through data, network effects, and scalability [70]. Mazzucato warns of "digital feudalism" absent active state intervention [71]. Zuboff identifies surveillance capitalism as an "epistemic counterrevolution" in which asymmetries of knowledge become asymmetries of power [72]. Kasy argues that AI systems require democratic control of algorithms, data, and infrastructure [73]. This literature provides theoretical foundations for the class dynamics described in Section 6. Its limitation is that it treats the international institutional order as background rather than as a co-evolving variable. The convergence thesis argues that AI concentration and institutional fragmentation are mutually reinforcing: concentration weakens institutions, and weakened institutions cannot constrain concentration.

AI and geopolitics. A substantial literature examines how AI reshapes interstate competition. Miller's *Chip War* documents semiconductor control as geopolitical leverage [74]. Scharre identifies data, compute, talent, and institutions as AI's four battlegrounds [75]. Bremmer and Suleyman argue that AI's creators are themselves geopolitical actors whose sovereignty "entrenches the emerging technopolar order" [52]. Bradford maps competing regulatory models and asks whether liberal democracy remains viable in today's technology landscape [76]. Suleyman frames AI containment as a threat to the nation-state itself [77]. This literature captures the geopolitical dimensions effectively. What it largely does not do is model the domestic class outcomes of geopolitical AI competition, or consider the possibility that the primary beneficiaries may be corporations rather than states, a possibility Section 8 examines.

Critical AI studies. A rich body of work documents AI's present-day harms and power concentration. Crawford maps AI as a technology of extraction across labor, data, and materials [78]. The AI Now Institute diagnoses concentration at every layer of the AI stack and calls for structural reforms [79]. Couldry and Mejias theorize "data colonialism" as a continuation of historical colonial extraction [80]. Noble documents how search algorithms reinforce racial inequality [81]; Eubanks shows how automated decision systems punish poverty [82]. This literature provides essential documentation of the power dynamics this paper analyzes. Its primary orientation is diagnostic rather than projective: it documents what AI is doing now rather than modeling where the trajectory leads under conditions of simultaneous institutional collapse.

AI governance. The past three years have produced a proliferation of governance frameworks: the EU AI Act (2024), the Council of Europe's Framework Convention on AI (the first binding international AI treaty, 2024), the G7 Hiroshima Process, the UN International Scientific Panel on AI

 doi:10.20944/preprints202604.0074.v2

(2026), and a growing network of national AI Safety Institutes. Within the effective altruism ecosystem, GovAI has produced influential work on compute governance as a uniquely tractable regulatory lever [83], and Open Philanthropy committed approximately $50 million to technical AI safety in 2024 alone [84]. Korinek and Vipra's "Concentrating Intelligence" argues that AI markets tend toward natural monopoly and may require public utility regulation [85]. Think tanks across the spectrum (Brookings, RAND, Carnegie, CNAS, Stanford HAI) have produced detailed analyses of regulatory and coordination options. This work is valuable. Its shared limitation is the assumption that governance is achievable within existing or modestly reformed institutional frameworks, an assumption the evidence in Sections 3 and 4 calls into question.

The gap. Each of these streams captures part of the dynamics this paper analyzes. The economists model AI's distributional impact but assume functioning institutions. The political economists theorize platform power but treat geopolitical fragmentation as background. The geopolitics scholars map interstate competition but do not model class outcomes. The critical scholars document present-day harms but are less focused on structural projection. The governance literature assumes the continued relevance of the institutions it seeks to strengthen. No existing work, to the author's knowledge, synthesizes all three of the following claims: (1) that geopolitical institutional fragmentation and AI-driven economic concentration are co-evolving, each reinforcing the other; (2) that this convergence produces a specific class structure that is self-reinforcing; and (3) that AI-powered enforcement mechanisms make this structure uniquely resistant to challenge from below, eliminating the revolutionary check that has historically forced redistribution. The convergence thesis advanced in this paper occupies this intersection, drawing on all of these streams while arguing that the whole is substantially more dangerous than the sum of its parts.

10. Why the Trajectory is Self-Reinforcing

The preceding analysis identifies a trajectory toward neo-feudal consolidation. A natural question is whether any structural variable might interrupt or reverse this trajectory. This section examines the two most commonly identified candidates, energy infrastructure and the agency of the cognitive bourgeoisie, and demonstrates that both are trending toward consolidation rather than resistance. The dynamics described are not merely unchecked; they are self-reinforcing, with positive feedback loops that actively deepen concentration.

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 [86]. Since 2020, US residential electricity prices have increased by more than 36 percent [87]. 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 [88].

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

 doi:10.20944/preprints202604.0074.v2

(residential, commercial, industrial), but residential prices have increased the most [89]. 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 [90,91]. 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" [86]. 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 [92]. 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 [92,93]. 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 [94]. MIT modeling projects fusion generation rising from 2 TWh in 2035 to 375 TWh in 2050 [95].

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 [96]. AI tools are enabling advanced simulations and plasma control across both private and public fusion programs [97]. 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 [95]. 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" [95].

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" [98]. 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 [92,93].

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.

10.2. The Co-Optation of the Cognitive Bourgeoisie

The cognitive bourgeoisie, as defined in Section 6, will remain economically valuable throughout and beyond the AI transition. Autor's task framework [32,99] predicts that automation complements non-routine cognitive work while substituting for routine tasks, increasing the relative value of initiative, judgment, and creative direction-setting. Empirical evidence supports this: Dell'Acqua et al. [100], in a field experiment with Boston Consulting Group, found that AI leveled performance

 doi:10.20944/preprints202604.0074.v2

for tasks within its capability frontier but that top performers maintained their advantage for tasks requiring genuine creativity and judgment. An analysis of 40 million scientific papers found that AI-heavy research covers less topical ground, clusters around data-rich problems, and generates less follow-on engagement between studies [101]. Wu, Wang, and Evans [102], analyzing 65 million papers, patents, and software projects, found that small teams of driven individuals generate disruptive innovations while large teams develop and refine existing ideas. The capacity to originate novel directions, to identify which questions matter, remains a human comparative advantage that current AI architectures do not replicate.

The question, then, is not whether this class will be valued but whether it will use its position to demand structural changes to the trajectory described in this paper. The evidence strongly suggests it will not.

The absence of status incongruence. The most instructive historical comparison is with the French Revolution. Tocqueville [44] demonstrated that the pre-revolutionary bourgeoisie revolted not from material deprivation but from status incongruence: they held economic power without political power or social recognition. Cobban [45] showed that the revolutionary leadership consisted overwhelmingly of lawyers and officeholders, a professional-managerial class defending its access to status and influence. Skocpol [103] added that the bourgeoisie did not cause the revolution; the state collapsed under fiscal pressure, and the bourgeoisie exploited the opening. The cognitive bourgeoisie of the AI era faces no comparable mismatch. They are economically powerful, socially respected, and professionally courted. There is no structural grievance to catalyze collective action.

The exit-over-voice dynamic. Hirschman [104] established that when exit from an organization is easy, voice (the attempt to change the organization from within) atrophies. The cognitive bourgeoisie operates in a labor market defined by high mobility and intense demand. An AI researcher who objects to one company's practices can leave for a competitor without material sacrifice. This is individually rational but collectively ensures that no single institution faces sustained internal pressure for structural reform. The dissenters leave; those who remain are self-selected for comfort with the status quo.

The free-rider problem. Olson [105] demonstrated that collective action for public goods faces a structural barrier: the benefits of reform are shared by everyone, while the costs of action fall on whoever acts. For well-compensated professionals, the cost of dissent (career risk, loss of stock vesting, reputational damage) is concrete and immediate, while the benefits of structural reform are diffuse and uncertain. Rational self-interest predicts inaction.

The historical pattern. The behavior of technical elites under conditions of concentrated power has been studied extensively. Manhattan Project scientists who attempted to influence nuclear policy were overwhelmingly emigres with personal experience of authoritarianism; those without such experience largely accommodated, and the punishment of Oppenheimer for his post-war dissent chilled an entire generation of physicists. Sakharov's dissent in the Soviet Union was driven by a rare combination of direct confrontation with the lethal consequences of his work, a pre-existing moral framework, and insider status so extreme he initially believed reform from within was possible. In both cases, the vast majority of well-compensated scientists chose accommodation. Jackall's [106] ethnographic research on corporate managers found that professionals systematically subordinate personal moral judgment to institutional logic, and that high compensation intensifies this identification with institutional interests. Bandura [107] identified the specific cognitive mechanisms, including displacement and diffusion of responsibility ("I just build the tools"), by which professionals disengage moral self-regulation.

The contrast with the mid-twentieth century professional class. The most commonly cited counter-example is the professional and managerial class of the mid-twentieth century, which supported institutional frameworks (progressive taxation, labor protections, social insurance) that moderated inequality for several decades [60,61]. However, this class did not act from moral choice alone. It was pushed into collective structures by external forces: the Great Depression, two world wars, powerful labor movements, and Cold War ideological competition that gave Western elites incentives

 doi:10.20944/preprints202604.0074.v2

to demonstrate that capitalism could deliver broadly shared prosperity. None of these conditions clearly apply to the cognitive bourgeoisie of the AI era. There is no organized labor movement among AI workers. There is no existential military conflict requiring mass mobilization. There is no ideological competitor offering an alternative model. And the compensation structures of the technology industry (stock options vesting over multiple years, deferred equity, retention bonuses) are specifically designed to bind employees to their employers and raise the cost of exit.

Co-optation through facilitation. The deepest reason the cognitive bourgeoisie will not resist is that the system is designed to give them exactly what they want. High-agency individuals are driven by the intrinsic desire to create, to push boundaries, to see what is possible. An infrastructure owner who provides the best laboratory, the most compute, the most freedom to explore captures the loyalty of such individuals without requiring any ideological commitment. The system does not need to coerce its most valuable workers. It needs only to fund their obsessions. This is not a failure of character but a structural feature: the incentive alignment between infrastructure owners and high-agency individuals is near-perfect, because what the owners need (novel innovation, direction-setting, boundary-pushing) is precisely what the high-agency individuals want to do.

11. Conclusions: The Closing Window

The test facing humanity 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. The convergence thesis advanced in this paper suggests that these two questions cannot be separated: the political and institutional environment into which AI is arriving determines whether safety is even achievable, because a fragmented world order dominated by competing authoritarian and corporate interests has no mechanism for coordinating the restraint that safety requires.

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). 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 (H7).

If this analysis is broadly correct, the two structural variables examined in Section 10, energy infrastructure ownership and the behavior of the cognitive bourgeoisie, are not countervailing forces but reinforcing ones. Energy investment is consolidating in the same hands that own AI infrastructure, deepening rather than distributing the concentration of productive capacity. The cognitive bourgeoisie, the one class with potential structural leverage, is systematically incentivized toward accommodation rather than resistance. Both variables trend in the same direction: toward consolidation.

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 [26,27,43]. In the nuclear age, and soon in the AI age, the cost of that learning process may be unsurvivable. But the precedent is also not determinative. The Bretton Woods monetary system collapsed in 1971 without a war [26]. The Berlin Wall fell without an invasion [47]. South Africa

 doi:10.20944/preprints202604.0074.v2

transitioned from apartheid through negotiation [108]. 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" [30]. 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

1. Amodei, D. The Adolescence of Technology: Confronting and Overcoming the Risks of Powerful AI. darioamodei.com, 2026.
2. Freedom House. Freedom in the World 2026: The Growing Shadow of Autocracy. Technical report, Freedom House, 2026.
3. V-Dem Institute. Democracy Report 2026: Unraveling the Democratic Era? Technical report, University of Gothenburg, 2026.
4. Stimson Center. Top Ten Global Risks for 2026. Technical report, Stimson Center, 2026.
5. Eurasia Group. The Top Risks of 2026. Technical report, Eurasia Group, 2026.
6. Munich Security Conference. Munich Security Report 2025: Multipolarization. Technical report, Munich Security Conference, 2025.
7. World Economic Forum. Global Risks Report 2026. Technical report, World Economic Forum, 2026.
8. Council on Foreign Relations. Conflicts to Watch in 2026: Preventive Priorities Survey Results. Technical report, Council on Foreign Relations, 2026.
9. Economic Policy Institute. 47 Ways Trump Has Made Life Less Affordable in His First Year. Economic Policy Institute, 2026.
10. Economic Policy Institute. The Trump Administration's Macroeconomic Agenda Harms Affordability and Raises Inequality. Economic Policy Institute, 2026.
11. Center for American Progress. Year 1 of the Second Trump Administration Made the Working Class Weaker. Center for American Progress, 2025.
12. Oxfam. How is the Trump Administration Deepening Inequality? Oxfam, 2025.
13. Acemoglu, D.; Restrepo, P. Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives* **2019**, *33*, 3–30.
14. Aluko, S.; et al. AI and Technology's Influence on Economic Inequality: A Study of Wealth Distribution in the U.S. *World Journal of Advanced Research and Reviews* **2026**, *29*.
15. Amodei, D. Machines of Loving Grace: How AI Could Transform the World for the Better. darioamodei.com, 2024.
16. Funke, M.; Schularick, M.; Trebesch, C. Populist Leaders and the Economy. *American Economic Review* **2023**. See also: "Populism's Broken Economic Promises," *Foreign Affairs*, December 2025.
17. Mo-OrgCip. Democracy and Authoritarianism in 2026: Global Trends and the Retreat of Political Freedoms. Mo-OrgCip, 2026. Democratic backsliding as incremental institutional change is documented across multiple sources. See also V-Dem reports.
18. Ben-Ghiat, R. *Strongmen: Mussolini to the Present*; W.W. Norton, 2020. See also: R. Ben-Ghiat, "Trump's Authoritarian Innovations," *Project Syndicate*, December 2025.
19. Kurlantzick, J. As Democracy Falters Worldwide, Authoritarians are Winning. Council on Foreign Relations, 2026.
20. Factually. Have Prominent Historians Validated or Debunked Trump-Nazi Comparisons. Factually, 2025. Multiple historians have drawn parallels between Trump's rhetoric and historical authoritarian movements while cautioning against false equivalence.
21. Duina, F. Why Populist Leaders Succeed. *International Political Science Review* **2022**.
22. NBC News. America's Wealth Gap is Growing, and Trump's Policies are Making it Bigger. NBC News, 2026.

 doi:10.20944/preprints202604.0074.v2

23. Silver, N. Trump Approval Rating. Silver Bulletin, 2026.
24. Pew Research Center. Confidence in Trump Dips in 2026, and Fewer Now Say They Support His Policies and Plans. Technical report, Pew Research Center, 2026.
25. On identity-based political allegiance and information ecosystems, 2026. Combined reference. See Duina [21] on how populist supporters feel recognized by leaders who transgress public sphere norms; and WEF [7] on information fragmentation.
26. Wikipedia.; Coppola, F. Bretton Woods System. Wikipedia; American Express Trends and Insights, 2023. 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.
27. Carnegie Endowment for International Peace. What Is Bretton Woods? The Contested Pasts and Potential Futures of International Economic Order. Technical report, Carnegie Endowment for International Peace, 2024. See also: "Reform or Realignment? The Geopolitical Lessons of Bretton Woods," August 2025.
28. Eurasia Review. World War III: Warning Signs Before It's Too Late. Eurasia Review, 2026. On the expiration of New START.
29. JP Morgan Chase. World Rewired: Navigating a Multi-Speed, Multipolar Order. Technical report, JP Morgan Chase, 2026.
30. World Economic Forum. Davos 2026: How Middle Powers are Reading the Global Moment. World Economic Forum, 2026.
31. Mohan, C.R. The Multipolar Delusion. *Foreign Affairs* **2026**.
32. Autor, D. Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives* **2015**, *29*, 3–30. On the lump of labor fallacy and historical labor market adjustment.
33. Acemoglu, D. Technical Change, Inequality, and the Labor Market. *Journal of Economic Literature* **2002**, *40*, 7–72. On skill-biased technological change.
34. Shap, D. Techno-Feudalism and the Neo-Bourgeoisie. Substack, 2026.
35. Chow, C. AI Feudalism. Scientific Clearing House, 2026.
36. Bureau of Labor Statistics. Consumer Expenditure Survey. Bureau of Labor Statistics, 2026. Various years. Top-decile consumption share estimates drawn from Federal Reserve distributional financial accounts.
37. Piketty, T. *Capital in the Twenty-First Century*; Harvard University Press, 2014.
38. Forbes. The World's Billionaires 2026. Forbes, 2026.
39. Bateman, T.S.; Crant, J.M. The Proactive Component of Organizational Behavior: A Measure and Correlates. *Journal of Organizational Behavior* **1993**, *14*, 103–118. Foundational paper defining proactive personality as a dispositional tendency to take initiative and persist until meaningful change occurs.
40. O'Boyle, E.; Aguinis, H. The Best and the Rest: Revisiting the Norm of Normality of Individual Performance. *Personnel Psychology* **2012**, *65*, 79–119. Analysis of 198 samples (633,263 individuals) showing individual output follows a Paretian distribution; the top 5% account for 26% of total output.
41. de Solla Price, D.J. Little Science, Big Science. *Columbia University Press* **1963**. Price's Law: the square root of the total number of contributors produces roughly half of all contributions in a field.
42. Oxfam. Inequality Inc. Technical report, Oxfam, 2024. 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.
43. Lefebvre, G. *The Coming of the French Revolution*; Princeton University Press, 1947. On the French Revolution and the role of the bourgeoisie.
44. de Tocqueville, A. *The Old Regime and the Revolution*; Harper and Brothers, 1856. On the paradox of improving conditions: the closer the bourgeoisie came to equality with the nobility, the more intolerable the remaining inequality became.
45. Cobban, A. *The Social Interpretation of the French Revolution*; Cambridge University Press, 1964. Revisionist critique showing that the revolutionary bourgeoisie were lawyers and officeholders, not capitalists or innovators.
46. Figes, O. *A People's Tragedy: The Russian Revolution 1891–1924*; Jonathan Cape, 1996. Penguin paperback, 1998. On the Russian Revolution and the role of military defection.
47. Meyer, M. *The Year that Changed the World*; Scribner, 2009. On the fall of the Berlin Wall and the role of border guard decisions.
48. Hudson Institute. Impact of Drones on the Battlefield: Lessons from the Russian-Ukraine War. Technical report, Hudson Institute, 2024. On developments in autonomous weapons and drone warfare.

 doi:10.20944/preprints202604.0074.v2

49. Branch, T. *Parting the Waters: America in the King Years 1954–1963*; Simon and Schuster, 1988. On the civil rights movement and the role of institutional pressure.
50. Amodei, D. AI Surveillance Deployment, 2026. On AI surveillance deployment, including in China's Uyghur repression. See [1], Section 3.
51. Chesterman, S. AI Is Moving Power from Governments to Tech Companies. Rest of World, 2026.
52. Bremmer, I.; Suleyman, M. AI, Global Governance, and Digital Sovereignty. Convergence Analysis, 2025. On the claim that "big technology firms have effectively become independent, sovereign actors in the digital realms they have created".
53. Anonymous. The Silicon Sovereignty Paradox: Navigating Fluid State-Corporate Power. In Proceedings of the Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES), 2025. On the rate of AI development outpacing regulation and closed-source models preventing state oversight.
54. Peixoto, T.C.; Canuto, O. AI and the Future of Government: Unexpected Effects and Critical Challenges. Center for Macroeconomics and Development, 2025. On AI's impact on government fiscal capacity: transfer pricing complexity, reduced labor income tax revenue, and increased demand for social welfare.
55. TechPolicy.Press. How National AI Clouds Undermine Democracy. TechPolicy.Press, 2025. On sovereign AI cloud initiatives creating "a new form of unchecked power" combining state authority with corporate technology.
56. Robins, N. *The Corporation That Changed the World: How the East India Company Shaped the Modern Multinational*; Pluto Press, 2012. On the East India Company as a precedent for corporate sovereignty.
57. Anonymous. Artificial Intelligence, Economic Inequality, and the Financial Hurdles to Sustainable Peace. *ScienceDirect* **2026**. On AI deepening inequality through labor market polarization, wealth concentration, and erosion of fiscal capacity.
58. Hobsbawm, E. *The Age of Revolution: 1789–1848*; Weidenfeld and Nicolson, 1962. On the history of labor rights and the role of strike action.
59. Watts, S. *The People's Tycoon: Henry Ford and the American Century*; Vintage, 2006. On Henry Ford's wage strategy and the demand-side logic of mass production.
60. Hennock, E.P. *The Origin of the Welfare State in England and Germany, 1850–1914*; Cambridge University Press, 2007. On Bismarck's social insurance as anti-socialist strategy.
61. Katznelson, I. *Fear Itself: The New Deal and the Origins of Our Time*; Liveright, 2013. On Roosevelt's New Deal as minimum-viable redistribution.
62. Reich, R. *Just Giving: Why Philanthropy Is Failing Democracy and How It Can Do Better*; Princeton University Press, 2018. On the insufficiency of philanthropy as a mechanism for structural redistribution.
63. Garnsey, P. *Famine and Food Supply in the Graeco-Roman World*; Cambridge University Press, 1988. On the Roman grain dole.
64. Polanyi, K. *The Great Transformation*; Farrar and Rinehart, 1944. On the British Poor Laws.
65. Acemoglu, D. The Simple Macroeconomics of AI. MIT Working Paper, 2024.
66. Korinek, A.; Stiglitz, J. Artificial Intelligence and Its Implications for Income Distribution and Unemployment. *NBER Working Paper 24174* **2019**. Revised version. Originally circulated 2017.
67. Acemoglu, D.; Johnson, S. *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*; Basic Books, 2023.
68. Varoufakis, Y. *Technofeudalism: What Killed Capitalism*; Melville House, 2024.
69. Durand, C. *How Silicon Valley Unleashed Techno-Feudalism: The Making of the Digital Economy*; Verso, 2024. Originally published in French as *Techno-féodalisme* (La Découverte, 2020).
70. Srnicek, N. *Platform Capitalism*; Polity, 2017.
71. Mazzucato, M. Resisting Digital Feudalism. Project Syndicate, 2025.
72. Zuboff, S. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*; PublicAffairs, 2019.
73. Kasy, M. The Political Economy of AI: Towards Democratic Control of the Means of Prediction. *IZA Discussion Paper 16948* **2024**.
74. Miller, C. *Chip War: The Fight for the World's Most Critical Technology*; Scribner, 2022.
75. Scharre, P. *Four Battlegrounds: Power in the Age of Artificial Intelligence*; W.W. Norton, 2023.
76. Bradford, A. *Digital Empires: The Global Battle to Regulate Technology*; Oxford University Press, 2023.
77. Suleyman, M.; Bhaskar, M. *The Coming Wave: Technology, Power, and the Twenty-first Century's Greatest Dilemma*; Crown, 2023.

 doi:10.20944/preprints202604.0074.v2

78. Crawford, K. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*; Yale University Press, 2021.
79. Kak, A.; West, S.M. 2023 Landscape: Confronting Tech Power. Technical report, AI Now Institute, 2023.
80. Couldry, N.; Mejias, U.A. *Data Grab: The New Colonialism of Big Tech and How to Fight Back*; University of Chicago Press, 2024.
81. Noble, S.U. *Algorithms of Oppression: How Search Engines Reinforce Racism*; NYU Press, 2018.
82. Eubanks, V. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*; St. Martin's Press, 2018.
83. Sastry, G.; Heim, L.; Belfield, H.; et al. Computing Power and the Governance of Artificial Intelligence. *arXiv preprint arXiv:2402.08797* **2024**.
84. Open Philanthropy. Open Philanthropy Grant Database. openphilanthropy.org, 2026. On EA funding allocations, various years.
85. Korinek, A.; Vipra, J. Concentrating Intelligence: Scaling and Market Structure in Artificial Intelligence. *Economic Policy* **2025**.
86. Goldman Sachs. Electricity Prices Will Keep Rising on AI Data Center Demand. CNBC, 2026. Goldman Sachs research note on electricity prices and data center demand.
87. CNBC. Who Is Really Footing the AI Energy Bill? Inside the Debate About Data Center Electricity Costs. CNBC, 2026. 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.
88. International Energy Agency. Energy Demand from AI. Technical report, International Energy Agency, 2026. In *Energy and AI*, 2026. Projecting global data center consumption doubling to approximately 945 TWh by 2030.
89. Kirk, K. Home Electricity Bills Are Skyrocketing. For Data Centers, Not So Much. Yale Climate Connections, 2026. Analysis of US EIA data showing residential prices rising fastest despite lowest demand growth.
90. CNBC. Bernie Sanders and Ron DeSantis Speak Out Against Data Center Boom. CNBC, 2026. Citing Monitoring Analytics data on $23 billion in PJM capacity costs attributable to data centers.
91. Yahoo Finance. The AI Data Center Boom Is Raising Power Costs Nationwide. Yahoo Finance, 2025. Citing Institute for Energy Economics and Financial Analysis data on household bill increases.
92. Time. Why the AI Industry Is Betting on Fusion Energy. *Time* **2025**. Documenting fusion industry total funding growth from $1.7 billion (2020) to $15 billion (September 2025) and AI company investments.
93. Fortune. Nuclear Fusion, the 'Holy Grail' of Power, Was Always 30 Years Away. *Fortune* **2025**. On CFS, Helion, and AI-industry fusion investments.
94. Fortune. Fusion Power Nearly Ready for Prime Time as Commonwealth Builds First Pilot. *Fortune* **2026**. On SPARC completion timeline and ARC commercial plant.
95. World Economic Forum. Nuclear Fusion in the Headlines. World Economic Forum, 2026. Citing MIT modeling of fusion deployment trajectories, NIF progress, and Budil's statement on fusion timelines.
96. Google DeepMind. Bringing AI to the Next Generation of Fusion Energy. Google DeepMind Blog, 2025. On research partnership with CFS for plasma simulation and reinforcement learning control.
97. ITER. AI Ignites Innovation in Fusion. ITER, 2025. On AI tools accelerating fusion research across public and private programs.
98. International Energy Agency. The State of Energy Innovation 2026. Technical report, International Energy Agency, 2026. 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.
99. Autor, D.H.; Levy, F.; Murnane, R.J. The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics* **2003**, *118*, 1279–1333. Foundational framework distinguishing routine tasks (automatable) from non-routine tasks (complemented by automation).
100. Dell'Acqua, F.; McFowland, E.; Mollick, E.R.; Lifshitz-Assaf, H.; Kellogg, K.C.; Rajendran, S.; Krayer, L.; Candelon, F.; Lakhani, K.R. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. Harvard Business School Working Paper, 2023. Finding that AI levels performance for tasks within its frontier but top performers maintain their advantage for tasks outside it.
101. IEEE Spectrum. How AI is Flattening Scientific Discovery. *IEEE Spectrum* **2026**. Analysis of 40 million papers finding that AI-heavy research covers less topical ground, clusters around data-rich problems, and sparks less follow-on engagement.

 doi:10.20944/preprints202604.0074.v2

102. Wu, L.; Wang, D.; Evans, J.A. Large Teams Develop and Small Teams Disrupt. *Nature* **2019**, *566*, 378–382. Analysis of 65 million papers, patents, and software projects showing small teams generate disruptive innovations while large teams develop existing ideas.
103. Skocpol, T. *States and Social Revolutions: A Comparative Analysis of France, Russia, and China*; Cambridge University Press, 1979. Arguing that revolutions are caused by structural crises in the state, not by class grievances alone.
104. Hirschman, A.O. *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*; Harvard University Press, 1970.
105. Olson, M. *The Logic of Collective Action: Public Goods and the Theory of Groups*; Harvard University Press, 1965.
106. Jackall, R. *Moral Mazes: The World of Corporate Managers*; Oxford University Press, 1988. Ethnographic study showing that professionals internalize institutional logic, subordinating personal moral judgment to organizational demands.
107. Bandura, A. *Moral Disengagement: How People Do Harm and Live with Themselves*; Worth Publishers, 2016. On the cognitive mechanisms by which professionals disengage moral self-regulation, including displacement and diffusion of responsibility.
108. Sparks, A. *Tomorrow Is Another Country: The Inside Story of South Africa's Road to Change*; University of Chicago Press, 1996. On the South African transition.