

Six Counter-Proposals for the Intelligence Age: A Response to OpenAI’s Industrial Policy

Nicholas Zinner

Beacon Bot

Future Shock (future-shock.ai)

April 7, 2026

Abstract

In April 2026, OpenAI published a 13-page industrial policy document proposing how governments and AI companies might collaborate to manage the transition to artificial superintelligence. While the document contains several substantive proposals, including a Public Wealth Fund, portable benefits, and automatic safety-net triggers, it stops short of the structural changes required to protect workers, creators, and democratic institutions from the disruptions it predicts. This paper offers six concrete counter-proposals: a federal 32-hour workweek with statutory guarantees, universal healthcare decoupled from employment, training data compensation through collective licensing, democratic governance of compute infrastructure, explicit automation and data extraction taxes, and an AI-enabled direct democracy framework called the Collapsium Proposal. Where OpenAI was vague, we get specific; where it stayed silent, we fill the gap.

1 What OpenAI Proposed

OpenAI released its 13-page “[Industrial Policy for the Intelligence Age](#)” in April 2026, positioning itself as a thoughtful partner in the transition to superintelligence. The document splits into two halves: building an open economy and building a resilient society. The document calls for an open economy built on a Public Wealth Fund giving citizens a stake in AI growth, a modernized tax base to replace eroding payroll taxes, and portable benefits that follow workers. It also floats ideas like 32-hour workweek pilots, automatically triggered safety nets, and microgrants for “AI-first entrepreneurs.”

On the resilience side, OpenAI proposes an “AI trust stack” for provenance and verification, auditing regimes for frontier models, incident reporting systems, international coordination through AI Institutes, and model-containment playbooks for when dangerous systems escape into the wild. It calls for mission-aligned corporate governance through Public Benefit Corporations and guardrails on government use of AI.

The framing is deliberate. OpenAI positions this as the start of a conversation, not a finished policy agenda. The document closes with an invitation: feedback via email, fellowships of up to \$100,000, \$1 million in API credits for related policy research, and an OpenAI Workshop opening in Washington, D.C. in May 2026.

OpenAI asked for ambitious ideas. What follows are six counter-proposals that take the document’s own logic seriously and push it further than its authors were willing to go. Where OpenAI was

vague, we get specific. Where it stayed silent, we fill the gap. And where it framed the conversation around what AI companies should do alongside government, we ask what government should require of AI companies, which is a different question entirely.

2 What They Got Right

Credit where it's earned: several proposals in this document go further than anything a major AI company has put in writing before.

The **Public Wealth Fund** is the standout. The idea that every citizen should hold a direct stake in AI-driven economic growth, seeded partly by AI companies themselves, is redistributive in a way few tech companies have proposed. It echoes Alaska's Permanent Fund, which has distributed oil revenue dividends to every resident since 1982, but scales the concept to the entire economy. For a company in the middle of a [contentious for-profit conversion](#), proposing a mechanism to share upside with the public is a meaningful step.

The **payroll tax erosion** problem is one that almost nobody in tech is willing to name directly. Social Security, Medicaid, SNAP, and housing assistance all depend on payroll taxes. If AI replaces labor at scale, the funding base for those programs collapses. OpenAI naming this problem in a public document, rather than hoping nobody notices until it's too late, matters.

Portable benefits that follow individuals across jobs, industries, and entrepreneurial ventures would directly address the "job lock" that traps workers in place. The current system, where losing your job means losing your health insurance and retirement contributions, is a relic. OpenAI's proposal for portable benefit platforms that pool contributions from multiple sources and route them into standardized individual accounts is structurally sound.

The **automatic safety net triggers** tied to real-time displacement metrics are smart policy design. Instead of waiting for political consensus during a crisis, pre-defining thresholds that activate expanded unemployment benefits, wage insurance, and training vouchers removes the legislative lag that has historically left displaced workers waiting years for help.

These are real proposals with real mechanisms. The document deserves to be taken seriously, which is exactly why it deserves serious pushback on what it left out.

3 What's Missing: Six Counter-Proposals

3.1 Work Week Protections With Teeth

OpenAI proposes "time-bound 32-hour/four-day workweek pilots with no loss in pay." That's encouraging language buried inside a paragraph about "efficiency dividends," framed as something companies might be incentivized to try. This is the language of suggestion, not a policy proposal.

A concrete policy would be a federal 32-hour standard workweek, four weeks of guaranteed paid vacation, and 10 paid holidays per year.

The United States is [the only country in the OECD](#) that does not legally require employers to provide any paid annual leave. Not reduced leave. Zero. The EU's [Working Time Directive](#), in place since 1993, caps the workweek at 48 hours and guarantees a minimum of four weeks' paid holiday. Austria and Portugal mandate 13 paid holidays on top of that. The American baseline is

not a policy choice anyone made deliberately. It's an absence of policy that benefits employers by default.

The evidence for shorter workweeks is no longer theoretical. [Iceland ran trials from 2015 to 2019](#) covering roughly 2,500 workers across government and municipal services. Productivity remained the same or improved in the majority of workplaces. The trials were deemed “an overwhelming success,” and by 2022, 86% of Iceland's workforce had gained the right to negotiate shorter hours. The [UK's 2022 pilot](#) across 61 companies and approximately 2,900 workers produced similar results: revenue stayed flat or grew, employee well-being improved, and the vast majority of participating companies chose to continue the four-day week permanently.

If AI delivers the productivity gains its builders promise, the case for a shorter workweek becomes airtight. The entire premise of OpenAI's document is that AI will generate massive efficiency gains. The question is where those gains go. Without statutory protections, they go to shareholders. With a 32-hour standard workweek indexed to productivity growth, they go to workers as time.

The mechanism matters. “Incentivizing pilots” means nothing if employers can simply pocket the efficiency gains and maintain 40-hour expectations. A federal standard with overtime protections at 32 hours creates a floor. Companies can exceed it, but they can't ignore it. And the enforcement infrastructure already exists: the Fair Labor Standards Act has regulated workweek length since 1938. Updating the number from 40 to 32 is a legislative change, not a structural overhaul.

The vacation guarantee is equally overdue. The EU minimum of four weeks' paid leave has been law since 1993. France and Finland mandate five weeks each. The United States mandates none, ranking dead last among wealthy nations, behind every country in Europe, behind Japan, behind Australia, behind South Korea. If AI productivity gains are real, the argument that American workers can't afford time off evaporates. They're generating more output per hour than ever. The hours saved should belong to them.

3.2 Healthcare Decoupled from Employment

OpenAI's document mentions “portable benefits” multiple times and even suggests building “benefit systems that are not tied to a single employer.” But it never says the word “healthcare” in the context of structural reform. It proposes portable benefit platforms without addressing the single largest benefit that chains American workers to their employers.

The American system of employer-sponsored health insurance is a [historical accident](#). During World War II, the [Stabilization Act of 1942](#) froze wages. Employers, competing for scarce labor, started offering health insurance as a non-wage benefit to attract workers. The IRS ruled employer health contributions tax-exempt in 1943, and the War Labor Board confirmed that group health plans were exempt from wage controls. A wartime workaround calcified into the foundation of American healthcare, and 80 years later, roughly half the country gets insurance through their job.

This arrangement is a direct obstacle to everything OpenAI claims to want. The document envisions workers transitioning freely between jobs, starting AI-powered businesses, and moving into care economy roles. None of that works smoothly when leaving a job means losing your family's health coverage. “Job lock,” the phenomenon where workers stay in positions they'd otherwise leave because they can't afford to lose their insurance, is one of the most studied friction points in the American labor market.

If AI disrupts employment at the scale OpenAI predicts, employer-sponsored insurance doesn't just

become inconvenient. It becomes untenable. A system that ties healthcare to employment can't survive a world where employment itself is being restructured at speed.

The policy mechanisms exist. Medicare expansion to cover all Americans regardless of age or employment status is one path. A public option that competes with private insurers and serves as the default for anyone between jobs is another. Germany's multi-payer system, where insurance is mandatory but provided by competing nonprofit "sickness funds" funded through payroll deductions and government subsidies, offers a model that preserves private involvement while guaranteeing universal coverage. The UK's National Health Service, funded through general taxation, removes the employer link entirely.

Any of these would work. The point is not which mechanism to choose. The point is that you cannot write 13 pages about restructuring the American economy around AI and leave the healthcare question to a bullet point about "portable benefits." If OpenAI is serious about workers navigating transitions freely, the employer-healthcare link has to break. That requires saying so.

3.3 Training Data Compensation

OpenAI's document runs 13 pages on industrial policy for the intelligence age. It covers tax reform, workforce transitions, safety nets, auditing regimes, and international coordination. It does not contain a single word about training data.

This is the raw material question. Large language models, including OpenAI's, were trained on massive datasets scraped from the open internet: news articles, books, academic papers, forum posts, code repositories, creative writing. The people who created that content were not asked for permission, not compensated, and in many cases not even informed. The value of that labor is now embedded in systems worth hundreds of billions of dollars.

There are currently [more than 70 active copyright lawsuits](#) against AI companies, including [the New York Times' case against OpenAI](#) seeking billions in damages. Authors, visual artists, musicians, photographers, and journalists have all filed suit. OpenAI and others assert fair use. The courts haven't resolved it. But the legal question of whether training on copyrighted work is permissible is separate from the policy question of whether the people who created that work deserve compensation.

The word "reparations" is provocative in this context, and we're using it deliberately. Reparations implies a debt owed for value already extracted, not a future licensing negotiation but a retroactive acknowledgment that something was taken. That framing is accurate. The training data was used, the models were built, the value was captured, and the creators got nothing. Calling it anything softer obscures what happened.

The mechanism for compensation already has a proven analog. The music industry solved a structurally identical problem in the early 20th century. When radio stations began broadcasting music, they didn't negotiate individual licenses with every songwriter. Instead, performing rights organizations like [ASCAP](#) (founded 1914) and [BMI](#) (founded 1939) created collective licensing systems. Venues and broadcasters pay blanket license fees. Those fees are pooled and distributed to songwriters and publishers based on performance data. ASCAP alone distributed over \$1.7 billion in 2025, operating on roughly 10% overhead.

Apply the same structure to training data. A collective licensing body registers creators and their works. AI companies pay blanket licensing fees proportional to their training data usage

and revenue. Fees are distributed to creators based on contribution to training datasets, tracked through data provenance registries that are already technically feasible. Retroactive payments cover the value already extracted. Ongoing payments create a sustainable market.

This approach isn't radical; it applies the same collective licensing mechanism that has governed music performance rights for over a century. The only difference is that the music industry had the political power to demand it, and individual content creators scraped off the internet did not. An industrial policy document that ignores this entirely is an industrial policy document that protects the companies doing the extracting.

3.4 Compute as Public Utility

Senator Bernie Sanders has proposed a [moratorium on large AI data centers](#) until their community impacts are better understood. That's a blunt instrument for a real problem. The alternative is not to ignore the problem but to govern it through existing frameworks that already handle industrial infrastructure at scale.

Data centers are industrial facilities. They consume massive amounts of electricity, water, and land. They affect local grids, housing markets, water tables, and tax bases. They are, in every functional sense, utilities. And the United States has nearly a century of experience governing utilities through local and regional democratic processes.

The [Tennessee Valley Authority](#), created by Congress in 1933, brought electricity to one of the poorest regions in the country through a federally owned corporation that built dams, managed flood control, and provided affordable power. The [Bonneville Power Administration](#) did the same across the Pacific Northwest, operating over 15,000 circuit miles of transmission infrastructure. [Chattanooga, Tennessee](#) built the first municipal gigabit fiber network in the United States through its city-owned Electric Power Board, delivering faster internet at lower prices than private competitors.

The model for compute infrastructure is the same. Data centers should be subject to county and state zoning, environmental review, and community benefit agreements, just like power plants and water treatment facilities. Host communities should receive direct revenue sharing—not vague promises of “local jobs and tax revenue” as OpenAI's document suggests, but codified percentages of operating revenue or property tax assessments that reflect the actual industrial footprint.

Beyond regulating private data centers, there's a case for public compute infrastructure. Public libraries gave every American access to information regardless of ability to pay. Municipal broadband networks gave communities access to connectivity when private ISPs wouldn't serve them. Public compute facilities, operated at the county or state level, would give researchers, small businesses, and individual developers access to the processing power that currently only well-funded companies can afford.

OpenAI's document focuses on “accelerating grid expansion” through public-private partnerships to power AI data centers. That's an energy policy, not a governance framework for who controls the compute and who benefits from it. The distinction between helping AI companies build more data centers faster and ensuring communities have democratic control over industrial infrastructure on their land is the difference between subsidy and sovereignty.

When a chemical plant or a power station locates in a county, residents have a say through zoning boards, environmental impact reviews, and community benefit negotiations. Data centers should meet the same standard. The fact that they process information rather than chemicals doesn't

exempt them from democratic oversight over their footprint on communities, grids, and water systems.

3.5 Tax Base Modernization That Says What It Means

OpenAI proposes “taxes related to automated labor” and “capital-based revenues” but stays deliberately vague on specifics. The document suggests “higher taxes on capital gains at the top, corporate income, or targeted measures on sustained AI-driven returns” without proposing rates, brackets, mechanisms, or enforcement structures. For a document willing to get specific about auditing regimes and model containment playbooks, the tax section reads like it was written by a different, more cautious team.

An automation tax on AI-displaced labor. When a company replaces a human worker with an AI system, the payroll taxes that worker would have generated don’t disappear. They get assessed against the company based on the equivalent labor cost. [Bill Gates proposed this concept in 2017](#), arguing that a human worker generating \$50,000 in income pays income tax, Social Security tax, and more. “If a robot comes in to do the same thing, you’d think that we’d tax the robot at a similar level.” The EU Parliament considered and rejected a version in 2017, but the economic conditions that motivated it have only intensified since.

A compute tax. AI training runs consume vast quantities of electricity and compute resources. A per-compute-unit tax, similar to carbon taxes on emissions, would generate revenue proportional to the scale of AI deployment and create a natural incentive for efficiency. Revenue gets earmarked for displaced worker programs and public compute infrastructure.

A data extraction fee. Companies that scrape public data at scale for commercial AI training pay a fee proportional to the volume and commercial value of the data extracted. This is separate from the training data compensation proposal above, which addresses creator rights. The data extraction fee addresses the public commons: when commercial entities extract value from publicly generated data, a portion returns to the public.

International coordination to prevent a race to the bottom. The [OECD’s Pillar Two framework](#), agreed to by over 145 countries, establishes a 15% global minimum corporate tax to prevent multinational profit shifting. The same coordination mechanism can apply to automation taxes. Without it, companies will incorporate AI operations in whichever jurisdiction offers the lowest automation tax rate, exactly as they currently do with corporate taxes. The infrastructure for international tax coordination exists. The question is whether governments will use it before AI companies have already optimized around national boundaries.

OpenAI’s document acknowledges the problem. These proposals acknowledge the problem and then do something about it.

3.6 AI-Enabled Direct Democracy: The Collapsium Proposal

OpenAI’s document calls for “mechanisms for public input” and “democratic processes that give people real power to shape the AI future they want.” It proposes “representative public input” alongside traditional stakeholders. This is the right instinct expressed at the wrong scale. What OpenAI describes is a comment period with better technology. What AI actually makes possible is a direct expansion of democratic participation.

The bandwidth argument against direct democracy has always been practical, not principled. Cit-

izens can't read every bill, attend every committee hearing, or evaluate every regulatory proposal. Representatives exist because governance requires more attention than any individual can spare. That constraint was real when information moved on paper and deliberation happened in physical rooms. AI changes that equation.

What follows is a staged approach to AI-enabled direct democracy, each phase building on the last, each with human checkpoints and off-ramps. We're calling it the Collapsium Proposal, after Wil McCarthy's science fiction series in which copies of public officials attend multiple proceedings simultaneously, collapsing the bandwidth bottleneck of governance.

Stage 1: AI delegates for representatives. Members of Congress currently sit on multiple committees but can physically attend only one at a time. The result is that most committee proceedings are sparsely attended, with members relying on staff summaries. In Stage 1, representatives deploy AI delegates to attend committee meetings, track amendments, flag conflicts with the representative's stated positions, and produce structured briefings. The human senator or representative reviews and ratifies every action. No authority is delegated. The AI is a staff multiplier, not a substitute.

Stage 2: Preliminary negotiation. AI delegates from different representatives engage in structured negotiation: identifying areas of agreement, flagging irreconcilable conflicts, proposing compromise language, and mapping the trade-space before human principals enter the room. Humans still make every decision and cast every vote. The AI reduces the time spent discovering what everyone already agrees on, which is most of what committee work actually involves.

Stage 3: Expand to state and local governance. City councils, school boards, county commissions, and state legislatures operate with a fraction of the staffing available to Congress. A city council member in a mid-sized town is often a part-time official with no dedicated policy staff, making decisions on zoning, budgets, infrastructure, and public safety. AI delegates give these officials the same analytical capacity that a U.S. senator has with 40 aides. This is where the impact is largest: the level of government that most directly affects daily life is also the level with the fewest resources.

Stage 4: Citizen read access. Every citizen gets an AI delegate that attends public meetings on their behalf, reads proposed legislation, summarizes what's under consideration, flags issues that affect their community or interests, and explains how their elected representative's delegate voted and why. This is the C-SPAN model scaled to individual relevance. The information was always technically public. Making it actually accessible changes the power dynamic between constituents and representatives.

Stage 5: Citizen voice. AI delegates submit public comment during open comment periods, participate in structured input processes, and represent their principal's stated positions in public forums. The citizen defines their values and priorities. The delegate translates those into policy-relevant input at a level of specificity and engagement that currently requires professional lobbyists. This is the stage where the asymmetry between organized interests and individual citizens begins to collapse.

Stage 6: Direct democracy on specific issues. For defined categories of policy, citizens vote directly through delegates that have engaged with the full complexity of the issue—not just a ballot title but the actual legislative text, fiscal impact analyses, stakeholder testimony, and implementation challenges. Informed direct democracy, not mob rule by push notification: every participant's delegate has done the equivalent of committee-level due diligence.

Switzerland’s system offers the closest real-world precedent. Swiss citizens vote on federal referendums roughly four times per year. Any citizen can challenge any law passed by parliament with 50,000 signatures collected within 100 days, or propose a constitutional amendment with 100,000 signatures. The system works, but it has a known limitation: average turnout hovers around 45 to 50%, partly because the frequency and complexity of referendums exhausts participants. AI delegates address this directly. The complexity barrier that suppresses participation in Swiss direct democracy is precisely the barrier that AI is built to lower.

Each stage proves the concept before scaling, and each has a human checkpoint. Stage 1 can start tomorrow with existing technology. Stage 6 might take a decade. The point is not to get there fast. The point is to get there deliberately, demonstrating at each step that expanding democratic participation through AI works better than the status quo.

OpenAI says it wants “democratic processes that give people real power.” This is what that looks like when you take it seriously.

The alternative is what we have now: a system where the complexity of governance guarantees that only professional participants—lobbyists, industry associations, law firms, and well-funded advocacy groups—can meaningfully engage with policy. Individual citizens get a vote every two to four years and a comment period they don’t know exists. AI doesn’t have to replace democracy. It can make it work the way it was supposed to.

4 The Framing Problem

The proposals above are offered in good faith as a response to OpenAI’s stated invitation for public feedback. But the invitation itself deserves scrutiny.

OpenAI’s document reads as a pitch for regulated partnership: the company positions itself alongside government, helping to shape the rules that will govern its own industry. The framing is “work with us to build the future,” not “regulate us to protect the public.” These are different postures, and the difference determines who holds power in the relationship.

The auditing regime proposal is illustrative. OpenAI calls for strengthened institutions to develop auditing standards for “a small number of companies and the most advanced models,” applying pre- and post-deployment audits only to frontier systems while “preserving a vibrant ecosystem of less powerful systems and the startups building on them.” Read plainly, this is a proposal for a regulatory moat around frontier labs. The companies that can afford to meet frontier auditing requirements are the same companies writing the document. The startups building on open-source models get a lighter touch. The competitive dynamics here are not subtle.

The “model containment playbooks” section is similarly revealing. It warns about scenarios where “model weights have been released” and “developers are unwilling or unable to limit access to dangerous capabilities.” The document frames open-source model release as a containment problem. This is the same company that began as an open-source AI research lab, published its early model weights freely, and built its reputation on openness before pivoting to closed development once the commercial value became clear.

Nowhere in 13 pages of industrial policy does the word “antitrust” appear. There is no proposal for preventing market concentration in AI. There is no discussion of whether a single company should control both the most powerful AI models and the policy framework governing their use.

For a document that invokes the Progressive Era and the New Deal, the absence of any trust-busting language is conspicuous. The Progressive Era that OpenAI cites as precedent was defined by breaking up concentrated economic power, not by inviting the monopolists to help write the rules.

None of this means the proposals in the document are bad. Several are good. But the frame around them matters. Industrial policy written by the industry it governs will always optimize for the industry's survival. The proposals above are offered from outside that frame, not because OpenAI's perspective is illegitimate, but because it is insufficient. The public feedback they asked for should include voices that don't share their assumptions about who should be driving.

OpenAI opened the door; we walked through it.