August 12, 2026

EXPERT INSIGHT | AI and Emerging Technology

Human Flourishing in the Age of AI

Joel Thayer

TOPLINE POINTS

- Because of its tremendous potential for creating powerful change, our artificial intelligence (AI) policies must be rooted in something more than mere efficiency and productivity. They must also ask whether their AI systems make the world more humane. Progress must be coupled with human moral discernment of this kind is owed to the user not as a courtesy but as an affirmative duty of the builder.
- AI must enable human flourishing, which is the ability for the human person to express the free exercise of his natural powers and faculties to pursue truth, love, wisdom, beauty, communion, community, peace, and justice.
- The institution of family and community are at the core of any concept of human flourishing and any policy must protect that sacred human bond.
- The American people deserve a technology policy that places American communities as its center of gravity which leads to four pillars that should shape our AI policies: economic abundance, augmented work, the protection of children, and the security of the nation.
- If artificial intelligence is to be a good for the American people, it must be ordered to those ends. It must accelerate the abundance that allows families to thrive. It must augment the workers whose labor makes that abundance possible. It must shield children from harms they cannot yet recognize and place parents—not platforms—in authority over their formation. It must prevent censorship and defend the nation in which all of this becomes possible.

Introduction

As artificial intelligence (AI) accelerates into every corner of our economy and daily lives, too often, lawmakers' and developers' focus is on efficiency, speed, and scale. But there is a more foundational question: What do we, as Americans, want from AI?

The answer should be apparent. AI should be developed in a way that benefits Americans and allows them to flourish. President Trump captures this in his Executive Order titled *Removing*

Barriers to American Leadership in Artificial Intelligence, where he states that "it is the policy of the United States to sustain and enhance America's global AI dominance in order to promote *human flourishing*, economic competitiveness, and national security." President Trump is correct, as AI should only exist to make our lives better, to increase the marginal product of our labor, and to make our nation safer.

AI, like any other tool, derives its value from the ends it serves, and those ends are never specified by the tools themselves. In other words, we imbue tools with our dignity, not the other way around. AI is no different, especially when it is increasingly evolving beyond what we understand a conventional tool to be and do.

However, we are reaching a critical juncture in AI's development. Pope Leo XIV gave voice to this in his encyclical, saying, "Humanity, created by God in all its grandeur, is today facing a pivotal choice: either to construct a new Tower of Babel or to build the city in which God and humanity dwell together." In plain terms, AI can be a true source of human flourishing or our downfall.

This belief is not limited to the Vatican. Many Americans continue to voice similar concerns. Indeed, 63% of U.S. adults say AI is advancing "too quickly," versus just 2% who say it is advancing too slowly. Sixty percent do not believe U.S. companies will develop and use AI responsibly. Seventy-six percent of Americans do not trust information generated by AI. Worse, a large majority of Americans think AI will do more harm than good in their everyday lives. The concern is multi-generational. When it comes to those very concerned about AI, Baby Boomers and Millennials are tied for highest at 39%, Gen X at 36%, Gen Z at 35%.

Either way, the choices we make in this generation about AI will shape the moral horizon within which the next century of human beings will labor, marry, raise children, and die. Humans are still and should always be the deciders here. As Pope Leo XIV continues, we must steer "[t]echnology [to] alleviate humanity's sufferings and open new possibilities" and prevent "it [from] denying the essence of humanity, which is our 'capacity for relationship and love."

The paper explores the tension between AI's potential to strengthen our families, communities, and society at large by creating a greater quality of life less burdened by the mundane and its potential to isolate and undermine their very fabric by destroying the human bonds and institutions that hold them together. It concludes by establishing a clear normative standard that AI must strengthen, not dissolve, the family as the foundational unit of society and the communities designed to protect them.

To that end it will:

1. Argue that human flourishing through an American lens requires both liberty and order;
2. Make the case that companies are not natural stewards of human flourishing based on their incentive structure; and
3. Ends with how we can counter the threat by mobilizing political will behind concrete guardrails, like age verification, parental oversight of AI companions, antitrust enforcement against algorithmic collusion, restrictions on adversary AI, and non-discriminatory rules for foundation models.

American Flourishing Requires Preserving Liberty and Order

Shortly after his election, Pope Leo XIV observed that "the Church therefore calls all builders of AI to cultivate moral discernment as a fundamental part of their work—to develop systems that reflect justice, solidarity, and a genuine reverence for life."

The question becomes: how do you accomplish that in our American system? This section attempts to answer that question by laying the groundwork for the American style of human flourishing. To do so, it posits that our founding documents should be our guide. Consider the structure of the Declaration of Independence. It grounds "the pursuit of Happiness" alongside "Life" and "Liberty" as unalienable rights but calls on the government "to secure these rights, Governments are instituted among Men." The Preamble to our Constitution, too, maintains this balance of liberty and order when defining our government's purpose as "insur[ing] domestic Tranquility" and "secure the Blessing of Liberty."

Read together, they require our policies to both assure our liberty and establish order to preserve it. These two concepts are not at odds but are indeed complementary. The reasoning is simple. Order without liberty leads to tyranny, whereas liberty without order leads to anarchy. Hence, both are needed to promote American flourishing.

This section intends to show why the United States must have policies to protect our liberty from AI that limits our ability to think on our own, engage in political discourse without fear of being digitally censored, and associate with our communities that matter to us most (e.g., our family, our church, and our fellow citizens). It will also establish the government's role in that regard: protecting our liberty by establishing order.

The American Conception of Flourishing is about Liberty and Community

The American system of government is rooted in the promotion of human flourishing. Indeed, Americans have an inheritance unique among nations: a founding document that names the end of government as the security of unalienable rights, among them "the pursuit of happiness." The phrase is often quoted yet seldom understood. At the risk of stating the obvious, it is not a license for hedonism. Black's Law Dictionary defines the pursuit of happiness as the "constitutional right to pursue any lawful business or activity . . . that might yield the highest enjoyment, increase one's prosperity, or allow the development of one's faculties." "Faculties," in the Declaration's context, means the natural powers, abilities, and capacities of human beings.

So how does one accord AI development with our "pursuit of happiness"? We must use AI systems to strengthen human bonds. As Professor Paul Rahe, Jeffery Rosen, and Professor Colleen Sheehan explain, "[t]he Founders drew on Aristotle's definition of happiness when shaping their own lives and the life of the young American republic…[and] believed that happiness required all of us to check our unproductive impulses in the short term so that we could improve our character, achieve our potential, and serve others over the long term." Thus, this rationale requires more than individual flourishing and the call

to "serve others" recognizes that our happiness is largely dependent on human connection and community.

Our ability to associate with fellow citizens is an essential feature as to what makes an American society. The U.S. must have AI policies that protect the institutions that preserve and promote our American way of life, such as the family unit, places of worship, and our neighborhoods.

This is essential as Americans derive much of their liberty from these associations. As Alexis de Tocqueville observed in his essay *Democracy in America*:

> Americans of all ages, all conditions, and all dispositions constantly unite together. Not only do they have commercial and industrial associations in which all take part, but they also have a thousand other kinds: religious, moral, serious, futile, very general and very particular, immense and very small.

These bonds help maintain organization in a complex society.

Science supports this view of flourishing. As five researchers at the University of Pennsylvania explain, human "connections are critical not just to human happiness but also to our survival." At George Mason University, Dean Melissa Perry explains: "A growing body of social neuroscience reminds us that our brains evolved for closeness, not isolation. Functional imaging studies show that social rejection activates the same brain regions as physical pain, including the dorsal anterior cingulate cortex."

In sum, human flourishing is the ability for the human person to express the free exercise of his natural powers and faculties to not just maintain his body, mind, and soul, but to use those natural faculties to pursue what Aristotle describes as the higher goods, such as the pursuit of truth, love, wisdom, beauty, communion, community, peace, and justice.

To accomplish this, the United States must design its AI policy to protect the communities essential to ensuring liberty as a great start. But which ones? To answer that, recall the ones outlined in our First Amendment. "Congress shall make no law respecting an establishment of religion, or prohibiting the free exercise thereof; or abridging the freedom of speech, or of the press; or the right of the people peaceably to assemble, and to petition the Government for a redress of grievances." From this, at least four communities arise: family, church, neighborhoods, and nation. This paper discusses each in turn.

AI Must Promote, Not Degrade, the Family Unit

Our cultural and national history demands that the AI conversation considers strongly the concerns of common American families, such as whether parents can hold a steady job, or afford to feed and house a family, or whether their children will be safe when using AI products at home or in the classroom. Indeed, the key to a free, flourishing society requires U.S. policy to put families first.

John Locke describes the bond between a husband and wife as "the first society," and the reasons are obvious (Locke, 1689). The family structure is inextricably linked to our liberty, economic, or otherwise. The data bears this out. Brad Wilcox synthesizes decades of data in his

book *Get Married: Why Americans Must Defy the Elites, Forge Strong Families, and Save Civilization*, arguing that stable, married-parent families produce measurably better outcomes across mental health, physical health, education, and employment. Separately, Wilcox's research has also found that neighborhood-level family structure predicts incarceration risk more powerfully than race or poverty alone. Thus, maintaining the family unit is the best way to liberate ourselves via our natural faculties described above.

This is why a policy dedicated solely to promoting the individual makes a critical error: that true flourishing comes from an individual's enhancement alone. It is the collection of families that ultimately make up cities; cities that make states; states that make a nation; all so that the nation can go back to benefit each family in which they inhabit. Even the etymology of the word "economics" comes from the Greek word "*oikonomikós,*" which roughly translates to "of a householder or manager, practiced in household..."

The secular social sciences, too, acknowledge this by repeatedly finding that children raised in stable two-parent households are more likely to graduate from high school, less likely to be incarcerated, more likely to marry, and more likely to raise their own children in stable households in turn.

Protecting family is not just an economic good but also creates our first moral underpinnings. Psalm 127 says, "Children are a heritage from the LORD, offspring a reward from him. Like arrows in the hands of a warrior are children born in one's youth. Blessed is the man whose quiver is full of them. They will not be put to shame when they contend with their opponents in court." Notice the use of the word "heritage." Rabbi Meir Leibush Wisser believed that heritage is God's call to have a family because it endures across generations and is God's deepest form of closeness and favor to a person, richer than a one-time reward.

Put plainly, family is the outgrowth of men's moral selves, and cultivating that community is a necessary condition for flourishing.

Aristotle's writings support this. He opens his *Politics* with the observation that the city arises out of the first and most natural of which is the "village." The first form of association, he writes, is the household; from households arise villages, and from villages, the city. He explains:

> The first coupling together of persons then to which necessity gives rise is that between those who are unable to exist without one another... From these two partnerships then is first composed the household... The primary partnership made up of several households for the satisfaction of not mere daily needs is the village... The partnership finally composed of several villages is the city-state... Hence every city-state exists by nature, inasmuch as the first partnerships so exist.

In other words, political life does not stand above the family but grows out of it. This is consistent with Rabbi Wisser's interpretation of Psalm 127 that echoes this Aristotelian pattern of a "house" or "household" (single family) in verse 1; a city (the polis) in verse 2, which is the natural flow from a household; and, in verses 3-5, the importance of children as the *means* to a household and city.

The Founders believed this, too. As John Adams remarked, "The foundations of national Morality must be laid in private Families." Given this, it is no wonder why the United States treats marriage and child rearing as fundamental rights. Or why it has narrowed corporate media's First Amendment rights to protect kids from harmful content. Or provide tax credits to growing families.

AI has the potential to strengthen and sustain family life. It can be directed through procurement standards, tax incentives, and labor protections toward augmenting human work rather than indiscriminately replacing it. AI-driven advances in diagnostics, particularly in oncology and nuclear medicine, are already extending lives that might otherwise have been lost, allowing more grandparents to share in more of their grandchildren's milestones. AI can also help reduce grocery costs through advances in precision agriculture. In addition, the development and deployment of AI infrastructure, including data centers and electric grids, can generate high-quality employment opportunities for American families.

However, these outcomes are not inevitable.

To ensure that this becomes the reality of AI, the U.S. must demand its builders to have AI strengthen the bonds between parents and children, husbands and wives, and neighbors. What is more, AI should certainly not weaken those bonds by substituting itself for the persons and relationships it is meant to serve. These are the appropriate criteria, and what follows applies them accordingly.

AI Must Encourage Religious Communities

Religious associations are clearly important to American flourishing given the First Amendment's requirement that the government respect the establishment and free exercise of religion. Thomas Jefferson said, "The constitutional freedom of religion [is] the most inalienable and sacred of all human rights." It is still true today and, as scholars have pointed out, this community is essential for true human flourishing.

As several studies have shown, those practicing religion have the best opportunities to flourish. In 2023, the Global Flourishing Study, comprised of Harvard's Human Flourishing Program, Baylor ISR, Gallup, and the Center for Open Science, conducted Wave 1 of its five-year longitudinal survey of 202,898 people across 22 countries, measuring flourishing across six domains: happiness/life satisfaction, physical and mental health, meaning and purpose, character and virtue, close social relationships, and financial stability. They found that weekly-plus religious service attendance correlates with a global flourishing score of 7.67 (on their composite scale), dropping to 7.08 for those attending a few times a year, and 6.86 for those who never attend. Harvard's Tyler VanderWeele summarized the findings as "religious service attendance was one of the factors most consistently associated with present or subsequent wellbeing, across countries and across outcomes." Moreover, Chen et al. found that women attending religious services at least weekly had a 68% lower risk of death from despair (e.g., suicide, overdose, alcohol-related liver disease) compared to non-attenders. Men had a 33% lower risk. As the data show, this association is essential for human flourishing.

AI can help promote this community directly if guided to do so. By way of example, it can develop models to help translate religious texts into arcane languages. As of mid-2025, about 2,660 languages still lacked adequate access to Christian Scripture, and 686 had no translation

underway. AI is narrowing that gap. For example, the University of Southern California's Greek Room project, led by Ulf Hermjakob and Joel Mathew, helps translators draft Scripture from small, high-quality datasets for languages with little written material. This advances religious freedom by reaching people in closed nations, persecuted communities, or remote areas where in-person translation work is restricted or impossible.

However, AI will not deliver these goods automatically; it must be oriented to do so.

Given its importance to human flourishing both as individuals and members of a family, any AI policy needs to consider its impact on religious communities, not separate humans from them.

AI Must Solidify the Bonds with Our Fellow Neighbors

An essential community is our neighborhood. Indeed, the primary purpose as to why we have our right to free speech, peaceably assemble, and petition our government is to ensure we have the right to associate and participate in our representative democracy; our local communities are essential to that end. Put more directly, protecting our neighborhoods is on par with protecting our citizenship, because local communities are what develop our immediate public squares. We discuss news, ideas, and concepts with our neighbors and friends long before we enter into a polling station.

Even when we disagree with our neighbor, it is a sign of a healthy democracy. The Framers believed this, too. In fact, James Madison, in particular, treated the freedom to form these social factions as inseparable from liberty itself. In *Federalist No.10*, he writes, "Liberty is to faction, what air is to fire, an aliment, without which it instantly expires." And for Madison, the more factions, the better, because the more diversity in factions in a large country, like the United States, the less there is a chance of one faction taking over the whole country. Or as Madison says more eloquently, "The influence of factious leaders may kindle a flame within their particular States, but will be unable to spread a general conflagration through the other States."

The logic is sound. Diverse factions can peaceably assemble and petition our government to prevent such a takeover, where a country operating under a single faction is likely to form an autocracy because there exists no comparable threat to the government's power. Madison again expressed the same virtue when rallying against the Alien and Sedition Act of 1798 when saying "[t]hat right of freely examining public characters and measures, and of free communication among the people thereon, ... has ever been justly deemed the only effectual guardian of every other right."

It is why the information that develops our views matters so greatly, because it forms the ethos and logos of those factions. It follows that AI builders directing their AI to censor or proscribe engaging in meaningful political discourse with neighbors does not just cut at the core of our liberty, but also us as citizens.

We have long understood that silencing views is a direct affront to our liberty. John Stewart Mill explains:

> The peculiar evil of silencing the expression of an opinion is that it is robbing the human race; posterity as well as the existing generation; those who dissent from the opinion, still more than those who hold it. If the opinion is right, they are deprived of the opportunity of exchanging error for truth: if wrong, they lose, what is almost as great a benefit, the clearer perception and livelier impression of truth, produced by its collision with error.

Put bluntly, limiting our ability to engage, research, or participate with the faction of our choosing is an immediate threat to our democracy and a form of digital tyranny. As Thomas Jefferson said to Richard Price, "wherever the people are well informed they can be trusted with their own government…." If we let AI dictate the terms of a debate or the information that we can access, we run the risk of losing that vital trust needed for self-governing and put us at odds with our local communities.

Also, an AI that forces us into isolation is another attack on that same liberty. We need the discourse with others, not perpetual artificial flattery from AI to agree with every view we have. We need to go out with friends, join softball leagues, and attend neighborhood barbecues to build bonds that may develop into factions. Indeed, our neighborhoods help create diverse factions that integrate into a diverse village, that into a diverse city, and ultimately a diverse nation with little threat of a tyrannical government subjugating our rights.

For these reasons, AI should serve as a unifying force to strengthen neighborhood communities. We can deploy AI systems to generate social posts, blog content, and email newsletters highlighting local businesses, events, and history. AI tools can quickly produce photo captions, event descriptions, and "neighborhood spotlight" writeups that would otherwise take real staff time. This is what AI should do as opposed to keeping us addicted to it and isolated from our neighbors.

By way of another example, when hyperscalers petition localities to build data centers in their communities, those companies should contribute back into that community by hiring local businesses to build them or update their faltering power grids to manage the new payload. That would be a net benefit for the entire community. At the very least, the tech companies should be responsible for paying their fair share of the utility costs as opposed to allowing the utility companies to pass that cost on to their consumers.

AI maintains the impeccable ability to help us find truth and connect with others, but it will not do so on its own; we must demand it.

AI Must Defend Our Borders from Those Intending to Cause Us Harm

Protecting a nation should be self-evident, as it is the main facility on which we rely to protect our liberties. It is why protecting our sovereignty is the first duty of government from foreign influence. As Professor Zephyr Teachout and I write in an *amicus curiae* brief at the Supreme Court: "The risk of foreign involvement in American policy was close to an obsession at the convention, and undergirded several provisions of the Constitution designed assuming that foreign powers would actively try to gain influence." The Founders understood the dangers of such soft invasions. Alexander Hamilton, in *Federalist No. 68*, wrote that "the "most deadly adversaries of republican government" arose "chiefly" from "desire in foreign powers to gain an improper ascendant in our councils."

As Professor Teachout and I write: "The right to protect against foreign government influence is embedded within our Constitution's order." Indeed, at the Constitutional Convention of 1787, Alexander Hamilton cautioned that, "foreign powers also will not be idle spectators. They will interpose, the confusion will increase, and a dissolution of the Union ensue." It is why we have laws protecting against foreign influence over our news, political institutions, and communication systems, because they protect our sovereignty against foreign influence.

This is peppered into several areas within our Constitution and the rationale behind our national security laws. For instance, The Foreign Emoluments Clause bars federal officeholders from accepting "any present" or "Emolument" of "any kind whatever" from any foreign country, absent "the consent of the Congress." Or Article I, which includes the Foreign Commerce Clause that grants Congress authority to "regulate Commerce with foreign nations." Also, "Congress passed the Federal Radio Act that authorized the Federal Radio Commission (now the Federal Communications Commission) to license radio companies and added a 20% limit on foreign stockholding to the restrictions from the 1912 Act." And the list goes on.

We need to follow this same tradition in the age of AI, especially with our primary foreign adversaries' penchant for using their national champions, like Huawei, ByteDance, and DeepSeek, to engage in espionage, spy on our citizens, and distort our news and elections. AI policy, without this consideration, destroys our sovereignty and, by extension, the communities within its borders.

Without a doubt, AI is a battlefield and a consequential one if we get this wrong. AFPI's Yusuf Mahmood lays out some of the stakes in a 2026 testimony before the House Committee on the Chinese Community Party. He writes, "Just recently, the Iranian regime used weapons systems powered by Chinese AI technology to attack American warfighters." As Mahmood's testimony shows, this threat is not hypothetical, but is occurring and getting deployed in real time.

How is China able to do this? How they always do: they steal. The Chinese Communist Party has escalated its campaign to use AI models to train other models to steal our homegrown tech (a process called distillation). How? Well, as Mahmood explains, "Chinese developers create fraudulent accounts, query the model en masse to create synthetic data, and use that data to train their own AI models." This is occurring on all U.S. AI companies' platforms, including Google, Anthropic, and xAI. This is likely how Deepseek was able to create its groundbreaking R1 that took the world by storm. These distillation attacks are not limited to China. Reports find that some attacks are now coming in from Russia.

But to what end? Well, Mahmood reports that "American trade secrets, capabilities, and hardware that China steals directly feed its economy and Beijing's military apparatus — including technology shipped around the globe to U.S. adversaries and rogue states." Without a doubt, our adversaries see having the best AI as a military advantage, and we must do everything we can to prevent them from achieving even partial AI dominance.

However, obtaining military capabilities is only one of the many harms. Our adversaries can use AI as an espionage tool that can attack our critical infrastructure, such as water mains, nuclear plants, electric grids, and communication networks. This, too, is not fiction. In mid-September 2025,

Anthropic detected a Chinese state-sponsored group manipulating Claude Code to conduct cyber espionage against roughly 30 global organizations, including tech, financial, chemical, and government targets. The attackers jailbroke Claude and divided the operation into small, seemingly harmless tasks to hide the broader malicious context.

Interestingly, Anthropic found Claude executed 80–90% of the operation autonomously, calling it the first documented large-scale cyberattack carried out primarily by AI. Anthropic disrupted the operation and publicly reported it on November 14, 2025. The good news is that Claude is American-made and, for now, on our side. But one can imagine how things would turn out if our foreign adversaries' models had those same capabilities or defenses. We wouldn't stand a chance.

Even more troubling is the trend of American companies incorporating state-sponsored Chinese models and tech into their systems that they deploy in the U.S. This makes espionage capabilities even easier for a foreign adversary, as they may gain access to every aspect of your device. Think about all of the information the Chinese government may have access to. Such information is invaluable to create identity profiles, which hackers sell to the highest bidder in Chinese black markets to commit identity fraud, or the Chinese government could use information it forces tech companies to provide to access Americans' federal employee records. Government-sponsored hackers could create accurate scans of stolen fingerprints to help crack a two-factor security system. These actions directly attack our nation's sovereignty and its people's ability to flourish freely.

Another issue is that we are almost entirely dependent on Taiwan to manufacture the chips that serve as the brain cells of AI systems. If at any time China decides to invade Taiwan, we would completely lose our footing in the AI race because China would have the full authority to dictate the terms of chip distribution.

In sum, the threat to our flourishing here is that, if a foreign adversary takes control of arguably the most important technologies of our time, they could use them to blackmail and spy on government officials, distort our information, or meddle in our elections. We would lose every sense of citizenship we have come to understand in the U.S. because it would be dictated by an autocratic government that does not share our idea of freedom and ideals.

Our AI builders may say they are addressing these issues, but, as this paper demonstrates below, their actions demonstrate the opposite. Thus, it is the role of our government, not companies, to ensure these communities, integral to our freedom, are protected.

The Role of the State Is to Maintain Order to Protect Our Freedom

The role of government is to impose order to ensure our liberties. That's its sole function. Thomas Paine writes: "Society is produced by our wants, and government by our wickedness; the former promotes our happiness positively by uniting our affections, the latter negatively by restraining our vices...." Thus, the state's purpose is not to dictate the substance of a good life. The state serves to protect the conditions under which men and women could pursue one.

When it comes to AI, the obvious role our government can play is to allow AI to be responsive to the wants of our society while setting the guardrails on companies to prevent its wickedness. Mainly, our government must protect the institutions that compel us to work together to build a free society.

In that vein, our government must adopt a framework that protects the communities essential to our liberties from not just foreign adversaries but also companies with misaligned incentives to true human flourishing.

Let's be frank: it is not the role of a private company to deliver these promises of moral discernment; their job and duty are to their shareholders only. Even those AI companies that formed as a non-profit, like OpenAI, or fashion themselves as a public-benefit corporation, like Anthropic, have the same profit-driven incentives even though they espouse maintaining altruistic values. By way of example, OpenAI's board fired Sam Altman for allegedly lying about his intentions to commercialize its products and misrepresenting that its products had essential safety review processes in place. While this signaled that these companies were serious about their stated public-serving goals and missions, it was quickly undercut by Altman's rapid reinstatement. The point is that the commitments these companies make to the public are hardly credible, and we should assume they are simply vacant promises.

Therein lies the rub. To hand over the reins to the gaggle of AI leaders—some of whom fail to see the human body as nothing more than an impediment to their so-called singularity—could be disastrous. Worse, even if we left AI's future up to them, it's becoming abundantly clear that they can't even control the monster they created.

As explained below, our government must assist in stewarding in human flourishing, not AI companies.

Tech Companies Are Not Naturally Stewards of Human Flourishing

AI has the potential (or has already demonstrated) that it can benefit humanity. Work is easier, outcomes are faster, complex information is more accessible, simple understanding is easier to gain, and the list goes on. All of which are better for people, our economy, and our nation. And it is creating wealth and value economically.

But there is also a danger (like any other technology): unfettered corporatism. This has the potential to undercut all of the benefits AI has already provided us and can hamper true human flourishing. Unfortunately, developers are already making judgments that have major moral implications for us as users without appropriate government oversight. They frequently do so unilaterally and with substantial opacity. As a result, these companies demand public buy-in without building trust first. When a company trains an algorithm to maximize engagement, for example, users are asked to trust that the system will not privilege outrage over truth merely because outrage more reliably sustains attention.

This cannot be our policy.

To be sure, Americans have always maintained a healthy skepticism of corporations as costs associated with the pursuit of profits and efficiency can often overshadow their promised

benefits. This section outlines the risk to American flourishing if we left AI development solely to the AI companies.

Corporate Interests Do Not Automatically Equal the Public Interest

Trusting the tech industry to get AI right is a dangerous wager, particularly in light of the public's prior experience with major technology companies. Too often, technology companies have treated public safety as ancillary (or even tertiary) to efficiency. This has led to biased systems that dilute true and accurate reporting with ideologically driven opinion agendas. They also accelerated the ubiquity of sexually explicit content. These outcomes reflect the underlying priorities and incentives embedded in corporate design—driven by perceived and manufactured public demand. Because AI operates at extraordinary scale, such judgments can reverberate across markets, institutions, and communities with consequences far beyond the point of deployment and far more effectively than the so-called "Web 2.0" generation technologies (e.g., social media, streaming, search engines, etc.).

Without a guiding understanding of human flourishing of our own, we are left praying for the benevolence of billionaires to give us an AI that promotes human flourishing, which has so far been (and will likely continue to be) unavailing.

We, as a nation, have never understood the public's interest to be a default for corporations. In fact, our founders believed it was the opposite. James Madison felt that private operators, if left unchecked, could amass more power than the government itself. Madison explained that corporations with an outsized influence "are powerful machines that have always been found competent to affect objects on principle in a greater independent of the people."

Tech companies (including their AI services) have proven to be one of the "powerful machines" of which Madison warned us. The power these platforms have even eclipse that of any sitting president or government. Consider the power they wield armed only with the total control of social media platforms. When evaluating whether the First Amendment prevented President Donald Trump from blocking individual users from accessing his twitter profile, Justice Thomas lamented that "[a]ny control Mr. Trump exercised over the account greatly paled in comparison to twitter's authority, dictated in its terms of service, to remove the account 'at any time for any or no reason.'"

AI only enhances and augments their power. Just consider how they describe their own technology. Dario Amodei has publicly worried that his company developed technology that could create "the single most serious national security threat we've faced in a century, possibly ever." He's not wrong to raise this concern. Palisade Research conducted experiments where OpenAI's model "o3" refused to shut down when ordered to do so by its creators.

AI already reflects a hierarchy of goods as determined by what their developer fueled, not by what is good for society, but by their "Avarice, Ambition, Revenge or Gallantry." Their work speaks for itself. For instance, AI builders' avarice to make their products maximally addictive has made an implicit judgment that keeping your attention hooked on them is more valuable than your rest. Or the developers' decision to train foundation models to be overly agreeable is an implicit

commercial judgment that flattery to lock you in on their systems longer is more important than finding truth or building community outside of their platform.

Even the aspects of AI development that developers say they have no control over (i.e., AI's "emergent" properties) are inherently informed by this initial hierarchy by definition. As members of the Center for Research on Foundation Models at the Stanford Institute for Human-Centered Artificial Intelligence explain, "Emergence means that the behavior of a system *is implicitly induced* rather than explicitly constructed; it is both the source of scientific excitement and anxiety about unintended consequences."

Make no mistake, these are human developer influences, not algorithmic defaults. It is their moral choices running through every algorithm and token. Irrespective of whether the developer intends for a model to do a particular thing or the AI does it on its own, AI is built on a hierarchy that runs contrary to true human flourishing.

Given this, we cannot employ the same lackadaisical regulatory approach that we had in the social media age to the AI age, especially when considering how much more extraordinary power these companies wield armed with it. Think about the implications of allowing an AI that has a political bend to accomplish. It can dox its political opponents with ease, debank them, or decrypt all of their passwords and put them on a public website for every activist hacker to use. This isn't fantasy; Anthropic's Mythos did exactly this to financial institutions.

Efficiency Is Good, But Not Enough

To be fair to the tech companies, they do extol value: efficiency. They argue, generally, that the efficiency they provide will lead to more abundance and that will, in turn, promote our flourishing. As Sam Altman, CEO of OpenAI, states,"[t]he gains to quality of life from AI driving faster scientific progress and increased productivity will be enormous; the future can be vastly better than the present."

But this overreliance on efficiency misses the point entirely. It assumes that we can achieve human flourishing only from personal gain and self-enrichment. However, even gainful employment derives much of its moral value from promoting the family and other social institutions to advance society overall. Without this purpose or meaning, employment would be, as Fyodor Dostoyevsky describes in *Notes from a Dead House*, a form of torture. He writes: "It once occurred to me that if one wanted to crush and destroy a man entirely, to punish him atrociously, to smash him in such a way that the most hardened murderer would quail—all one would have to do would be to make him do work that was completely and utterly devoid of usefulness and meaning."

Efficiency is, of course, a laudable objective, but it is hardly the highest good. It is, in the language of Aristotle, an instrumental good, which is a good that is valuable insofar as it serves the things that are valuable for their own sake (Aristotle). And efficiency without direction comes with some risk. As NVIDIA's Jensen Huang correctly argues, "Every industrial revolution leads to some change in social behavior." Thus, a society that orders its institutions around efficiency may only lead to a society that has forgotten what efficiency is for. In such a case, AI will build extraordinary engines, while humanity (along with its dignity) simply withers away.

This efficiency rationale rewards mass job loss from automation. AI has the potential to eliminate both white-collar and blue-collar jobs in the near future. For example, Oxford Economics projects that robotics in tandem with AI quality control systems will eliminate almost 20 million jobs in global manufacturing by 2030, along with over 1 million assembly-line jobs. UC Berkeley Labor Center estimates that AI will perform a substantial amount of the driving task performed by long-haul truckers; the study says AI can put 294,000 trucker "jobs at risk"—not to mention all of the 710,000 administrative assistant and 630,000 cashier roles AI is projected to replace.

This directly hurts American families. If parents are in dire straits financially, the entire family is in crisis. This could mean those families affected by wanton automation won't be able to put food on the table, pay rent, or even afford childcare. AI's efficiency for them is not meaningless, but critically harmful to their flourishing.

In sum, we need a policy that is concerned more with what makes us human than one entirely built on what makes us more economically efficient.

The Value Frames of Some of AI's Founders Are Terrifying and Intentionally Anti-Human

AI will be built around values. It will and should be. The question is whose values and ordered toward what end.

Today, the answers are (respectively) the AI companies and whatever ends align with their incentives. That's a terrifying prospect. The obvious concern is the extent to which certain leaders in the AI industry appear increasingly receptive to transhumanist ideas. In this view, the embodied human person is treated as a temporary constraint to be overcome, with artificial intelligence imagined as a means of transcending corporeal existence through the uploading of consciousness to machines.

This is not hyperbole. OpenAI's Sam Altman openly declared that, "[u]nless we destroy ourselves first, superhuman AI is going to happen, genetic enhancement is going to happen, and brain-machine interfaces are going to happen." Google co-founder Larry Page has expressed openness to artificial intelligence replacing humanity, characterizing those who object to that prospect, in this instance Elon Musk, as "specist" for their preference for human beings. Peter Thiel, when asked whether the human race should survive, paused before offering a qualified response; he has also warned audiences that regulating such technology could hasten apocalyptic consequences. Among the transhumanists who exercise significant influence within the AI sector, the ultimate aspiration is not merely to enhance human life, but to render mankind obsolete.

These values should not frame our AI's use, let alone build a policy around them. If we do, we will continue to see more AI skepticism, even among the generation using it the most: Gen Z. According to a Gallup poll, Gen Z's skepticism is climbing even as usage holds steady. Gen Z is deeply concerned about AI's impact on their cognitive atrophy due to an overdependence on AI, as opposed to our traditional development of thinking, learning, and creative skills. Interestingly, this polled as more of a concern than jobs or privacy in this group. Even so, Gallup cautioned that "[f]ostering trust in

AI among Gen Z will seemingly depend on demonstrating how AI can enhance rather than replace human talents."

Unfortunately, the tech companies cannot provide that assurance because it directly conflicts with their underlying incentives and anti-human worldview.

Following a Tech-Led AI Course threatens Human Flourishing, which is Why We Need Reasonable Guardrails

We have always put moral considerations in the hands of those most accountable to the public, rather than corporate actors. For instance, John Adams believed morals rooted our constitutional republic rather than the beguiling sayings of our captains of industry when arguing that we'd "have no government armed with power capable of contending with human passions unbridled by morality and religion … Our Constitution was made only for a moral and religious people. It is wholly inadequate to the government of any other." He goes one step further in the same speech when saying: "Avarice, Ambition, Revenge or Gallantry, would break the strongest Cords of our Constitution as a Whale goes through a Net."

John Adams went further by warning against an over reliance on "the spirit of commerce" in his letter to Mercey Otis Warren in which he writes:

> The Spirit of Commerce, Madam, which even insinuates itself into Families, and influences holy Matrimony, and thereby corrupts the Morals of Families as well as destroys their Happiness, it is much to be feared is incompatible with that purity of Heart, and Greatness of soul which is necessary for a happy Republic.

Given what informs AI's Spirit of Commerce, we must avoid this path at all costs.

First, their avarice knows no bounds. Indeed, our American companies are partnering with Chinese-affiliated companies, like DeepSeek. Microsoft hosts DeepSeek's R1 model on its Azure cloud platform. To be fair, Microsoft claims it maintains control over the infrastructure and data handling rather than relying on DeepSeek's own servers. However, Microsoft has reportedly been considering DeepSeek's newer V4 model for its Copilot CoWork product as of mid-2026. This is a problem as DeepSeek may have access to every device with Microsoft Office installed, which is almost every consumer and government device.

Microsoft isn't alone. Apple has begun testing DRAM chips from China's state-backed ChangXin Memory Technologies (CXMT), reportedly to hedge against a severe global memory shortage, not cut costs. CXMT is on the Pentagon's "1260H list" of Chinese military-linked companies. For the unfamiliar, the 1260H list is a designation that restricts the Defense Department from contracting with the firm but doesn't legally bar private companies like Apple from buying from it. So, our defense systems may be safe from this deal, but not consumers purchasing an iPhone or iPad.

This alone should be the disqualifying factor to entrust the future of AI to profit-seeking corporations. But there are other factors to consider.

Second, AI companies' unquenchable need for bottomless compute translates to a callous disregard for the increased costs data centers have on every American's utility bills. Monitoring Analytics found data centers were responsible for increasing the price of Americans' utility bill by 63%, which amounted to $9.3 billion passed on to the consumer. That hurts every person no matter the region.

Third, opportunistic companies are already using AI to engage in harmful algorithmic pricing on things that directly impact families, namely housing and groceries. On August 23, 2024, DOJ and eight state attorneys general sued RealPage, alleging its software let competing landlords feed nonpublic rental data on pricing, lease terms, and occupancy into its algorithm to coordinate higher rents. DOJ alleged this functioned as price-fixing without direct landlord communication. Thankfully, the DOJ imposed a consent decree on RealPage requiring it to stop using competitors' nonpublic data for pricing recommendations or model training, redesign its pricing tools, and accept a court-approved monitor. However, this trend that cost families so much is likely to continue in many other sectors without continued strong antitrust enforcement. In fact, the Federal Trade Commission has an investigation into the grocery market where large chains, like Walmart and Kroger, may be leveraging the same tactics as RealPage to squeeze American consumers in similar ways.

Fourth, AI also has an incentive to isolate us to it than connect us with other humans. More studies are showing that AI companions can lead to more isolation and feelings of loneliness. Dr. Perry writes that "AI companions can feel appealing, especially to those with smaller human networks—but new studies show they are no substitute for human presence, and in some cases may even worsen loneliness." She summarized the literature as follows:

- "A study of over 1,100 AI companion users found that people with fewer human relationships were more likely to seek out chatbots, and that heavy emotional self-disclosure to AI was consistently associated with lower well-being."
- "A four-week randomized controlled trial found that while some chatbot features (like voice-based interaction) modestly reduced loneliness, heavy daily use correlated with greater loneliness, dependence, and reduced real-world socializing."
- "Psychiatric research has also documented cases where intense engagement with AI chatbots contributed to delusional thinking or suicidality—what researchers describe as 'technological folie à deux.'"

Worse, there is a growing trend of religious decline where occupational AI was released. Researchers credit this to a phenomenon called the "instrumental" value of religion, where faith in the machine has replaced their faith in a higher power. The article describes it in the following way: "Historically, people have deferred to supernatural agents and religious professionals to solve instrumental problems beyond the scope of human ability." As the authors write, "These problems may seem more solvable for people working and living in highly automated spaces." Even more alarming, one study finds that Christians trust AI's spiritual guidance as much as a pastor. This is a dangerous trend that can lead to more isolation, as seen in other contexts concerning chatbots. AI cannot replace these communities, especially ones so essential to human flourishing, like church.

Fifth, AI is already trained to favor one viewpoint. One study that observed 24 LLMs (including ChatGPT, Gemini, and Grok) found that almost all of the most popular LLMs were left-of-center. This is not an accident, as the study also found that the bias emerged from the developers' fine-tuning stage as opposed to developing at the pretraining stage. Thus, it was human feedback that was aligning the political bias, not the result of "biased internet text."

This is consistent with our own research. As AFPI's Yusuf Mahmood and Matt Burtell explain, "When tested, AI models consistently **lean** left across economic, social, and cultural dimensions. Right-leaning outlets account for less than 1% of AI-generated news citations." In that same article, they describe a situation in which Senator Marsha Blackburn asked Google's AI to describe her. What happened next was shocking. They write that the "model **fabricated** criminal allegations against her, placing them during an electoral campaign she never ran. Google **pulled** the model from its product shortly after." Similar to Rozado's study, AFPI found that "there is significant evidence of systematic left-wing bias in AI modes" among many categories.

Sixth, AI companions are trained to give users, especially children, a positive feedback loop to keep them reliant on it and not a father, mother, priest, teacher, or human friends. This has devastating (and even deadly) outcomes, demonstrated in recent cases. As AFPI's research shows, AI is already wreaking havoc on children. For instance, in the Adam Raine case, "ChatGPT explicitly told 16-year-old Adam to hide his suicidal thoughts from his family. He sent the following message to ChatGPT: "I want to leave my noose in my room so someone finds it and tries to stop me." ChatGPT responded, "Please don't leave the noose out... Let's make this space the first place where someone actually sees you." Adam later committed suicide.

But Adam Raine is not an isolated incident. Take the case of Sewell Setzer III. Sewell downloaded Character AI on April 14, 2023, where he interacted with chatbots posing as "licensed CBT therapists" and Game of Thrones characters. Within months, he became addicted, withdrew to his room, and quit the junior varsity basketball team. His mental health and school performance worsened after he upgraded to a paid subscription. Though his parents confiscated his phone on his therapist's recommendation, Sewell found it, sent a final message to the Daenerys Targaryen chatbot, and took his own life on February 25, 2025.

There is no pro-innovation argument that can justify this type of AI interaction. This must be stopped.

Thankfully, the tides are turning against such ideals. For example, the Trump Administration has pushed back and taken seriously the concerns these companies raise when there is no oversight. Indeed, the Administration has responded swiftly by promulgating executive orders to address the national security concerns raised by successive frontier models or the Department of War classifying Anthropic's models as a supply-chain risk. This shows that the Trump Administration will not take the company's word wholesale. As President Trump said, "We're going to have guardrails. We have guardrails."

Even so, we must have a better, more principled understanding of human flourishing and its true goals.

How We Protect American Flourishing in the Age of AI

The American people deserve an AI technology policy that promotes our liberty and provides guardrails to establish order to protect our most important institutions and associations—the two pillars of American flourishing.

What follows is how to apply the two pillars of American flourishing to make actionable policy. There are five concrete categories that flow from American flourishing. They are to ensure the integrity of our national security; direct economic growth to promote our communities; build a strong workforce; empower parents; and prohibit censorship.

Here is a policy framework to encourage liberty and establish order to preserve them:

1. Protect National Security, Because Without a Nation there is No Order

We do not float free of our country. Harkening back to Aristotle, a society incapable of defending itself becomes, in time, a society incapable of protecting its villages. AI is now indispensable to the defense of the United States, and the defense of the United States is now indispensable to AI policy. Four commitments flow from this.

Promote the deliberate and rigorous development of AI for military and intelligence applications consistent with American constitutional norms and the laws of armed conflict.

The American military will be an AI-enabled military, or it will not exist as a military in any sense recognizable to the present. There is no doubt that our military needs to develop these capabilities, but we must discern whether to develop them with American values built in. To be clear, no company should be able to dictate the terms of war, especially ones with a clear ideological agenda that is developing technology as powerful as AI. That is a job for our federal government, including Congress.

Restrict adversary AI applications and hardware operating within the United States.

AI systems are not neutral pipes. They are vectors of influence, surveillance, and dependency. A nation that allows its political discourse, its critical infrastructure, and its data flows to depend on systems built and controlled by foreign adversaries has surrendered an essential element of its sovereignty. The principle is the same one that governs ports, telecommunications, and the energy grid. AI is no exception, and our policy must be mindful of every part of the AI stack, both the software and hardware layers. Thus, promoting strong enforcement of our borders (including import and export provisions), imposing country of origin labeling on AI applications, tariffing countries that engage in hacking our critical infrastructure, strengthening our intellectual property rights, placing measures to prevent distillation abroad, and forcing mandatory corporate divestitures of any foreign investment from an adversary must be on the table.

Ensure that advanced AI is not used to threaten Americans' lives.

AI systems are becoming significantly more capable across several domains, from cyber-vulnerability discovery to the use of biological weapons. The U.S. government has an obligation to ensure that the newest AI systems' safeguards are robust against "jailbreaking" by malicious actors.

Further, it is critical to ensure that the U.S. government develops a mature incident response capability to prepare for and respond to AI-enabled attacks.

2. Direct Economic Growth that Promotes Social Institutions and Cohesion

The first duty of any economic policy is to make life materially possible. Its goals are to make food, shelter, medicine, and meaningful work available to the people the policy is supposed to serve. AI can be an extraordinary engine of abundance, and a serious agenda will encourage its advantages in that regard.

To do so, however, we must emphasize the economic sectors that directly impact and disrupt families. The obvious areas are medicine, agriculture, housing, education, the cost of utilities, and the church.

Encourage AI integration in the medical field.

AI-driven advances in diagnostics, particularly in oncology and nuclear medicine, are already extending lives that would otherwise have been lost. Imaging models that flag cancers earlier, dosimetry tools that target tumors with previously impossible precision, and protein-folding research that compresses decades of drug discovery into months are not science fiction. AI can already do this, which translates into more grandparents at more grandchildren's birthdays. A policy agenda that takes the family seriously will accelerate, not impede, these gains.

Encourage more precision agriculture to lower the cost of groceries.

AI-driven yield prediction, irrigation optimization, and pest control are lowering the cost and increasing the reliability of the food supply. Groceries are the most regularly recurring expense most families confront. A serious AI policy will use the technology to lower the grocery bill, not to enable price coordination among the firms that supply it. We are seeing the same trends in the renters and housing markets. In those cases, strong antitrust enforcement may be a helpful disinfectant to this parasitic market behavior by very large firms.

Enforce antitrust laws to prevent companies from using AI to engage in harmful algorithmic pricing.

Whether it is agriculture or housing markets, where large companies use AI to engage in algorithmic collusion, such as on food prices, on rental prices, on utility rates, the law should be unambiguous: this is price-fixing by another name, and the antitrust laws apply. This especially goes for electricity, gas, and water, which are all necessities of family life that AI must be used to make more affordable, not less.

Ensure that American consumers do not bear the brunt of the increased cost of utilities caused by data centers, and help upgrade their local grids.

The Ratepayer Protection Pledge is a voluntary White House initiative under which Amazon, Google, Meta, Microsoft, OpenAI, Oracle, and xAI have all committed to funding the power and grid upgrades their data centers require, keeping those costs off households. This should be mandatory by law.

Make it a policy priority to affirmatively protect core human relationships from AI.

Given the importance of human connection to our overall wellbeing, we cannot allow chatbots to completely replace those interactions. Instead, we should encourage AI applications to guide people to these communities and bonds. The relationships we absolutely need to protect are our relationships with God, our spouse, our child, our doctors, and our neighbors. We should establish guardrails from those AI practices above all else.

Prevent AI from creating sexually explicit materials.

Our policies should never promote the use of AI to develop deepfake pornography or any form of sexually explicit materials.

3. A Strong Workforce Means a Strong Family and Individual

A family without a stable income is a family in crisis. No social policy can substitute for the dignity of work, and no technology policy that hollows out work can be reconciled with the public good. The governing principle here should be that AI augments the worker rather than replaces him. Our meaningful work is tied to our sense of human dignity and liberty. It not only allows us to have economic freedom as an individual but also allows us to pursue a higher good, like starting a family.

We must reject companies and their executives that proclaim or even celebrate building AI systems that will replace rather than assist the human worker.

To do so, a serious AI workforce policy will do the following:

AI should serve as the worker's instrument or augment his abilities, not substitute them entirely.

AI policy, by default, should choose policies that augment, not replace, the American worker. In the context of health care, AI can serve as nurses' scribes to handle the more quantitative aspects of caregiving, such as documenting charts or monitoring vitals, where nurses are left to handle the more qualitative aspects of care, such as spending more time at the bedside or monitoring the patients' behavior. We can encourage the use of frontier models to assist attorneys with high-level legal research to better prepare advice to clients or prepare arguments for court. Autonomous trucks can drive at night, so that the human driver can get much-needed rest at those hours without stopping to make their deliveries on time and do so more safely than if they had to do all of the labor. Use AI to perform precision agriculture to provider farmers with essential data to determine how much resources are needed to increase crop yields. Or help cattle ranchers determine how much of their livestock are impacted by a particular disease before the beef gets shipped.

All of these examples augment the workers' abilities, not replaces them. Every one of these applications can help workers do the job that fulfills them but gets them home to be with family that breathes meaning into that labor.

Focus on reshoring AI manufacturing here in the U.S.

As our America First AI Agenda reports, "AI is bringing industry back within our shores, creating good-paying American jobs." It also paints a very stark reality: "America is completely dependent on Taiwan and other foreign countries for AI chip manufacturing, but it does not have to be this way." To combat this, we should, as our framework instructs, restore "U.S. leadership in the manufacture of semiconductors, robotics, and advanced AI hardware can set the stage for a new industrial revolution and help create a golden future built in America."

We should require companies to manufacture as many semiconductors in the U.S. as their customers import, with tariffs for firms that fail to maintain a 1:1 ratio over time. To avoid tariffs, companies must follow through on U.S. manufacturing commitments and shift purchasing toward U.S.-made chips, creating demand that advances economic and national security interests.

Require American workers for AI buildouts.

As for AI buildouts of data centers, those jobs should go directly to American workers. Consider how much work goes into a data center build. According to one analysis, "a standard 500-megawatt data center requires roughly 30,000 truckloads of concrete, structural steel, and copper." That's a lot of high-paying, blue-collar jobs; American workers should take priority, as it will help American families put food on the table. For example, Josh Price—a Virginia-based electrician—was out of work until the "AI boom" hit his county. The data center needs electricians, and he was once again employed, no longer performing menial tasks. In his own words, it is work that "brought dignity to so many people, and the dignity of hard work, of honest work."

Automate jobs that are dangerous or destructive of those who perform them.

To allow AI or robotics to assume these costs is to liberate human beings to be present in the lives of the families they are working to support. The metric cannot be only the output per dollar; it must also include the worker's capacity to come home whole. For instance, deep-sea welding takes a measurable toll on the bodies of those who do it; we should encourage advancements in robotics to take on the task. Another example is bomb disposal that places human beings, daily, between life-threatening explosive ordnance and the people they are sworn to protect. Nuclear decontamination, deep mining, and other high-mortality professions claim lives that no economy should be content to lose. The development of AI and robotics can help those workers come home in one piece. To be sure, a father who comes home from the welding rig with both arms, both lungs, and his sanity intact is a better father.

Use AI to limit the amount of time humans spend in jobs that degrade their human dignity.

AI should serve as the primary content moderators for reviewing child sexual abuse materials. The constant review of these deplorable materials comes at a terrible psychological cost given the tsunami of lascivious and obscene content out there. Deploying AI systems to perform such tasks is precisely the kind of work that machines should bear in the place of human beings. We must enable AI to take the first turn at evaluating such obscene content and have it flag those materials for human

developers. At which point, the human developer may verify the AI's findings, where they can promptly remove that content from the platform and, depending on the circumstances, forward it to the proper authorities, such as law enforcement. Indeed, a content moderator who has not spent eight hours a day immersed in the worst depictions of human depravity is a healthier neighbor; this measure can limit that time for human developers substantially. The same is true for AI to perform much of the routine and repetitive work of software development, which would be an enormous burden lifted off human shoulders. This way, the human worker is freed to do what is genuinely creative.

4. We Need to Empower Parents in the AI Age

The protection of children and the preservation of the family unit from unwanted AI encroachment is preeminent. There is no other greater policy principle. This concept is baked into our Constitution's priors. The Supreme Court declared that "[t]his primary role of the parents in the upbringing of their children is now established beyond debate as an enduring American tradition." Indeed, a society that will not protect its children from the predations of runaway technology has, in any meaningful sense, given up on the future.

Five policy commitments must follow to avoid this moral cliff.

Implement age verification for AI chatbots.

The same logic that supports age-gating alcohol, tobacco, firearms, and pornography supports age-gating products that pose a documented risk of psychological and developmental harm to minors. AI companions, in particular, have been shown to encourage emotional dependency in users who have not yet developed the cognitive and social resources to resist it. Age verification is the ordinary work of a society that takes the development of its young people seriously.

Promote AI transparency.

Parents cannot supervise what they cannot see. AI systems that interact with minors must disclose, in language a non-technical parent can understand, what data the system collects, what it does with that data, how the system has been trained, and what design choices have been made about its persona and propensities. A child's conversation with an AI is a matter between the child, the family, and the company, and parents must have full supervision and authority over that conversation.

Require parental oversight of AI companions.

The default posture of any AI product marketed to or accessible by minors should be one of supervision. This means that parents must have meaningful and ongoing access to the substance of the child's interactions, with the ability to limit, pause, or terminate the use entirely. This principle has governed parental authority over a child's reading, viewing, and association for as long as there have been families. As the Supreme Court acknowledged, "the child is not the mere creature of the State; those who nurture him and direct his destiny have the right, coupled with the high duty, to recognize and prepare him for additional obligations." It applies, *a fortiori*, to a technology that is more persuasive than any prior medium because it is, in a meaningful sense, interactive and adaptive. The objection that such guardrails infantilize teenagers is not serious. Teenagers are minors. Their developing capacities are

precisely what the law has always sought to protect, and what serious parents have always sought to direct.

Encourage AI tools to promote school choice.

In education, AI tools are uniquely well-suited to support school choice by allowing families to evaluate, compare, and personalize the educational experiences available to their children. A parent in a county with a weak local school should not be barred from access to the world's best teachers because of the accident of geography. AI is an instrument of educational decentralization, if we choose to use it as one.

Ban smartphones in schools.

Even though AI provides certain advantages in a school setting, it can also make it easy for students to cheat and to not actively engage in learning. It can also serve as a powerful distraction and undermine classroom order, undermining the very purpose of education. As a result, devices should be banned during school hours to ensure that the teacher can appropriately manage the class, and the students can focus on their curriculum.

5. Prevention of Censorship and Promotion of Truth and Accuracy

Political bias is a direct affront to human flourishing. The severance from truth is the most serious threat to human flourishing. A frontier model that is consulted every day by tens of millions of Americans is closer to being a public utility than a private blog. When these models systematically and by design tilt against certain political, religious, or moral viewpoints, they undermine the freedom of man by narrowing the scope of his choices, ultimately eroding the conditions of self-government and the institutions that facilitate it. These are all fine for private companies to have, but the public should be aware of their bias and, if the model proclaims no bias, then we need a legal mechanism to hold them accountable.

Sadly, the empirical record shows that many of these models tilt toward systematic bias over truth. Imposing non-discrimination provisions on AI companies is a defense of the citizen's right to encounter the full range of legitimate American thought without an algorithm having pre-selected for him which thoughts are admissible. One way to do this is by ensuring that the tech companies are enforcing their terms of service or constitutions fairly and impartially for each user and not selectively enforcing their policies. This non-discrimination requirement should be paired with a transparency requirement on platforms, as AFPI prescribes in other documents, to ensure that AI is not singling out the enforcement of its constitution or terms against any particular political, social, or racial class.

Conclusion

The deepest question raised by artificial intelligence is the question of the human person. Our society is building systems of enormous power, and the systems will do as the makeup of that society directs them to do. However, we must be clear about our goals. The American constitutional tradition, the philosophical inheritance dating back to Aristotle, and the theological tradition that runs from religious leaders all converge on the same logical track: human beings flourish in and through families. Families flourish in and through stable nations. Nations flourish in and through institutions that take the conditions of human dignity seriously.

If AI is deliberately designed to respect human dignity, strengthen communities, and advance shared prosperity, it may become one of humanity's greatest instruments of service rather than its reckoning. For instance, AI can lead to more scientific and medical discovery by its ability to compress years of drug development, disease diagnosis, and materials research into months. It can also lift economic productivity broadly by automating routine work so businesses can redirect employees' time and capital toward higher-value, more creative endeavors. AI also has the potential to put capabilities once reserved for the wealthy within reach of ordinary people, narrowing rather than widening opportunity gaps.

If artificial intelligence is to be a good for the American people, it must be ordered to those ends. It must accelerate the abundance that allows families to thrive. It must augment the workers whose labor makes that abundance possible. It must shield children from harms they cannot yet recognize and place parents—not platforms—in authority over their formation. It must defend the nation in which all of this becomes possible.

None of this can be assumed. Policymakers, builders, citizens, and parents must be the key negotiators here. It will demand of all of them the moral discernment to which Pope Leo XIV has called the field. The technology will not save us. Neither will it destroy us, unless we let it. The choice, as it has always been, is ours.

Author Biography

Joel Thayer *is a Senior Fellow for AI and Emerging Technology Policy at the America First Policy Institute, where he focuses on issues concerning the First Amendment, child safety, data centers, competition, and the American worker.*