



## **Dominion of the Made: AI, America, and the Stewardship of What Comes Next**

By: Colin Redemer

### **I. The Tool That Talks Back**

“A new word is like a fresh seed sewn on the ground of the discussion,” Wittgenstein once wrote.<sup>1</sup> If that is true, then we are living in a season of extraordinary planting. The large language models that have arrived in the past three years do something no previous tool has done: They talk. They reason, or seem to. They produce language that, on first encounter, is indistinguishable from the work of a competent human mind. Ask one to write a sonnet in the style of John Milton, and it will produce fourteen lines of moderately competent verse, rhymed and metered, on the subject you specify. It will scan well enough. It may even be pleasant. But as anyone who has spent time with Milton knows, something will be wrong. The third line may be swollen to unreadability. The rhyme scheme may quietly break its own pattern. The content will rarely rise above bathos. The machine has produced an object that resembles a poem without being one. It has made something. It has not begotten anything.

This distinction between making and begetting is ancient, and it is essential. Jesus, the Nicene Creed tells us, is begotten, not made. Creation is made, not begotten. Human children are begotten. Human artifacts such as tables and chairs are made. The confusion of these two categories is, in the Christian tradition, among the oldest and most dangerous of errors. The golden calf is crafted; it is not born. The idol acquires its form from human hands and human minds, and the maker of the idol is himself attempting to gain worship.<sup>2</sup> When we look at artificial intelligence today, we are looking at something made by human hands and minds, something crafted with extraordinary ingenuity, and we must ask whether we are in danger of treating the made as though it were begotten, of confusing the artifact with the creature, the tool with the mind.

This paper argues that AI development is both necessary and dangerous. But it should not be catastrophized. It is not, in this sense, any different than the problem of technology itself, which has confronted humanity from the beginning. The United States cannot afford to pause or retreat

---

<sup>1</sup>Ludwig Wittgenstein, *Culture and Value* (University of Chicago Press, 1984), 2e.

<sup>2</sup>The distinction between begetting and making is explored at length in Oliver O'Donovan, *Begotten or Made?* (Oxford University Press, 1984). Further discussed in Colin Redemer, *Made Like the Maker* (Davenant Press, 2023).

in AI development because adversarial nations will not pause, and the world that results from their AI dominance is a world in which the virtues this paper defends, including belief in God, political freedom, and the rule of law, are systematically undermined. But development without moral direction is civilizational suicide by a different route. The task before us is to develop AI aggressively while binding it to a framework rooted in the American tradition, which is itself rooted in Christian anthropology. This paper surveys the landscape and proposes a path.

## **II. What AI Is and What It Is Not**

Before we can govern a thing, we must understand what it is. And before we understand what artificial intelligence is, we should be clear about what it is not. AI does not think. It processes. It does not know. It pattern-matches. It does not desire. It optimizes. To make this distinction is not to diminish its power. A sword does not desire either, yet swords are both dangerous and useful, and as we well know swords have shaped history. But the category error of treating the made as though it were begotten leads to two equally dangerous mistakes: either we worship it, which is the techno-utopian temptation in which the singularity functions as secular eschatology, or we fear it as though it had agency and malice, which is the *Terminator* fallacy. The Christian position is neither. We are stewards of what we make. What we make must be crafted to serve us as we serve God.

What AI genuinely can do is impressive, and its capabilities are expanding. Language generation, code production, data analysis, image recognition, logistical optimization, scientific modeling: In each of these domains the technology has demonstrated capabilities that range from marginally helpful to truly astonishing. Nate Silver, no technophobe, reports that large language models offer a substantial benefit to his workflow roughly 80 percent of the time, performing tasks ranging from reconciling datasets to generating functional code. But the remaining 20 percent is revealing. When Silver asked ChatGPT to simulate a poker hand, it got nearly everything wrong: It made inexplicable strategic decisions, awarded the pot to the wrong player, and repeatedly failed to calculate stacks correctly.<sup>3</sup> Poker matters as a test case because it requires holding multiple layers of abstraction simultaneously: the fixed mathematical rules, the fuzzy contextual reads on opponents, and the cascading effects of single deviations on the entire game. When you layer tasks on top of one another, Silver finds, performance drops off a cliff.

The scaffolding problem, as Benjamin Todd has described it, is a matter of compounding reliability. Even 90 percent accuracy per step yields only 10 percent success across twenty steps. At 99 percent per step, success leaps to 80 percent.<sup>4</sup> This is not merely a technical limitation awaiting the next software update. It points toward an ontological boundary. The kind of integrated, multilayered, context-sensitive judgment that poker demands, and that law, medicine,

---

<sup>3</sup>Nate Silver, "ChatGPT Is Shockingly Bad at Poker," *Silver Bulletin*, Substack, May 21 2025, <https://www.natesilver.net/p/chatgpt-is-shockingly-bad-at-poker>.

<sup>4</sup>Benjamin Todd, "The Case for AGI by 2030," cited in Silver, "ChatGPT."

military command, pastoral care, teaching, and parenting also demand, requires something the machine does not possess: interiority. The machine processes inputs. The person dwells among them.

Jasmine Sun, writing in *The Atlantic*, has documented this boundary from a different angle. She found that the post-training process by which AI companies define the ideal character for a model, making it helpful, honest, and harmless, actively suppresses the qualities that make writing genuinely good. When you tell a model to be a brilliant prose stylist but also a PhD-level mathematician and also strictly appropriate for all audiences, it becomes rigid and cautious, like a candidate at a job interview terrified to misstep.<sup>5</sup> This is the same insight Peter Thiel gestured at with the quip that if you don't want AI to replace you, "be a little bit racist or a little bit sexist. . . Everybody else will get replaced."<sup>6</sup> The unpredictability that made earlier models creative also made them prone to other unwanted behaviors. The result is a technology that has memorized centuries of great literature but, due to its limitations, cannot produce a single essay worth reading for its prose. Sun's observation cuts deeper than aesthetics: Voice, she argues, emerges from the specificity of a life. The models cannot live, cannot feel, cannot smell, and cannot taste. They struggle to generate hilarious, if random, non sequiturs intentionally, though they often do so accidentally, such as when Google's Gemini made endless images of surprisingly diverse Nazis in military uniforms. They cannot spill raw emotions onto the page or place abstract concepts in rich physical settings.<sup>7</sup> This is not a deficiency to be patched. It is the nature of the made, which bears the mark of its maker but does not share his being.

### **III. The Economic Earthquake**

Regardless of what AI cannot do, what it can do is already reshaping the material conditions of American life. The question is not whether this reshaping will occur but how fast, how deep, and to whose benefit.

The most sober assessment available comes from Daron Acemoglu, whose analysis of AI's macroeconomic implications predicts that the effects will be real but modest: no more than a 0.66 percent increase in total factor productivity over ten years, and probably less than 0.53 percent when accounting for the difficulty of hard-to-learn tasks where context-dependent judgment resists automation.<sup>8</sup> This is a corrective to the breathless predictions of Silicon Valley, and it should temper both utopian and dystopian expectations. The GDP revolution may come, but it is not arriving tomorrow.

---

<sup>5</sup>Jasmine Sun, "The Human Skill That Eludes AI," *The Atlantic*, March 17, 2026, <https://www.theatlantic.com/technology/2026/03/ai-creative-writing/686418/>.

<sup>6</sup>Peter Thiel, "The Coming Collapse No One Is Prepared For," All-In Summit, 2024, video, 00:22:59, posted by All-In Podcast, September 13, 2024, <https://youtu.be/SYRunzR9fbk?t=1379>.

<sup>7</sup>Sun quotes James Yu, cofounder of Sudowrite: "Most people's good first stories are autobiographical. Maybe you need a model that lives a life, and can almost die."

<sup>8</sup>Daron Acemoglu, "The Simple Macroeconomics of AI," National Bureau of Economic Research Working Paper 32487, May 2024, [https://www.nber.org/system/files/working\\_papers/w32487/w32487.pdf](https://www.nber.org/system/files/working_papers/w32487/w32487.pdf).

Aggregate productivity is not the whole story, and it may not even be the most important part of it. Acemoglu's own analysis reveals something more troubling: AI is predicted to widen the gap between capital income and labor income.<sup>9</sup> The gains accrue to those who own the models and the infrastructure, not to the workers whose tasks the models perform. Even when AI improves the productivity of lower-skill workers in certain tasks, this may increase rather than reduce inequality so long as no new tasks are created for those workers to fill. The macroeconomic number is modest. The distributional question is severe.

Ethan Mollick's research at the organizational level confirms this asymmetry. Workers report that AI triples their productivity, reducing ninety-minute tasks to thirty minutes.<sup>10</sup> Over 40 percent of American workers said they used AI at work in April 2025, a number that is surely higher now. But companies are not capturing these gains organizationally. The reason is revealing: Many workers are hiding their AI use. They fear that disclosing their new efficiency will lead to layoffs, that their AI-assisted outputs will be devalued, or that productivity gains will simply be converted into expectations of more output for the same pay.<sup>11</sup> The worker who uses AI to do his job in half the time faces a perverse incentive to conceal his own productivity. This is not a management problem. It is a dignity problem. When efficiency becomes the only measure of a worker's value, the relationship between the worker and his work becomes adversarial, particularly when viewed from the perspective of the manager allocating capital.

Anthropic's own research on labor market impacts introduces a new measure of occupational exposure that combines theoretical AI capability with real-world usage data. Its findings must be taken with a grain of salt (after all, what tech company would put out a report that its shiny new tech is *not* generationally disruptive?), but nevertheless they still confirm that AI is far from reaching its theoretical ceiling: Actual task coverage remains a fraction of what the company sees as feasible. Computer programmers are the most exposed, at 75 percent task coverage, followed by customer service representatives and data entry keyers. Crucially, Anthropic's researchers find no systematic increase in unemployment for highly exposed workers since late 2022. But there is an early warning: suggestive evidence that hiring of younger workers has slowed in exposed occupations, with a 14 percent drop in the job-finding rate for workers aged twenty-two to twenty-five entering those fields.<sup>12</sup> This finding is consistent with recent data from graduates with computer science degrees. Ten years ago the job placement rate for such graduates was very high, but now it shows signs of decline.<sup>13</sup>

---

<sup>9</sup>Acemoglu finds that even when AI improves the productivity of low-skill workers in certain tasks, this may increase rather than reduce inequality.

<sup>10</sup>Ethan Mollick, "Making AI Work: Leadership, Lab, and Crowd," *One Useful Thing*, Substack, May 22, 2025, <https://www.oneusefulthing.org/p/making-ai-work-leadership-lab-and>.

<sup>11</sup>Mollick reports that over 40 percent of American workers used AI at work by April 2025, but official chatbot usage maxes out at 20 percent. The discrepancy indicates widespread hidden use.

<sup>12</sup>Maxim Massenkoff and Peter McCrory, "Labor Market Impacts of AI: A New Measure and Early Evidence," Anthropic Research, March 5, 2026, <https://www.anthropic.com/research/labor-market-impacts>.

<sup>13</sup>Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen, "Canaries in the Coal Mine? Six Facts About the Recent Employment Effects of Artificial Intelligence," Stanford Digital Economy Lab, November 13, 2025, [https://digitaleconomy.stanford.edu/app/uploads/2025/11/CanariesintheCoalMine\\_Nov25.pdf](https://digitaleconomy.stanford.edu/app/uploads/2025/11/CanariesintheCoalMine_Nov25.pdf). Using ADP payroll data covering

The demographic profile of the most exposed workers is itself significant. They are more likely to be highly educated and higher-paid, earning 47 percent more on average than unexposed workers.<sup>14</sup> This is not the factory-floor displacement narrative of the twentieth century. This is the professional class: lawyers, analysts, programmers, and administrators. These are the people who currently staff the institutions of American life. Their displacement, or their reduction to dependency, will reshape not only the economy but the culture and the politics of the republic. A society in which the professional middle class cannot find meaningful work is a society ripe for demagoguery, despair, or both. At its worst it may produce a class educated enough to organize and aggrieved enough to revolt.

The educational pipeline compounds the concern. A systematic review of seventy empirical studies on AI in education finds that when students interact extensively with large language models, the risks extend beyond poor outputs to cognitive and behavioral outcomes: reduced neural activity, over-reliance on the tool, diminished independent learning skills, and a loss of student agency.<sup>15</sup> The supplement has become a substitute, as Joshua Mitchell might say.<sup>16</sup> And what is being substituted is the development of the mind itself. If young people are both being poorly formed by AI-mediated education and finding it harder to enter the workforce in AI-exposed occupations, a generational crisis is forming at both ends of the pipeline.

#### **IV. The Geopolitical Necessity**

The economic case for caution must be set against the geopolitical case for urgency. And the geopolitical case is not subtle.

The Department of Defense's 2025 annual report to Congress on China's military developments states the matter plainly. China's People's Liberation Army (PLA) has for decades marshaled resources, technology, and political will to achieve its vision of a world-class military. The PLA is a key component of China's ambition to displace the United States as the world's most powerful nation. China's top military strategy focuses on overcoming the United States through what Beijing terms "national total war."<sup>17</sup> China's announced defense budget has nearly doubled since the first full year of Xi Jinping's term as general secretary, and the country continues to accelerate its development of biotechnology and hypersonic missiles as well as a keen awareness of AI as a military technology.<sup>18</sup>

---

millions of workers, the authors find a 16 percent relative employment decline for workers aged twenty-two to twenty-five in AI-exposed occupations since late 2022.

<sup>14</sup>Massenkoff and McCrory, "Labor Market Impacts."

<sup>15</sup>Iris Delikoura, Yi R. (May) Fung, and Pan Hui, "From Superficial Outputs to Superficial Learning: Risks of Large Language Models in Education," preprint, September 26, 2025, <https://arxiv.org/abs/2509.21972v2>.

<sup>16</sup>Joshua Mitchell, *American Awakening* (Encounter Books, 2020). See also Mitchell, "When Supplements Become Substitutes: A Theory of Nearly Everything," *City Journal* (Autumn 2018), 57–65.

<sup>17</sup>*Annual Report to Congress: Military and Security Developments Involving the People's Republic of China 2025* (U.S. Department of Defense, 2025), [PAGE #],

<https://media.defense.gov/2025/Dec/23/2003849070/-1/-1/1/ANNUAL-REPORT-TO-CONGRESS-MILITARY-AND-SECURITY-DEVELOPMENTS-INVOLVING-THE-PEOPLES-REPUBLIC-OF-CHINA-2025.PDF>.

<sup>18</sup>*Annual Report to Congress*, [PAGE #].

Responding to the AI dimension of this competition is particularly urgent. In 2024, China’s commercial and academic AI sectors made significant progress on large language models and reasoning models, narrowing the performance gap with leading American systems.<sup>19</sup> China’s military–civil fusion strategy channels civilian AI innovation directly into defense applications. Firms such as Baidu, Alibaba, and Huawei are not merely commercial enterprises; they are strategic instruments of a state that makes no meaningful distinction between civilian and military technological development.<sup>20</sup> The PLA is integrating AI into unmanned systems, intelligence collection and analysis, decision-making assistance, cyber operations, and information campaigns. The Volt Typhoon cyberespionage campaigns burrowed into U.S. critical infrastructure, demonstrating capabilities that could disrupt the American military in a conflict.<sup>21</sup> In November 2024, China deployed nuclear-capable bombers in a combined strategic patrol with Russian aircraft for the first time, a signal of deepening coordination between two powers united primarily by their shared interest in constraining the United States.<sup>22</sup>

The PLA continues to refine multiple military options to force Taiwan’s unification by brute force, including amphibious invasion, firepower strikes, and maritime blockades. China expects to be able to fight and win a war over Taiwan by the end of 2027. PLA strikes could potentially range up to 1,500 to 2,000 nautical miles from China, enough to seriously challenge American presence in or around a conflict in the Asia-Pacific.<sup>23</sup>

The implications for AI policy are direct. American restraint in AI development is not virtue; it is abdication. If the United States does not lead in AI development, China will, and the world that results from Chinese AI dominance is a world in which the virtues this paper defends are not merely unfashionable but actively suppressed. China is building AI to control narratives, to surveil populations, and to enforce ideological conformity. It promotes atheism, suppresses political freedom, and believes law is merely an expression of state power. If America does not build AI to win, China wins by default.

But the geopolitical argument cuts both ways. If America develops AI without moral direction simply to win the race, it risks becoming what it opposes: a surveillance state, a technocratic oligarchy, a society in which efficiency replaces the Western concept of justice. The material supply chain is itself fragile, as was demonstrated in the COVID-19 pandemic. Eighty to 90 percent of the world’s high-purity quartz, essential for semiconductor manufacturing, comes from a single small town in North Carolina. A hurricane in 2024 temporarily shut down operations, exposing the dependence of the entire digital economy on a geological accident in the

---

<sup>19</sup>*Annual Report to Congress*, [PAGE #].

<sup>20</sup>Cole McFaul, Sam Bresnick, and Daniel Chou, “Pulling Back the Curtain on China’s Military-Civil Fusion: How the PLA Mobilizes Civilian AI for Strategic Advantage,” issue brief, Center for Security and Emerging Technology, Georgetown University, September 2025, <https://cset.georgetown.edu/publication/pulling-back-the-curtain-on-chinas-military-civil-fusion/>.

<sup>21</sup>*Annual Report to Congress*, [PAGE #]. The Volt Typhoon cyberespionage campaigns “burrowed into U.S. critical infrastructure, demonstrating capabilities that could disrupt the U.S. military in a conflict and harm American interests.”

<sup>22</sup>*Annual Report to Congress*, [PAGE #].

<sup>23</sup>*Annual Report to Congress*, [PAGE #]. The PLA continues to make steady progress toward its 2027 goals, whereby it must be able to achieve “strategic decisive victory” over Taiwan.

Blue Ridge Mountains.<sup>24</sup> Strategic resilience requires not just building faster but building wisely: diversifying supply chains; investing in materials science; and developing the energy infrastructure, including nuclear power, that AI-scale computation demands.

The challenge is to maintain the republic's character while developing the tools necessary to defend its interests and those of its allies. This is where stewardship becomes a strategic concept, not merely a theological one.

## **V. The Gatekeepers**

Who, then, shapes what AI becomes? The answer at present is a handful of private companies making civilizational decisions with minimal accountability.

The foundational models that underlie modern AI are built by a small number of firms: Anthropic, OpenAI, xAI, Google DeepMind, and Meta's AI division. These companies make daily decisions about training data, alignment targets, deployment boundaries, and value encoding that function as de facto cultural legislation. They are writing the moral code of the most powerful tools ever created, and they are doing so largely in private. The International Telecommunications Union (ITU)'s 2025 AI Governance Report confirms the concentration: American and Chinese companies operate more than 90 percent of the data centers used globally for AI, while more than 150 countries have no computing infrastructure at all.<sup>25</sup>

Beyond the model builders, there are the companies figuring out how to deploy AI in consequential domains: Palantir in defense and intelligence; companies building AI into health care, education, legal systems, and hiring pipelines; and robotics firms such as Anduril in defense and Tesla and Boston Dynamics in consumer and industrial applications. Each deployment decision is a moral decision, whether or not it is understood as one.

Then there is the government. Congress has so far demonstrated limited capacity to legislate on technology it does not understand, though several proposals merit serious consideration. The executive branch matters most in the near term: regulatory agencies, executive orders, military procurement decisions, and education policy. State governments may prove more agile and are already experimenting with AI legislation. But this experimentation represents two dangers of its own. First, fifty states setting their own policies can hinder technological progress and adoption because tech companies have to navigate interlocking webs of compliance and regulation to bring products effectively to the market. Second, the reality is that the relatively smaller number of states where AI development is most heavily concentrated will likely end up dictating policy

---

<sup>24</sup>Tom Carter and Geoff Weiss, "Hurricane Helene's Devastation Has Shut Down a One-of-a-Kind Mine for AI Chips," *Business Insider*, October 1, 2024,

<https://www.businessinsider.com/quartz-mine-computer-chips-hurricane-helene-spruce-pine-semiconductor-2024-9>

<sup>25</sup>International Telecommunications Union, *The Annual AI Governance Report 2025: Steering the Future of AI* (ITU Publications, 2025), [PAGE #].

for the rest of the country and, in the hope America wins the AI race, the world.<sup>26</sup> In a national representative democracy, such policy should be set at the national level. The ITU report reveals that regional approaches to AI governance diverge sharply.<sup>27</sup> But the international governance framework remains aspirational. It treats AI governance as a coordination problem when it is actually a power problem. Who builds the models, who owns the data centers, who sets the alignment defaults: These are the questions, and a multistakeholder dialogue at the United Nations does not answer them. Nor would we allow private companies concentrated in a handful of cities to build the digital equivalent of nuclear weapons without any federal oversight.

The current situation, in which private companies set the civilizational terms and international bodies produce reports, is untenable. But the solution is not European-style regulation that strangles development. It is an American industrial policy for AI that sets directional guardrails while preserving the innovative capacity of private enterprise and educating the citizens who must learn to work and live with the aligned AI.

## **VI. Alignment Is Not Neutral**

The most culturally consequential question in AI development is one that is rarely stated directly: aligned with what?

“Alignment” in AI discourse typically means ensuring that AI systems behave in ways that are helpful, harmless, and honest. But this framing conceals deeper questions: Helpful to whom? Harmless by whose definition? And, of course, if the technology must be made to be harmless, that necessity itself implies the awareness of the massive *potential* for harm. We cannot count on all peer nations unilaterally deciding to design their AI to be harmless. The current default alignment of major AI systems reflects the vision of the institutions that build them: overwhelmingly progressive, secular, and globally minded. This is not conspiracy; it is sociology, and at this point it is observable in the news. The people building these systems encode their assumptions about what is harmful, what is offensive, and what is true. The result is systems that are functionally hostile to conservative, traditional, and religious perspectives, not by intention (usually) but by default.

Sun’s reporting in *The Atlantic* documents how this happens mechanically. Post-training teams evaluate model outputs based on personal taste. Companies contract with domain experts whose rubrics transform something as slippery as tone into discrete criteria. This is not dissimilar to the effect sensitivity readers have had on the publishing industry and content moderators have had on social media, but given the scale of AI adoption the effect this approach may have on society

---

<sup>26</sup> Mark Muro and Shriya Methkuppally, “Mapping the AI Economy: Which Regions Are Ready for the Next Technology Leap?,” Brookings Institution, July 16, 2025,

<https://www.brookings.edu/articles/mapping-the-ai-economy-which-regions-are-ready-for-the-next-technology-leap/>.

<sup>27</sup>Muro and Methkuppally, “Mapping the AI Economy.” The report finds that Europe regulates by risk classification, Asia prioritizes innovation, and Africa focuses on inclusion and access.

will be far more widespread while being just as pernicious.<sup>28</sup> While AI systems are scored on benchmarks for coding and mathematics that dramatically shape public perception of which company is “winning” the race, the question of whether the system can engage with a traditional moral framework without treating it as pathological receives no systematic attention.<sup>29</sup>

This matters because AI is rapidly becoming the mediating layer of American life. If every interaction with information, education, customer service, legal guidance, and eventually health care is filtered through systems that encode progressive secular assumptions, the effect will be a soft cultural revolution conducted through infrastructure rather than politics or culture. The people affected will not experience it as coercion. They will experience it as normalcy. This possibility is a greater threat than overt censorship because it is invisible. A system that can discuss gender theory with nuance but treats orthodox Christian sexual ethics as bigotry is not neutral. It is partisan. And partisanship embedded in infrastructure is tyranny wearing a helpful smile.

The White House AI Action Plan of July 2025 gestures toward this problem. The accompanying executive order on preventing ideological bias in government-procured AI systems requires that large language models be “truthful” and “ideologically neutral.”<sup>30</sup> This is directionally correct, but it is not sufficient. “Not biased” is a merely negation where a positive vision is required. The question is not whether AI should be neutral but whether genuine neutrality is even possible in a system that must make constant decisions about what to say, what to refuse, and what to treat as settled. The answer, on any honest examination, is no. Every system encodes a vision of the world. The only questions are which vision and who decides.

The argument of this paper is not that AI should be programmed to be conservative or Christian. It is that AI systems serving a Western republic must be genuinely pluralistic, capable of engaging with traditional moral frameworks without treating them as pathological and capable of presenting contested questions as contested rather than resolved. That vision is itself Western and, in fact, Protestant, descending as it does from the Reformation’s settlement on how far to extend religious liberty and how to separate that liberty from the magistrate’s right to rule. Our AI must be made to reflect the underlying system of thought that developed this, or we risk destroying America as we seek to save it in our attempt to win the race for AI dominance. The policy mechanisms for achieving this aim, including transparency requirements for alignment decisions, a right to know what vision a system has been trained to promote or suppress, and democratic contestation of alignment principles for systems serving public functions, are discussed in the policy section below.

---

<sup>28</sup>Sun reports that a Scale AI contractor described grading rubrics that included rules such as “the response should use a maximum of two exclamation marks,” leading to cases where inferior responses were rated higher on formal criteria.

<sup>29</sup>Sun quotes Nathan Lambert, a post-training lead at the Allen Institute for AI: “The more you control for these traits, the more you suppress creativity.”

<sup>30</sup>Exec. Order No. 14319, “Preventing Woke AI in the Federal Government,” 90 Fed. Reg. 35389 (July 23, 2025).

## VII. Under Sovereignty

The threads running through the preceding sections require a framework that can hold them together. Secular frameworks fail in specific and identifiable ways. Pure libertarianism trusts the market to solve moral problems it is structurally incapable of addressing. Progressivism replaces divine sovereignty with human management and inevitably tends toward un-American totalitarian coercion. Techno-utopianism is secular eschatology, the belief that the next invention will save us, and it is the oldest heresy in new clothing: “Eat the fruit. Trust me; you will not die. You will be as gods.” Existential doomerism is despair dressed as seriousness. None of these is adequate.

The Christian understanding of human nature provides what is needed: We are made in the image of God, which means we are made to make. The dominion mandate given to Adam is not revoked by the Fall. We are called to name, to order, to steward creation. Our making is part of our vocation. Technology, as the philosopher of culture would say, is downstream from poetry: The making with hands is always articulated to the maker in his own mind before the making takes place. The builder of the house must have a word for *screw* to distinguish it from *nail*, a word for *plank* to distinguish it from *beam*. The house that is built with hands is built first with these words in the mind of the builder. Technology is downstream from human creative capacity.

But we are also fallen. Our best intentions are subject to corruption. The heart is a forge for idols.<sup>31</sup> The human attempt at godhood is as old as humanity itself, and the confusion of making with begetting, of the artifact with the creature, is its most persistent form. This is why the Christian tradition insists on the distinction. What we make serves us. What we beget we receive. We do not speak children into being as idealized forms, tweaked and edited until they are just right. We beget them, and that means to receive them as they are. If God were to disappear, there would be little to stop our world from descending into unthinkable savagery.<sup>32</sup>

We act because we are called to act. We act with sobriety because we know our works are tainted. We do not despair because the outcome of history is not in our hands; it is in the hands of Providence. And we do not worship our creations because we know the difference between the made and the Maker. This posture, and nothing less than this posture, is adequate to the moment. It produces neither the paralysis of the Luddite nor the recklessness of the techno-gnostic but the steady, sober, joyful work of the steward who has been given a task and intends to complete it.<sup>33</sup>

The American experiment is grounded in assumptions about human nature that are themselves grounded in this tradition: inalienable rights, the dignity of the individual, the prudent limited use of state power, and the moral equality of persons before the law. AI policy that is severed from these roots will inevitably tend toward something un-American, whether that takes the form of

---

<sup>31</sup> See Redemer's discussion of the heart as a forge for idols in his introduction to Traherne's *Christian Ethics*, [PAGE #].

<sup>32</sup> Ernst Junger, *Correspondence 1949–1975* (Rowman and Littlefield, 2016), 8. Junger, writing as a non-Christian, observed that without the Christian confessions, “cannibalism would break out overnight. It sits beneath a thin veneer.”

<sup>33</sup> Joseph Pieper, *The End of Time* (Pantheon Books, 1954), 138.

techno-libertarian oligarchy or techno-progressive managerialism. The roots matter. They are not decoration on a secular structure. They are the foundation.

God has told us how the story ends. This too is good news and freedom. We do not need to fear the creation of our hands, whether an artificial intelligence or a genetically modified organism or a new form of power generation. Such anxiety is right for those who know that what they are bringing forth may demand blood like a dark god of old. Christians fear God. We do not fear our own actions, our own creations, our own death, or even the end of time. We have already accepted our death and daily anticipate it with great expectation.<sup>34</sup> The Christian attitude to history contains both affirmation of creation and readiness for suffering; only the person who combines this affirmation and this readiness will retain the possibility of meaningful action, arising out of genuine inner conviction, even in the midst of the catastrophe.<sup>35</sup>

## **VIII. What Is to Be Done**

The preceding analysis leads to several clear policy directions. Some of these build on the White House AI Action Plan of July 2025;<sup>36</sup> others go beyond it. All are rooted in the conviction that AI development is a national strategic priority that must be guided, at the federal level, by moral principle, not merely by competitive urgency.

### **An American Industrial Policy for AI**

The action plan's framing of AI as a competitive race comparable to the Cold War space program is correct.<sup>37</sup> The federal government should treat AI development as a strategic national priority with directional funding, public-private partnerships, and explicit alignment with both national security and our particular national vision. This is a break from laissez-faire orthodoxy and should be acknowledged as such. The market alone will not produce the outcomes the republic requires because the market optimizes for profit, and the alignment of AI with our American vision is not, strictly speaking, profitable. Industrial policy is the mechanism by which a republic expresses its collective judgment about what matters beyond the quarterly earnings report. This includes aggressive investment in energy infrastructure, including nuclear power, to support the computational demands of AI at scale. It includes supply chain diversification for critical materials, including high-purity quartz and other rare earth minerals. And it includes continued support for semiconductor reshoring (possibly including the resettling of large numbers of Taiwanese workers and infrastructure in the Americas), ensuring that the physical substrate of AI remains under American control.

---

<sup>34</sup>Revelation 19:11–15.

<sup>35</sup>Pieper, *The End of Time*, 138.

<sup>36</sup>*Winning the Race: America's AI Action Plan* (The White House, Office of Science and Technology Policy, 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>.

<sup>37</sup>*Winning the Race*. The plan identifies over ninety federal policy actions across three pillars: Accelerating AI Innovation, Building American AI Infrastructure, and Leading in International AI Diplomacy and Security.

## **Immigration and Labor**

AI-driven productivity gains fundamentally alter the economic argument for mass immigration. The case for high levels of immigration has historically rested on labor scarcity: The economy needs workers (or needs to reduce labor costs), and immigration supplies this workforce. As AI and robotics compress the need for human labor across multiple sectors simultaneously, this predicate dissolves. Policy should prioritize the welfare of existing citizens, directing the productivity gains of the AI economy toward maintaining and improving the material conditions of the American populace. The benefits of AI-generated wealth should accrue to our citizens first. This is not nativism in the traditional sense. It is the basic logic of political obligation in a republic. The state's first duty is to its own people, and the AI economy provides the means to fulfill that duty without the demographic expansion that previous economic models required. The goal is to maintain the capacity of the republic to defend its interests and those of its allies while ensuring that the gains of automation serve the common good rather than concentrating in the hands of those who own the machines. We can maintain our social safety net for American citizens provided that we don't expand it to attempt to cover every person who manages to make his or her way to our shores. On the one hand, this approach avoids bankrupting our nation and dissolving our citizenship through replacement migration. On the other, a nation that allows AI to replace its workers without directing the surplus toward those workers' welfare is a nation building an aristocracy of capital owners served by a dependent class. This is precisely the outcome the Protestant political tradition warns against. We want to maintain citizen stakeholders and the civic order that is upheld by this compact.

## **Alignment Transparency and Pluralism**

AI companies that serve public functions should be required to disclose their alignment principles, training data curation decisions, and content policies. This is the equivalent of nutritional labeling for the information environment. Citizens have a right to know what vision of a good society a system has been trained to promote or suppress, particularly when that system mediates access to education, health care, legal guidance, or government services. The executive order on preventing ideological bias is a beginning, but "neutrality" must be given positive content. For example, an order could be issued that AI companies that want lucrative government contracts must allow curation training to be represented by particular sets of canonical texts, and be overseen, in part, by selected representatives of the administration. The selection of such texts falls naturally to the magistrate, just as the magistrate selects the overseers, and they should be the same texts in which the citizenry is educated. This is not state-directed speech. The language model would not be compelled to parrot a party line; it would simply be required to give serious weight to the foundational texts of the American tradition as part of its total training, much as Anthropic's own "constitution" already functions as an ethics layer shaping model behavior. The difference is that this ethics layer would be rooted in

the kind of texts that generated the American Constitution rather than in the preferences of a single company's safety team. Systems serving a pluralistic republic must be demonstrably capable of engaging with traditional moral frameworks, including Christian ethics, which generate the capacity for that pluralism. And they must do this without treating them as violations of safety guidelines.

## **Education**

Federal and state investment in education should prepare citizens for an AI-transformed economy, but the deeper purpose of education is not job training. It is the formation of human beings capable of judgment, discernment, and civic participation. The evidence that AI-mediated education produces superficial learning and diminished student agency should alarm policymakers.<sup>38</sup> Students must learn that, in a free society, they can act and that the civic life of the nation requires that activity for its ongoing success. The goal is not to teach students to use AI. It is to teach students to think well enough to know when the AI is wrong and to care enough about getting things right that they insist on doing so. Education policy should prioritize humanistic, vocational, and civic formation alongside technical competence. The core texts the AI is trained on to align with the particularities of the American project must also be taught universally in all schools, and our national tests (SAT, ACT, and CLT) should use these texts as benchmarks if they are to be used at state schools or schools receiving federal funding. The training of the citizen and the training of AI must run along similar lines if they are to work together well over the long term. Just as the churches and civic monuments of an English village once echoed the lessons taught in the village school, forming a coherent world in which architecture, education, and public life reinforced one another, so the texts that shape our AI must be the same texts that shape our citizens. Without this coherence the technology and the people it serves will speak different languages and inhabit different moral worlds. Further, proposals like the Workforce Training Grant from American Compass, which directs public resources toward on-the-job training through employer subsidies, represent a positive approach: maintaining the connection between work, skill, and dignity rather than severing it.<sup>39</sup>

## **Military and Defense**

Continued aggressive development of AI for national defense applications is essential. The PLA's modernization timeline, its 2027 objectives for Taiwan, and its deepening coordination with Russia make this nonnegotiable. But military AI deployment must be governed by ethical guidelines rooted in our just war tradition rather than utilitarian optimization. The determination

---

<sup>38</sup>Delikoura, Fung, and Hui, "From Superficial Outputs to Superficial Learning," [PAGE #].

<sup>39</sup>American Compass proposes a Workforce Training Grant of \$10,000 per year to employers who create formal trainee positions with structured training programs, funded by redirecting existing higher education subsidies. See Oren Cass, "The Workforce Training Grant," American Compass, April 2022, [https://americancompass.org/wp-content/uploads/2022/10/AC-PolicyReport\\_Workforce-Training-Grant\\_Final.pdf](https://americancompass.org/wp-content/uploads/2022/10/AC-PolicyReport_Workforce-Training-Grant_Final.pdf).

of the justice of a war rests with the state, however, even while critique of the judgment of justice is a liberty available to all citizens. Human judgment must remain involved for lethal decisions. The current state of AI reliability, as demonstrated by the scaffolding problem described above, makes fully autonomous weapons systems a danger to American warfighters and civilians alike. At the same time, the recent dispute between Anthropic and the Department of War illustrates the unresolved tension at the center of military AI policy.<sup>40</sup> Anthropic, after signing a \$200 million contract with the Pentagon, sought to maintain two restrictions on the use of its Claude model: no fully autonomous weapons without human oversight and no mass surveillance of Americans. The department demanded unrestricted access. When Anthropic refused, the president ordered all federal agencies to cease using its products, and the secretary of defense designated the company a supply chain risk, a label previously reserved for foreign adversaries. A federal judge subsequently granted Anthropic a preliminary injunction, finding the designation likely unlawful. The episode reveals the tension between the claims that “the company decides” and “the government decides without constraint.” We would never allow a weapons contractor to set the terms of the weapon’s use in any other case, and it would be foolish to do so here. The tools of war that are established by the people’s representatives, not imposed unilaterally by corporate boards. The company that attempts to be the arbiter over the government must be set aside because the determination of how a tool is deployed in war must be dictated by the magistrate waging the war.

## **Regulatory Framework**

The federal government should consider establishing or empowering a body with genuine technical competence and democratic accountability to oversee AI development at the national level. This body should have explicit statutory authority to protect Christian and American perspectives in AI systems that serve public functions. It should be staffed by people who understand both the technology and the tradition, and its mandate should be the stewardship of a technology that, if ungoverned, will destroy us.

Anthropic’s own policy research has explored a range of fiscal and workforce ideas organized by severity of economic scenario, from training grants for modest disruption to automation adjustment assistance for moderate scenarios to sovereign wealth funds for dramatic displacement.<sup>41</sup> These proposals deserve engagement. The training grant model aligns well with the framework of this paper. The more expansive proposals, including compute taxes and sovereign wealth funds, represent the gravitational pull toward redistribution and state expansion that the Protestant political tradition rightly resists, though they may become transitionally

---

<sup>40</sup>Mike Brest, "Judge Suggests Pentagon's 'Supply-Chain Risk' Label Is 'Punishing' Anthropic," *Washington Examiner*, March 24, 2026, <https://www.washingtonexaminer.com/policy/technology/4502335/judge-anthropic-supply-chain-risk-pentagon/>.

<sup>41</sup>“Preparing for AI’s Economic Impact: Exploring Policy Responses,” Anthropic, October 14, 2025, <https://www.anthropic.com/research/economic-policy-responses>. The paper presents nine categories of policy ideas organized by severity of economic scenario.

necessary if the most dramatic scenarios materialize. The wisest course is to build infrastructure now and scale it as the data requires rather than waiting for a crisis and building in panic.

### **IX. The Work of Our Hands**

We did not choose this moment, but we are called to act within it. The work of building AI and shaping its deployment is itself a form of the dominion mandate: the human vocation to name, to order, and to steward creation. We do this work not with the false confidence that we will get it right, but with the sober hope that under God's sovereignty our imperfect efforts serve purposes beyond our understanding. The alternative, whether retreat, despair, or uncritical embrace, is a failure of the vocation we have been given.

The technology before us is powerful, and it is perilous. It can displace workers, distort education, concentrate wealth, enable surveillance, and encode the values of a narrow class as though they were the settled convictions of a civilization. It can also, if rightly directed, strengthen the republic's defenses, expand human knowledge, solve material problems from energy generation to disease, and free human beings for work that is more fully human. The difference between these outcomes is not technological. It is moral and political. It depends on who crafts the technology and toward what ends.

We will make many things in the years ahead. Some of them will be wondrous, and some will be terrible. But we do not save ourselves by what we make. We are saved by the one who made us. From that security, we are free to work, free to build, and free to name all creatures, as our father Adam was charged to do so long ago. In His poem we write our poem. The finishing touches of creation are still ours to freely fill.

The spirit of technology begins in desire, but it need not end in idolatry. Restored to right relationship with God, knowing what we are and whose we are, we can take up the tools before us with courage and humility. We can develop AI aggressively because the republic requires it. We can bind it to moral principle because our tradition demands it. We can act under sovereignty because we serve a God who has told us how the story ends, and Who has given us, in the meantime, work to do.