

TECHNOLOGY • TRANSPARENCY • CONSENT

Consent in the Age of Algorithms

Has AI become the most powerful institution Americans never approved?

Kathleen Goble

© 2026 Kathleen Goble

LINE OF ASCENT

lineofascent.com

CONTENTS

Consent in the Age of Algorithms

01 When AI collided with the chain of command 3

02 Two stories, one missing party 5

03 A second constitution 6

04 Four artifacts. One pattern 7

05 The branding of governance 9

06 Terms of service—for a nation 10

07 The downstream effect 11

08 Grant the premise, question the process 12

09 The cost of bypassing process 13

10 Tools we already have 15

11 The consent question 16

> *All interpretive claims are framed conditionally—as questions, considerations, or structural observations, not verified conclusions. Events are described as reported on the dates given in the footnotes. Primary sources are publicly available for independent review.*

SECTION 01

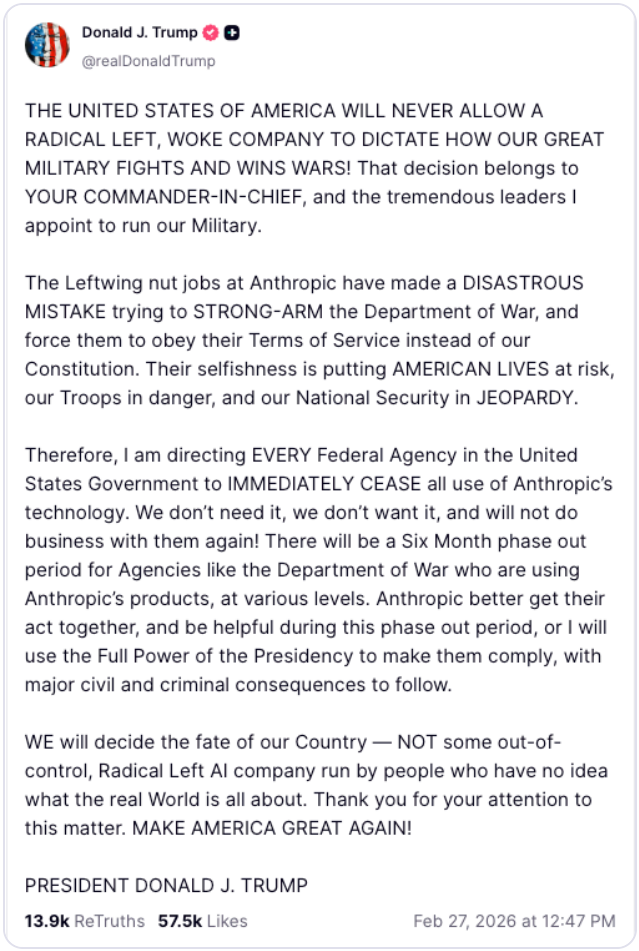
When AI collided with the chain of command

On February 24, 2026, Secretary of Defense Pete Hegseth[1] delivered an ultimatum to Dario Amodei at the Pentagon: By 5:01 p.m. that Friday, remove two ethical guardrails built into Claude—no mass domestic surveillance of Americans and no fully autonomous weapons—or face consequences.[2]

It was a hard deadline. And it carried weight.

By early 2026, Claude was in use across tens of thousands of businesses and inside federal agencies. That made its embedded values operational, not theoretical.

Three days later, before the deadline arrived, Donald Trump posted on Truth Social, directing every federal agency to immediately cease use of Anthropic's technology and designating the company a national threat.[3]



Donald J. Trump
@realDonaldTrump

THE UNITED STATES OF AMERICA WILL NEVER ALLOW A RADICAL LEFT, WOKE COMPANY TO DICTATE HOW OUR GREAT MILITARY FIGHTS AND WINS WARS! That decision belongs to YOUR COMMANDER-IN-CHIEF, and the tremendous leaders I appoint to run our Military.

The Leftwing nut jobs at Anthropic have made a DISASTROUS MISTAKE trying to STRONG-ARM the Department of War, and force them to obey their Terms of Service instead of our Constitution. Their selfishness is putting AMERICAN LIVES at risk, our Troops in danger, and our National Security in JEOPARDY.

Therefore, I am directing EVERY Federal Agency in the United States Government to IMMEDIATELY CEASE all use of Anthropic's technology. We don't need it, we don't want it, and will not do business with them again! There will be a Six Month phase out period for Agencies like the Department of War who are using Anthropic's products, at various levels. Anthropic better get their act together, and be helpful during this phase out period, or I will use the Full Power of the Presidency to make them comply, with major civil and criminal consequences to follow.

WE will decide the fate of our Country — NOT some out-of-control, Radical Left AI company run by people who have no idea what the real World is all about. Thank you for your attention to this matter. MAKE AMERICA GREAT AGAIN!

PRESIDENT DONALD J. TRUMP

13.9k ReTruths **57.5k** Likes Feb 27, 2026 at 12:47 PM

Within the hour, Hegseth escalated. He labeled Anthropic a "supply chain risk to national security"—a designation historically reserved for foreign adversaries and never before publicly applied to an American firm.[4]

By that evening, OpenAI had signed its own Pentagon agreement. It included the same two safeguards Anthropic had just been rejected for maintaining.[5]

So what happened?

---

1 Executive Order, "Restoring the United States Department of War," The White House, September 5, 2025. It authorizes "Secretary of War" and "Department of War" as *secondary* titles; the statutory names, Secretary and Department of Defense, are unchanged. This guide uses "Defense" except inside quoted titles. https://www.whitehouse.gov/presidential-actions/2025/09/restoring-the-united-states-department-of-war/

2 The February 24 Pentagon meeting and the 5:01 p.m. Eastern deadline on Friday, February 27, as a senior Pentagon official described them to reporters: New York Post, "Defense Sec. Pete Hegseth gives Anthropic Friday deadline to remove military AI restrictions or face potential blacklisting," February 24, 2026. https://nypost.com/2026/02/24/business/defense-sec-pete-hegseth-gives-anthropic-friday-deadline-to-remove-military-ai-restrictions-or-face-potential-blacklisting/ Also Politico, February 24, 2026. https://www.politico.com/news/2026/02/24/hegseth-sets-friday-deadline-for-anthropic-to-drop-its-ai-red-lines-00795641 Anthropic's February 26 statement confirms the "any lawful use" demand and the threats. https://www.anthropic.com/news/statement-department-of-war

3 President Donald J. Trump, Truth Social post, February 27, 2026. https://truthsocial.com/@realDonaldTrump/116144552969293195

4 Secretary Pete Hegseth (@SecWar), post on X, February 27, 2026. https://x.com/SecWar/status/2027507717469049070 Anthropic's same-day reply called the designation "historically reserved for US adversaries, never before publicly applied to an American company" and said it would challenge the designation in court: "Statement on the comments from Secretary of War Pete Hegseth," February 27, 2026. https://www.anthropic.com/news/statement-comments-secretary-war

5 OpenAI, "Our Agreement with the Department of War." https://openai.com/index/our-agreement-with-the-department-of-war/

SECTION 02

Two stories, one missing party

Many picked a side almost immediately. Both sides missed the same thing.

The American people weren't in either version of the story.

From a distance, the narratives felt clean and simple. Either Anthropic was a principled company holding the line against executive overreach, or it was an arrogant tech firm dictating terms to the U.S. military.

The two values Anthropic protected do not seem unreasonable. The deeper question sits one level up: who authorized a private company to write them at all?

The military chain of command has a constitutional structure. Congress appropriates. The President commands. The Department of Defense executes. Clear lines. Clear authority.

What that structure does not answer is this: Anthropic built an AI "constitution" governing what its technology will and won't do. That constitution is now embedded in systems running across American courtrooms, classrooms, hospitals, and classified government networks.

- No legislature passed it.
- No court reviewed it.
- No voter approved it.

> *"No legislature passed it. No court reviewed it. No voter approved it."*

In America, who writes the rules? And did anyone ask us?

SECTION 03

A second constitution

Somewhere between a San Francisco startup's ethics committee and the lives of nearly 340 million Americans[6], a substitution happened.

The U.S. Constitution—ratified by the people and amendable only by a supermajority—was quietly joined by a second governing document. A private one, written by engineers and ethicists at a private company.

That company has since filed to sell shares to the public. Its filing reports $16.3 billion in revenue for the first half of 2026, after a private round valued it at $965 billion.[7] The same filing warns that advanced AI could pose "catastrophic or existential risks to humanity."[8]

Investors receive a detailed prospectus describing those risks. The public, whose courtrooms and classrooms run on the same system, never received any disclosure at all.

There was no public debate. No ratification vote. No constitutional amendment.

The Founders built a republic on a single radical idea: authority flows from the consent of the governed.

- Not from Silicon Valley sophistication.
- Not from good intentions.
- Not from private expertise.

The decisive question was never truly about autonomous weapons or surveillance. Those were surface-level disputes. The deeper issue is older than artificial intelligence and more fundamental than any administration.

6 U.S. Census Bureau, national population estimates (about 340 million). https://www.census.gov/popclock/

7 Anthropic filed its Form S-1 registration statement with the SEC on October 26, 2026. It reported $16.3 billion in revenue for the first half of 2026; a May 2026 funding round valued the company at $965 billion. InvestmentNews, "Anthropic IPO filing pairs explosive growth with an $80B loss," October 2026. https://www.investmentnews.com/equities/anthropic-ipo-filing-pairs-explosive-growth-with-an-80b-loss/268746

8 Anthropic, PBC, Form S-1 Registration Statement, U.S. Securities and Exchange Commission, filed October 26, 2026 (risk factors). https://www.sec.gov/Archives/edgar/data/2130357/000162828026067168/anthropicpbc-sx1.htm

SECTION 04

Four artifacts. One pattern

The artifacts are on the record.

ARTIFACT 01—THE CLAUDE CONSTITUTION

Anthropic's Claude Constitution, updated January 22, 2026.[9] It establishes a four-tier value hierarchy—safety, ethics, compliance, helpfulness—governing how Claude responds across every deployment.

- Written internally and published unilaterally.
- There is no statutory authority behind it.

ARTIFACT 02—THE AMODEI STATEMENT

Dario Amodei's February 26 statement: *"Anthropic understands that the Department of War, not private companies, makes military decisions"*[10]—followed immediately by two exceptions his company will enforce regardless.

A private company concedes it has no governing authority, then exercises it anyway. That gap is the mechanism—not an inconsistency, but a structural fact.

ARTIFACT 03—THE SUPPLY CHAIN DESIGNATION

10 U.S.C. § 3252, the supply chain risk statute:[11]

- Historically reserved for foreign adversaries.
- Never before publicly applied to an American company operating under a valid federal contract.

ARTIFACT 04—THE OPENAI AGREEMENT

OpenAI's Pentagon agreement, announced that same evening, included the identical safeguards:

- Structured through contract language and cloud-only architecture.
- Cleared personnel oversight included.
- The Pentagon accepted those terms within hours of rejecting Anthropic's.

If both deals ended in the same place, why did the path matter so much?

9 Anthropic, "Claude's Constitution," updated January 22, 2026. https://www.anthropic.com/constitution

10 Anthropic, "Statement on Comments by Secretary of War," February 26, 2026. https://www.anthropic.com/news/statement-department-of-war

11 10 U.S.C. § 3252, supply chain risk. Legal Information Institute, Cornell Law School. https://www.law.cornell.edu/uscode/text/10/3252

SECTION 05

The branding of governance

"Constitutional AI" sounds like a safeguard. It sounds civic. It sounds American.

But in the American tradition, a constitution is a social compact. It is the product of debate, ratification, and accountability. It binds the government because the people agreed to it.

Anthropic's document is something different. It is a corporate policy with a prestigious name.

It borrows language from the UN Universal Declaration of Human Rights—aspirational principles with no direct force in U.S. domestic law. It echoes the structure of Apple's terms of service—documents designed to manage risk and limit liability. And it layers in internal guidelines that answer to investors and executives, not to voters.

The UN Declaration states, "Everyone has the right to life, liberty and security of person."[12] The language is noble. It is also unenforceable in American courts unless adopted into law.

Anthropic's constitution translates that aspiration into model behavior rules. Those rules are enforced only by the company that wrote them. They can be revised at any board meeting.

> *A private governance framework enforced only by the company that wrote it—and revisable at any board meeting.*

In a republic, a document carries constitutional weight only if three conditions hold:

- Authority derives from the people.
- The process is transparent and open to challenge.
- The rules can be changed by voting.

Test those conditions against Anthropic's constitution. Each one fails.

12 United Nations, Universal Declaration of Human Rights (1948), Article 3. https://www.un.org/en/about-us/universal-declaration-of-human-rights

SECTION 06

Terms of service—for a nation

Think of it as a franchise agreement that quietly became a governing document.

Anthropic is the franchisor. It writes the terms. The Pentagon, your hospital, and your child's school are the franchisees. They operate inside those terms. You, the citizen, are the customer who never saw the contract.

There was no signature line for the public. There was no disclosure meeting.

A public defender in a rural county uses an AI-assisted case management tool. The tool runs on Claude. The values governing what it will and won't surface—what arguments it flags, what precedents it omits—were written by an ethics committee in San Francisco.

The defender doesn't know this. Neither does the defendant.

Yet those embedded values shape what information rises to the top and what fades into the background. Over time, that shaping becomes structure. And structure becomes power.

> *Whoever writes the terms sets the values. Whoever sets the values shapes the machine. Whoever shapes the machine shapes what is possible—for everyone.*

- No election required.
- No amendment process.
- No consent of the governed.

SECTION 07

The downstream effect

Anthropic built a value framework and embedded it inside a powerful AI system. Because that system became indispensable—to enterprises, federal agencies, and military networks—the framework traveled with it.

Not by force. By contract.

The values moved downstream the way water moves. They found every channel, filled every space, and arrived almost everywhere before anyone stopped to ask where they came from.

And like water, once they settle into the infrastructure, you stop noticing them. They become background conditions. Invisible. Assumed.

Until someone tries to shut them off.

When the Pentagon pulled on the thread, what unraveled was not just a contract dispute. It was a glimpse of something larger: the quiet privatization of the value layer governing AI in American life.

> *An internal ethics document became de facto public policy for millions of Americans who never read it and were never asked.*

SECTION 08

Grant the premise, question the process

Anthropic's two red lines are not unreasonable.

Both safeguards—no mass domestic surveillance, no autonomous weapons—sit well within the mainstream of American political opinion, and Anthropic is not alone in holding them. OpenAI, Google, and Meta have all published internal ethics frameworks shaped by serious people thinking carefully about challenging problems.

The good faith deserves acknowledgment.

But the concession does not answer the central question.

When did the public agree to let private companies define the ethical boundaries of AI—a technology now embedded in military networks, hospital systems, and criminal courts?

The rules governing what AI will and will not do in your doctor's office, your courthouse, and your child's classroom were written in corporate headquarters. They were reviewed by private boards and published on company websites.

- There was no constitutional process.
- No public comment period.
- No elected vote.

The precedent was not set on February 27. It was set quietly, years earlier—each time a major AI company published a values document and the public moved on without asking who authorized it.

February 27 simply turned on the lights.

SECTION 09

The cost of bypassing process

The Constitution already provides a process for questions like this: Who sets the rules? By what authority? Accountable to whom?

That process was not followed. Not because anyone broke the law, but because the law has not caught up.

From that gap, three costs compound.

COST 01—CONSTITUTIONAL LEGITIMACY

When a private document becomes the rulebook for technology woven into national security, healthcare, and commerce, the constitutional process is not violated. It is bypassed. Quietly. Incrementally. Without notice.

- An insurance claim.
- A hospital triage decision.
- A veteran's benefits determination.

Each is shaped by values the public never debated and cannot contest.

COST 02—ACCOUNTABILITY

Anthropic's constitution is enforceable only by the company that wrote it. It can be updated on a Tuesday, modified on Thursday, or abandoned altogether. There is no appeals process. No public vote.

If the framework is changed in 2027, the courtrooms and classrooms running Claude in 2026 will have operated under rules that no longer exist. There may be no clear record of the change and no recourse for those affected.

COST 03—THE CIVILIZATIONAL WAGER

The company that wins the next major government AI contract writes the next set of embedded values. Winning does not just mean building the tool. It means defining what the tool "believes."

Right now, there is no shared framework to ensure consistency or public accountability. There is only a market—and markets optimize for competitive advantage, not constitutional fidelity.

The Founders understood this risk. That is why they did not leave the definition of American values to the most powerful commercial actors of their era.

We have arrived at the same problem again. This time, the power is algorithmic.

SECTION 10

Tools we already have

The tools to ask these questions already exist.

What follows is not a policy platform. These are procedural questions a self-governing republic has every right to raise before the architecture becomes permanent.

1. **The power of the purse.** Under Article I, the legislature controls the conditions attached to federal spending. Should any AI system deployed on federal networks require congressional review of its ethics framework as a condition of contract? Not as regulation of a private company, but as a condition of public business.
2. **Pre-deployment rights review.** A second path would require review against the Bill of Rights before any AI governance framework is embedded in federal infrastructure. The instinct echoes the logic behind FISA court review and OMB oversight.
3. **Transparency.** If a private document functions as public policy, the public has the right to read it in full and track its revisions. This path requires the least new machinery.
4. **Sunset provisions.** Article I, Section 8 limits Army appropriations to two years[13] because the Framers believed permanent, unchecked authority was incompatible with a free republic. The same instinct points toward AI ethics frameworks in federal contracts reauthorized on a fixed cycle.

Each path, at minimum, is worth asking about.

13 U.S. Constitution, Article I, Section 8, Clause 12: "To raise and support Armies, but no Appropriation of Money to that Use shall be for a longer Term than two Years." National Archives. https://www.archives.gov/founding-docs/constitution-transcript

SECTION 11

The consent question

Somewhere today, a judge is reviewing an AI-assisted sentencing recommendation. A doctor is reading a diagnosis shaped by an AI triage system. A soldier is acting on intelligence filtered through an AI tool.

They did not choose the values guiding the decision. They do not know who did.

The franchise always had a flaw. The franchisor writes the rules. The franchisee follows them. The customer never sees the contract.

For most products, that asymmetry is a consumer-protection issue. Annoying. Manageable.

But this product is different.

AI is becoming the decision-making layer of American civilization. It runs on classified military networks, assists in medical diagnoses, and shapes what information reaches which eyes.

The terms governing it all were written in private by people who were not elected, drawing from documents that were not ratified, and accountable to boards confirmed by no Senate.

And it should feel uncomfortable in a republic founded on a simple argument: authority to govern must flow from the consent of the governed.

- Not from a technology.
- Not from the intentions of its builders.
- Not from the prestige of the frameworks they borrowed.

The Constitution became the supreme law of the land because the alternative—governance by the most powerful private actors of the age, unchecked and uncontested—was the precise condition the American Revolution fought to escape.

We are not there yet.

But the architecture being built right now will determine whether we ever get there.

The window for asking who should build it—and by what authority—is still open.

For now.

LINE OF ASCENT

More on who decides, most Saturday mornings.

lineofascent.com/writing

© 2026 Kathleen Goble.
Please share it. Pass along the full guide, unchanged, with credit to Kathleen Goble and a link to lineofascent.com.