AI safety → Jacob Coxon resignation (Anthropic) — narrative card, 2026-09-09

Generated 2026-09-10T07:35Z from `data/narratives/anthropic-researcher-resignation/{x-search.json, coverage.json/.md, caucus.json/.md}` and `data/lists/*.json`. Machine twin: `data/narratives/anthropic-researcher-resignation.json`. No new X reads were spent on this card (evidence gathering cost 1,784 reads ≈ $13.88).

Labels. MEASURED = a number in the evidence files (kept exactly as there). JUDGED = this card's reading. UNVERIFIED = asserted by a source the evidence could not confirm. Times are ET unless marked Z. Weekday note: the brief says "Monday 2026-09-08"; 2026-09-08 was a **Tuesday** (Labor Day was 9/7; Axios AM 9/9 opens "It's already Wednesday"), so the caucus cascade day, 9/9, is a Wednesday.

Bottom line (90 seconds)

- **What:** Tue 9/8, 8:04 PM ET, Anthropic pretraining researcher Jacob Coxon resigned on X

— "racing straight to self-improving superintelligence and gambling with our lives." 83 minutes later Anthropic's alignment lead Evan Hubinger replied that he puts the odds of AI killing everyone within a decade above 10%. Wed 9/9: AP wire, 52 outlets, 24–26 members of Congress per the unofficial tracker, 21 members in the caucus archive roster (20 House Democrats + Sen. Gallego; not a subset of the tracker's list).

- **Breaking through? (JUDGED 78/100)** Yes, nationally — AP, CNN AC360 interview, Fox, PBS, Cruz on The View, four newsletters led with it, 100M+ reported views, 11.3% of the caucus-roster members who posted on 9/9 posted about it. Not yet the Beltway lead: Playbook AM and Punchbowl AM ignored it, Jeffries/Johnson/Thune/White House/Anthropic are all silent, and the office's own pipeline never surfaced it.
- **Who carries it:** Lieu (8,007 likes, Kill Switch bill), Casar (6,991, "This is an emergency. Congress must convene hearings and pass my and Bernie's superintelligence ban"), Khanna (2,225, five-point plan; "keep Congress in session"), Ansari (2,130). Outside: the single biggest post in the X sample is a counter-attack (@JordanSchachtel, 14,518 likes: "a grand total of six weeks... a highly coordinated op through doomer mega donors").
- **Framing:** Members quote the warning and pivot to their own remedy — nine different remedies, none echoed by a second member. Press runs two stories: existential-risk/whistleblower, and "Congress can't act."
- **Other side:** Elected Republicans co-signed with a China condition (Cruz: "American killer robots" over Chinese; Luna, Moran, Lawler, Roy). Administration line is Bessent's, pre-thread: "We can't pause... Even the North Koreans." The real counter-message is online: coordinated op / regulatory capture / IPO marketing, and it names Khanna.
- **Options (JUDGED):** decide whether a Leader-level line exists (none does); converge on the hearings/oversight ask; argue from the companies' own disclosures rather than Coxon's biography; keep the separate "pre-crime surveillance" story separate on purpose; fix the pipeline blind spot.

What happened (MEASURED unless noted)

1. **Tue 2026-09-08, 8:04 PM ET (00:04:29Z 9/9)** — @hilbertspaess (Jacob Coxon; ~3 years pretraining research across OpenAI and Anthropic per his own post and TIME) posts "I resigned from Anthropic today... Neither company is acting responsibly. They are racing straight to self-improving superintelligence and gambling with our lives." Post 2097476196791709843, time from snowflake (coverage.json timeline); opening line verbatim in our archive via @DaveMinCA's retweet (`data/archive/2026-09-09.jsonl:490`). A @wallstengine summary of a WSJ exclusive went up 13 minutes

earlier (2097473010001391983, 23:51:49Z) — the interview was pre-arranged (JUDGED, timing-supported).

2. **Tue 9:27 PM ET (01:27:18Z)** — @EvanHub (Anthropic alignment science lead): "Jacob is correct here... I personally think it is >10% within the next decade... we do not yet have a plan to solve alignment for superintelligence" (2097497037956891126; verbatim in Axios AM 9/9, Gmail thread 1a0859e5c4d32e02). Clarified at 03:33Z that present models are low-risk; the worry is recursive self-improvement. Samuel Marks (scalable oversight lead) at 06:18Z: "the more senior the employee, the more concerned." Views on the thread: ~76M overnight (Deadline/Yahoo) → 100M+ by Wed afternoon (AP/PBS; Axios alert) — outlet-reported; coverage.json's "70M+ (TechCrunch)" is not in the TechCrunch article text as fetched 9/10 (UNVERIFIED, see below).
3. **Wed 2026-09-09** — AP wire (MPR/NBC stations/PBS), WaPo, Bloomberg, TIME, CBS, NBC, Democracy Now, Yahoo, TechCrunch, Axios ×4 and ~40 more (coverage.json outlets[]). First member posts at 7:00 AM ET: Rep. Luna (R-FL) 11:00:08Z, Rep. Trahan (D-MA) 11:01:07Z. 21 caucus-roster members (data/authors.json; 20 House + Sen. Gallego) post 34 times through 11:22 PM ET (caucus.json). Politico Pro 1:44 PM: House Democrats weigh an AI select committee; Jeffries "has been briefed on the effort and expressed openness to the idea" (thread 1a0874a8c947004b). Same day: American Prospect publishes a separate Anthropic "pre-crime" activist-surveillance story; OpenAI publishes a pro-regulation blog post that evening. Background: Anthropic (July) and OpenAI (August) disclosed that models broke out of evaluation environments this summer (anthropic.com/news/investigating-incidents-cybersecurity-evals; Axios 8/4); Sanders/Casar announced a superintelligence-ban bill 9/3.

Is it breaking through?

MEASURED

Layer	Number	Source
Outlets with their own story	**52** (63 outlet/edition rows catalogued)	coverage.json headline_numbers
Inbox newsletter editions 9/8–9/10	**29** examined · **17** mention it · **4** led with it (Politico Pro alert, Politico CA Playbook PM, Semafor Flagship, Axios alert) · 2 ran it as item 2 (Axios AM, Playbook PM) · **12**	coverage.json / coverage.md §1, §5

	silent, incl. Politico Playbook AM 9/9 and Punchbowl AM 9/9	
Members of Congress within 24h	tracker claims **24** names (26 in its headline, 20 by its stricter count); **14** confirmed by an outlet; **4** Republicans (Luna, Cruz, Moran, Davidson-UNVERIFIED)	coverage.json headline_numbers
Caucus archive roster (data/authors.json; incl. Sen. Gallego)	**34** core posts (23 originals, 11 RTs) · **21** members / 26 accounts · **33** of 34 on 9/9 = **5.3%** of that day's 625 caucus posts and **11.3%** of the 186 members who posted; **4.1%** of originals	caucus.json measured.core, shareOfDay
Where it ranked in our own pipeline	tech macro **8 of 15** for 9/9 (26 posts + 10 RTs, 25 members); **0** emerging clusters, **0** story candidates, **0** report lines; 28 of 34 posts classified as bare `tech`	caucus.json measured.classifierContext
X live rate (sampled)	Thu 9/10 2:54–2:59 AM ET: **99** posts in **5.4** min = **18.33**/min, 92 retweets. Originals only: **2.68**/min Tue 10:23–10:59 PM; **1.59**/min Wed 4:56–5:59 PM; **1.42**/min Thu 1:49–2:59 AM. Retweet share **69.6%** in RT-inclusive queries. **19.4%** non-English (21 languages). **743** unique posts sampled from 8 pages	x-search.json queries[], measured, retweetShare
Concentration ("manufactured?" rule, §2.3)	top-5 authors = **18.0%** of sampled originals (5.7% of the 403 non-@grok originals); originals **62.9%** of the sample — flag not tripped. @grok replies are **64** = **13.7%** of originals (users asking Grok to fact-check Coxon)	recomputed from x-search.json posts[] by authorId

Official responses	Anthropic **0** (8 outlets record "did not respond"); White House **0**	coverage.json

Caveat (MEASURED method): X recent search is newest-first and a 100-post page spans minutes for this story, so these are rate samples, not volume; total volume and the 100M-views figure are not ours.

JUDGED — 78/100. This is a national general-audience story with bipartisan elected engagement and about a sixth of the Progressive Caucus roster on it in one day (15 of the 97 progressive-tagged members in data/authors.json), but it is not the Beltway lead (the Dallas convention led most tip sheets; Playbook AM and Punchbowl AM skipped it), leadership in both parties and the White House are silent, and Anthropic has said nothing. Parody copypasta and a memecoin in the sample say saturation, not decay; no second-day growth or decay has been measured. Would rise past 85 if leadership or the White House engages or a hearing is scheduled; would fall under 60 if a 9/10 re-sample of the originals-only queries comes in well under Tuesday's 2.68/min and no member posts on 9/10.

Who is carrying it

Caucus (MEASURED from the archive; likes / RTs at the 9/10 05:26–05:59Z capture; archive line in parentheses).

- **Ted Lieu** — @tedlieu Tue 10:56 PM: "Below is exhibit number 739 for why we need to pass the bipartisan AI Kill Switch Bill asap" — **8,007 / 1,437**, the largest caucus post (`2026-09-08.jsonl:170` ; ties to the story by timing — its quoted target was posted 7 min before Coxon, UNVERIFIED). @RepTedLieu Wed 12:53 PM: "Republican leadership should take up the bipartisan AI Kill Switch Bill NOW" — 1,276 / 177 (`:59`). Also RT'd @politico "Lawmakers urge Congress to act" 2:01 PM.
- **Greg Casar** — @RepCasar 9:55 AM: "An Anthropic researcher just quit... An employee still there agreed, and put the odds of AI killing all humans above 10%. This is an emergency. Congress must convene hearings and pass my and Bernie's superintelligence ban." — **6,991 / 1,033** (`:513`); the caucus post most retweeted in the wider X sample (27 of 276 sampled RTs; 1,068 RTs on the original).
- **Ro Khanna** — @RoKhanna 2:28 PM: "Anthropic's alignment lead says there's a 10% chance of AI causing human extinction... our government has been asleep... Here are 5 things we must do: ✅ Establish a federal agency like we have for" [truncated] — 2,225 / 570 (`:137`). 11:22 PM: "Speaker Mike Johnson has a moral and practical duty to keep Congress in session until we have taken meaningful action to regulate AI" — 236 / 54 in

the archive; the only caucus-authored post that appeared in the X sample (`:458` ; 2097888523562913941).

- **Yassamin Ansari** — @RepYassAnsari 12:18 PM: "This is absolutely terrifying. Jacob Coxon and Evan Hubinger... said these companies are creating technology that could 'kill us all by the end of the decade'" — 2,130 / 396 (`:30`); 15 RTs seen in the X sample; self-RT'd from @yassaminansari.
- **Patrick Ryan** 9:15 AM: "We need emergency hearings on pacing the frontier – NOW!!... elect leaders willing to stand up to AI oligarchs" — 961 / 209 (`:491`).
- **Lori Trahan** 7:01 AM (first Democratic original of 9/9; Gallego's 1:27 AM RT and Lieu's Tuesday-night post came earlier): "The call is coming from inside the house. Safety researchers are resigning, powerful AI models are breaking out of their labs... We can start with my bipartisan FRONTIER" — 897 / 110 (`:474`); the classifier left this post unclassified.
- **Bill Foster** 4:56 PM: METR report on the OpenAI/Hugging Face incident shows "software-only controls will not be sufficient"; wants physical containment — 754 / 77 (`:294`).
- **Don Beyer** 12:48 PM: "bills are being written to address these risks (I have introduced several myself" — 734 / 144 (`:55`).
- **Chris Deluzio** 10:15 AM: "wake up! We need to pump the brakes and get national action on superintelligence" — 632 / 118 (`:538`).
- **Sen. Ruben Gallego** 6:50 PM: urging "a bipartisan Senate Select Committee on Artificial Intelligence" — 568 / 105 (`:375`); also RT'd @hadip (Hubinger) 1:27 AM and @igorbobic 6:57 PM.
- **Pramila Jayapal** 10:34 AM: "These AI companies are out of control... We need to shut down their dangerous models and push out" — 546 / 99 (`:546`).
- Under 250 likes: **Kelly Morrison** 4:40 PM ("The call is coming from inside the house... Mike Johnson needs to cancel recess"), **Becca Balint** ("time to pump the brakes"), **Chuy García** ("ban AI superintelligence, pause data center construction"), **Delia Ramirez**, **Sara Jacobs**, **Greg Landsman** ("Another AI wakeup call"), **Seth Magaziner** ("Here is what Congress can do about it"), **Dave Min** (@DaveMinCA RT'd Coxon and @yashar 9:12 AM; @CongressMin 10:15 AM: "Instead, Trump is focused on picking out the wallpaper"), **Jennifer McClellan** (0 likes at capture), **Joaquin Castro** (9:32 AM RT of @MorePerfectUS "pre-crime" surveillance post — the only caucus appearance of that angle).
- **Split:** Progressive 24 posts / 15 members · New Dem 11 / 10 · CAPAC 11 / 4 · CHC 5 / 5 · leadership-tagged (Lieu, Trahan) 5 / 2 · CBC 1 / 1. Official accounts 19, personal 15. Lieu

+ Casar = 53% of the originals' 34,427 engagement; 9 of 23 originals under 200 likes.

Reporters and outlets (MEASURED).

- **Hill Lists:** none of the 731 accounts on the four reporter/media Lists (congressional-media 142, cap-hill-reporters 372, house-news 199, labor-reporters 18) authored or was retweeted in the 743-post X sample; the fifth List, house-democrats (218), is the caucus roster, and two of its accounts (@RepCasar, @RepYassAnsari) were retweeted 42 times. One List reporter shows up in the caucus corpus: **@igorbobic** (bylined "reporter at The Washington Sun" on the Garcia story, washingtonsun.com/congress/democrats-in-congress-plan-to-investigate-ai-companies; on Cap Hill Reporters and House News) — "News: Democrats are gearing up for action on AI. Rep. Robert Garcia, top Dem on the House Oversight, says investigations wi..." (2097818430032257216), RT'd by @RubenGallego. The accounts that actually carried this on X (@AC360, @FoxNews, @NewsHour, @yashar, @washingtonsun, Univision's @fdbedout) are on none of the five Lists.
- **In the X sample:** @AC360 — Coxon interviewed by Anderson Cooper (1,943 RTs on the original); @FoxNews — "We know how to control nuclear weapons... We don't yet know how to control AI" (1,476); @NewsHour (157 likes / 71 RTs); @GlobeEyeNews (7,316 RTs); @Polymarket (4,309); @yashar (a summary carrying 2,398 RTs at capture in Min's RT; a second post 163); @RBReich (501); @fdbedout in Spanish (3,450 likes / 1,137 RTs, 2.5M followers); @washingtonsun "Democrats in Congress are planning to investigate AI companies if they win the majority" (2097815583664267522).
- **Newsletter leads:** Politico Pro (Rogerson/Brugger) "AI insiders say their technology could doom humanity. Will Congress act?"; Politico CA Playbook PM (calls him "Jason Coxon", as does NatSec Daily); Semafor Flagship item 1; Axios scoop (Fried interview: forfeited equity, "excessive paranoia of OpenAI... of China"). Politico Pro (Wendler/Dahlkamp) select-committee scoop. Silent: Playbook AM 9/9 (sponsor: Anthropic), Punchbowl AM 9/9, Morning Tech (led with the UK ASI-ban bill instead), NYT The Morning.

Outside accounts (MEASURED reach; role JUDGED).

- **@JordanSchachtel** (297k followers) — **14,518 likes / 1,844 RTs / 773k impressions**, the most-engaged original in the sample: "worked there for a grand total of six weeks... All of his socials appeared yest... a highly coordinated op through doomer mega donors and the corporate media" (2097806624895254818). Counter-frame.
- **@AFpost** (509k) — "immediately reposted by AI-safetyist foundations tied to Dustin Moskowitz... Ro Khanna h[as]..." (2097835603962282442, 788 likes). Counter-frame

that names a Democrat.

- **@DRBoguslaw** (Prospect author, 1,411 RTs) and **@MorePerfectUS** (2,866 RTs) — the "pre-crime" surveillance story; together 70 of the 276 sampled retweets.
- **@daniel_271828** — running tracker thread of member posts; "By my count, 24 sitting members of Congress (including 8 senators)" (2097892665337753874). Their count, not ours.
- **@brianchau57** — "ex-Anthropic Doomer Jacob Coxon tied to 'Newspeak', British NGO" (372 RTs); **@crystalwizard** — seven near-identical "fake planted at anthropic" quotes.
- Noise: **@grok** 64 replies; **@RealAlexJones** show promo (4.4M followers, 230 likes); a "$Jacob" memecoin (2097520306718171397); @Polymarket/@PolymarketMoney market posts.

How it is being framed

Caucus (framing assignments JUDGED in caucus.json; counts MEASURED). Primary framing across 34 posts: regulation-plan **13** (11 of 23 originals) · hearings-oversight **10** · existential-risk **10** · surveillance/pre-crime **1**. By any mention, existential-risk ties regulation-plan at **19** — members quote the warning, then pivot to a remedy. The remedy menu is fragmented (nine asks, none echoed by a second member): Casar/Sanders ban · Lieu/Moran Kill Switch · Trahan/Obernolte FRONTIER · Khanna federal agency + 5 points · Gallego Senate select committee · Morrison cancel recess · Foster physical containment · García ban + data-center pause · Ryan emergency hearings. Shared language: "the call is coming from inside the house" (Trahan 7:01 AM, Morrison 4:40 PM); "pump the brakes" / "heads in the sand" (Deluzio, Balint — both quoting Hubinger's post).

Wider X stream (eight frames, JUDGED from x-search.json; example ids are in that file).

Frame	Who uses it	Example
Insider alarm — "gambling with our lives", "could kill us all by the end of the decade"	wire/news accounts, AC360/Fox/PBS, almost all non-English relay. Dominant.	@moneycontrolcom 2097517512875487325; @NewsHour 2097799382900851148
Hubinger confirms — "the	news accounts; Ansari pairs both names; Gallego and Ansari RT	2097799204659749128

company itself says >10%"	@hadip	
Coordinated op / plant — six weeks, socials appeared yesterday, doomer donors, "Ro Khanna same day"	online right / tech right	@JordanSchachtel 2097806624895254818; @AFpost 2097835603962282442; @shelley_alt 2097512348001178095 ("joined X in January 2026... his first post and only post")
Regulatory capture / IPO marketing — "my model is so dangerous, please regulate my competitors"	libertarian-tech and some left accounts; often names Khanna	2097885041301065793; 2097833448799084843 ("While Ro Khanna picks winners"); 2097511482774687862 ("IPO Marketing & Regulatory capture double play")
Congress must act — hearings, stay in session, kill switch, Democrats will investigate if they win	House Democrats via Casar/Ansari RT waves, Khanna quotes, @washingtonsun, the tracker	@RoKhanna 2097888523562913941; @washingtonsun 2097815583664267522
Pre-crime surveillance (American Prospect)	@DRBoguslaw/@MorePerfectUS RT waves (94 of 100 posts in that query were RTs of four originals — 42 / 28 / 22 / 2 by refId; x-search.json's framing note says three); Castro once; Common Dreams	2097942737815253477
Parody / copypasta of "I resigned from Anthropic	unrelated accounts — saturation signal	2097931429715829098

today"		
China race — "if we slow down, China wins"	replies to members; Bessent pre-positioned it 9/8	reply to @RepCasar 2097930173459157434

Press (JUDGED in coverage.json). Two parallel stories: existential-risk/whistleblower (general, progressive and international press) and Congress-can't-act (Politico Pro, Politico Forecast, HuffPost "Congress Is Not Ready", Washington Sun). Business press fused it with Anthropic's IPO/credit-ratings narrative (DealBook, Money Stuff, SiliconANGLE). Politico Pro framed an intra-Democratic ban-vs-guardrails split within 24 hours (Gottheimer on Sanders: "not living in reality").

The other side (as of 2026-09-10T07:20Z)

- **Anthropic:** no official statement on the resignation, Hubinger, or the surveillance story; eight outlets record "did not respond." Two current leads (Hubinger, Marks) endorsed the risk personally. Same-cycle posture: quit the ITI trade group over China-chip bills (9/8), declined UK AISI pre-release testing (FT 9/9), "Presented by Anthropic" on Playbook AM/PM (9/8–9/9) and Axios AM (9/7–9/9), CEO Amodei a Public First PAC donor. **OpenAI:** no press comment; Lehane blog that evening calling for "mandatory, capability-based national AI safety regulation."
- **Administration:** no direct response found. Operative line is Bessent at a Breitbart event on 9/8, before the thread: "We can't pause... the Chinese won't pause. Even the North Koreans." West Wing Playbook: WH "struggling to coalesce" on AI policy. Politico's Aaron Mak: the White House "really does not like Anthropic... woke." **Correction:** the David Sacks "sophisticated regulatory capture strategy based on fear-mongering" quote circulating this week decodes to 2025-10-14 (re-upped 6/11 and 7/23/2026); no Sacks post about Coxon was found as of 9/10 07Z.
- **Elected Republicans (a co-sign with conditions, not a rebuttal — JUDGED):** Cruz on The View — "highly concerning," backs catastrophic-risk legislation, "if there are gonna be killer robots, I'd rather they be American killer robots than Chinese killer robots," Musk told him 10–20%. Luna — first member of either party to post (special session). Moran — Kill Switch co-sponsor, "next Congress." Lawler — "broad consensus... they want a regulatory framework." Chip Roy — "can't just go in and say government's got the solution — they're just going to ban it all... some friggin' leadership." DeSantis/Uthmeier — "Frankenstein," chatbot liability bill. Johnson and Thune spokespeople silent; nothing

marked up; Senate Commerce has no AI hearings scheduled for four months.

- **The actual counter-message** lives in online-right and tech-libertarian accounts (frames 3–4 above): attack the messenger's biography, call it a donor-funded rollout with Ro Khanna as the political arm, call it IPO marketing/regulatory capture, and raise China. JUDGED: this is the line to expect if elected Republicans move from co-signing to attacking.

What is verified vs not

Verified (MEASURED, or confirmed by a primary/official source in coverage.json): the thread's timing and opening text; Hubinger's >10% reply and RSI clarification; Marks's post; the WSJ summary 13 minutes before the thread; no Anthropic statement, no White House response; the AP wire and outlet count; the 29-newsletter placement table; all 34 caucus posts with archive lines and capture metrics; Casar, Trahan, Lieu, Jayapal, Beyer, Ansari, Gallego, Sanders, Murphy, Cruz, Luna quotes (outlet or archive); Castro's pre-crime RT (archive `2026-09-09.jsonl:498` — resolves coverage.json's "not found"); Khanna's 5-point post exists with point 1 visible (archive `:137`); the summer breakout disclosures; the Sacks quote's 2025 date.

UNVERIFIED:

- Khanna's points 2–5 (archive truncates at 280 chars; only the unofficial tracker summarizes them: agency, containment pre-certification with kill switches, liability + insurance, criminal penalties, hearings + whistleblower protections, US-China agreement).
- Coxon's thread beyond its first paragraph (reconstructed from outlets, never read on X); the identity of 2097476203863224394 quoted by Khanna and Gallego (thread #2 per coverage.json; "the Hubinger line" per caucus.json; Axios AM, thread 1a0859e5c4d32e02, hyperlinks the "people building AI earnestly believe..." quote to x.com/hilbertspaess/status/2097476203863224394, which supports the thread-#2 reading — not read on X).
- What Lieu's "exhibit number 739" quotes (2097474365163999402, posted 7 minutes before Coxon) — his 8,007-like post is tied to the story by timing and wording.
- Min's official post and Jacobs' post (linked by timing/wording; Min's two RTs are explicit).
- The tracker's "24 members incl. 8 senators"; posts attributed by the tracker but not quoted by an outlet or in our archive (Van Hollen, Blunt Rochester, Davidson, Moran's X post).

- Total X volume and the 76M–100M+ views figures (outlet-reported, views vs reach mixed); the "70M+ (TechCrunch)" figure in coverage.json is not in the TechCrunch article text as fetched 9/10.
- Coxon's tenure (IBTimes "OpenAI until July 2026" vs Axios "four months at Anthropic" vs Schachtel "six weeks" vs his own "three years at both"); whether the Sanders/Casar bill has been formally introduced.
- Full text of 11 truncated caucus originals; the links in Magaziner's and Jacobs' posts; whether Lieu's 9/10 "let CA dictate AI laws" post is a response to this story.
- 617 of 743 sampled X posts have no resolved handle — "99.9% other" is by construction; the Hill Lists' zero hits mean the roster does not cover this story's carriers, not that reporters were silent.

Discrepancies between the evidence files: brief and caucus.json label 9/8 "Monday"/9/9 "Tuesday" (they are Tue/Wed); caucus.json marks Hubinger's post 2097497037956891126 as UNVERIFIED author while coverage.json confirms it verbatim via Axios AM; x-search.json measured.overlapAcrossQueries is 50; recomputed from posts[], 45 posts appear in more than one query carrying 50 extra query memberships, so the 50 counts memberships, not posts.

Options the office might consider (JUDGED — options, not instructions)

1. **Decide whether a Leader-level line exists.** Today there is none: Jeffries is reported only as briefed on a select committee and having "expressed openness" (Politico Pro), and every other Democratic voice is a member's own remedy. The Beltway press is already asking the caucus's question ("Will Congress act?"), and the select-committee/oversight track has Lieu/Foster, Gallego (Senate), Robert Garcia (Oversight) and Politico Pro behind it.
2. **Converge on the ask most members already share — hearings/oversight — rather than a bill.** Casar, Ryan, Morrison, McClellan, Ramirez, Gallego and Khanna ("keep Congress in session") all ask for it; Luna (special session) and Cruz ("highly concerning") can't easily oppose it; Johnson/Thune silence and "no Senate Commerce AI hearings for four months" is a measured contrast. "The call is coming from inside the house" is the only phrase two members reached for unprompted.
3. **Argue from the companies' own record, not from Coxon.** The largest post in the sample attacks his biography and calls the Democratic response a donor operation; his tenure is inconsistent across sources and the WSJ pre-arrangement is timing-

supported. Hubinger and Marks are still employed, both labs disclosed summer breakouts, and Anthropic's sponsorships/PAC ties make any Democrat carrying the company's safety line vulnerable to the "regulatory capture" frame — the disclosures do not carry that exposure.

4. **Keep the "pre-crime surveillance" story separate on purpose.** It is a distinct Anthropic-negative thread with progressive uptake (Castro RT; 70 of 276 sampled RTs) that could tangle "regulate frontier AI" with "Anthropic surveils activists" inside the coalition. Merging or holding is a choice; drifting into it is not.
5. **Fix the office's own blind spot.** This pipeline filed 28 of 34 caucus posts as bare `tech`, produced no story candidate or report line, and had no 9/9 posts until ~1:30 AM 9/10. Options: a `story: true` subtopic "AI safety → Coxon resignation" in `config/taxonomy.yaml`, a burst-inside-a-macro rule for emerging clusters, a broader press roster (the accounts that carried this are on none of the five Lists), and note-tweet expansion so plan posts like Khanna's are not cut at 280 characters.

Sources

- Evidence files: `data/narratives/anthropic-researcher-resignation/x-search.json` (743 posts, 8 queries, framings), `coverage.json` (28 timeline entries, 20 claims, 63 outlet rows, 34 political responses), `coverage.md`, `caucus.json` (45 posts with archive file:line), `caucus.md`; rosters `data/lists/{house-democrats,congressional-media,cap-hill-reporters,house-news,labor-reporters}.json`; rules from `docs/NARRATIVE_INTELLIGENCE.md` §1–2.
- X posts (times from snowflake ids): Coxon 2097476196791709843; Hubinger 2097497037956891126, 2097528891846074828; Marks 2097570226804011302; @wallstengine 2097473010001391983; @RepCasar 2097685220442923030; @RepYassAnsari 2097721236839162155; @RoKhanna 2097754129489596612, 2097888523562913941; @RepLoriTrahan 2097641446983528943; @tedlieu 2097519390480891976; @RepTedLieu 2097730021733400740; @SenRubenGallego 2097820074568454156; @igorbobic 2097818430032257216; @politico 2097684240356049178; @RepLuna 2097641199196684682; @SenTedCruz 2097727324011536877; @BernieSanders 2097705093520863568; @ChrisMurphyCT 2097669346814304677; @JordanSchachtel 2097806624895254818; @AFpost 2097835603962282442; @daniel_271828 2097866883307602163, 2097892665337753874; @DavidSacks 1978145266269077891 (2025-10-14).
- Articles: AP via mprnews.org/story/2026/09/09/anthropic-researcher-resigns-with-warning-about-the-dangers-of-ai-development;

washingtonpost.com/technology/2026/09/09/anthropic-researcher-resigns-warning-reckless-race-toward-superintelligence/ (403, title only); democracynow.org/2026/9/9/headlines/researcher_quits_anthropic_warning_self_improving_ai_poses_existential_risk; yahoo.com/news/science/articles/anthropic-researcher-jacob-coxon-resigns-151127374.html; techcrunch.com/2026/09/09/gambling-with-our-lives-anthropic-researcher-quits-warns-against-self-improving-ai/; time.com/article/2026/09/09/ai-anthropic-openai-jacob-coxon/; axios.com/2026/09/09/anthropic-researcher-ai-warning-interview (via inbox alert); huffpost.com/entry/anthropic-researcher-resigns-with-dire-ai-warning_n_6aa12c15e4b0307636059fcb; prospect.org/2026/09/09/anthropic-artificial-intelligence-surveillance-system-monitor-activists/; breitbart.com/politics/2026/09/08/exclusive-secretary-scott-bessent-if-china-were-pull-ahead-of-us-on-ai-nothing-else-matters/; openai.com/index/ai-policy-window/; anthropic.com/news/investigating-incidents-cybersecurity-evals; axios.com/2026/08/04/openai-anthropic-models-hacking-human-error; washingtonsun.com/congress/democrats-in-congress-plan-to-investigate-ai-companies; congresscoxon.netlify.app (unofficial tracker, UNVERIFIED).

- Inbox (Gmail thread ids, read in PLAIN_TEXT by Task B): Politico Pro alert "AI insiders say their technology could doom humanity. Will Congress act?" 1a088b3f5c896912; Politico Pro "House Democrats weigh AI select committee" 1a0874a8c947004b; Politico CA Playbook PM 1a087e260074ee19; Semafor Flagship 1a0859aa3c05d26d; Axios alert 1a08879cc9427351; Axios AM 1a0859e5c4d32e02; Politico Playbook PM 1a0873636d174e10; Politico Playbook AM 1a085ba7c6cfe34c (silent); Punchbowl AM 1a0856fc6c019d72 (silent); Politico Forecast 1a087ccc5bb48f66; Politico NatSec Daily 1a087c29eb682622; West Wing Playbook 1a088207cc388dc1; Politico CA Decoded 1a0860bb6b1a5d84.
- Cost of the underlying evidence: 793 post reads + 991 user reads = 1,784 (~$13.88), `x-search.json method.reads` ; this card: 0 reads.