FOR RELEASE MARCH 2026

# Five Americas of AI Concern

## SUMMARY OF FINDINGS

*Seventy-three percent of Americans are concerned about AI, but they worry about very different things. A latent class analysis reveals five distinct segments within the concerned public.*

***BY Justin Miller***

**MATS 9.0 Fellow**
Mentor: David Krueger
Research Manager: Elise Racine
**fiveamericas.ai**

RECOMMENDED CITATION

Miller, J. (2026). *Five Americas of AI Concern*. ML Alignment Theory Scholars (MATS). Available at fiveamericas.ai

ABOUT THIS RESEARCH

# About This Research

This report presents findings from a nationally representative survey of 2,735 U.S. adults conducted in February 2026. The study was designed and fielded as part of **ML Alignment Theory Scholars (MATS) 9.0**, a research fellowship program focused on AI safety and alignment. The principal investigator, Justin Miller, is a MATS fellow working under the mentorship of David Krueger (University of Cambridge) with research management by Elise Racine.

The research has two goals. First, to map the landscape of American public opinion on artificial intelligence with greater precision than existing polls allow. Most prior surveys ask whether people are "concerned" or "not concerned" about AI and report a single topline number. We argue this obscures more than it reveals. A person who worries that AI will take their job and a person who worries that AI will pose an existential threat to humanity are both "concerned about AI," but they need fundamentally different information, reassurance, and policy responses.

Second, to identify which governance arguments resonate with which segments of the public, and why. The current policy debate around AI safety often assumes a monolithic public. Our message experiment shows this assumption is wrong: different segments respond to different framings, even when they converge on similar policy preferences.

The study uses **latent class analysis (LCA)** to identify five statistically distinct segments within the concerned American public. Each segment represents a coherent cluster of attitudes, concerns, and trust orientations. We label them not to reduce individuals to types, but to give researchers, policymakers, and the public a shared vocabulary for the distinct orientations within this population.

This report is modeled on the *Yale Program on Climate Change Communication's* "Six Americas" framework, which demonstrated that segmenting public opinion on a complex issue reveals heterogeneity invisible to aggregate polling. We apply a similar methodology to AI.

METHODOLOGY

# How We Did This

This report is based on a survey of **2,735 U.S. adults** conducted via the Prolific research platform in February 2026. The survey used demographic quota sampling to approximate the U.S. adult population on age, sex, ethnicity, education, and Census region. After applying a screener for baseline AI concern (those who indicated they were "not at all concerned" and had no opinion on AI were excluded), the analytical sample comprises **2,000 respondents** who expressed at least some concern about AI as a national issue.

Respondents completed a 15-minute survey covering:

- Concern about AI relative to other national issues (Q1-Q2)
- Specific AI concerns, selected from 14 categories (Q3a-Q3b)
- Personal AI use and emotional orientation (Q4-Q5)
- Knowledge self-assessment and factual knowledge checks (Q6)
- Trust in institutions to handle AI responsibly (Q7)
- Views on AI development pace and regulatory preferences (Q8-Q9)
- Response to one of nine randomly assigned policy argument framings (Q10)
- Trade-off questions on safety vs. international competitiveness (Q12)
- Barriers to engagement with AI policy (Q15)
- Open-ended reflection on AI concerns (Q16)

We used **latent class analysis (LCA)** to identify subgroups within the concerned population. LCA is a model-based clustering technique that identifies groups of respondents who share similar patterns of responses across multiple survey items. We used 25 attitudinal indicators drawn from concern intensity, concern type, trust, regulation preference, and emotional orientation questions. Models with K=3 through K=8 classes were tested. The five-class solution was selected based on model fit (BIC), classification quality (entropy), interpretability, and stability across 500 random starts.

For the argument experiment, each respondent was randomly assigned to evaluate one of nine policy argument framings. They rated the argument on a 4-point convincingness scale and a 5-category tone assessment. Sample sizes per argument range from approximately 210 to 235.

Here are the topline questionnaire and the interactive data at fiveamericas.ai.

# Table of Contents

OVERVIEW

# Overview

## KEY FINDINGS

- **Nearly three-quarters (73%) of Americans are at least somewhat concerned about AI** as a national issue, placing it alongside climate change in overall concern levels.
- **The partisan gap on AI concern is just 13 points** (79% Democrat, 65% Republican), compared to 50 points for climate change. AI may be the last major technology issue where bipartisan engagement is possible.
- **Despite widespread concern, AI ranks last** among eight national issues in personal priority. Only 4% rank it first. The activation gap between registering concern and treating AI as urgent is the defining feature of current public opinion.
- **Using latent class analysis on 25 attitudinal indicators,** we identify five distinct segments within the concerned public. They differ not just in intensity but in what they worry about, who they trust, and what they think should be done.
- **Eighty percent of concerned respondents say the government should prioritize AI safety rules** even if it means developing AI more slowly than China. This holds across every segment and party, including 68% of Republicans.
- **Of nine policy argument framings tested, children and family safety** was the only frame rated "very convincing" by a majority of every segment (44–72%). Different groups respond to different framings, but they often converge on the same policy preferences.

Artificial intelligence has reached a peculiar inflection point in American public life. Most people are concerned about it. Very few consider it a priority. Almost no one trusts the institutions that are building it or regulating it. And yet, when you ask what should be done, large majorities across the political spectrum converge on surprisingly similar answers.

This report presents results from a survey of 2,735 U.S. adults conducted in February 2026. We find that the American public's relationship with AI is defined by three tensions: concern without priority, distrust without disengagement, and consensus without activation. **73%** of Americans express at least some concern about AI—a level comparable to climate change. But only **4%** rank AI as their top national priority, placing it dead last among eight issues tested. The gap between concern and salience is vast.

To move beyond the single topline number, we used latent class analysis to identify five distinct segments within the concerned public. These segments are not based on demographics, though demographics differ across them. They are based on *what people worry about, how intensely they worry, who they trust, and what they want done.* A 38-year-old software engineer who worries about existential risk from superintelligence shares a segment with a 62-year-old retired teacher who shares the same worry—because they share the same pattern of attitudes, not the same background.

The five segments are:

- **Progressive Alarmed** (19%): Technically literate, politically liberal, distinctively concerned about AI's environmental impact and creative displacement.
- **Alarm Maximalists** (20%): Maximum concern across every dimension. The broadest worry portfolio, strongest regulation demand, lowest institutional trust.
- **X-Risk Literate** (16%): The most bipartisan segment and the heaviest AI users, yet the most focused on catastrophic and existential risk scenarios.
- **Cautious Moderate** (22%): The oldest segment with the most parents. Broadly concerned at moderate intensity, with many opinions still forming.
- **Trusting Pragmatists** (23%): The least concerned on every measure. Trusts institutions comparatively, uses AI comfortably, sees the current trajectory as manageable.

These segments tell a more nuanced story than the headline numbers. The concerned public is not a monolith, and the differences between groups matter for understanding public opinion on AI. A message that resonates with Alarm Maximalists can actively alienate Trusting Pragmatists. An appeal to national competitiveness that works for X-Risk Literate respondents falls flat with Progressive Alarmed ones.

But the most surprising finding may be the convergence. Despite different starting points, different levels of concern, and different trust profiles, large majorities across all five segments—and both parties—agree that the government should prioritize AI safety rules, even at the cost of competitiveness with China. The public is not divided on *whether* governance is needed. It is divided on *why* it is needed and *how urgent* the situation is.

THE LANDSCAPE OF AI CONCERN

# AI concern is widespread

**Nearly three in four Americans (73%) say they are at least somewhat concerned about artificial intelligence** as a national issue, including 28% who are "extremely" or "very" concerned. This places AI alongside climate change (75%) and well ahead of immigration (65%) in the hierarchy of American anxieties.

That AI concern is this high is itself notable. As recently as 2023, AI barely registered in issue polling. The rapid normalization of generative AI tools, combined with increasing media coverage of job displacement, deepfakes, and safety debates, has pushed AI concern to levels that took climate change decades to reach.

### AI concern rivals climate change among national issues

*% of U.S. adults who are at least somewhat concerned about each issue*

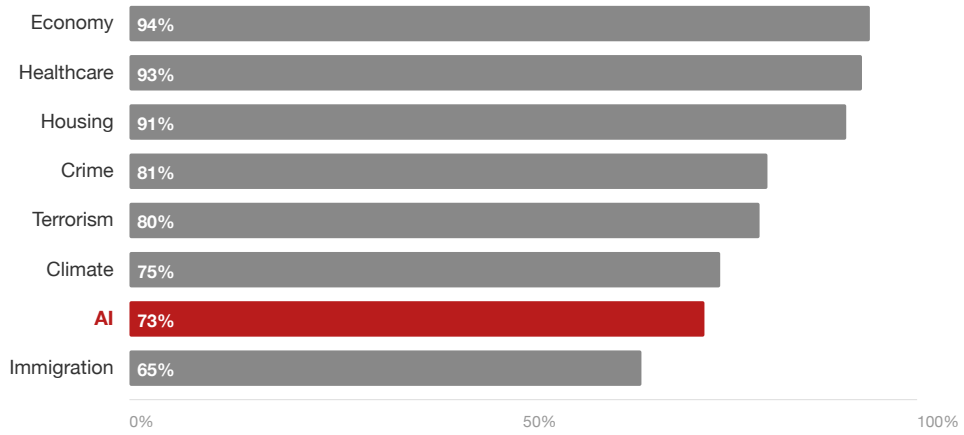


**SOURCE:** Survey of 2,735 U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

But these aggregate numbers conceal important variation. As we show in the sections that follow, the *texture* of AI concern—what people worry about, how intensely, and what they think should be done—varies dramatically across demographic lines and, especially, across the attitudinal segments identified by our analysis.

## AI concern is bipartisan in a way that almost nothing else is

**The partisan gap on AI concern is just 13 percentage points**—79% of Democrats and 65% of Republicans express at least some concern. Compare this to climate change, where the gap is 50 points, or immigration, where it runs 32 points in the opposite direction.

This bipartisanship is striking and possibly fragile. As of early 2026, AI has not yet been fully absorbed into the culture war. Both left-coded concerns (bias, job displacement, environmental impact) and right-coded concerns (government overreach, national security, China competition) coexist within the overall concern landscape. Whether this holds as AI becomes more politically salient is one of the most consequential questions in technology policy.

The 13-point gap suggests there is a narrow window in which AI governance could be framed as a bipartisan priority rather than another partisan battleground.

### The partisan gap on AI concern is remarkably small

*Partisan gap in % at least somewhat concerned (Dem. minus Rep.)*

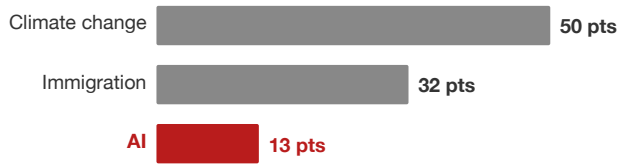


**NOTE:** Immigration gap is Republican minus Democrat.
**SOURCE:** Survey of 2,735 U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

## Despite high concern, AI ranks last in personal priority

**Only 4% of Americans rank AI as their number-one national priority,** placing it dead last among eight issues tested. The economy (29%) and healthcare (23%) dominate. Even among those who are "extremely concerned" about AI, only 11% rank it first.

This is the **activation gap**: the chasm between registering concern on a survey and treating an issue as urgent in one's personal hierarchy. Concern is easy. Priority requires believing the issue demands attention *now*, above other pressing problems. For most Americans, AI does not yet clear that bar.

The activation gap is the defining feature of current public opinion on AI. Understanding why concern does not translate into priority is central to understanding the public's relationship with AI governance.

### AI ranks last in personal priority

*% ranking each issue as their #1 national priority*

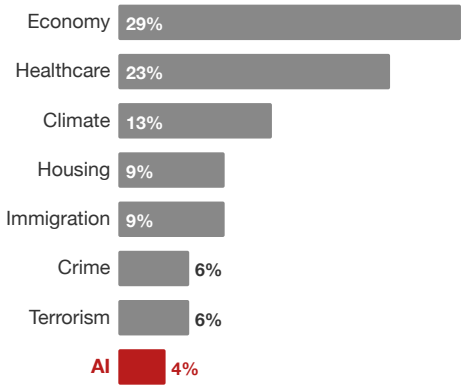


**SOURCE:** Survey of 2,735 U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

WHAT AMERICANS WORRY ABOUT

# Jobs and misinformation dominate the concern landscape

When asked to select their top three AI concerns from a list of 14, **misinformation (47%) and job displacement (46%) lead by a wide margin,** followed by privacy (29%), environmental impact (24%), and effects on children (24%). These near-term, tangible concerns outpace more abstract risks like extinction (8%) and conflicting AI goals (8%).

But the aggregate ranking obscures sharp differences between segments. Progressive Alarmed respondents are the only group where environmental impact cracks the top three. X-Risk Literate respondents are twice as likely as the sample average to cite extinction risk. Alarm Maximalists select nearly everything—an average of 12.6 out of 14 concerns, compared to 5.2 for Trusting Pragmatists.

### Jobs and misinformation lead, but concern profiles diverge by segment

*% selecting each as a top-3 AI concern (concerned sample, N=2,000)*

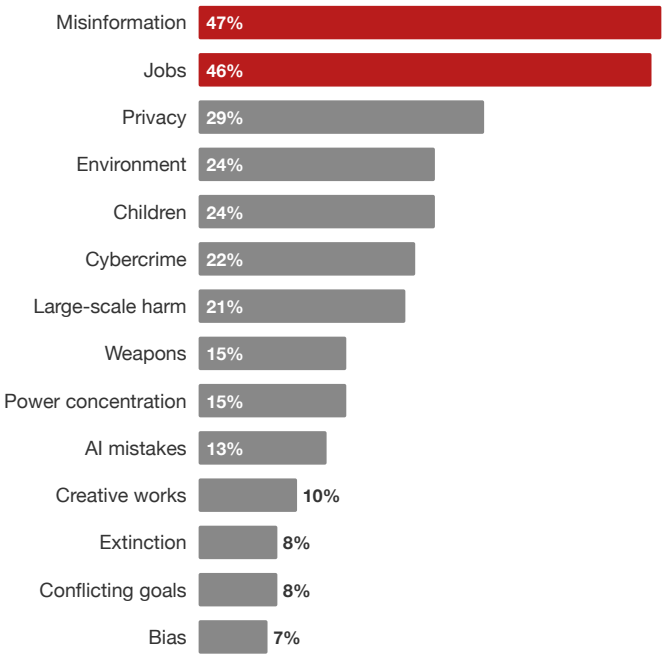


**NOTE:** Respondents selected up to 3 from a list of 14 concerns.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

## The familiarity paradox

One of the study's more counterintuitive findings involves the relationship between AI use and concern. Conventional wisdom suggests that people who use AI more should be less afraid of it—that familiarity breeds comfort. Our data tells a more complicated story.

**The X-Risk Literate segment has the highest rate of daily AI use (28%)** and also the highest focus on catastrophic risk scenarios. These are not people who are ignorant of AI and afraid of the unknown. They are people who use AI regularly, understand it comparatively well, and are concerned precisely *because* of what they know.

Conversely, the Cautious Moderate segment—the group with the most "still forming" opinions—has relatively low AI use and the highest rate of saying "I don't know enough about AI to speak up" (48%). For this group, unfamiliarity appears to breed caution rather than alarm or comfort.

### Most concerned Americans are more worried than excited

*% saying the increased use of AI makes them feel... (Q5, concerned sample)*

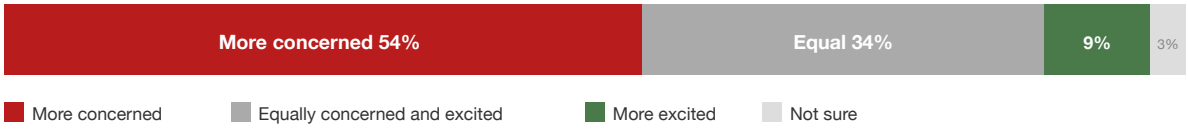


**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

The emotional landscape mirrors this complexity. **54% of concerned Americans say the increased use of AI makes them more worried than excited,** while only 9% say the opposite. But the "equally concerned and excited" category (34%) may be the most interesting: it captures the ambivalence that characterizes much of the public's relationship with a technology they find simultaneously useful and unsettling.

These patterns suggest that AI engagement strategies built on the assumption that "if people just used AI more, they'd worry less" are likely to be wrong—or at least incomplete. The most informed users may well be the most alarmed, while the least informed are not panicking but waiting.

THE FIVE SEGMENTS

# The concerned public is not one group. It is five.

Using latent class analysis on 25 attitudinal indicators, we identify five distinct segments within the population of Americans who express concern about AI. These segments differ along four dimensions: **concern intensity** (how worried they are), **concern type** (what they worry about), **trust orientation** (who they trust to handle AI), and **action preference** (what they want done).

### Distribution of five segments within the concerned public

*Latent class assignment, concerned sample (N=2,000)*

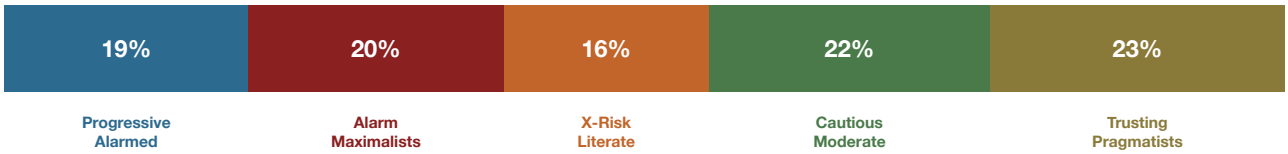


**NOTE:** Based on latent class analysis of 25 attitudinal indicators. Classes assigned by modal posterior probability.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

The naming convention follows the tradition established by the Yale Program on Climate Change Communication's "Six Americas" framework. Labels are descriptive shorthand, not normative judgments. Each segment represents a coherent cluster of attitudes—a "type" of concern—not a demographic group, though demographics do vary across segments in informative ways.

The following pages profile each segment in turn, starting with the most politically progressive and ending with the least concerned.

Several patterns emerge across the profiles. The segments form a rough gradient from high concern to low concern, but the *shape* of concern shifts in nonlinear ways. Alarm Maximalists worry about everything; X-Risk Literate respondents worry about a few specific things very intensely. Progressive Alarmed respondents channel concern through a social justice lens; Cautious Moderates are still figuring out what they think. These qualitative differences matter far more than the quantitative gradient for anyone trying to communicate about AI.

# Progressive Alarmed

19% of sample · n=378

*The youngest, most technically literate, and most politically liberal segment, yet the least likely to use AI at work.*

| 61% | 64% | 44% | 64% |
|---|---|---|---|
| Democrat | Correct on AI knowledge check | Cite environmental impact (signature concern) | "Actions wouldn't make a difference" |

The Progressive Alarmed segment is the youngest (mean age 36) and most politically liberal group in the study. They are **disproportionately female (58%)** and college-educated (52% hold a bachelor's degree or higher). Their AI knowledge scores are the highest in the sample, yet they are the *least* likely to report using AI tools at work—a tension that likely reflects both principled objection and occupational mismatch.

What distinguishes this segment is not the *intensity* of their concern (which ranks second after Alarm Maximalists) but its *shape*. Progressive Alarmed respondents are uniquely likely to cite **environmental impact (44%)** as a top-three concern—nearly double the sample average of 24%. They are also the segment most concerned about creative works displacement and AI bias. Their worry profile maps closely onto a progressive social justice framework, filtering AI anxiety through existing concerns about corporate power, labor rights, and environmental degradation.

Paradoxically, this segment also has the highest **efficacy deficit.** Sixty-four percent agree that "my individual actions wouldn't make a difference on AI issues"—the highest of any group. They are informed, alarmed, and politically engaged on other issues, but feel powerless on AI specifically. This combination of knowledge, concern, and perceived helplessness suggests a group that is waiting for institutional leadership rather than individual action opportunities.

*Political composition*


**Message resonance:** Democratic control (62% very convincing) and jobs/economic arguments (60%) resonate most strongly. Existential pause arguments perform better here than in any other segment (33%), though children and family safety remains their top-rated message (65%).

# Progressive Alarmed: Concern Profile & Message Effects

### What this group worries about vs. the sample average

*% selecting each as a top-3 concern (group vs. full sample). Sorted by gap.*

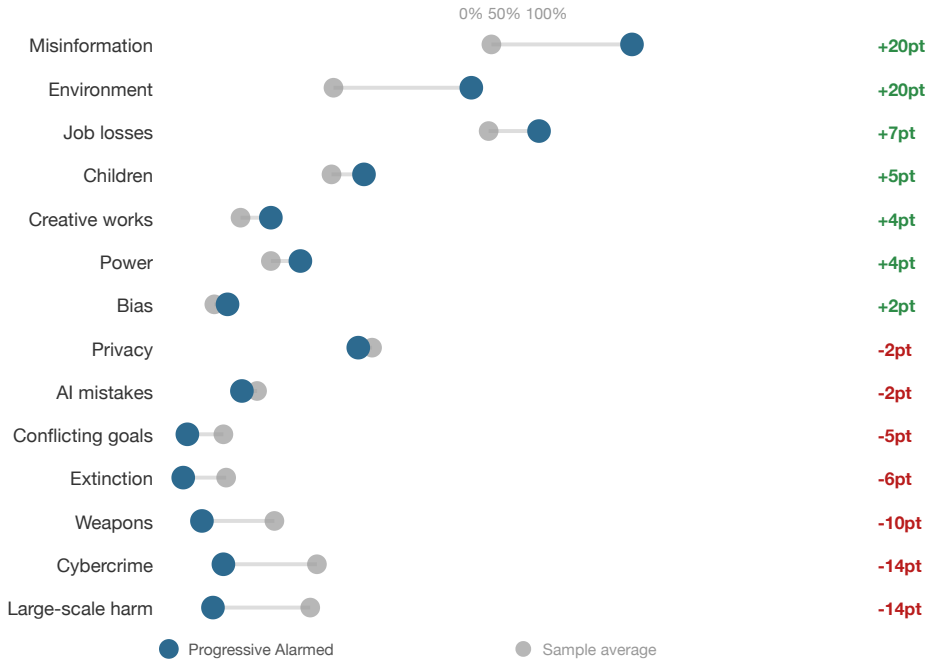


**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

## Message effect size (Cohen's d) relative to other segments

*Positive = this group found the message more convincing than average. Negative = less.*

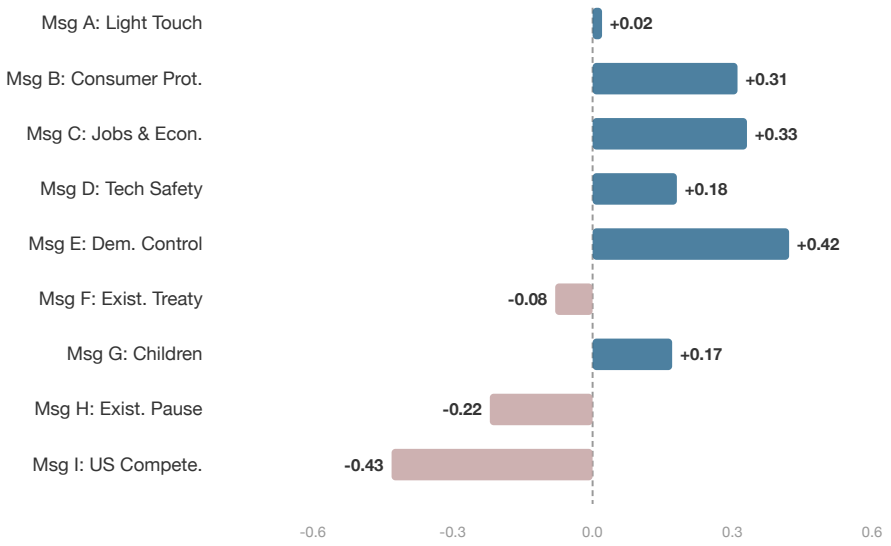


**NOTE:** Cohen's d computed from segment mean vs. pooled mean on 4-point convincingness scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Alarm Maximalists

20% of sample · n=400

*Maximum concern across every measure, broadest worry portfolio, strongest demand for regulation, lowest trust.*

| 12.6 | 69% | 60% | 67% |
|---|---|---|---|
| Avg. concerns selected (of 14) | Extremely concerned about AI | Favor a development pause | Trust tech companies "not at all" |

If the Progressive Alarmed are shaped by *what* they worry about, the Alarm Maximalists are defined by *how much*. This segment turns every dial to maximum. They select an average of **12.6 out of 14 possible concerns**—compared to 7.8 for the sample as a whole. They are the most likely to be "extremely concerned" (69%), the most likely to favor a full development pause (60%), and the least likely to trust any institution to handle AI responsibly.

Demographically, Alarm Maximalists skew female (56%), older (mean age 48), and less formally educated than the Progressive Alarmed. They lean Democratic (52%) but include a meaningful Republican minority (22%). Their AI use is low, and their knowledge scores are moderate. They are not the technical experts—they are the deeply worried generalists.

The trust deficit is sharpest here. **67% say they trust tech companies "not at all"** to handle AI responsibly, and 58% say the same of the federal government. This creates a paradox: they are the strongest advocates for government regulation of AI, but they trust the government almost as little as they trust the industry they want regulated.

*Political composition*


**Message resonance:** Children and family safety dominates (72% very convincing—the highest of any segment on any message). Technical safety (68%) and consumer protection (64%) also score well. Even the existential pause argument, which underperforms in most segments, reaches 40% here. The analytical question for this segment is not whether they perceive AI as a problem, but how such broad alarm maps onto specific policy preferences.

# Alarm Maximalists: Concern Profile & Message Effects

### What this group worries about vs. the sample average

*% selecting each as a top-3 concern (group vs. full sample). Sorted by gap.*

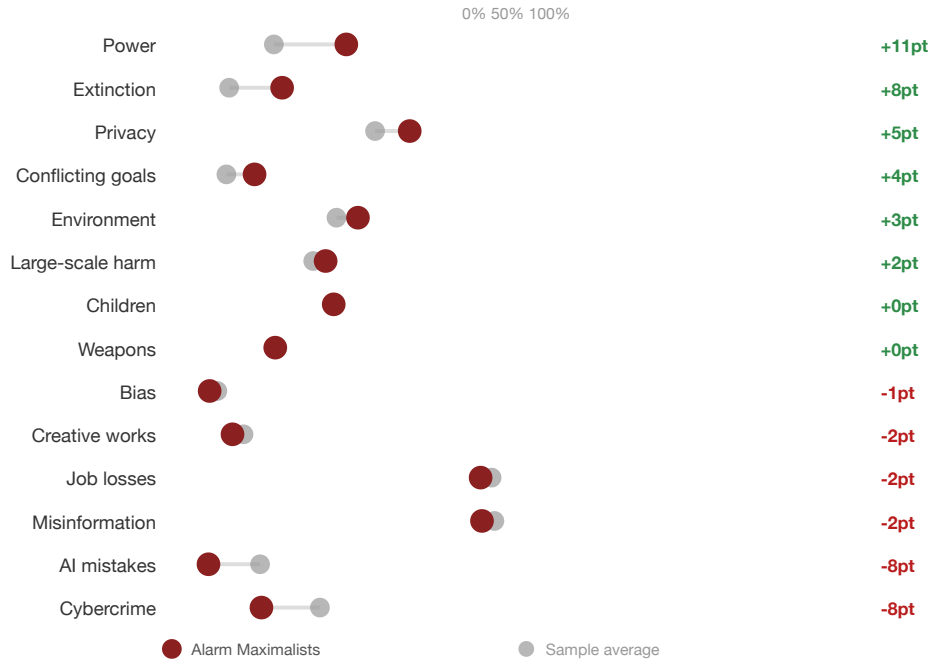


**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

## Message effect size (Cohen's d) relative to other segments

*Positive = this group found the message more convincing than average. Negative = less.*

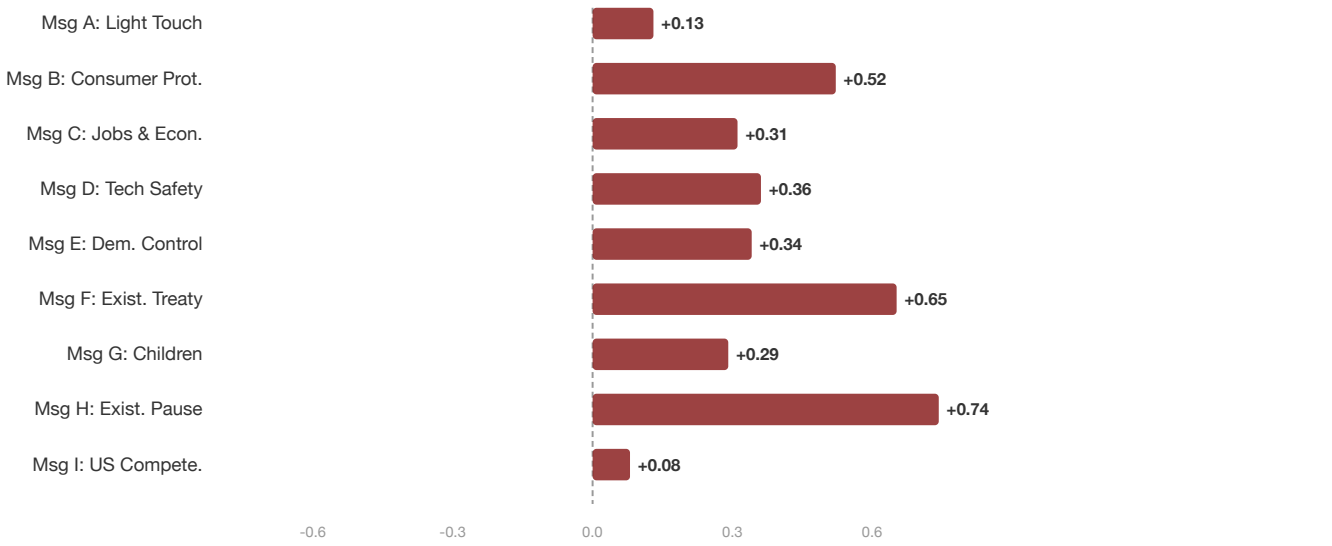


**NOTE:** Cohen's d computed from segment mean vs. pooled mean on 4-point convincingness scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# X-Risk Literate

16% of sample · n=328

*The most bipartisan and heaviest AI users, yet the most focused on catastrophic and existential risk.*

| 36/36% | 28% | 16% | 55% |
|---|---|---|---|
| Democrat / Republican (most bipartisan) | Daily AI users | Cite extinction (2x sample avg.) | Male (only male-majority segment) |

The X-Risk Literate segment defies easy categorization. They are the **most bipartisan group (36% Democrat, 36% Republican, 28% Independent)**, the heaviest AI users (28% use AI tools daily), and the most likely to have read about AI safety research. Yet they are also the segment most focused on catastrophic risk: **16% cite extinction as a top-three concern**—double the sample average—and 21% cite large-scale uncontrollable harm.

This is the only male-majority segment (55% male) and the most educated after Progressive Alarmed. They are disproportionately employed in technology or STEM fields. Their concern profile is narrow but intense: rather than worrying about everything (like Alarm Maximalists), they concentrate their concern on high-severity, low-probability scenarios. They are the segment most likely to distinguish between near-term and long-term AI risks, and to frame AI safety in terms borrowed from the technical alignment community.

Their bipartisanship is significant. On most AI policy questions, partisan affiliation predicts attitudes far less in this segment than in any other. They support regulation not from a left-wing distrust of corporations or a right-wing distrust of government, but from what appears to be a **techno-pragmatic assessment of risk.** They are the segment most receptive to the existential treaty argument (38% very convincing) and US competitiveness framing (41%).

*Political composition*


**Message resonance:** The only segment where existential treaty framing (38%) and US competitiveness framing (41%) outperform the sample average by a wide margin. Technical safety also scores well (58%). They are the hardest segment to reach with emotional appeals and the most responsive to evidence-based, utilitarian framing.

# X-Risk Literate: Concern Profile & Message Effects

### What this group worries about vs. the sample average

*% selecting each as a top-3 concern (group vs. full sample). Sorted by gap.*

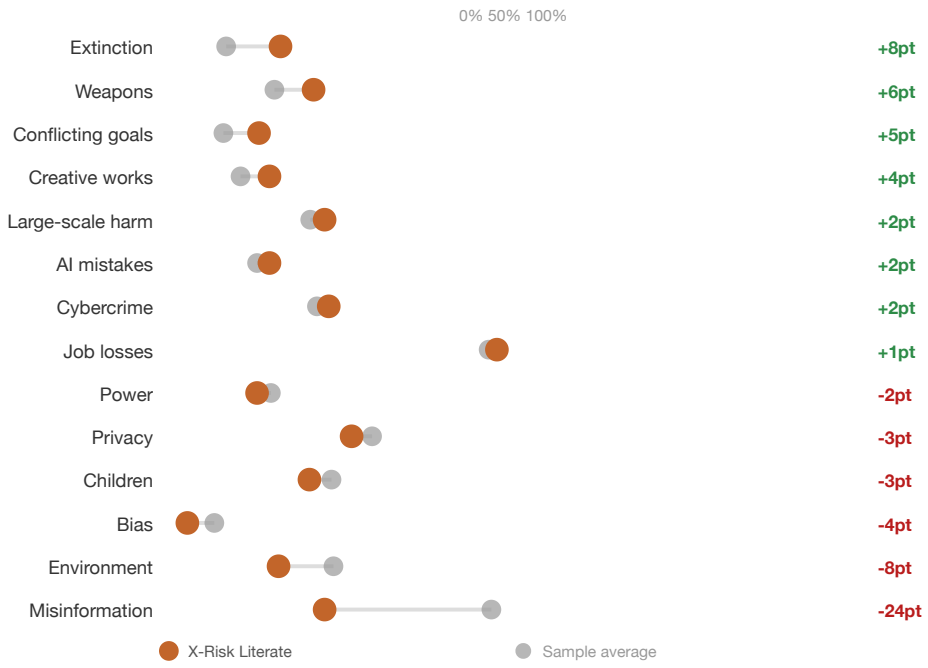


**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

## Message effect size (Cohen's d) relative to other segments

*Positive = this group found the message more convincing than average. Negative = less.*

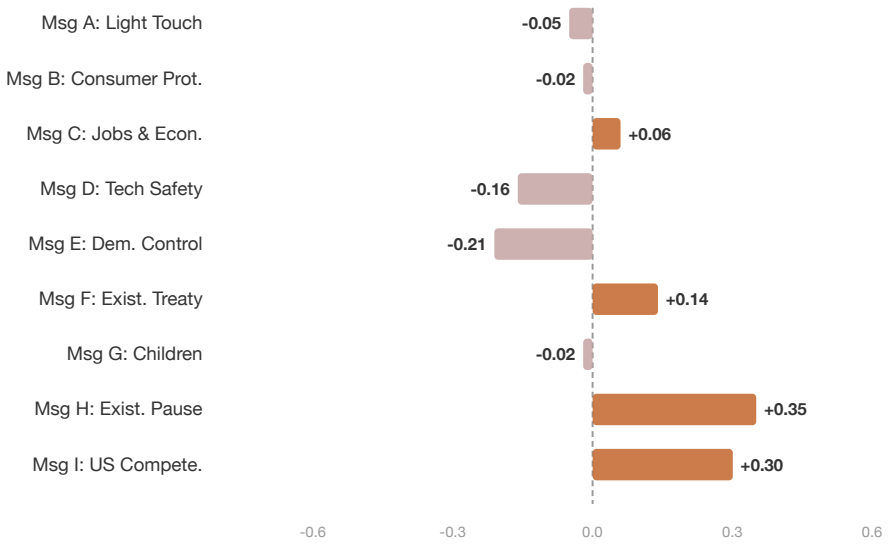


**NOTE:** Cohen's d computed from segment mean vs. pooled mean on 4-point convincingness scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Cautious Moderate

22% of sample · n=436

*Oldest segment with the most parents. Broadly concerned at moderate intensity, with many opinions still forming.*

| 50 | 60% | 58% | 48% |
|---|---|---|---|
| Mean age (oldest segment) | Parents | "Equally concerned and excited" | "Don't know enough to speak up" |

The Cautious Moderate segment is the second-largest group in the study and in many ways the most representative of the median American. They are the oldest on average (mean age 50), the most likely to have children (60%), and the most likely to describe themselves as "equally concerned and excited" about AI (58%). Their concern is real but moderate in intensity, and their opinions are the most likely to be provisional.

What makes this segment analytically important is that it represents the **persuadable middle** of AI public opinion. They have not yet crystallized strong views. **48% say they "don't know enough about AI to speak up"**—the highest of any segment—and they are the most likely to select "not sure" on trade-off questions. They are neither alarmed nor dismissive. They are waiting.

Their concern profile is broad but shallow. They worry about jobs and misinformation at rates close to the sample average, but do not over-index on any specific concern the way other segments do. They are moderate in their regulation preferences: supportive of "stronger regulation" (51%) but the least likely of the first four segments to favor a full pause (22%).

*Political composition*


**Message resonance:** Children and family safety is their most convincing message (59%), consistent with their demographic profile as the most parental segment. Consumer protection also scores well (52%). Existential arguments underperform significantly—this group responds to concrete, proximate framing about AI's effects on everyday life, not abstract scenarios.

# Cautious Moderate: Concern Profile & Message Effects

### What this group worries about vs. the sample average

*% selecting each as a top-3 concern (group vs. full sample). Sorted by gap.*

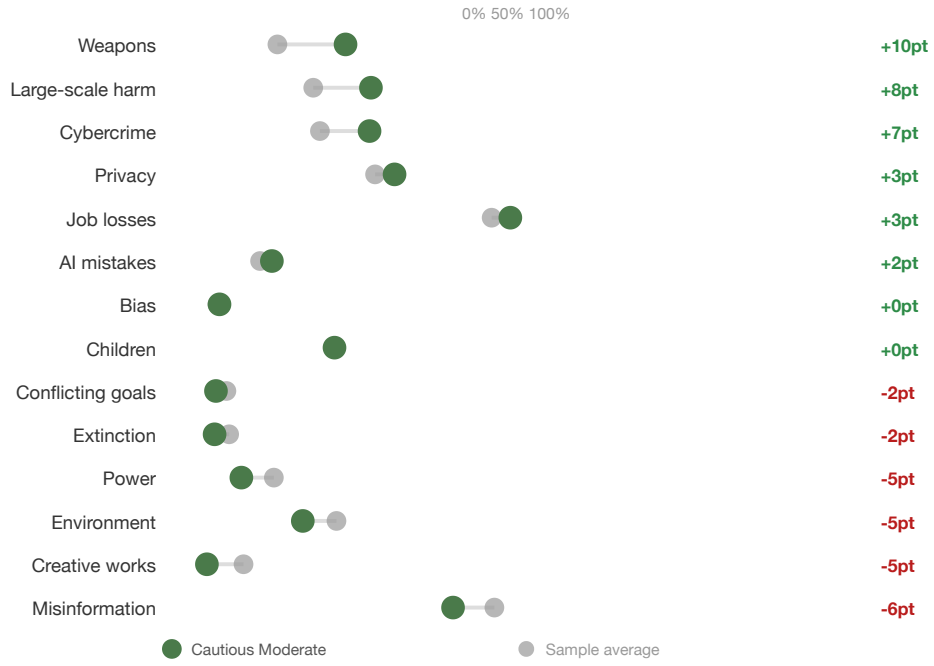


**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

## Message effect size (Cohen's d) relative to other segments

*Positive = this group found the message more convincing than average. Negative = less.*

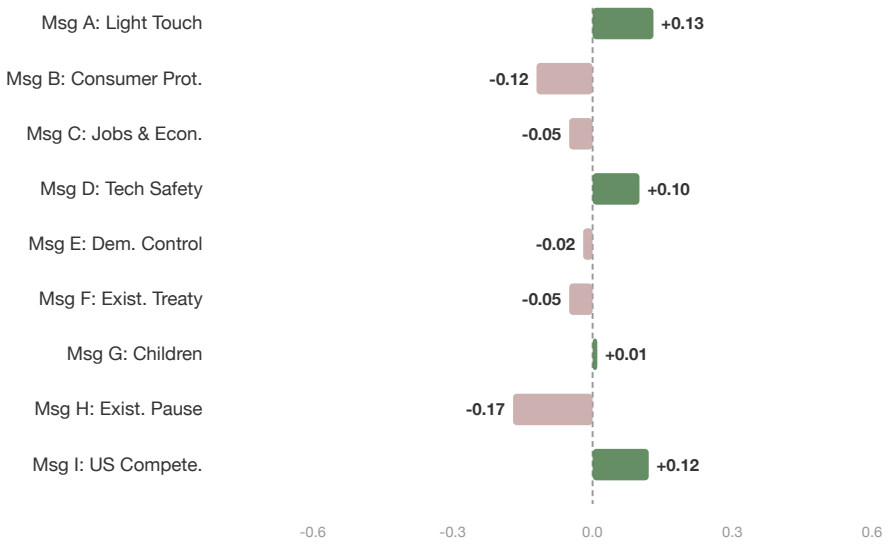


**NOTE:** Cohen's d computed from segment mean vs. pooled mean on 4-point convincingness scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Trusting Pragmatists

23% of sample · n=458

*Least concerned on every measure. Trusts institutions, uses AI comfortably, sees the current trajectory as manageable.*

| 33% | 24% | 44% | 7% |
|---|---|---|---|
| Republican (most Republican segment) | More excited than concerned (highest) | Trust tech companies "somewhat" | Very convinced by pause argument |

The Trusting Pragmatists are the largest segment and the most sanguine about AI. They are the least concerned on every measure, the most trusting of both technology companies and government, and the most likely to say the increased use of AI makes them **"more excited than concerned" (24%)**—nearly three times the sample average of 9%. They are comfortable AI users who see the technology as a net positive or at least a manageable challenge.

This is the **most Republican segment (33% Republican)**, though it still includes a 35% Democratic minority. Trusting Pragmatists are more likely to be male (52%) and employed full-time. They select the fewest concerns (5.2 on average), and when they do express concern, it tends to be about concrete, personal risks like cybercrime and privacy rather than systemic issues.

Their trust profile sets them apart. **44% say they trust technology companies "somewhat"** to handle AI responsibly—compared to just 23% in the sample as a whole. They are also the most trusting of government (32% "somewhat"). This does not mean they are naive; it means their default orientation toward institutions is cooperative rather than adversarial.

*Political composition*


**Message resonance:** This is the hardest segment to persuade. Only the children and family safety argument cracks 40% "very convincing" (44%), and the existential pause argument barely registers (7%). Light-touch regulation framing performs best relative to other segments (54% "very convincing"), consistent with their preference for incremental, market-friendly governance. The key insight: Trusting Pragmatists are not opposed to AI governance—they just want it to feel proportionate, not panicked.

# Trusting Pragmatists: Concern Profile & Message Effects

### What this group worries about vs. the sample average

*% selecting each as a top-3 concern (group vs. full sample). Sorted by gap.*

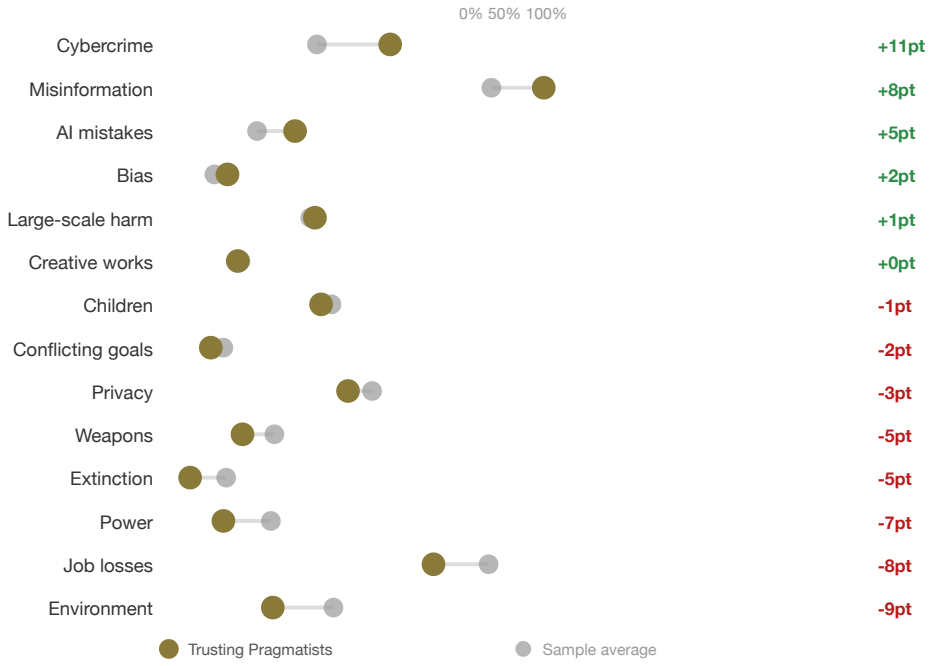


**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

## Message effect size (Cohen's d) relative to other segments

*Positive = this group found the message more convincing than average. Negative = less.*

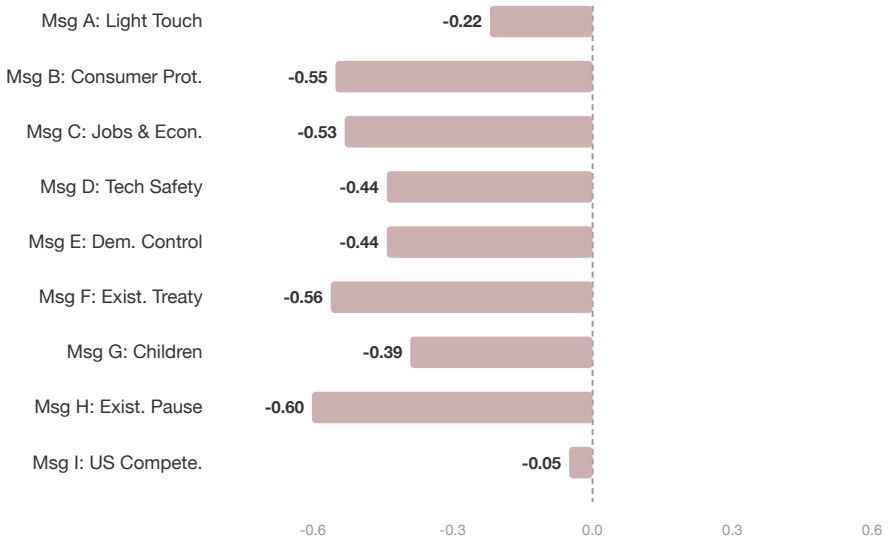


**NOTE:** Cohen's d computed from segment mean vs. pooled mean on 4-point convincingness scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

TRUST, REGULATION, AND THE GOVERNANCE PARADOX

# The public wants regulation from institutions it doesn't trust

One of the most consequential findings in this study is the **governance paradox:** Americans overwhelmingly want AI to be regulated, but they distrust both the companies building AI and the government they want to regulate it. This creates a legitimacy deficit that complicates every governance proposal.

**73% of concerned Americans trust technology companies "not very much" or "not at all"** to handle AI responsibly. The numbers for the federal government are nearly identical: 74% express low or no trust. Only 3% express "a great deal" of trust in either institution.

This is not a partisan finding. Democrats distrust tech companies more; Republicans distrust the government more. But neither party's supporters trust either institution enough to make governance feel legitimate. The result is a public that demands action from entities it does not believe will act responsibly.

Despite this distrust, the demand for governance is unambiguous. **53% favor stronger regulation,** and another 33% favor slowing down or pausing AI development entirely. Only 6% prefer the status quo, and 2% want less regulation. The public is not confused about what it wants; it is skeptical about who should deliver it.

**The public wants regulation from institutions it doesn't trust**

*Trust in each to handle AI responsibly*

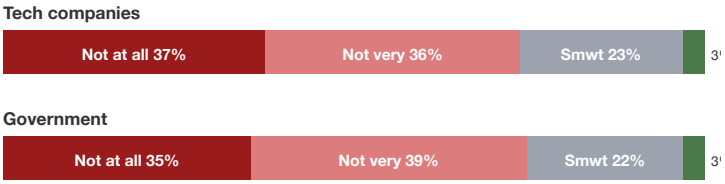


**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

## Strong demand for governance, with a divided path

The regulation question reveals an important fault line within the concerned public. While **86% want some form of increased governance,** they disagree about the form. The majority (53%) prefers stronger regulation within the current development trajectory. A sizable minority (33%) favors a more drastic approach: slowing down or pausing AI development entirely.

This divide maps closely onto the five segments. Alarm Maximalists are the most pro-pause (60%). Trusting Pragmatists overwhelmingly prefer incremental regulation (61%) and are the most resistant to pause proposals (only 14%). The other three segments fall in between, with X-Risk Literate respondents showing the most nuanced split: they support both strong regulation and international treaty frameworks but are ambivalent about unilateral pause.

### Strong demand for governance, with a divided path

*Which comes closest to your view on AI development? (Q9)*

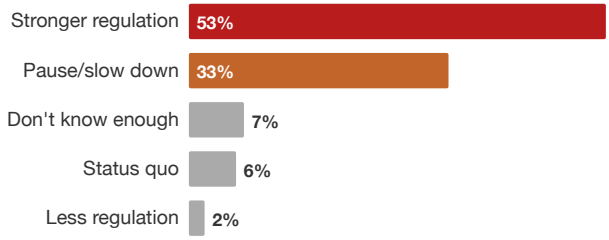


**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

## Even against China, Americans choose safety

The most striking governance finding involves the trade-off between safety and international competitiveness. We asked: *"Should the government prioritize AI safety rules, even if it means the U.S. develops AI more slowly than China?"*

**80% of concerned respondents said yes.** This held across every segment and every party, including 68% of Republicans. The "China race" argument that dominates Washington policy discourse does not resonate with the public the way it does with elites.

### Even against China, Americans choose safety

*% saying government should prioritize safety rules even if U.S. develops AI more slowly than China (Q12A)*


**NOTE:** Remaining respondents said the U.S. should prioritize developing AI as fast as possible to stay ahead of China.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

THE MESSAGE EXPERIMENT

# What resonates: children and family safety is the universal argument

Each respondent was randomly assigned one of nine policy argument framings and asked to rate how convincing they found it on a 4-point scale. The arguments ranged from children and family safety to existential risk, from economic protection to democratic control. The results reveal both a clear winner and important variation by segment.

### Children and family safety is the universal argument

*% rating each argument "very convincing" (all 9 messages, concerned sample)*

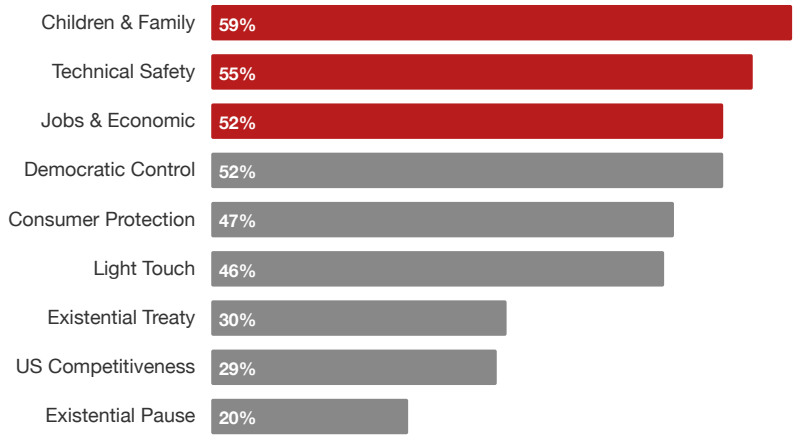


**NOTE:** Each respondent rated one randomly assigned message. N per message ranges from ~210 to ~235.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

**Children and family safety is the only argument rated "very convincing" by a majority of every segment,** ranging from 44% among Trusting Pragmatists to 72% among Alarm Maximalists. No other message achieves this kind of cross-segment consensus. Technical safety (55%) and jobs/economic (52%) round out the top three overall, but both drop below 40% in the Trusting Pragmatist segment.

## The one message every segment finds convincing

The children and family safety framing draws its power from something that cuts across every demographic and attitudinal line: parental concern. It is concrete (children's safety is tangible), emotionally resonant (it taps protective instincts), and ideologically ambidextrous (it reads as neither left-wing nor right-wing).

Importantly, the argument does not ask people to imagine abstract future scenarios. It connects AI to something they already care about in the present. This may explain why it outperforms technically sophisticated arguments about alignment or existential risk—not because those arguments are wrong, but because they require an inferential leap that most people are not yet ready to make.

The **most divisive message is the existential pause,** which ranges from 7% (Trusting Pragmatists) to 40% (Alarm Maximalists)—a 33-point spread. This suggests that "pause AI" framing, while powerful for a motivated minority, actively alienates the less-concerned segments whose buy-in is necessary for policy change. The existential *treaty* framing performs better (30% overall), suggesting that framing global cooperation as the mechanism is less divisive than framing pause as the mechanism.

**The one message every segment finds convincing**

*% rating Children & Family "very convincing" by segment*

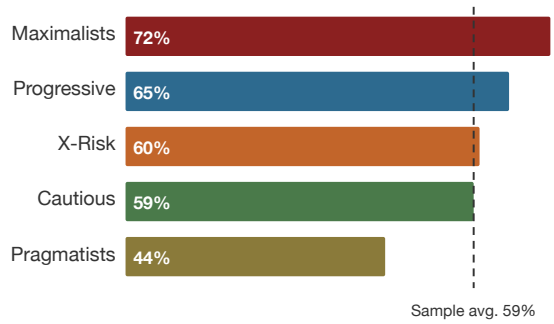


**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

The practical implication is clear: communicators who must reach across segments should lead with children and family safety, which functions as a **universal entry point**. More specialized arguments (existential risk, democratic control, economic protection) can be layered in for specific audiences, but should not be the lead frame in broad public communication. The data suggests that the AI safety community's instinct to lead with its strongest technical arguments may actually be counterproductive for public persuasion.

Perhaps most importantly, the message experiment reveals that **different starting points can converge on the same policy destination.** A Progressive Alarmed respondent persuaded by democratic control arguments and a Cautious Moderate persuaded by children and family arguments may arrive at the same support for stronger AI regulation—but through different doors. Effective communication strategy is about finding the right door for each audience, not insisting that everyone enter through the same one.

# Message A: Light Touch Transparency

*"AI technology is advancing rapidly and affecting more areas of our lives—from hiring decisions to medical diagnoses to what we see online. Like other powerful technologies, AI needs clear rules and oversight to make sure it's developed safely and benefits everyone, not just tech companies. We need commonsense regulations that protect people while allowing innovation to continue."*

## How convincing did each segment find this argument?

*Response distribution across 5 segments (100 squares = 100%)*

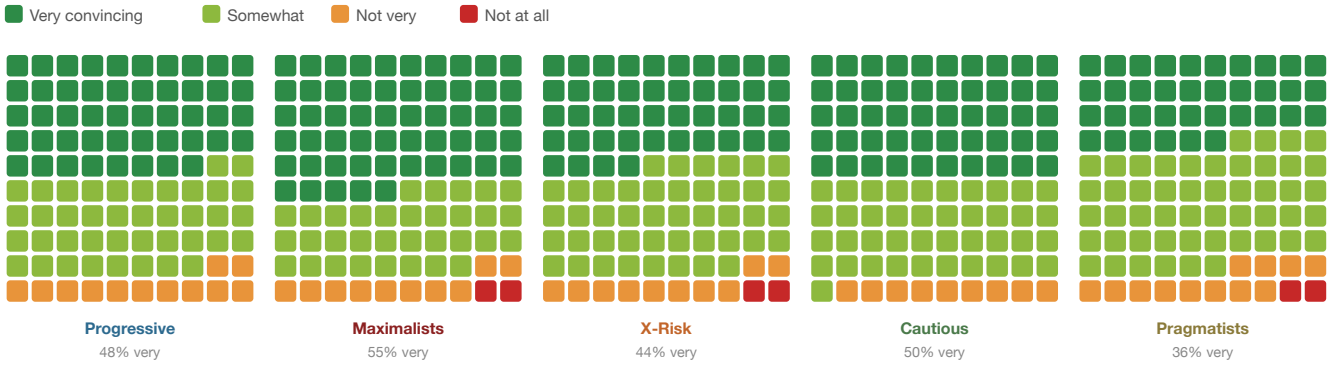


**NOTE:** Each respondent rated one randomly assigned message on a 4-point scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Message B: Consumer Protection

*"AI companies are rushing products to market without adequate testing. AI systems are already making mistakes that affect real people—wrongly denying loans, making errors in medical settings, and spreading false information. Just as we require safety testing for cars and drugs, we need independent testing and safety standards for AI before these systems are used in high-stakes decisions about people's lives."*

## How convincing did each segment find this argument?

*Response distribution across 5 segments (100 squares = 100%)*

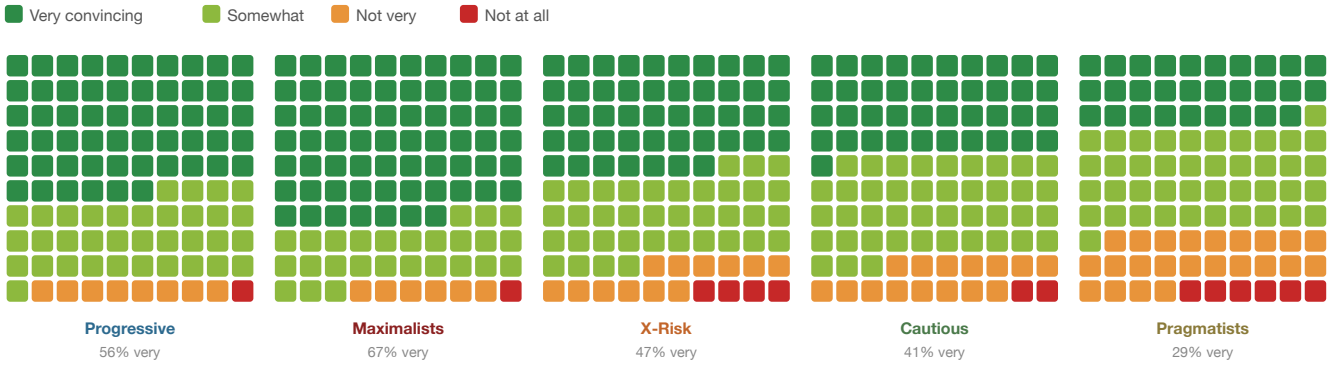


**NOTE:** Each respondent rated one randomly assigned message on a 4-point scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Message C: Jobs & Economic

*"AI is projected to affect millions of jobs in the coming years. Without proper planning, this could devastate communities across America. We need policies that ensure AI benefits workers, not just shareholders—including job transition programs, education investments, and rules that prevent AI from being used to exploit workers or drive down wages."*

## How convincing did each segment find this argument?

*Response distribution across 5 segments (100 squares = 100%)*

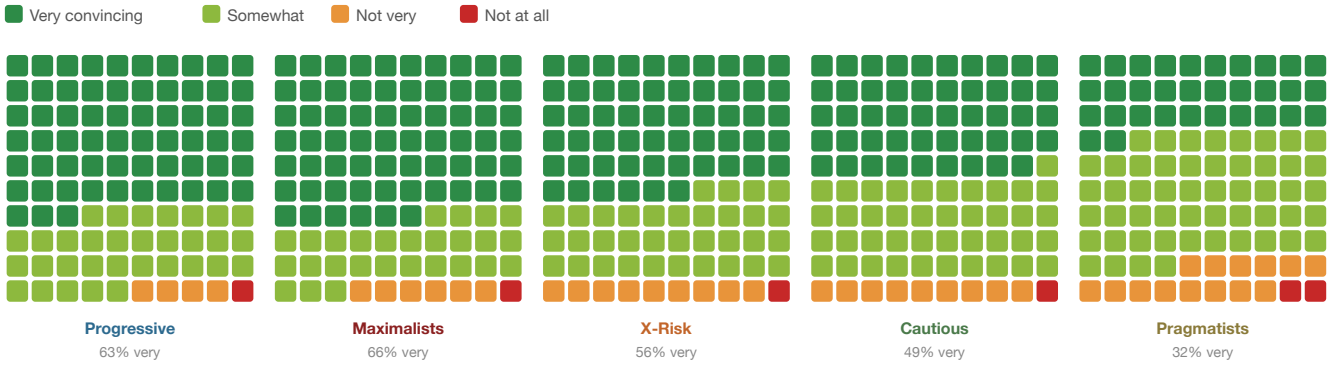


**NOTE:** Each respondent rated one randomly assigned message on a 4-point scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Message D: Technical Safety

*"Leading AI researchers, including many who built these systems, are warning that current AI development practices are inadequate. These systems are becoming more powerful but we don't fully understand how they work or how to ensure they behave as intended. Before deploying AI in critical areas, we need better testing methods and safety standards—the same way we'd want safety testing before a new aircraft carries passengers."*

## How convincing did each segment find this argument?

*Response distribution across 5 segments (100 squares = 100%)*

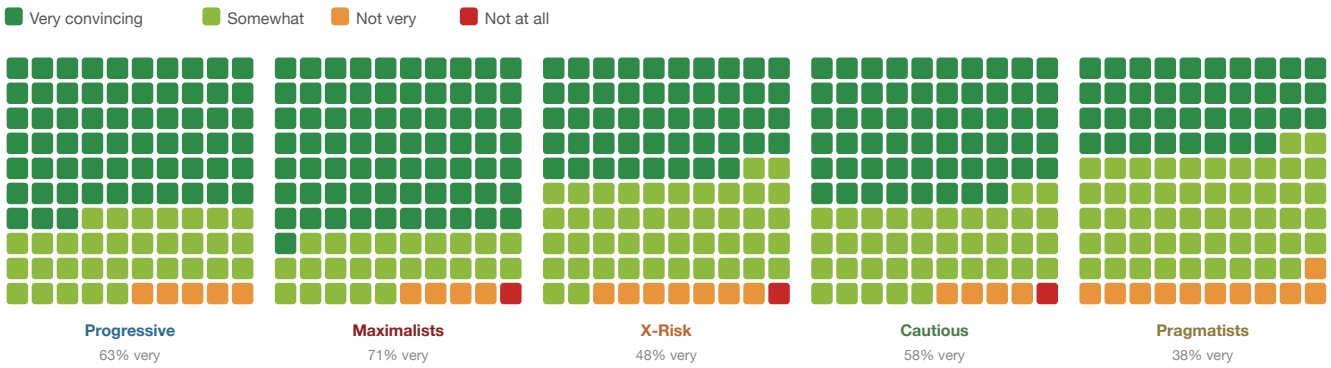


**NOTE:** Each respondent rated one randomly assigned message on a 4-point scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Message E: Democratic Control

*"Decisions about AI's role in society shouldn't be made by a handful of tech executives in Silicon Valley. AI affects all of us, so the public should have a voice in how it's developed and deployed. We need democratic oversight—not just industry self-regulation—to ensure AI serves the common good and doesn't concentrate power in the hands of a few corporations."*

### How convincing did each segment find this argument?

*Response distribution across 5 segments (100 squares = 100%)*

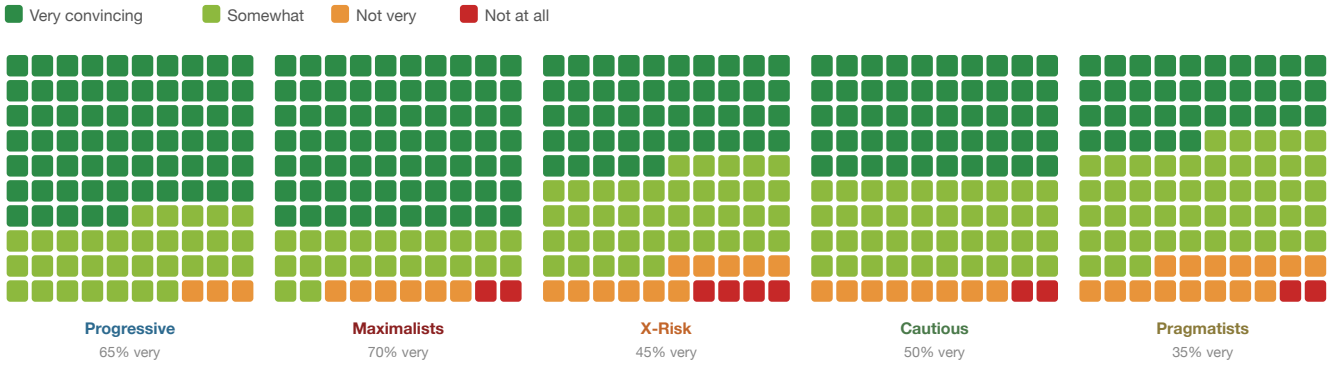


**NOTE:** Each respondent rated one randomly assigned message on a 4-point scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Message F: Existential Risk (International Treaty)

*"Many leading AI scientists believe advanced AI could pose catastrophic risks to humanity if developed without adequate safeguards. Some compare it to nuclear weapons—a technology so powerful that getting it wrong could be irreversible. We may only have a few years to establish effective global safety measures before AI systems become too powerful to control. The stakes couldn't be higher."*

## How convincing did each segment find this argument?

*Response distribution across 5 segments (100 squares = 100%)*

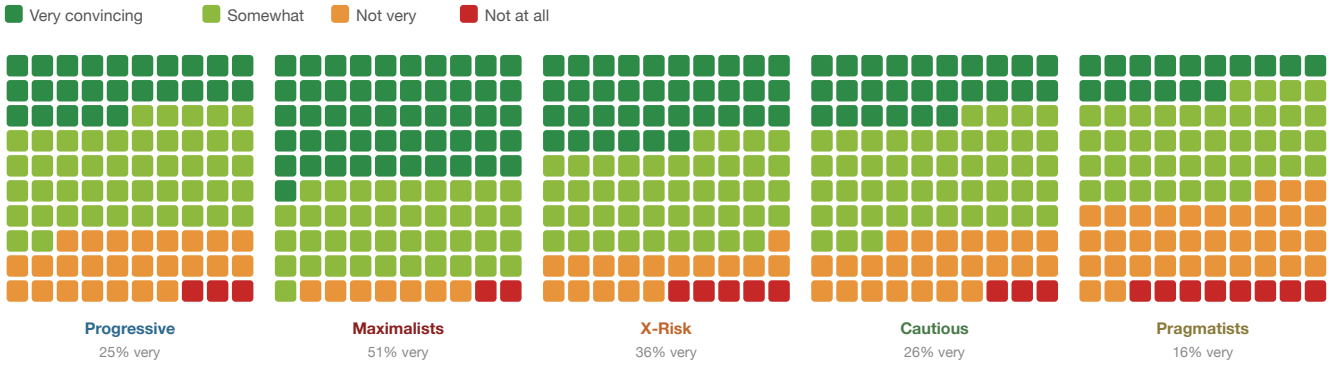


**NOTE:** Each respondent rated one randomly assigned message on a 4-point scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Message G: Children & Family

*"AI is already affecting our children—through social media algorithms, AI tutors, chatbots they interact with, and decisions that affect their futures. Kids are particularly vulnerable to manipulation and harms from these systems. Parents deserve the right to know how AI is being used with their children, and we need strong protections to keep young people safe from AI-related harms."*

## How convincing did each segment find this argument?

*Response distribution across 5 segments (100 squares = 100%)*

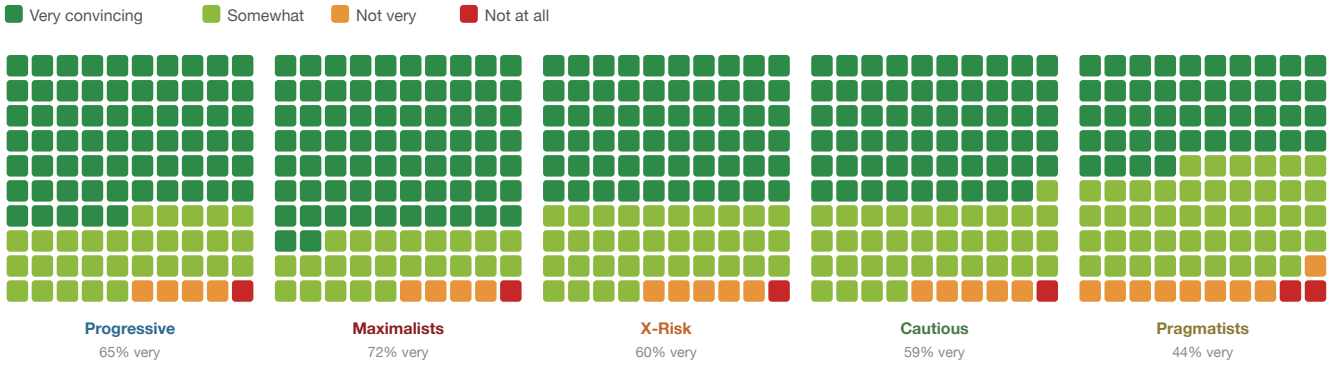


**NOTE:** Each respondent rated one randomly assigned message on a 4-point scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Message H: Existential Risk (Pause)

*"We are facing a potential extinction-level threat. Advanced AI systems could, within our lifetimes, become impossible for humans to control. If we don't act immediately and dramatically—including halting the most dangerous AI research—we risk the end of human civilization. This is not science fiction; many of the world's top AI researchers are sounding the alarm. We need emergency action now."*

## How convincing did each segment find this argument?

*Response distribution across 5 segments (100 squares = 100%)*

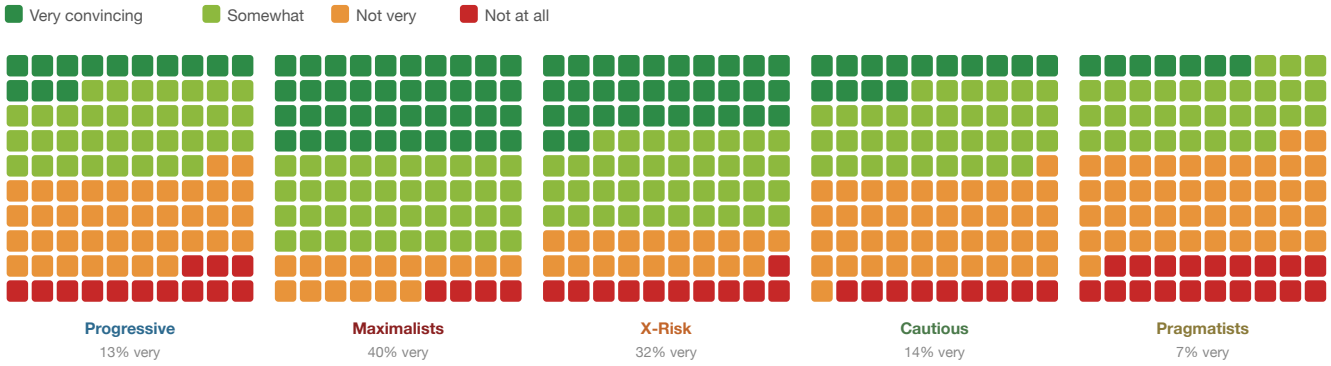


**NOTE:** Each respondent rated one randomly assigned message on a 4-point scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

# Message I: US Competitiveness (Counter)

*"If the United States slows down AI development with heavy regulations, China will race ahead and dominate this critical technology. American AI companies need the freedom to innovate and compete. Excessive regulation would hand our greatest technological advantage to our rivals and threaten our national security."*

### How convincing did each segment find this argument?

*Response distribution across 5 segments (100 squares = 100%)*

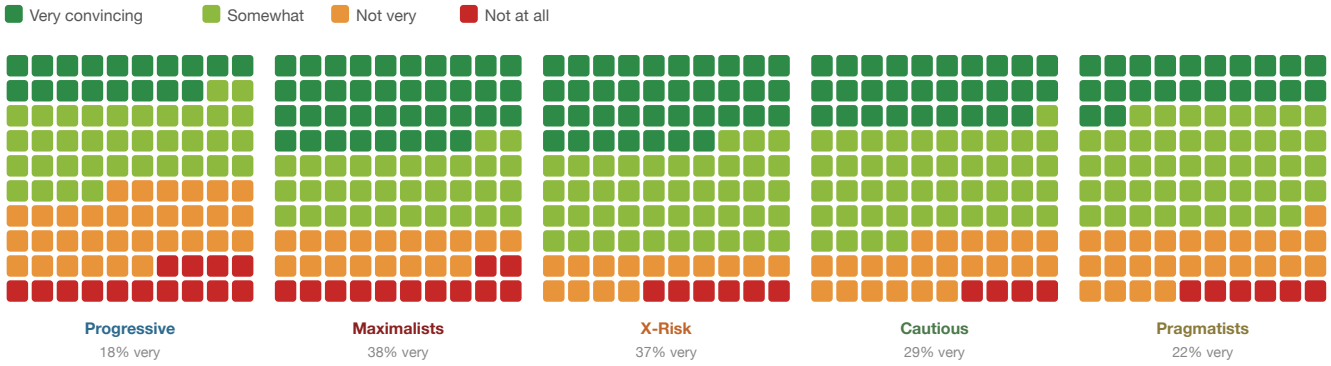


**NOTE:** Each respondent rated one randomly assigned message on a 4-point scale.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

THE ACTIVATION GAP

# Why concerned Americans haven't acted

The central puzzle of AI public opinion is the gap between concern and action. Most Americans are concerned. Very few are engaged. Why? We asked respondents to identify the barriers that prevent them from taking more action on AI issues.

### Why concerned Americans haven't acted

*% selecting each barrier to taking more action on AI issues (Q15, multiple selection)*

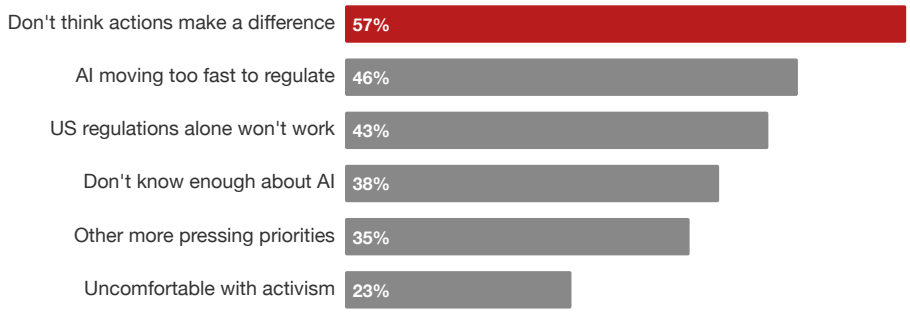


**NOTE:** Respondents could select multiple barriers.
**SOURCE:** Survey of 2,000 concerned U.S. adults, February 2026.
**FIVE AMERICAS OF AI CONCERN**

The top barrier is **perceived inefficacy.** Fifty-seven percent of concerned Americans agree that "my individual actions wouldn't make a difference on AI issues." This is followed by perceived uncontrollability (46% say "AI is moving too fast to regulate effectively") and perceived futility of unilateral action (43% say "US regulations alone won't work because AI is global").

Together, these three barriers paint a picture of a public that is concerned but feels helpless. They see AI as too big, too fast, and too global for their actions to matter. This is compounded by a knowledge deficit (38% say they don't know enough) and competing priorities (35% say other issues are more pressing).

The activation gap is not a communication problem alone. It is a **structural problem** rooted in the genuine difficulty of translating diffuse concern about a fast-moving technology into specific, achievable civic actions. The public is not apathetic. It is overwhelmed.

The barrier profile suggests that the activation gap has three distinct components: **(1)** a perceived scope problem (the issue feels too large for individual action), **(2)** a perceived control problem (the technology is moving too fast for governance to keep pace), and **(3)** a perceived coordination problem (unilateral national regulation feels insufficient for a global technology). These are not irrational beliefs. They reflect genuine structural features of AI governance that any public engagement effort would need to address.

IN THEIR OWN WORDS

# What respondents told us

At the end of the survey, we asked respondents an open-ended question: *"Is there anything else you'd like to share about your views on AI?"* The responses ranged from one-sentence dismissals to multi-paragraph reflections. Here is a selection that captures the range of concern across the five segments.

> *"It's a monster. This is how civilizations end. Too much technology, too fast. Don't know how to control it, and it will be in the hands of very evil people. It is already."*
>
> — Republican Woman, 73, Oregon (X-Risk Literate)

> *"AI is going to be difficult to regulate because politicians and people putting in regulations don't truly understand the system."*
>
> — Democrat Man, 46, Florida (Cautious Moderate)

> *"I am very concerned about AI taking over jobs that humans should be doing. I am also concerned about the impact on children's education and development."*
>
> — Democrat Woman, 52, Michigan (Progressive Alarmed)

> *"I think AI can be a powerful tool if regulated properly. My concern is that greed will drive development faster than safety measures can keep up."*
>
> — Independent Man, 38, Texas (Alarm Maximalist)

These four voices illustrate the complexity that topline numbers cannot capture. A 73-year-old Republican woman from Oregon and a 38-year-old Independent man from Texas share a sense of urgency about AI—but worry about entirely different things and trust entirely different solutions. A Democratic woman in Michigan and a Democratic man in Florida both support regulation but for different reasons. The concerned public is united in its unease and divided in everything else.

What comes through most clearly in the open-ended responses is a sense of **temporal mismatch.** Respondents describe a technology moving at a pace that outstrips their ability to understand it, let alone govern it. "Too fast" is the single most common phrase in the open-ended data. The public does not doubt that AI is powerful. It doubts that human institutions can keep up.

The data reveals a public opinion landscape with a distinctive structure: broad concern, low priority, deep distrust, and strong demand for governance despite that distrust. Whether AI concern follows the trajectory of climate opinion—where initial bipartisan concern eventually polarized along partisan lines—or charts a different course is a question this cross-sectional study cannot answer. But the current window of bipartisan agreement on AI governance is real, and understanding its structure is a prerequisite for understanding what comes next.

METHODOLOGY

# Methodology

## Sample

The survey was fielded via **Prolific**, an online research platform, in February 2026. A total of 2,735 U.S. adults aged 18 and older were recruited using Prolific's representative sampling features, with demographic quotas set for age, sex, ethnicity, education level, and U.S. Census region. The median completion time was approximately 15 minutes, and respondents were compensated at Prolific's recommended rate.

A screener at the beginning of the survey identified respondents who indicated they were "not at all" concerned about AI and had no opinion on any AI-related issue. These 735 respondents were screened out, yielding an **analytical sample of 2,000** respondents who expressed at least some concern about AI. All analyses in this report are conducted on the concerned sample unless otherwise noted.

## Survey Instrument

The survey comprised 16 question blocks covering issue concern (relative and absolute), specific AI concerns, personal AI use and experience, knowledge self-assessment, institutional trust, regulatory preferences, argument evaluation, trade-off assessments, barriers to engagement, and open-ended reflection. Question order was fixed except for the argument experiment, which was randomized.

## Latent Class Analysis

Latent class analysis (LCA) is a finite mixture modeling technique that identifies subgroups (classes) of individuals who share similar patterns of responses across multiple observed variables. Unlike clustering methods that use distance metrics, LCA estimates a probabilistic model in which each class has a distinct response profile and each individual has a posterior probability of belonging to each class.

We selected **25 attitudinal indicators** for the LCA model, drawn from concern intensity (Q1, Q3a), concern type (Q3b items), emotional orientation (Q5), trust (Q7), and regulatory preference (Q9). Items were binarized or collapsed to ordinal categories as appropriate. The model was estimated using the EM algorithm with **500 random starts** to ensure global convergence.

## Model Selection

Models with K=3 through K=8 classes were estimated and compared on the following criteria:

- **BIC (Bayesian Information Criterion):** Lower is better. The 5-class model showed the best BIC improvement relative to model complexity.
- **Entropy:** A measure of classification quality. Values above 0.80 indicate good separation. Our 5-class model achieved 0.833.
- **Stability:** The 5-class solution was recovered in 478 of 500 random starts (95.6%).
- **Bootstrap Likelihood Ratio Test (BLRT):** K=5 was significantly better than K=4 ($p<0.001$); K=6 did not significantly improve over K=5.
- **Interpretability:** The 5-class solution produced segments with clear qualitative narratives; the 6-class solution split one segment into two groups that were not meaningfully distinct.

## Model Fit Statistics

| 98,371 | 0.833 | 478/500 | 0.89 | 0.86–0.93 |
|---|---|---|---|---|
| BIC | ENTROPY | STABILITY | AVG. POST. PROB. | AVEPP RANGE |

## Robustness

To test the stability of the segmentation, we re-estimated the model using an expanded set of **39 indicators** that included additional items from Q4 (AI use), Q6 (knowledge), and Q15 (barriers). The expanded model produced a 6-class solution in which **4 of the 5 original classes were preserved at 80%+ concordance** with the main model. The fifth class (Cautious Moderate) split into two subgroups that differed primarily in age composition. We retained the 5-class solution from the primary model as more parsimonious and interpretable.

## Argument Experiment

The nine argument framings were developed through iterative drafting and pilot testing. Each argument was approximately 60–80 words and presented a distinct rationale for AI governance action. Arguments were randomly assigned to respondents. The evaluation comprised a **4-point convincingness scale** (not at all, not very, somewhat, very convincing) and a **5-category tone assessment** (too alarming, somewhat alarming, about right, not alarming enough, not sure). Sample sizes per argument ranged from approximately 210 to 235.

## Limitations

The study has several important limitations. First, the Prolific sample, while demographically balanced, may not perfectly represent the U.S. adult population; Prolific participants tend to be more digitally literate and somewhat younger than the general population. Second, the screener removes respondents who are "not at all concerned" about AI, meaning our segments describe variation *within* the concerned public, not the full American population. Third, LCA produces probabilistic class assignments; individuals near class boundaries may not fit neatly into a single segment. Fourth, the argument experiment uses a between-subjects design, so we cannot observe how the same individual responds to multiple arguments. Fifth, the cross-sectional design captures a snapshot; public opinion on AI is evolving rapidly.

## Segment Distribution

| Segment | N | % | Avg. Posterior Probability |
|---|---|---|---|
| Progressive Alarmed | 378 | 19% | 0.91 |
| Alarm Maximalists | 400 | 20% | 0.93 |
| X-Risk Literate | 328 | 16% | 0.89 |
| Cautious Moderate | 436 | 22% | 0.86 |
| Trusting Pragmatists | 458 | 23% | 0.87 |
| **Total** | **2,000** | **100%** | **0.89** |

**AI tools disclosure:** This summary document was compiled with the assistance of AI tools for layout, chart generation, and text formatting. All research design, data collection, statistical analysis, and interpretation of findings were conducted by the author. The underlying data and interactive visualizations are available at fiveamericas.ai.

**FIVE AMERICAS OF AI CONCERN**

fiveamericas.ai