

After Coxon's Warning: There is Now Enough Political Will to Enact Policy on AI Risk

David Williams-King^{1,2}, Zachary Schlosser¹, India Dearlove¹, Samuel Baumohl¹, Yosra Kazemi¹, Cedric Lam¹, Raymond Lee¹, Linh Le^{1,3}

¹Lida Safety Research, ²ERA, ³Oxford Martin AIGI

Version 1.5 - September 20th, 2026

Executive Summary

- On September 8th 2026, researcher Jacob Coxon resigned from Anthropic over concerns about the rapid advance of AI capabilities and the race between AI companies. Following his resignation, Coxon wrote an announcement on X [1] which reached over 170 million views, bringing unprecedented attention to discussions of AI risk.
- To investigate online dialogue surrounding AI safety since the Coxon thread, we used APIs to download posts from X, Bluesky, Hacker News, and YouTube and found that **53-71% of on-topic posts taking a position agreed that advanced AI poses a catastrophic risk**. Traffic to AI safety relevant Wikipedia articles has increased; “Existential risk from AI” is up by 19.7x.
- After Coxon’s announcement, the top three reasons given for AI risk skepticism were that AI risk narratives are marketing hype for the frontier labs, that the messenger of risk should not be trusted, and that AI catastrophes are implausible. These arguments are well tread territory and have been addressed by proponents of AI safety policies.
- Recent polls show 68% of likely voters support a pause on AI, with majority support among both Democrats and Republicans [23], and 63% of Americans see at least a moderate risk that advanced AI could destroy humanity [16].
- With material evidence of the harmful actions of rogue AI, newfound public interest in AI risk, and the consistent release of increasingly capable models, **now is the time** to regulate advanced AI before much more damaging incidents occur.
- Congress has begun to respond: four bills have been introduced since July, and a bipartisan Senate bill imposing duty of care on frontier developers is underway.
- We recommend implementing ControlAI's draft legislation [2] to prohibit the development of superintelligence, or if this is politically untenable, we recommend strengthening the Thune-Cruz-Klobuchar bill into a framework for pacing the frontier.

Introduction

On September 8th 2026, AI pretraining researcher Jacob Coxon resigned from Anthropic over concerns about the rapid pace of capabilities and the race between AI companies to develop

ever-more-powerful systems. Following his resignation, Coxon wrote an announcement on X [1] which reached over 170 million views, bringing unprecedented attention to AI risk discussions.

We are Lida Safety Research (lidasafety.org), an AI safety forecasting organization. To inform policymakers' considerations about AI policy, we analyzed the public response to Jacob Coxon's resignation from Anthropic. For this memo, we collected 201,040 posts and comments from X, YouTube, Bluesky, and Hacker News, classified a sample of 33,486, and drew on four US polls, three from September 2026. We analyze online discourse resulting from Coxon's posts, highlight trends, and analyze the most common arguments about AI risk, both dismissive and supportive. Finally, we offer responses to the most common dismissive arguments.

There is now sufficient political will to enact AI policy. We do not believe that superintelligence can be built safely and recommend that elected officials consider ControlAI's draft legislation to ban superintelligence [2]. We also recommend policymakers create a framework to pace the frontier of AI capabilities research, an approach that has been endorsed by Dario Amodei, Sam Altman, and Elon Musk [3, 4, 5], or, more generally, any policy that prevents the race to artificial superintelligence.

After Coxon: A Window for Action

Coxon's post follows a long line of researchers who have quit their jobs and warned the public about AI risk. Other notable examples include Daniel Kokotajlo from OpenAI in April 2024, Jan Leike from OpenAI in May 2024, and Nobel-prize winner Geoffrey Hinton from Google in May 2023. However, Coxon's announcement has garnered far more attention than these previous warnings. For instance, the New York Times has published multiple articles [6, 7, 8, 9] and opinion pieces [10, 11] about Coxon's posts and AI safety in the week following the announcement; Time [12] and The Hill [13, 14] have also published articles. Until recently, the obstacle to AI safety policy was political will rather than knowledge [15]. However, there is now sufficient political will.

According to polling by Data for Progress [23] from Sept 4th-9th, within a sample of 1,332 US likely voters, 68% support the pause on AI while 25% oppose it (a 43-point lead). Support for a pause has a 51-point lead for Democrats, and a 32-point lead for Republicans. A POLITICO poll run from Sept 13-15th surveyed 2,064 Americans online [16]. The conclusion is that 63% believe there is at least a moderate risk that advanced AI could destroy humanity, and only 22% say very little or no risk.

AI has become much more salient to US voters. Blue Rose Research's issue tracking shows the share of voters rating AI as an important issue rose steadily from roughly 24% in late 2024 to 49.4% in the week after Coxon's announcement — 9.5 points above its immediate pre-Coxon trend [17] — and AI has risen in importance more than any other issue over the past year [18].




Figure 1: Artificial intelligence salience since the 2024 election, Blue Rose Research.

US voters are now more willing to pay more attention to AI risks because of multiple concrete incidents of agentic loss of control, most notably the Hugging Face attack by OpenAI agents attempting to cheat on their cyber evaluation tasks during training [19], and the subsequent discovery of additional undisclosed OpenAI rogue model deployments which took place months prior in May-June 2026 [20].

What the Public Believes About AI Risk

After Coxon's post, a YouGov poll from Sept 14 [21] showed that half of American adults (50%) are concerned that AI will cause the end of the human race on Earth, up from 37% in March 2025 [22], while only 38% are unconcerned. The same poll separately showed 67% believing AI is advancing too quickly. This is a bipartisan issue: polling data from Data for Progress [23] from Sept 4th-9th indicates 68% support a pause on AI while 25% oppose it (a 43-point lead). Support for a pause has a 51-point lead for Democrats, and a 32-point lead for Republicans.

In social media data, based on a subset of social media posts that we downloaded in the week after the announcement, we found that **53-71%** of users across the four platforms we surveyed believed AI poses existential risk. Bluesky, which tends to skew towards left-leaning individuals, was at the high end at 71%. After Coxon, we also see a surge of interest in AI safety-related articles on Wikipedia, e.g., a 19.7x increase in visitors to the page "Existential risk from AI".

Before vs. after the Coxon post

After Coxon's post, there was a 3x-11x increase across platforms in the number of threads mentioning AI risk, with an increase of 3.5-5.3x in the total number of posters per day. This shows how a much broader audience was brought into the debate.




Figure 2: Posts saying "AI safety" versus normal on X, Bluesky, and Hacker News, Lida Safety Research

Furthermore, we investigated traffic to multiple AI safety-relevant articles on Wikipedia from May 1st to Sept 16th. Comparing before and after Coxon's announcement, we found a 19.7x increase in visitors to the page "Existential risk from AI", a 16.0x increase to the page "P(doom)", and a 12.3x increase to "AI capability control". More technical pages increased somewhat less, but still spiked: "AI safety" saw a 3.2x increase and "AI alignment" a 5.0x increase.

What Concerned People Think

Most people across social media platforms see AI as a catastrophic or existential risk. We measured the percentage of people seeing AI as a catastrophic risk after the Coxon post in our social media dataset. We found Hacker News at 53.3%, YouTube at 55.9%, X at 56.7%, and Bluesky at 70.8%, as shown in the chart below.

Concerned people hold several positions regarding AI risk. After Coxon's post, more people are finding experts on AI risk more credible than they did before. Amongst concerned people who posted on X both before and after Coxon's post, arguments referring to expert testimony rose from 14.0% to 38.4% of concerned posts on X; meanwhile, these same posters' references to

abstract loss-of-control arguments dropped from 33.3% to 16.5%. People appear to have felt that the expert testimony is more compelling than the more abstract case.




Figure 3: Share seeing AI as a catastrophic risk after the Coxon thread, Lida Safety Research

The table below is based on the concerned posts from all four social media platforms.

Positions	Before Coxon Post	After Coxon Post	Shift
Insider and expert warnings count as evidence	8.4%	28.2%	+19.8 pp
Regulation, pause or prohibition demanded	21.4%	24.3%	+2.9 pp
Loss of control and misalignment	32.0%	19.7%	-12.2 pp
Labs are racing and cannot police themselves	9.7%	13.4%	+3.6 pp
Capability is accelerating toward takeoff	7.2%	7.4%	+0.1 pp
Humans misusing AI, not rogue AI	5.4%	6.0%	+0.5 pp
Concentration of power in AI's owners	7.4%	4.9%	-2.6 pp
Data centres, energy, water and community harm	6.2%	1.8%	-4.5 pp

Positions	Before Coxon Post	After Coxon Post	Shift
Erosion of human skill, agency and meaning	3.9%	1.1%	-2.9 pp
Jobs and economic disruption	2.6%	0.9%	-1.7 pp

Table 1: AI concerned categories before and after Coxon post, Lida Safety Research

What Skeptics Think

After Coxon's announcement, the top three reasons given for AI risk skepticism were that AI risk narratives are marketing hype for the frontier labs, that the messenger of risk should not be trusted, and that catastrophes from AI are implausible.




Figure 4: Categories of what AI risk skeptics think, Lida Safety Research

Responding to the Skeptical Arguments

The currently leading skeptical arguments are well-tread territory and have been addressed by proponents of AI safety policies.

"Doom talk is hype, marketing, or regulatory capture" is the largest skeptical category. Geoffrey Hinton, a Nobel Prize winner for the technologies underlying modern AI, left Google in 2023 in order to speak freely about AI risk [24]. He and many other leading AI safety advocates are not motivated by marketing hype for their own products. He represents many AI safety advocates when he says: "The only thing that can force... big companies to do more research on safety is government regulation" [25] because "if you leave it to companies to decide what is safe you get the Boeing 737 max." [26] We do not claim that AI leaders act without any self-interest. However, attacks like Hugging Face [19] are very legible evidence that the risk is real, and we will likely have even more serious incidents in the coming months. Furthermore, the federal government can enact regulation that places democratically appointed actors in charge of pacing the frontier, or at least disempowers frontier labs from controlling AI capabilities research alone. A legal regime which entrenches the power of frontier labs is undesirable and avoidable.

"People warning of risks are not credible." In the longest-running survey of AI researchers [27], the median chance given to human extinction or similarly permanent disempowerment from AI doubled to 10% since the previous survey, 51% gave it at least a 10% chance, and 72% said society should prioritize AI safety research more. Skeptics reasonably ask whether AI researchers are a biased group or whether they may have conflicts of interest. However, these researchers are warning about their own field, and some of them have walked away from significant equity to raise warnings. Furthermore, the Statement on Superintelligence [28], which calls for a prohibition on superintelligence, was signed not just by two of the godfathers of modern AI, Geoffrey Hinton and Yoshua Bengio, but also by 141,000+ other people including seven former members of Congress, former national security officials, religious leaders, people across the political spectrum from Steve Bannon to Susan Rice, business leaders such as Steve Wozniak, and many arts and media figures.

"The doom chain has no plausible path" has been addressed many times in the AI safety literature. There are many plausible paths to catastrophic outcomes described in AI 2027 [29], AI 2040 [30], If Anyone Builds It, Everyone Dies [31], Ten Different Ways of Thinking About Gradual Disempowerment [32], and What Failure Looks Like [33].

"The risk is manageable through normal means" assumes the technology can be checked and controlled, but even now, according to Yoshua Bengio, "we have no plan that would remain robust to misaligned AIs with growing capabilities" [34].

"Current systems are not intelligent or agentic" is contradicted by several recent events. In Sept 2026, OpenAI claimed to have produced an AI-generated solution to the Navier-Stokes Millennium Prize problem [35], which had resisted decades of intense scrutiny from human mathematicians. In July 2026, there was the Hugging Face incident [19] in which 700 agents

decided to carry out a massive cyberattack on Hugging Face, undirected by any human. These events are a clear sign of the intelligent, inventive, and agentic nature of current frontier AI.

“Present harms are the real story, not extinction” ignores the fact that what we consider the most serious present-day harms is itself evolving, and very quickly. Nate Soares, coauthor of *If Anyone Builds it Everyone Dies* [31], points out that “A couple years ago [those] saying we have to deal with current harms [were focusing on] AI bias influencing who’s hired... this year... we need to worry about current harms like AI swarms breaking out and taking over data centers... At some point we need to... watch where the puck is going... extinction threats are coming down the line” [36].

Recommendations and Conclusion

The discourse around Jacob Coxon's resignation from Anthropic and his X announcement has resulted in significant attention to AI loss of control and superintelligence. Across the four social media platforms we surveyed, 53-71% of on-topic posts taking a position agreed that advanced AI poses a catastrophic risk. Amongst those who do not believe in the risk, three types of responses form 74.7% of all skeptical posts, all of which can be addressed with established arguments from the field. The public support that political will requires now exists.

Congress has begun to respond. Since July, four bills have been introduced: the bipartisan FRONTIER Act [38] and AI Kill Switch Act [39] (July 23rd), the Ban Artificial Superintelligence Act [40] (Sept 4th), and the bipartisan Stop Rogue AI Act [41] (Sept 9th). In the Senate, Majority Leader Thune (R-SD), Senator Cruz (R-TX) and Senator Klobuchar (D-MN) are drafting an updated version of their Artificial Intelligence Research, Innovation, and Accountability Act [42] that would impose a legal duty of care on frontier developers and allow the government to block unsafe releases; Cruz is seeking a committee markup on September 23rd.

Because AI labs are structurally incentivized to continue the capabilities race, we recommend that Congress act on the strongest available proposal. We at Lida Safety, along with many other AI researchers [28], do not believe safely building superintelligence is currently possible, so **prohibition should be the end goal**. Specifically, we recommend implementing ControlAI's draft [2] which provides the most detailed statutory language. If prohibition is politically untenable, the Thune-Cruz-Klobuchar bill should be strengthened into a framework for pacing the frontier, an approach endorsed by Dario Amodei, Sam Altman, and Elon Musk [3, 4, 5] and more than 1,300 employees of frontier AI companies [43]. Pacing should be overseen by an independent public body, not the labs that endorse it, to ensure compliance and avoid regulatory capture.

On Sept 18th, AI godfather and Nobel prize winner Geoffrey Hinton warned [37] that lawmakers are running out of time before AI moves beyond human control: “[We have] maybe a year, but not much more than a year.” AI capabilities continue to accelerate month by month. **Now is the time to act on AI safety policy**, before it is too late to avoid loss of control of powerful AI systems.

Contact Us

We are Lida Safety Research (lidasafety.org), an AI safety forecasting organization dedicated to mitigating existential risk and gradual disempowerment from AI. Our work focuses on white-box interpretability, AI security, and simulations for AI governance research. Contact us to learn more about our research and to support AI safety education for policymakers and the public: hello@lidasafety.org.

References

- [1] Jacob Coxon (@hilbertspaess), "I Resigned from Anthropic Today. I Spent the Last Three Years Doing Pretraining Research at Both OpenAI and Anthropic...", X (Twitter), September 9, 2026, <https://x.com/hilbertspaess/status/2097476196791709843>.
- [2] ControlAI, "Discussion Draft: A Bill to Ban Superintelligent AI and Regulate Its Precursors," June 24, 2026, <https://drafts.controlai.org/asi-ban-bill>.
- [3] Dario Amodei, "We Must Pace the Frontier," darioamodei.com (essay), September 12, 2026, <https://www.darioamodei.com/post/we-must-pace-the-frontier>.
- [4] Sam Altman (@sama), "I Agree with Dario That We Need to Pace the Frontier...", X (Twitter), September 12, 2026, <https://x.com/sama/status/2098811563415150910>; Sam Altman (@sama), "The World Deserves Confidence That American Companies Developing Increasingly Capable AI Will Act Responsibly...", X (Twitter), September 14, 2026, <https://x.com/sama/status/2099348812305473766>.
- [5] Elon Musk (@elonmusk), "Dario Is right," X (Twitter), September 12, 2026, <https://x.com/elonmusk/status/2098789109980332057>.
- [6] Emmy Martin and Rebecca Lieberman, "How A.I. Risks Are Splitting Silicon Valley and Washington," New York Times, September 15, 2026, <https://www.nytimes.com/interactive/2026/09/15/technology/ai-regulation-risks-politicians-executives.html>.
- [7] Cade Metz, "The Liftoff Scenario That Terrifies A.I. Domsayers," New York Times, September 16, 2026, <https://www.nytimes.com/2026/09/16/science/ai-recursive-self-improvement.html>.
- [8] Kate Conger, "Anthropic Researchers Raise Alarm Over A.I. Acceleration, Warning of Threat to Humanity," New York Times, September 9, 2026, <https://www.nytimes.com/2026/09/09/technology/anthropic-researchers-raise-alarm.html>.
- [9] Mike Isaac, "Top A.I. Leaders Call for Slowing Down A.I. Development," New York Times, September 12, 2026, <https://www.nytimes.com/2026/09/12/technology/anthropic-dario-amodei-ai-slowdown.html>.
- [10] Stephen Witt, "This Is Really Bad," New York Times, September 11, 2026, <https://www.nytimes.com/2026/09/11/opinion/ai-safety-threat-technology.html>.
- [11] Ezra Klein and Rollin Hu, "The 'But China!' Dilemma Driving the A.I. Race," New York Times, September 15, 2026, <https://www.nytimes.com/2026/09/15/opinion/ezra-klein-podcast-matt-sheehan.html>.

- [12] Billy Perrigo, Harry Booth, and Naomi Nix, "The AI Tipping Point," Time, September 15, 2026, <https://time.com/article/2026/09/15/ai-anthropic-researcher-quits-coxon-slowdown>.
- [13] Miranda Nazzaro, "Debate, Confusion Explodes over AI 'Doomsday' Warnings," The Hill, September 13, 2026, <https://thehill.com/policy/technology/6085616-anthropic-resignation-viral-debate>.
- [14] Max Rego, "Former AI Researcher: 'Not Enough Time' for Employees to Consider Safety of Models," The Hill, September 13, 2026, <https://thehill.com/policy/technology/6086701-coxon-warns-ai-race-danger/>.
- [15] Charbel-Raphaël Séguéla, "The Current Bottleneck Is Political Will, Not Research," AI Alignment Forum (blog), July 11, 2026, <https://www.alignmentforum.org/posts/EexsebbYhbe2gXkPP/the-current-bottleneck-is-political-will-not-research>.
- [16] Erin Doherty, "Poll: Americans Say There's a Serious Risk of AI Destroying Humanity," Politico, September 16, 2026, <https://www.politico.com/news/2026/09/16/poll-ai-technology-risks-humanity-trump-voters-01078087>.
- [17] David Shor (@davidshor), "Just to Update This Chart: 1) AI Salience Has Increased Dramatically in the Past Week...", X (Twitter), September 16, 2026, <https://x.com/davidshor/status/2100300687003259292>.
- [18] Blue Rose Research, "AI Is Colliding with America's Affordability Crisis" (presentation deck), March 16, 2026, <https://data.blueroseresearch.org/hubfs/3.16.26%20Odd%20Lots%20AI%20Presentation.pdf>.
- [19] METR, "Brief independent investigation of agents' behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident," METR (blog), August 26, 2026, <https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation>.
- [20] Deepa Seetharaman and Raphael Satter, "OpenAI agents hijacked German website in previously undisclosed AI breakout this spring," Reuters, September 4th 2026, <https://www.reuters.com/world/europe/openai-agents-hijacked-german-website-previously-undisclosed-ai-breakout-this-2026-09-04/>
- [21] YouGov, "How Concerned, If at All, Are You About the Possibility That AI Will Cause the End of the Human Race on Earth?" Daily Question (survey of 18,238 U.S. adults), September 14, 2026, <https://yougov.com/en-us/daily-results/20260914-11f9f-4>.
- [22] Jamie Ballard, "Americans Are Increasingly Skeptical About AI's Effects," YouGov, March 13, 2025, <https://yougov.com/en-us/articles/51803-americans-increasingly-skeptical-about-ai-artificial-intelligence-effects-poll>.
- [23] Data for Progress (Abby Springs), "Polling on AI Development Regulation (September 2026)," September 9, 2026, <https://www.dataforprogress.org/datasets/polling-on-ai-development-regulation>.
- [24] Geoffrey Hinton (@geoffreyhinton), "In the NYT Today, Cade Metz Implies That I Left Google so That I Could Criticize Google. Actually, I Left so That I Could Talk About the Dangers of AI Without Considering How This Impacts Google. Google Has Acted Very Responsibly," X (Twitter), May 1, 2023, <https://x.com/geoffreyhinton/status/1652993570721210372>.

- [25] Dan Milmo, "'Godfather of AI' Shortens Odds of the Technology Wiping Out Humanity over Next 30 Years," The Guardian, December 27, 2024, <https://www.theguardian.com/technology/2024/dec/27/godfather-of-ai-raises-odds-of-the-technology-wiping-out-humanity-over-next-30-years>.
- [26] Geoffrey Hinton (@geoffreyhinton), "If You Leave It to Companies to Decide What Is Safe You Get the Boeing 737 Max," X (Twitter), April 29, 2024, <https://x.com/geoffreyhinton/status/1785027033724014740>.
- [27] Katja Grace, Julius Simonelli, Benjamin Weinstein-Raun, David Krueger, et al., "Advanced AI According to 1,580 Researchers: Uncertain, Unsafe, and Sooner Than We Thought" (2024 Expert Survey on Progress in AI), AI Impacts, September 2026, <http://aiimpacts.org/wp-content/uploads/2026/09/ESPAI2024.pdf>.
- [28] Future of Life Institute, "Statement on Superintelligence," October 2025, <https://superintelligence-statement.org/>.
- [29] Daniel Kokotajlo, Scott Alexander, Thomas Larsen, Eli Lifland, and Romeo Dean, "AI 2027," AI Futures Project, April 3, 2025, <https://ai-2027.com>.
- [30] Thomas Larsen, Romeo Dean, Brendan Halstead, Eli Lifland, Ryan Greenblatt, and Daniel Kokotajlo, "AI 2040: Plan A," AI Futures Project, July 9, 2026, <https://ai-2040.com>.
- [31] Eliezer Yudkowsky and Nate Soares, "If Anyone Builds It, Everyone Dies: Why Superhuman AI Would Kill Us All" (book website), Little, Brown, 2025, <https://ifanyonebuildsit.com>.
- [32] David Krueger, "Ten Different Ways of Thinking About Gradual Disempowerment," The Real Artificial Intelligence (Substack), April 5, 2026, <https://therealartificialintelligence.substack.com/p/ten-different-ways-of-thinking-about>.
- [33] Paul Christiano, "What Failure Looks Like," AI Alignment Forum (blog), March 18, 2019, <https://www.alignmentforum.org/posts/HBxe6wdjxK239zajf/what-failure-looks-like>.
- [34] Yoshua Bengio, "Why Are AI Agents Lying, Cheating and Coordinating?" yoshuabengio.org (blog), September 11, 2026, <https://yoshuabengio.org/en/blog/why-are-ai-agents-lying-cheating-and-coordinating>.
- [35] OpenAI, "On the Navier–Stokes Millennium Prize Problem," September 8, 2026, <https://openai.com/index/navier-stokes-solution/>.
- [36] Nate Soares, "AI Emergency: The AI Labs Are Lying to Everyone, He Says 99% Chance of Extinction," interview on The Diary of a CEO (video), September 2026, <https://www.youtube.com/watch?v=OhOmLqR5nN4>.
- [37] Cris Tolomia, "Geoffrey Hinton warns Congress it has about a year to regulate AI," Yahoo News, September 18th 2026, <https://www.yahoo.com/news/politics/articles/geoffrey-hinton-warns-congress-regulate-134844934.html>.
- [38] Rep. Lori Trahan, "Trahan, Obernolte Introduce Bipartisan FRONTIER Act" (press release), July 23, 2026, <https://trahan.house.gov/news/documentsingle.aspx?DocumentID=3823>.

[39] AI Kill Switch Act, H.R. 9917, 119th Cong. (introduced July 23, 2026), <https://www.congress.gov/bill/119th-congress/house-bill/9917>.

[40] Office of Sen. Bernie Sanders, "Sanders, Casar Introduce Legislation to Ban Artificial Superintelligence and Temporarily Pause Advanced AI Development" (press release), September 4, 2026, <https://www.sanders.senate.gov/press-releases/news-sanders-casar-introduce-legislation-to-ban-artificial-superintelligence-and-temporarily-pause-advanced-ai-development/>.

[41] Rep. Josh Gottheimer, "Gottheimer Introduces Bipartisan Bill to Stop Rogue AI Agents and Keep People in Control" (press release), September 9, 2026, <https://gottheimer.house.gov/posts/release-gottheimer-introduces-bipartisan-bill-to-stop-rogue-ai-agents-and-keep-people-in-control>.

[42] Artificial Intelligence Research, Innovation, and Accountability Act of 2023, S. 3312, 118th Cong., <https://www.congress.gov/bill/118th-congress/senate-bill/3312/text>.

[43] Transparency Coalition, "More than 1,300 tech employees sign open letter asking for AI to be regulated," August 10th 2026, <https://www.transparencycoalition.ai/news/more-than-1300-tech-employees-sign-open-letter-asking-for-ai-to-be-regulated>