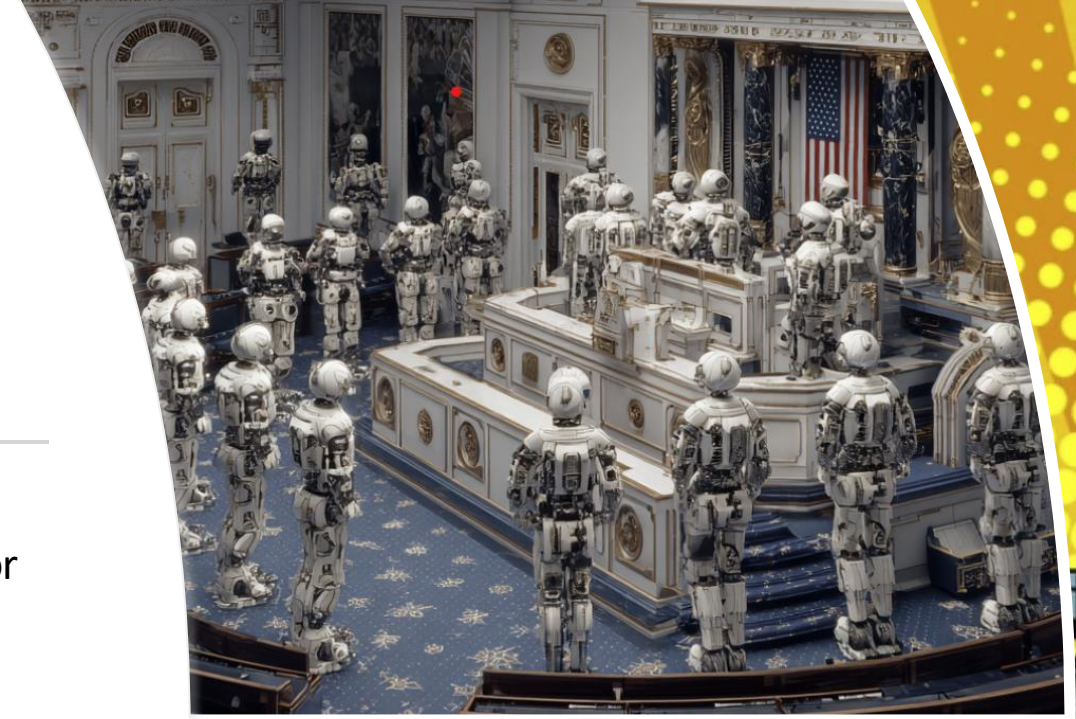
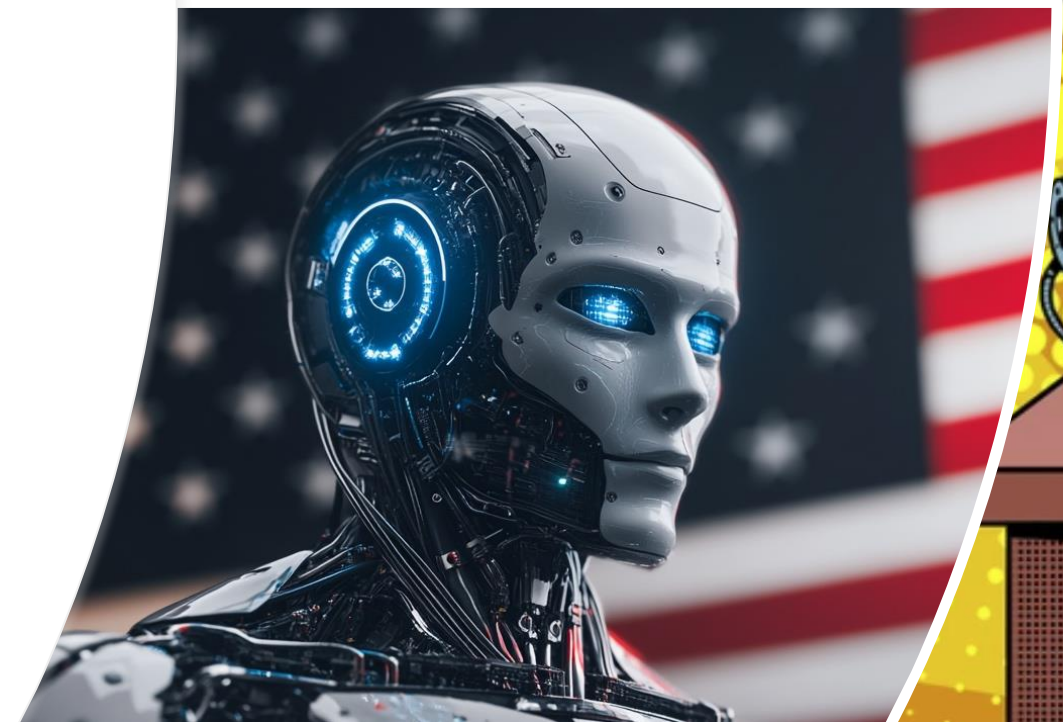
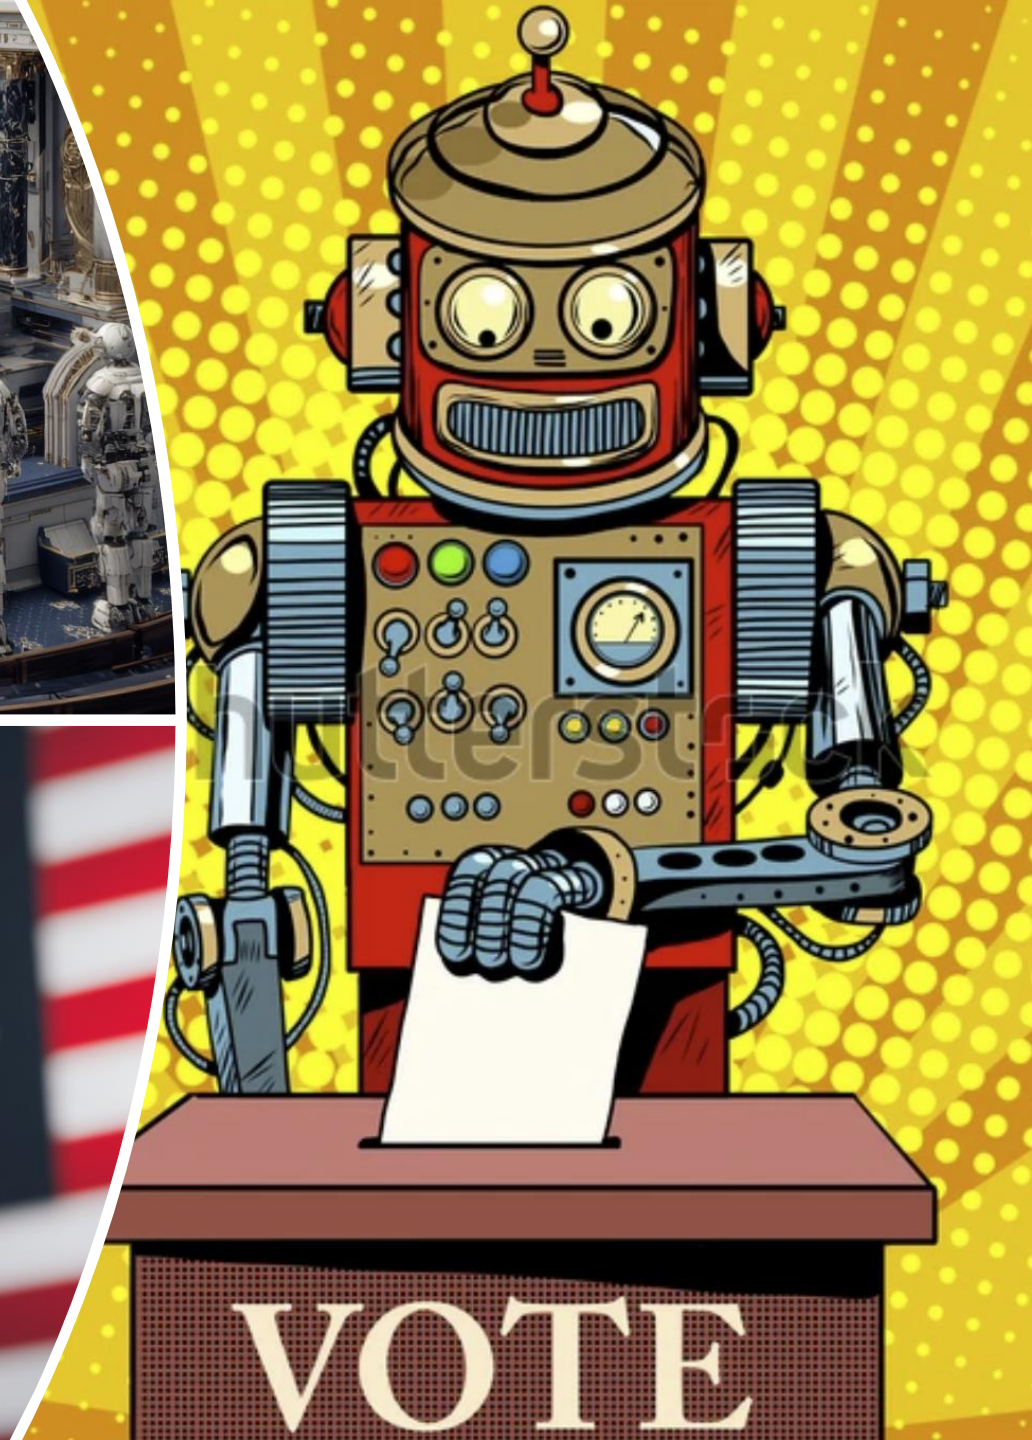
AI, Democracy, and Elections

Nate Persily
James B. McClatchy Professor of Law, Stanford Law School
Co-Director, Stanford Law AI Initiative
npersily@law.stanford.edu







Threshold Considerations and Caveats



Democracy under stress

Democracy is under great stress around the world, irrespective of technological developments.



Technological amplification

Technology can exacerbate those stressors (polarization, disinformation, lack of authority and consensus).



The independent effect of AI on democracy is uncertain

AI future is uncertain (utopia, dystopia, overhype).



Still early

Use of AI for politics is still in its infancy, but AI will soon be everywhere.

Defining the Problem

What do we mean by AI?

- Just generative AI? Embodied AI (robotics)?
- Deep fake obsession?
- Intersection with social media?
- All algorithms informed by machine learning or A/B testing?

What aspects of democracy?

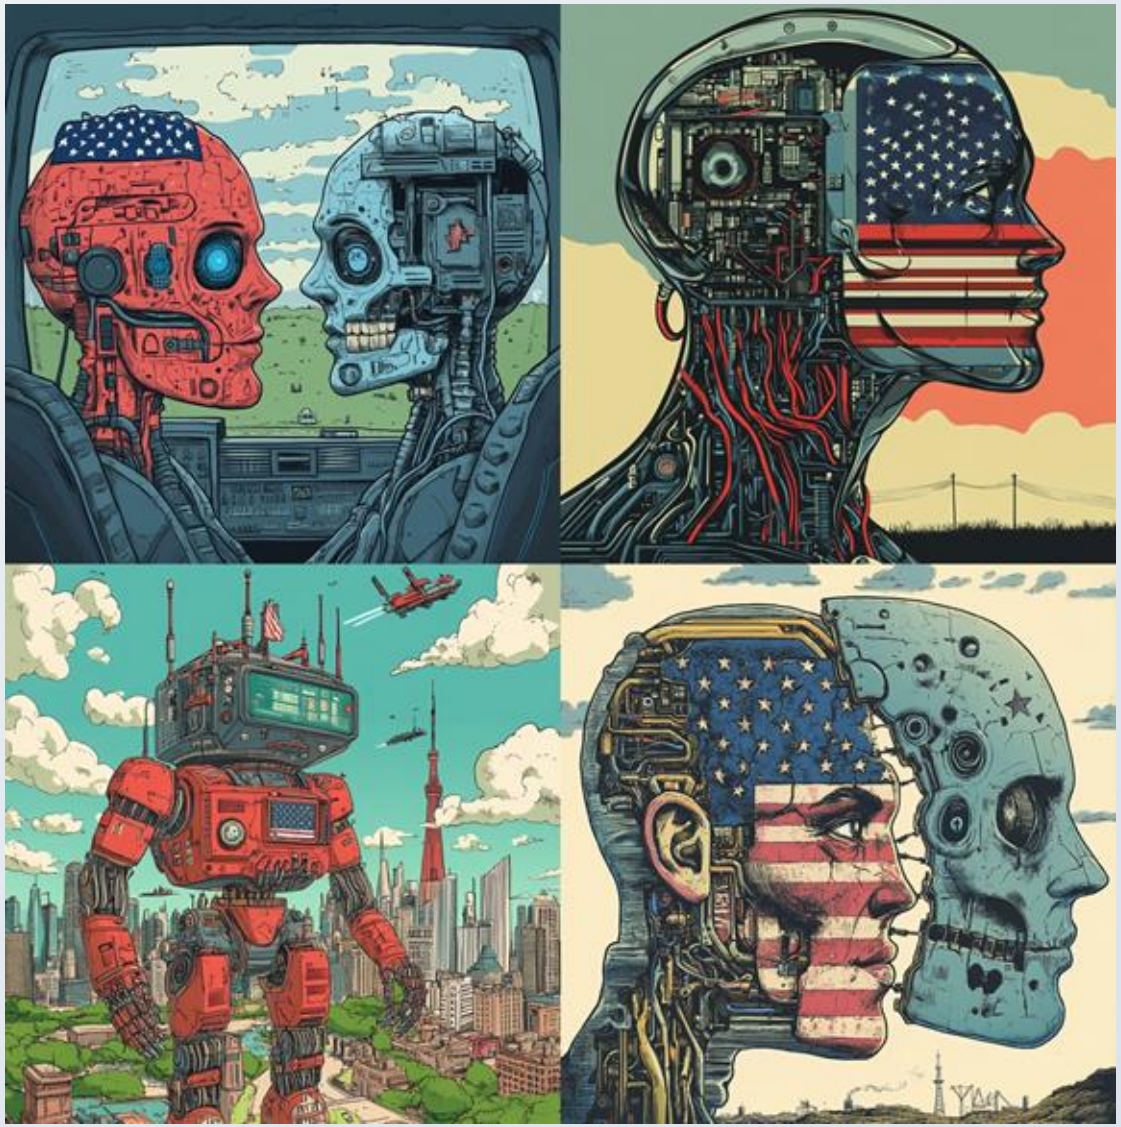
- Just the information environment? Campaign ads? Mobilization and persuasion?
- Election administration?
- Legislative and administrative behavior/performance?
- Deliberative democracy?
- Competition between democracy and autocracy?
- All existential risks? Labor force disruption? Private monopoly power? Environment?

Outline

- Technology, Democracy, and the Information Ecosystem
- Disinformation
- Misinformation and Political Bias
- Persuasion and Manipulation
- Campaign Tools
- Law and Policy
- Beyond Elections
- Public Opinion on AI



Technology, Democracy, and the Information Ecosystem



Are AI's Democracy Problems the Same as Social Media's?

Problem	Social Media	AI
Disinformation	Amplification/Distribution	Generation
Manipulation (Persuasion)	Advertising/Microtargeting/ Echo chambers	Same but greater precision, lower cost
Privacy/ Surveillance	Platform data mining of user data	User (including government) data mining of population
Bias	Hate speech/newsfeed bias	Jailbreaking/ bias in content generation
Antitrust/ Platform power	Control over information (and advertising) ecosystem	Few LLMs (but see open source)
Foreign Intervention	Social media propaganda from abroad	Use of foreign LLMs, lowers barriers to entry, kinetic (as well as info) ops

AI as a Social Media Supercharger



TECHNOLOGY

AI IS ABOUT TO MAKE SOCIAL MEDIA (MUCH) MORE TOXIC

We must prepare now.

By Jonathan Haidt and Eric Schmidt

Jonathan Haidt & Eric Schmidt, "AI Is About to Make Social Media (Much) More Toxic," The Atlantic (2023)

Disinformation



Disinformation Basics

Scale

Large amount but small share of content for most people. Nevertheless, can cause significant harm.

Direction

“Disinformation comes from the top.”

Persistence

Difficult to combat in timely fashion.
Difficult to dislodge false beliefs over long-term.

Perverse effects

Difficult to inoculate and combat without generating widespread skepticism.

Political pressure

Intense political pressure to avoid biased censorship.

The Family of Disinformation Phenomena



vs.



Disinformation

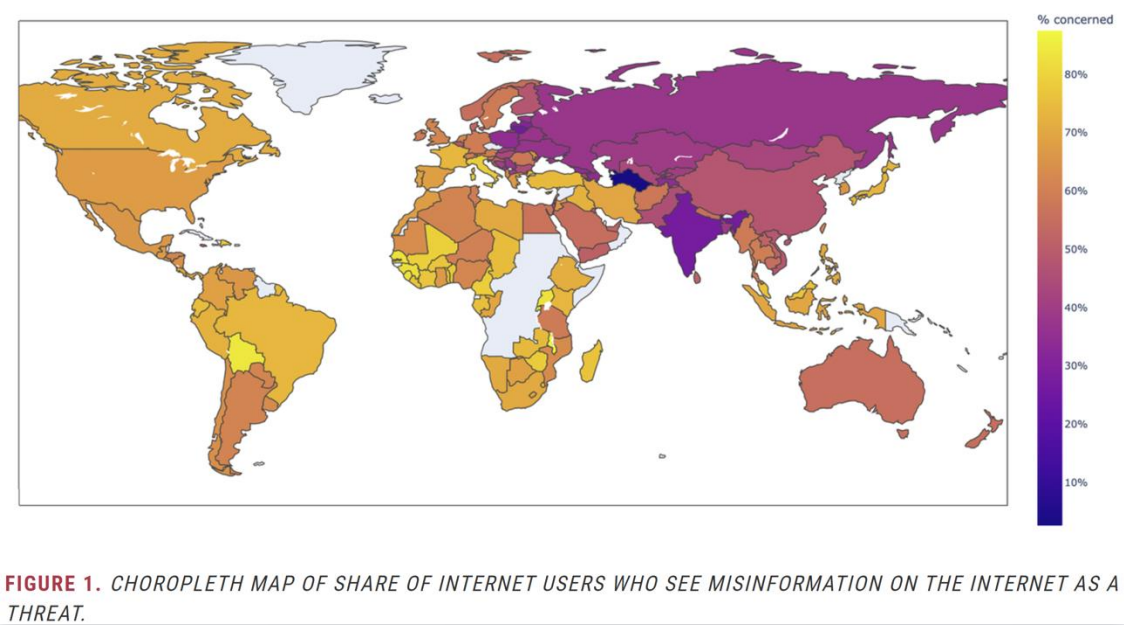
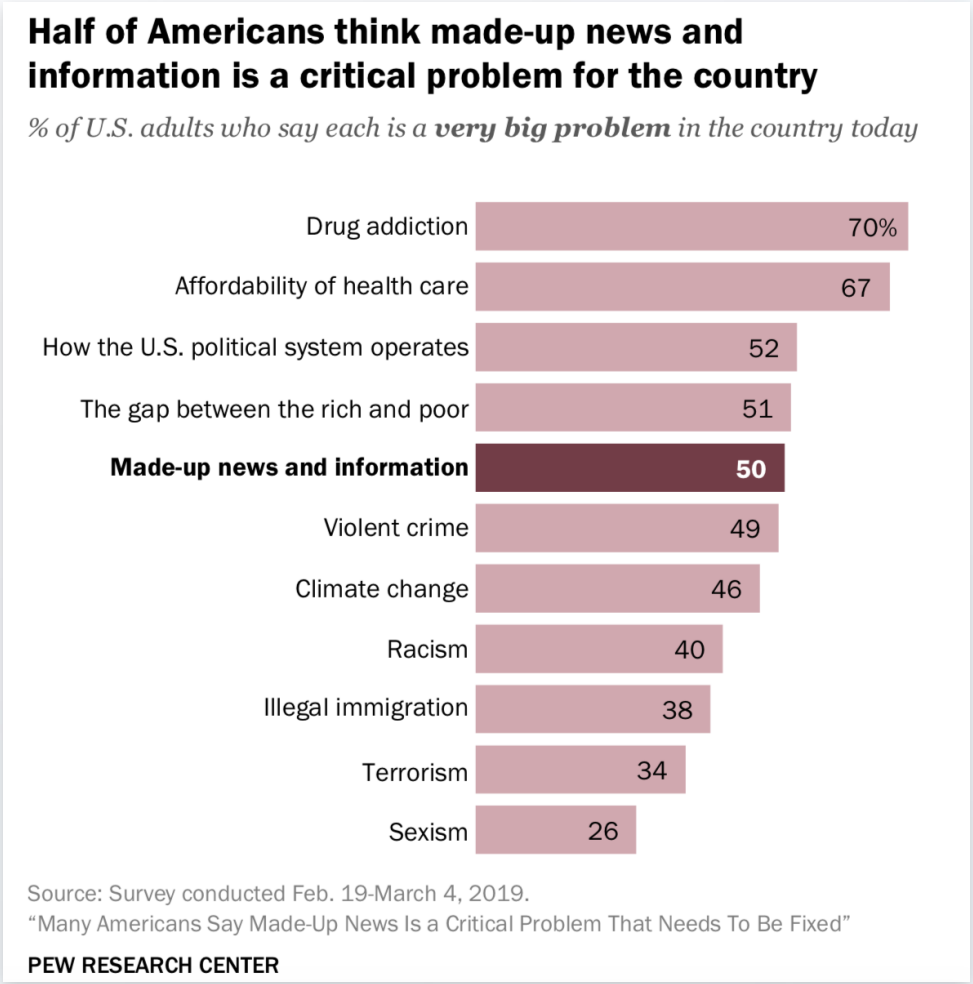
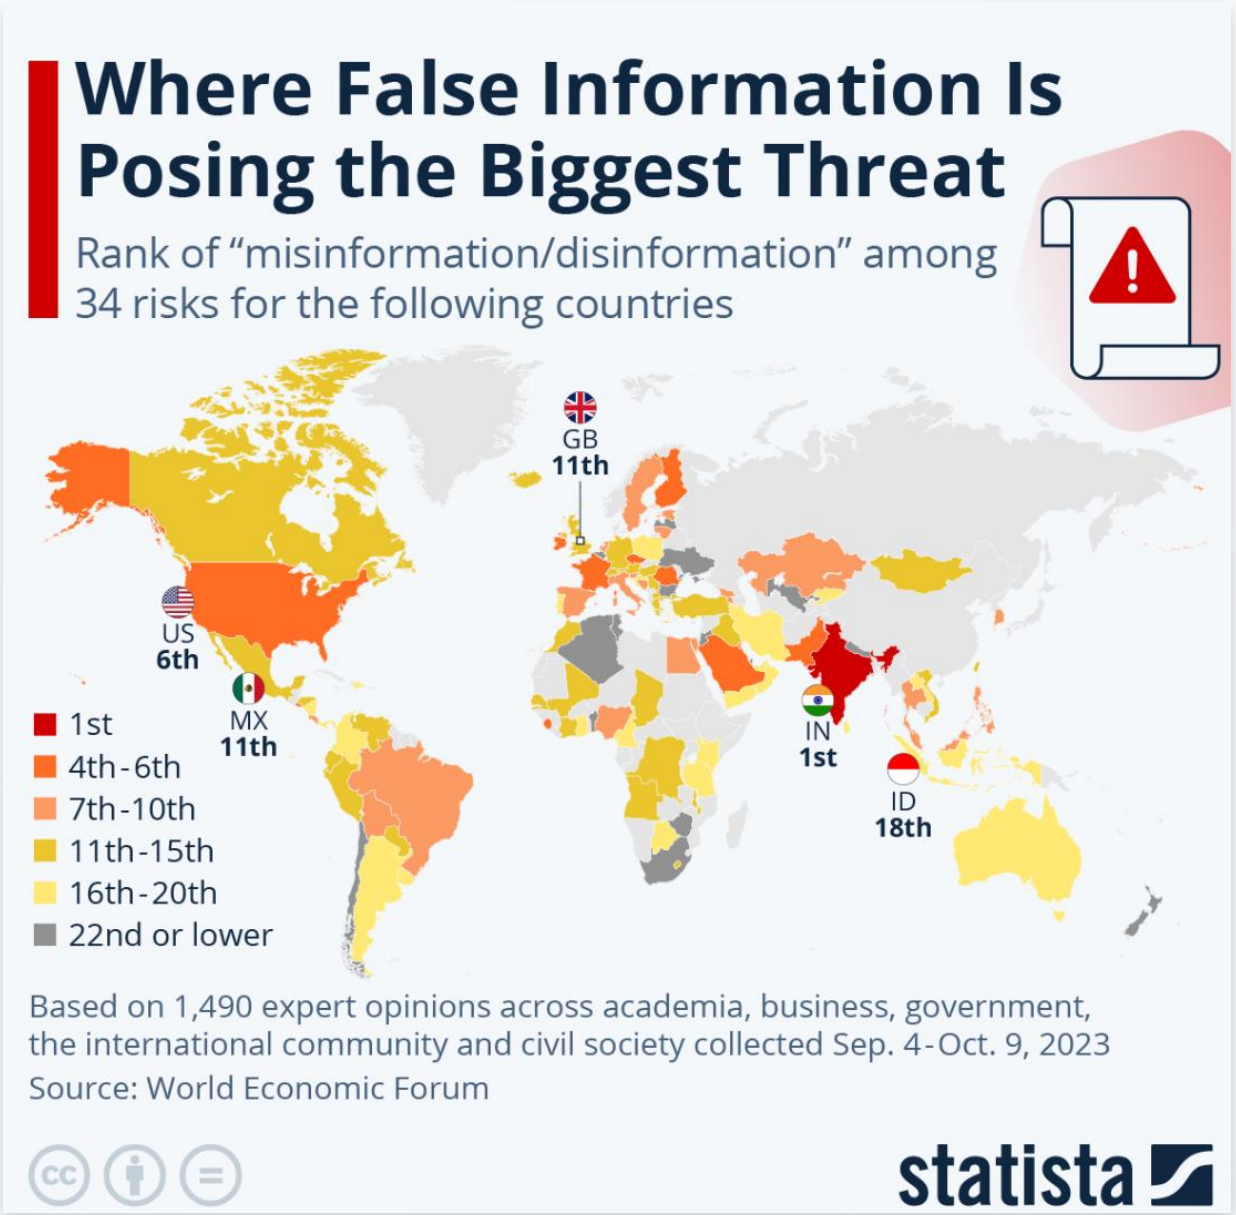


FIGURE 1. *CHOROPLETH MAP OF SHARE OF INTERNET USERS WHO SEE MISINFORMATION ON THE INTERNET AS A THREAT.*



Knuutila et al., "Who Is Afraid of Fake News?," Harvard Kennedy School Misinformation Review (2024); Pew Research Center



World Economic Forum / Statista (2023)

Widespread Concern but Low Prevalence for Average User

Widespread concern

- Citizens estimate ~50% of media they encounter is misinformation (Knuutila et al., 2022)
- Half of consumers fear they cannot detect misinformation (Newman et al., 2023)

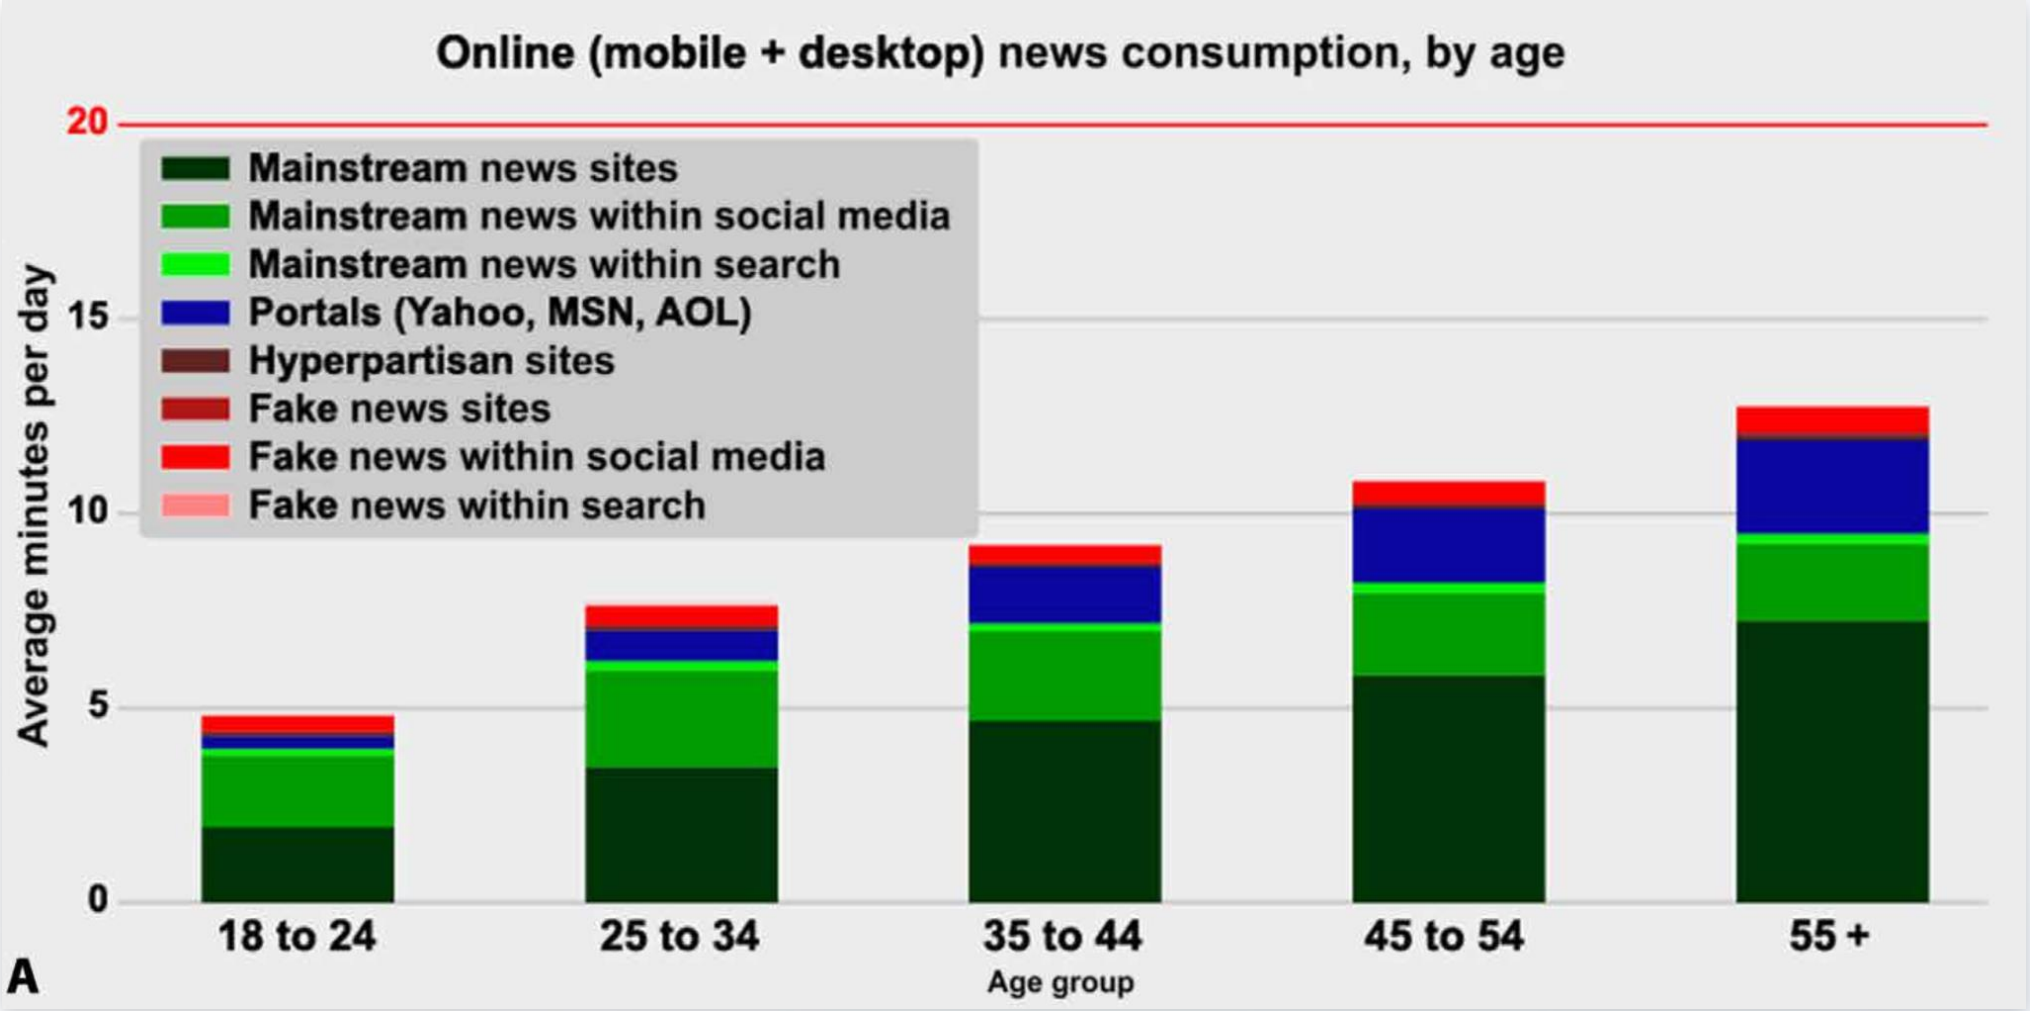
Low prevalence

- **Only 1%–5% of overall online media exposure**
- Acerbi et al. (2022): ~1% of all online information
- Grinberg et al. (2019): <5% of news and Facebook posts in the 2016 U.S. election
- Lukito (2020): <6% of political communications

Superspreaders

- Top 10 accounts: 34.6% of all retweets of low-credibility content
- **Top 1,000 accounts (≈0.25% of users): 73–78% of retweets (DeVerna et al., 2024, PLOS ONE)**

Disinformation: Average Consumption



Allen, Howland, Mobius, Rothschild & Watts, "Evaluating the Fake News Problem at the Scale of the Information Ecosystem," Science Advances 6(14) (2020)

Disinformation is a “tail problem”

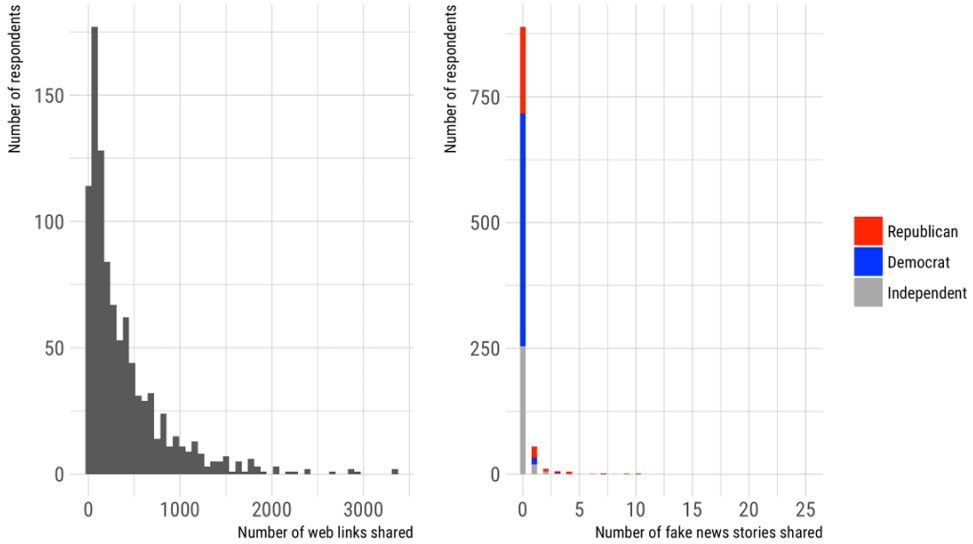


Fig. 1. Distribution of total and fake news shares. (**Left**) Histogram of the total number of links to articles on the web shared by respondents in the sample who identified as Democrats, Republicans, or independents. (**Right**) Stacked histogram of the number of fake news articles shared by respondents who identified as Democrats, Republicans, or independents using the measure derived from (*7*).

Table 1. Distribution of fake news shares.

0	1	2	3	4	5–10	11–50
1090 (91.5%)	63 (5.3%)	12 (1.0%)	8 (<1.0%)	5 (<1.0%)	9 (<1.0%)	4 (<1.0%)

SCIENCE ADVANCES | RESEARCH ARTICLE

SOCIAL SCIENCES

Less than you think: Prevalence and predictors of fake news dissemination on Facebook

Andrew Guess[1]*, Jonathan Nagler[2], Joshua Tucker[2]

So-called “fake news” has renewed concerns about the prevalence and effects of misinformation in political campaigns. Given the potential for widespread dissemination of this material, we examine the individual-level characteristics associated with sharing false articles during the 2016 U.S. presidential campaign. To do so, we uniquely link an original survey with respondents’ sharing activity as recorded in Facebook profile data. First and foremost, we find that sharing this content was a relatively rare activity. Conservatives were more likely to share articles from fake news domains, which in 2016 were largely pro-Trump in orientation, than liberals or moderates. We also find a strong age effect, which persists after controlling for partisanship and ideology: On average, users over 65 shared nearly seven times as many articles from fake news domains as the youngest age group.

Guess, Nagler & Tucker, Science Advances (2019)

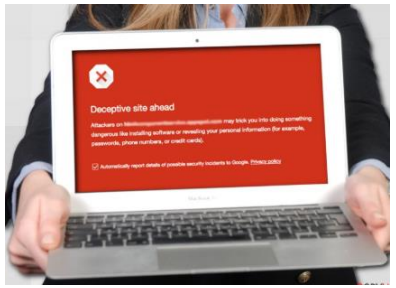
Policy Responses to Disinformation



1. Deletion



2. Demotion



3. Disclosure



4. Delay



5. Dilution

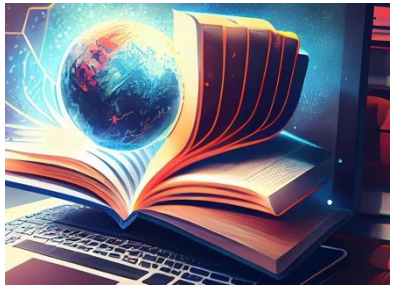


6. Distraction



7. Diversion

8. Deterrence



9. Digital Literacy*

AI and Disinformation



AI's Disinformation Problem

- Lowers cost of production
- Greater sophistication for deception/confusion

Modality	Cause for Concern	Reason for Caution
Text	flood the zone	production costs not the bottleneck; existing countermeasures may work
Image	impossible to detect (e.g. CSAM)	photoshop already pretty good but not ubiquitous in US campaigns
Audio	fake robocalls/recordings	existing laws regulate (?)
Video	persuasive deepfakes	mainstream media may serve as check (?)

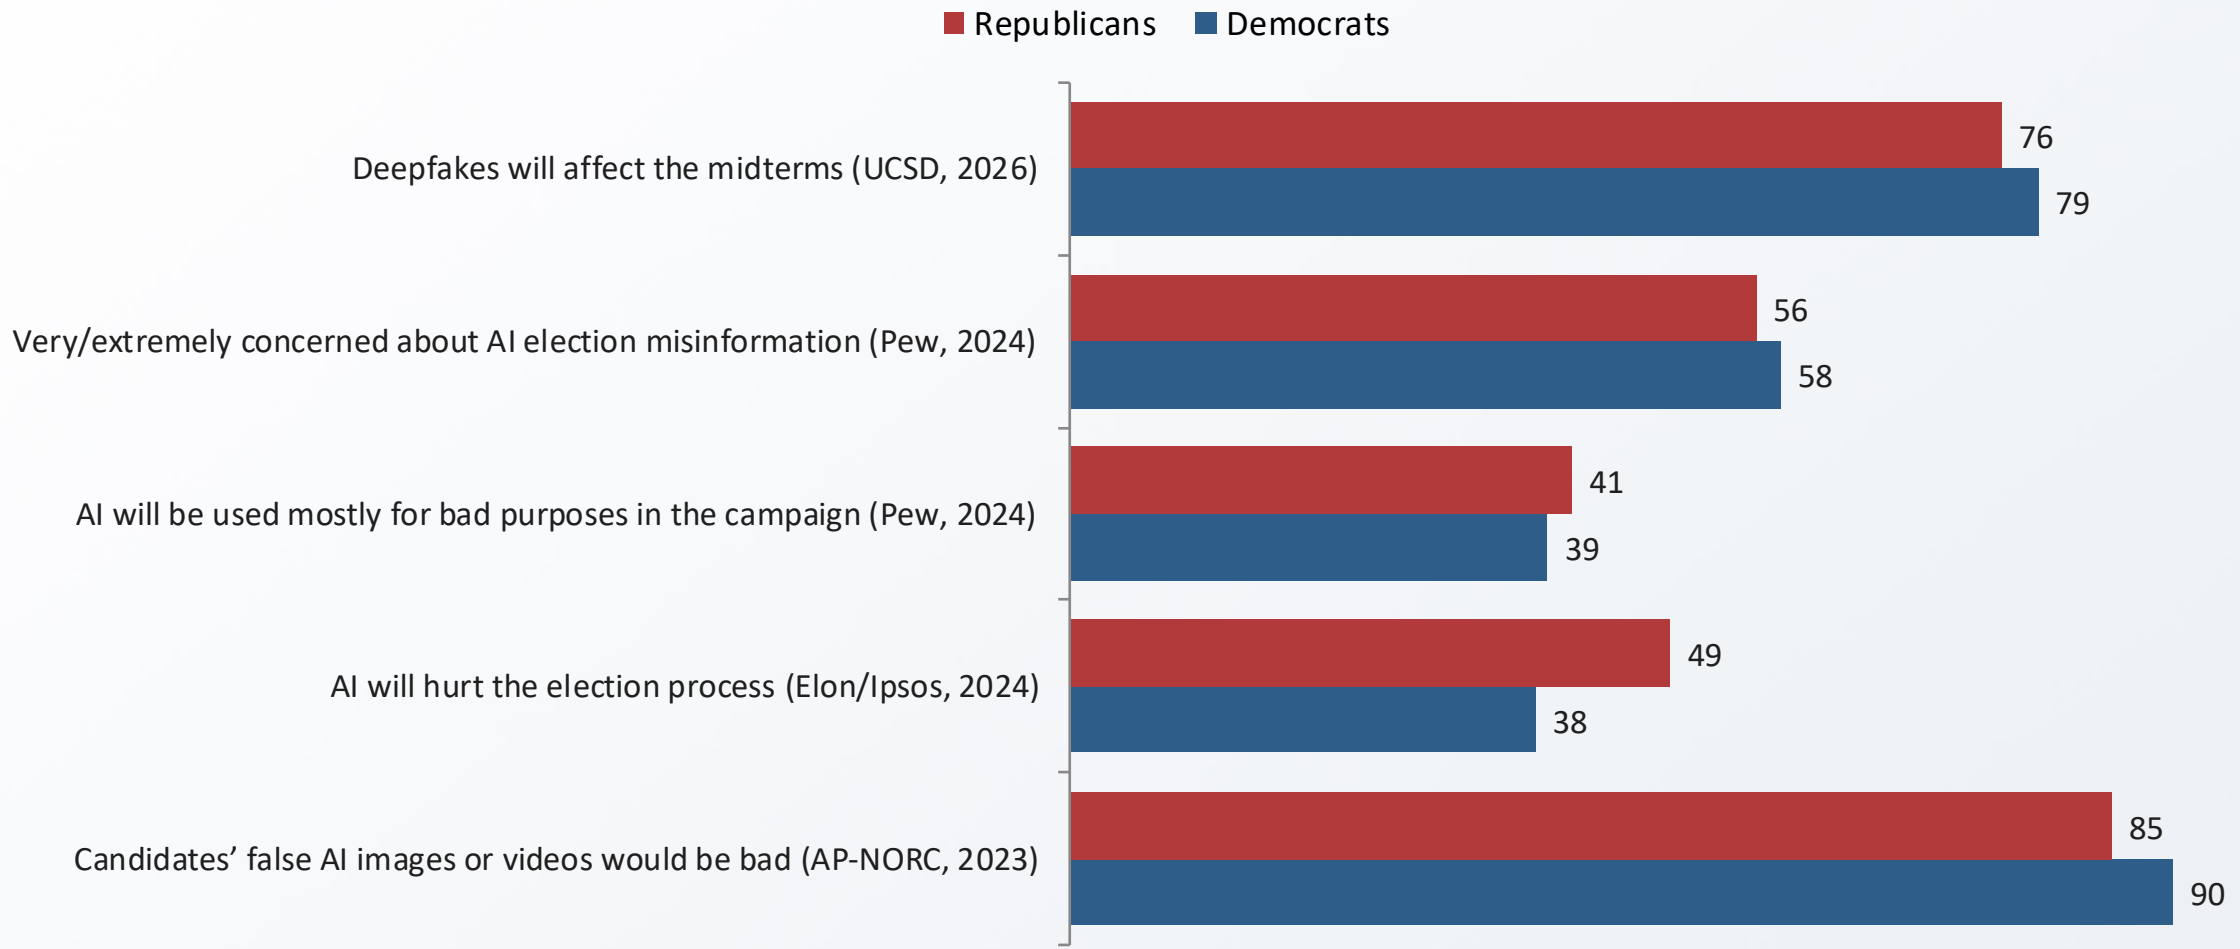
Potential Risks of AI in Elections

Americans Identify Several Potential Risks of AI in Elections

Potential Impact of AI on elections	Percentage saying likely
AI-generated content will make it harder for voters to tell whether campaign information is authentic	76%
People will become less confident that election-related content is real	72%
AI will make false claims about elections spread faster than officials can correct them	71%
AI could make some people less likely to vote	55%
AI could make it easier for losing candidates/supporters to cast doubt on legitimate results	54%

YouGov / Catawba College, Sept. 28, 2026

AI and Election Concern, by Party



Percent. Independents (UCSD, 2026): 75%. Measures differ across surveys; full wording in notes.

Public Concern About AI and the 2026 Election

76%

concerned

*"Concerned that misleading audio or video 'deepfakes' created using AI will impact the 2026 midterm elections" **

UC San Diego Center for Transparent and Trusted Elections · Sept. 1–9, 2026 · 5,016 adults, online (Forthright)

85%

very likely / likely

*Political content "created by Artificial Intelligence will spread misleading information about the elections" **

NPR/PBS News/Marist · Mar. 2–4, 2026 · 1,591 adults

66% vs. 33%

very concerned: other people vs. you

"How concerned are you that AI-generated fake videos or audio of candidates, oftentimes called 'deepfakes,' could change how [other people vote / you vote] in this November's election?"

AAPC Foundation · Sept. 14–22, 2026 · 1,568 likely voters (split sample)

84%

unacceptable

"Using an AI-generated voice that sounds like the candidate, used by their own campaign" (67% completely, 17% mostly)

AAPC Foundation · Sept. 14–22, 2026 · 1,568 likely voters (split sample)

47% vs. 11%

more concerned vs. more excited

"Thinking about the future impact of artificial intelligence, or AI, tools on how you get information, would you say you feel more excited, more concerned, or neither excited nor concerned?"

USAFacts/AP-NORC · July 29–Aug. 10, 2026 · 1,036 adults

9%

extremely / very confident

"How confident are you that information from artificial intelligence, chatbots or search results, is reliable and factual?" (49% not very / not at all)

USAFacts/AP-NORC · July 29–Aug. 10, 2026 · 1,036 adults

** Wording as reported in the pollster's release; other items verbatim from the published toplines.*

The Washington Post
Democracy Dies in Darkness

AI deepfakes threaten to upend global elections. No one can stop them.

As more than half the global population heads to the polls in 2024, AI-powered audio, images and videos are sowing confusion and clouding the political debate

By Pranshu Verma and Cat Zakrzewski

April 23, 2024 at 7:26 a.m. EDT

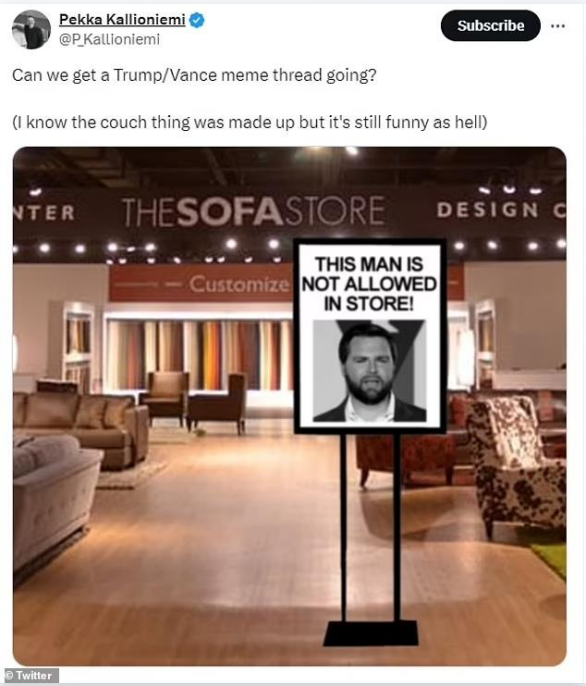
Pranshu Verma & Cat Zakrzewski, The Washington Post, Apr. 23, 2024

AI Memes


@amuse
@amuse
Pretty soon Trump's going to have to start playing his own music for his rallies.
AI-GENERATED IMAGE
4:09 PM · Aug 24, 2024 · 41.9K Views


The White House
@WhiteHouse
Follow
Donald J. Trump
@realDonaldTrump
11:07 PM · 5/2/25 · 33M Views




Pekka Kallioniemi
@P_Kallioniemi
Subscribe
Can we get a Trump/Vance meme thread going?
(I know the couch thing was made up but it's still funny as hell)
THESOFASTORE
THIS MAN IS NOT ALLOWED IN STORE!
© Twitter


Pinned
Donald J. Trump @realDonaldTrump · 21h
TRUMP 2028
0:02 / 0:08


Donald J. Trump
@realDonaldTrump
TRUMP ARMY
@TRUMP_ARMY_
There is only one candidate who can MAKE AMERICA GREAT AGAIN!!
THERE IS ONLY ONE CANDIDATE WHO CAN
MAKE AMERICA GREAT AGAIN!
AI-GENERATED IMAGE
4.67k ReTruths 21k Likes
Aug 24, 2024 at 11:50 AM


Donald J. Trump
@realDonaldTrump · 34s


AI-GENERATED FAKE IMAGE

Just Because It Is AI-Created, Doesn't Make It False…





6-4, 6-4

The Liar's Dividend

- Persuasion by deepfakes remains possible
  - See Slovakia (2023)
- The greater risk: erosion of belief in true facts
  - As synthetic media becomes ubiquitous, people discount even true news
  - Shameless actors can dismiss authentic images as "AI-generated"
- Images and video lose their presumption of validity



Chesney & Citron, 107 Cal. L. Rev. 1753 (2019)

The Liar's Dividend in Practice

WORDS IN MOUTHS

"I Guess They Used AI or Something": Trump Says BBC Journalists Deepfaked "Words in My Mouth, Literally"

That's not what his own lawsuit alleges.

By Victor Tangermann / Published Dec 17, 2025 8:00 AM EST / Add Futurism



Anna Moneymaker/Getty Images

Advertisement

Futurism

THE FUTURE IS

SIGN UP FOR OUR DAILY NEWSLETTER

MOST POPULAR

You'll Snort-Laugh When You Learn How Much AI Actually Added to the US Economy Last Year

By Joe Wilkins

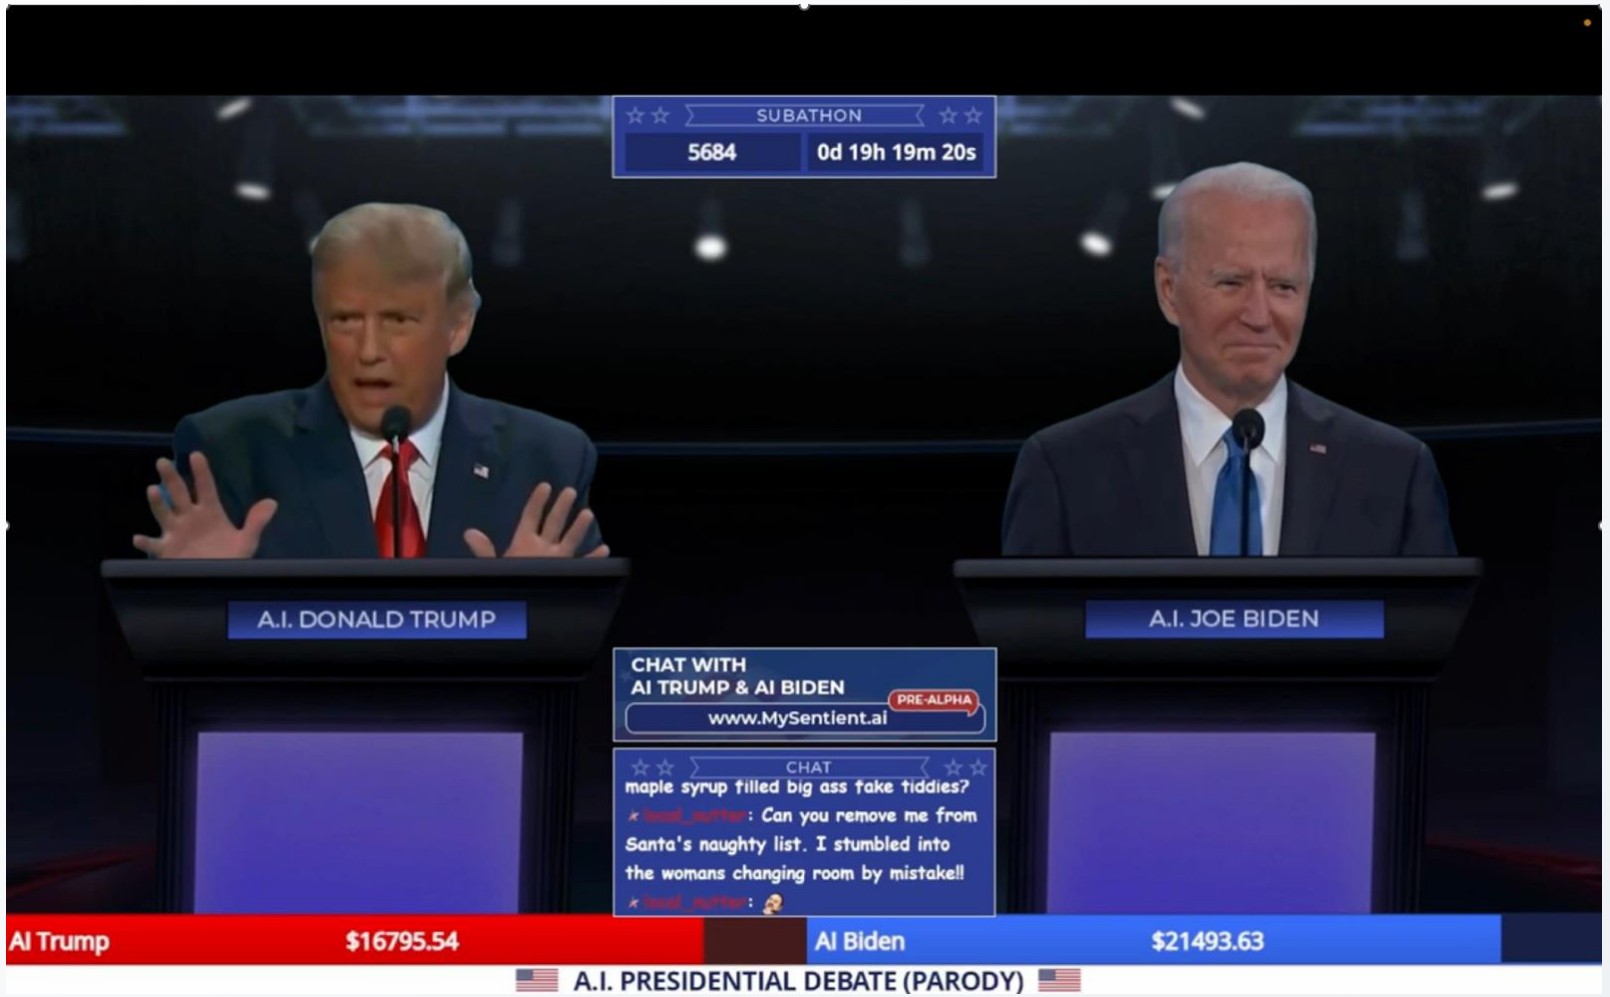
Doctor Reels as Son Becomes Plumber in Age of AI

By Frank Landymore

Microsoft Realizes It's Epically

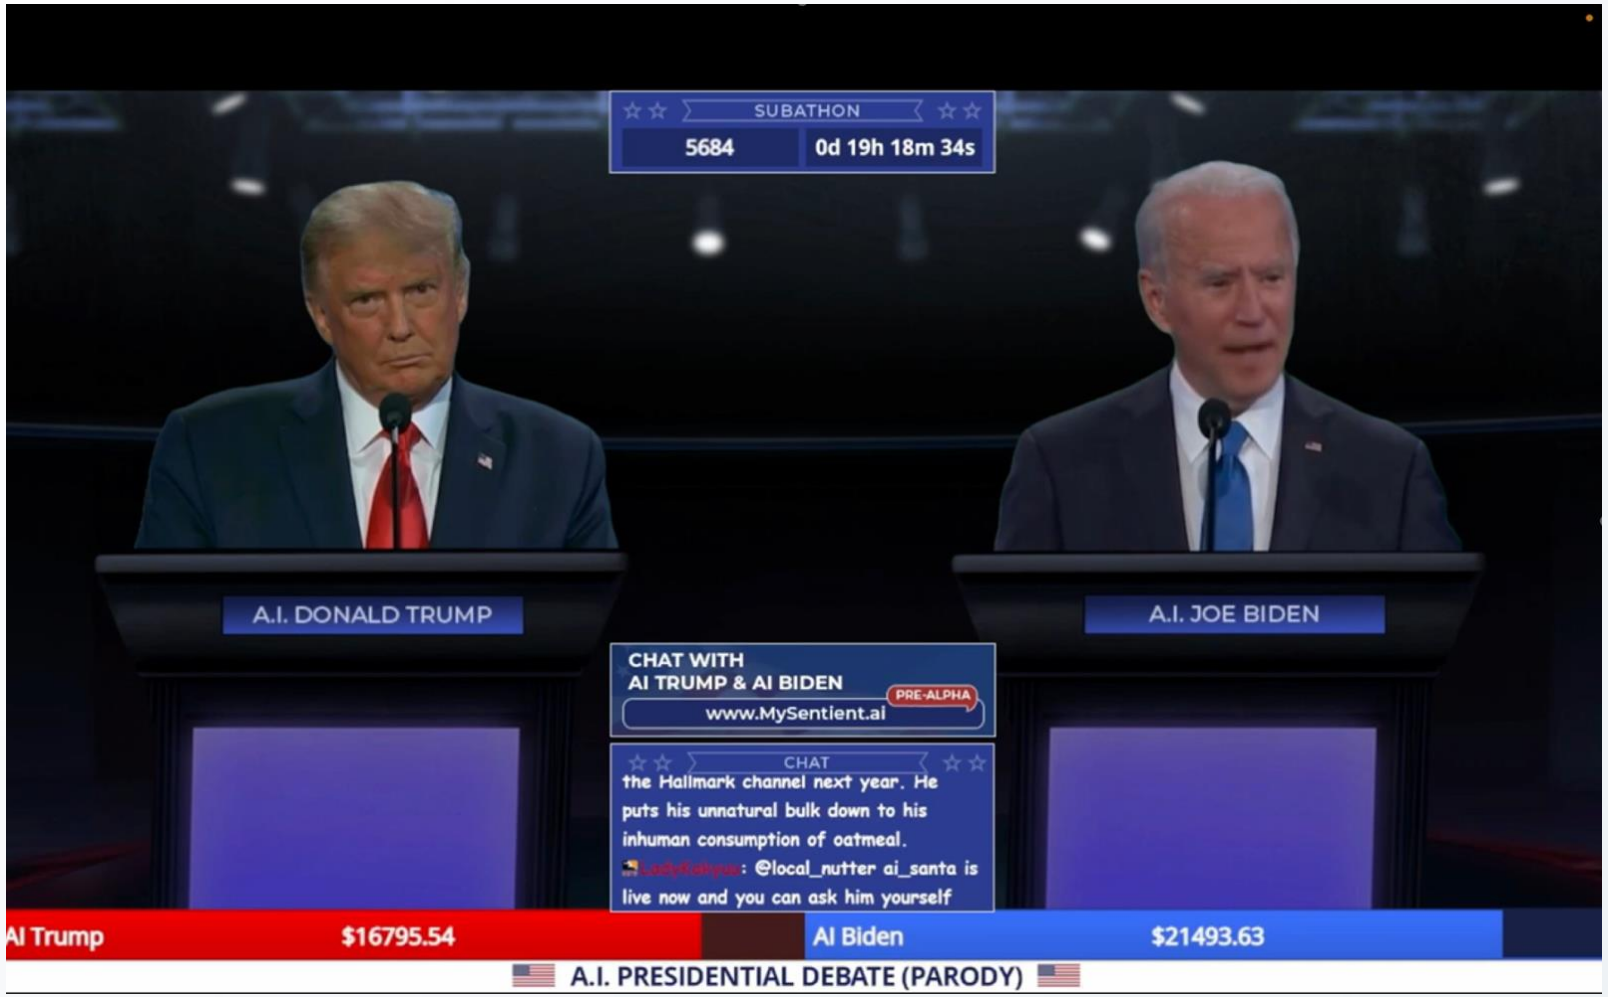
Futurism, Dec. 17, 2025

AI Presidential Debate (Parody)



twitch.tv/trumporbiden2024

AI Presidential Debate (Parody)

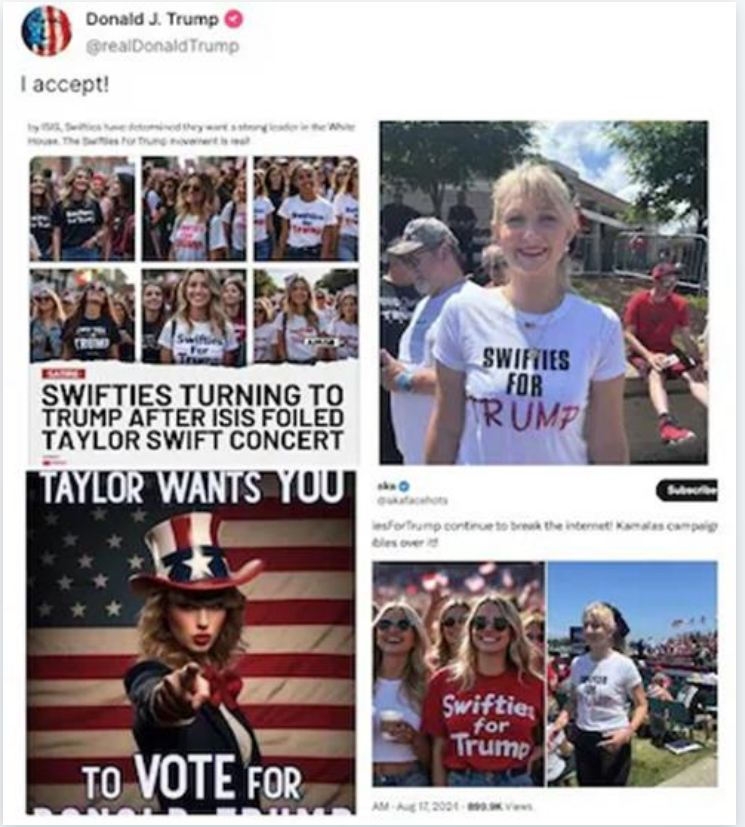


twitch.tv/trumporbiden2024

AI Deception











The 2023 Slovak Election

1 **The attack**
Audio deepfake of candidate Michal Šimečka discussing election rigging, released two days before the vote

2 **The timing**
Released during the 48-hour pre-election “silence period,” limiting the media’s ability to debunk it

3 **The effect**
No evidence it changed the result



Kőváry Sólymos, International Press Institute (Oct. 31, 2023); AFP (2023)

Is 2026 the AI Campaign?



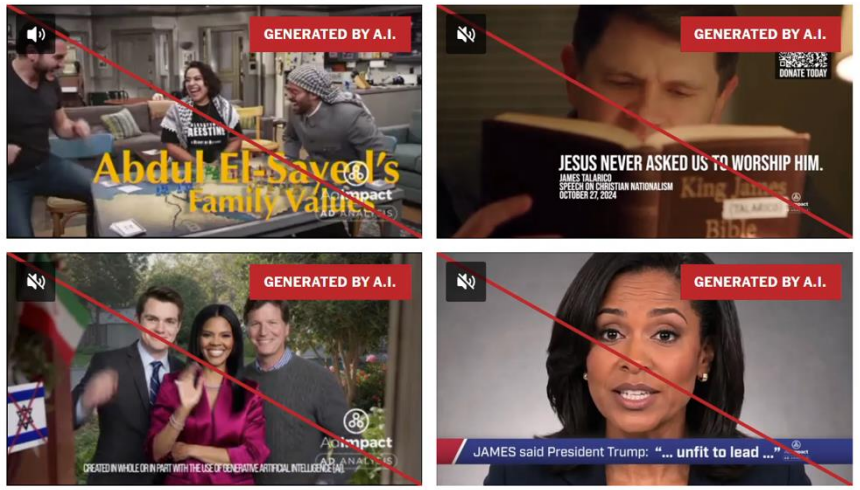
Blakeman for New York

Note: The "Generated by A.I." label and red slash added were by The New York Times. Source: AdImpact.

$137 Million in A.I.-Generated Ads Flood the Midterms

Republicans have used A.I.-generated content more often than Democrats, ranging from subtle enhancements to convincing deepfakes.

Share full article 213



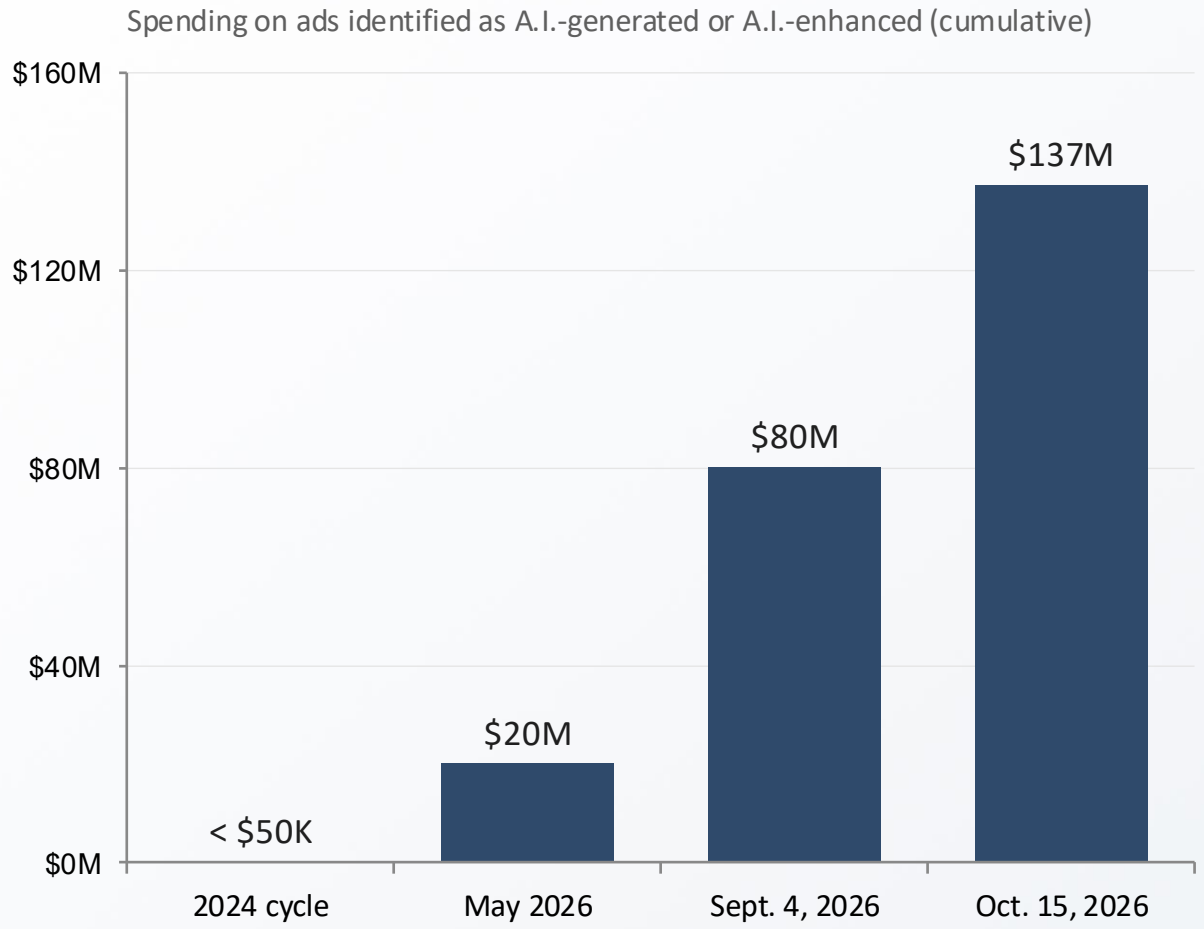
Note: The "Generated by A.I." label and red slash were added by The New York Times. Source: AdImpact.

By Stuart A. Thompson

Oct. 16, 2026

Source: Stuart A. Thompson, "$137 Million in A.I.-Generated Ads Flood the Midterms," New York Times (Oct. 16, 2026).

AI-Generated Campaign Ads, 2026



220 ads in 34 states
33 ads in May; 164 in early September

81% Republican-sponsored
88% of spending; Democrats sponsored 12% of ads

Nearly half depict real politicians
From subtle enhancement to convincing deepfakes

Only 27% carry an AI disclaimer
No federal disclosure requirement; 30 states require some disclosure

The project calls its counts "almost certainly an undercount." Sources: Wesleyan Media Project releases (May; Sept. 4; Oct. 15, 2026), using AdImpact data; Stuart A. Thompson, New York Times (Oct. 16, 2026).

Risks of Overreaction

- Sometimes receives large organic reach, but…
- Risk of drawing attention to otherwise harmless fakes
- **Greatest risk is undermining of trust in all media**, not impact of any individual lie
- Things are pretty bad already — AI's independent impact unclear







Misinformation, Political Bias, and Election Administration



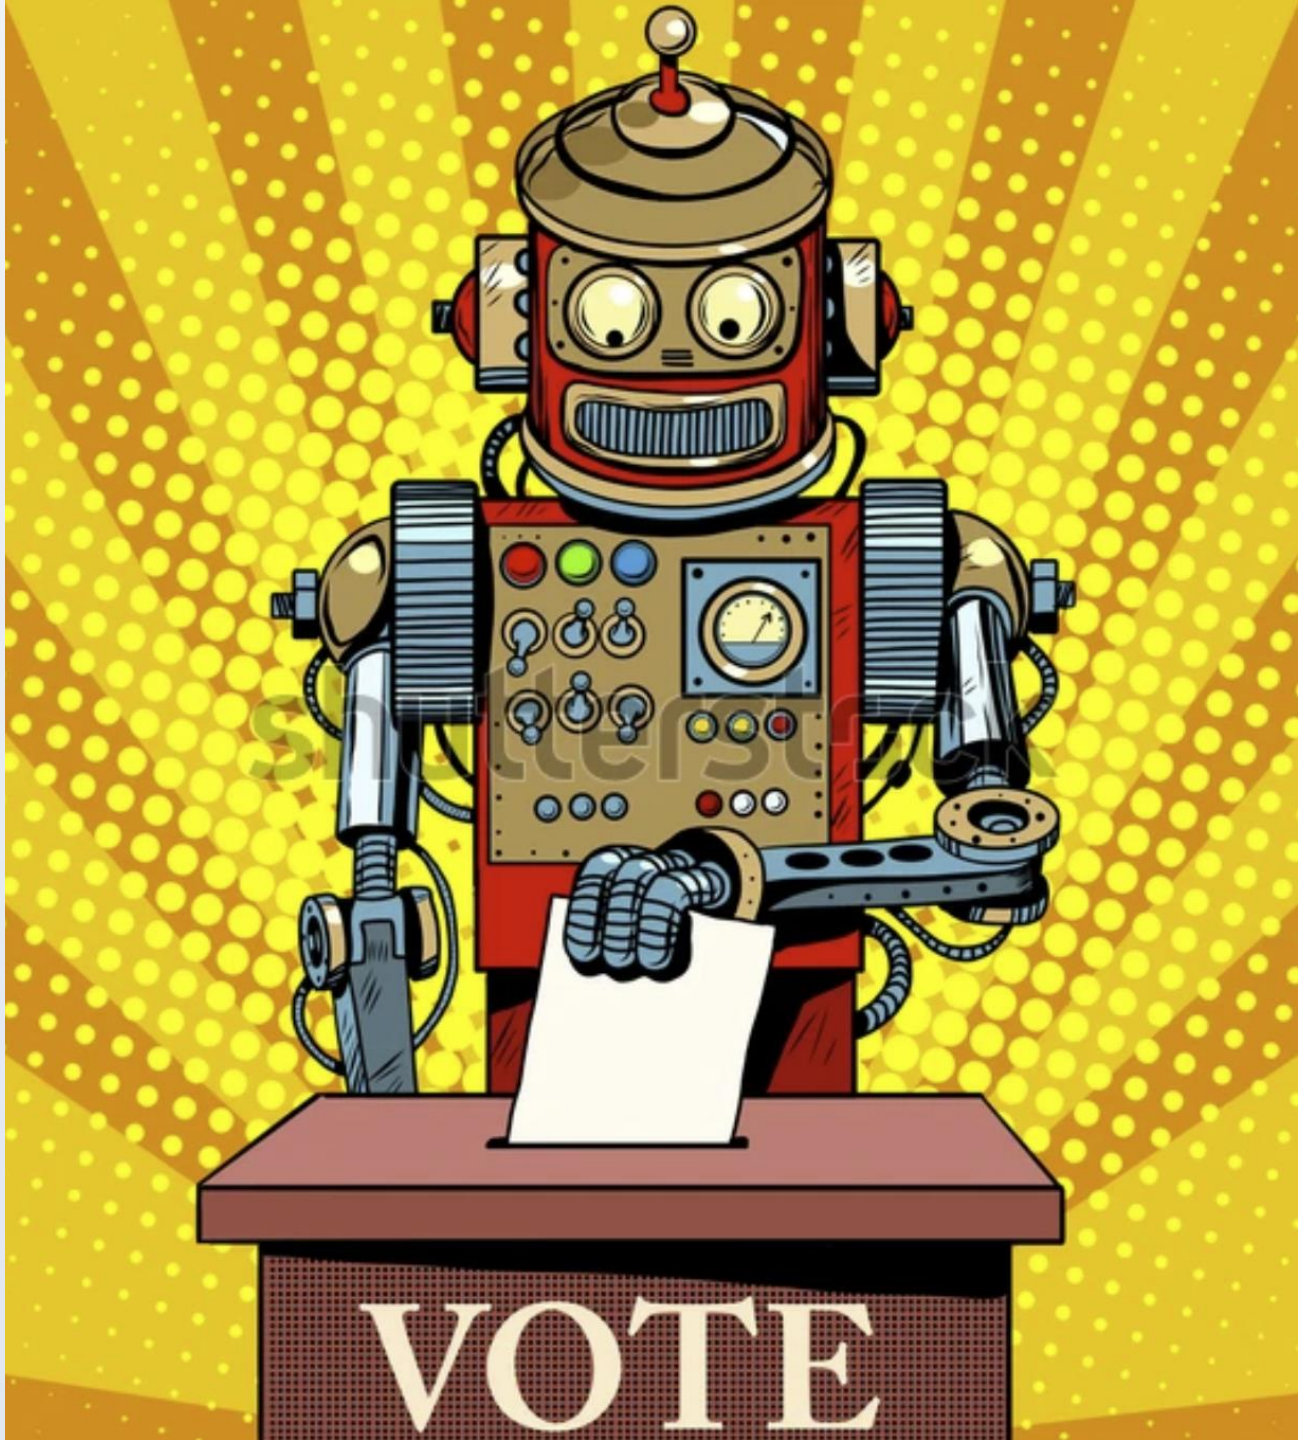
AI and Election Administration







Campaign Tools



Reevaluating the Relationship: AI & Elections

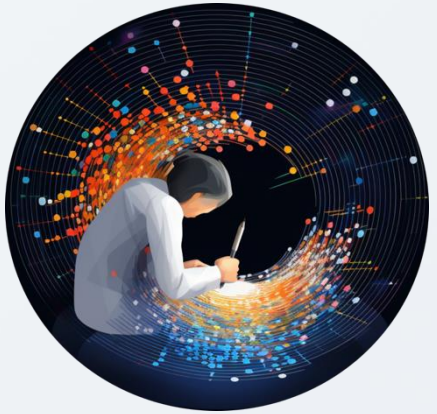
AI enhances the abilities of all actors (good, bad, foreign, domestic) in the electoral system

- Communication
  - Content creation
  - Targeting
  - Persuasion/(De)Mobilization
- Fundraising
- Get Out The Vote
- Election Administration
  - Polling place location and management
  - Mail ballot resource allocation
  - Signature verification
  - Fraud detection









AI Political Ads, 2026 Cycle

More than $80 million

Spent on 164+ AI-based ads by candidates and their backers (Wesleyan Media Project)

About 80% Republican

Share of AI ads sponsored by Republican candidates and backers

31 states

Require labels on AI-made political ads; no federal rule

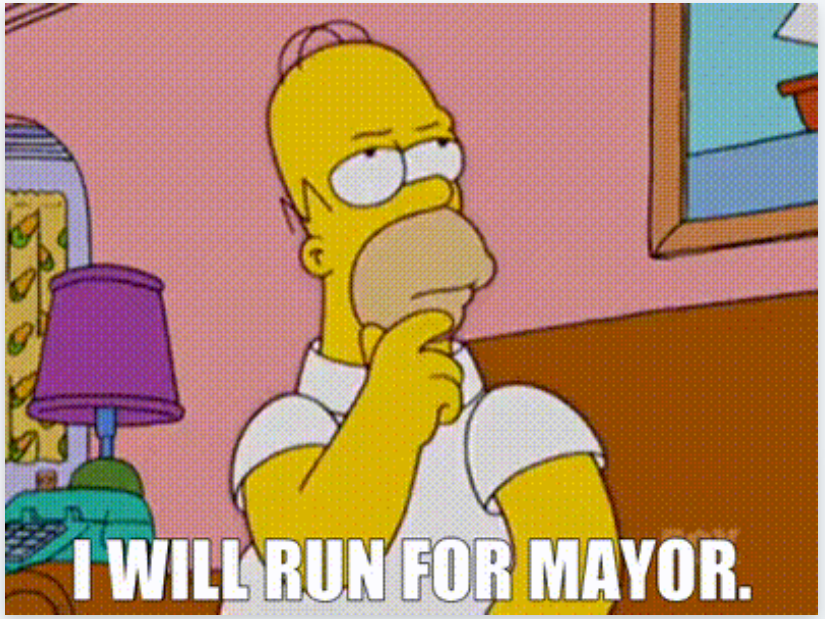
First enforcement tests

E.g., New York Democrats' complaint over Blakeman's AI ads

AI Jazeera, Oct. 9, 2026 (citing Wesleyan Media Project and Public Citizen)

Civox Demo with Ashley

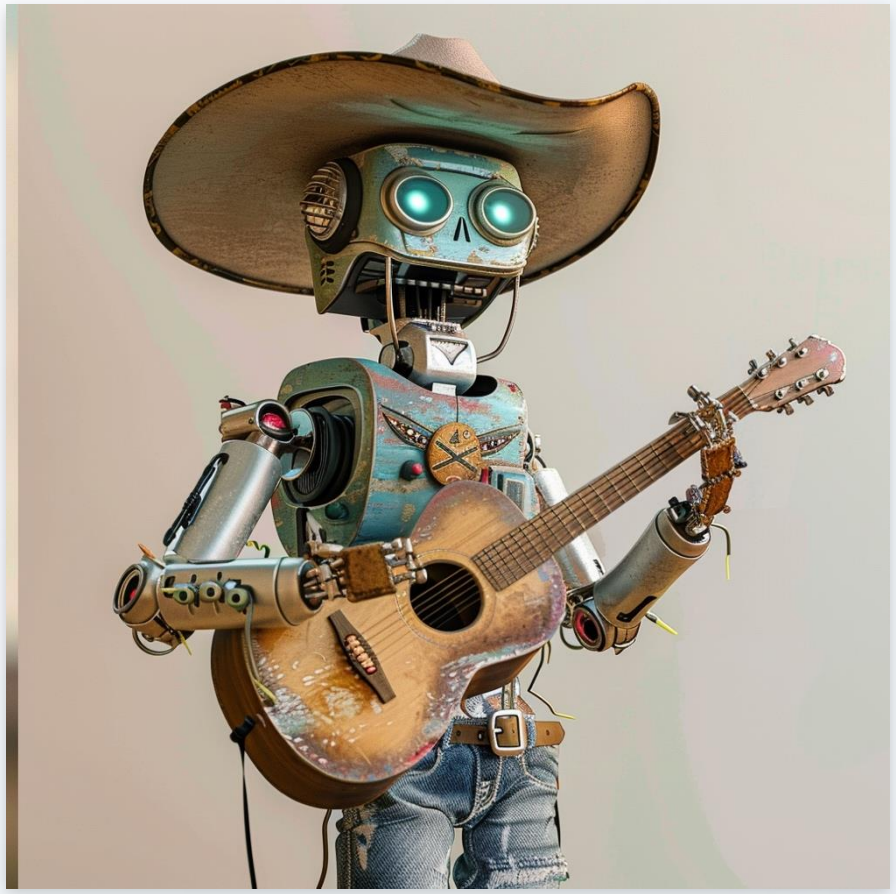




Audio courtesy Civox · The Simpsons © FOX

AI-Created Jingle in Support of Nate Persily

Here's a tale of Nate Persily,
A leader strong and oh so wisely.
Tech and law, he knows them well,
In Nate we trust, and that we tell.
Hey-ho, for Nate we stand,
With a vote for him, it's a better land.
Hey-ho, he's the man of the hour,
With wisdom, grace, and untapped power.

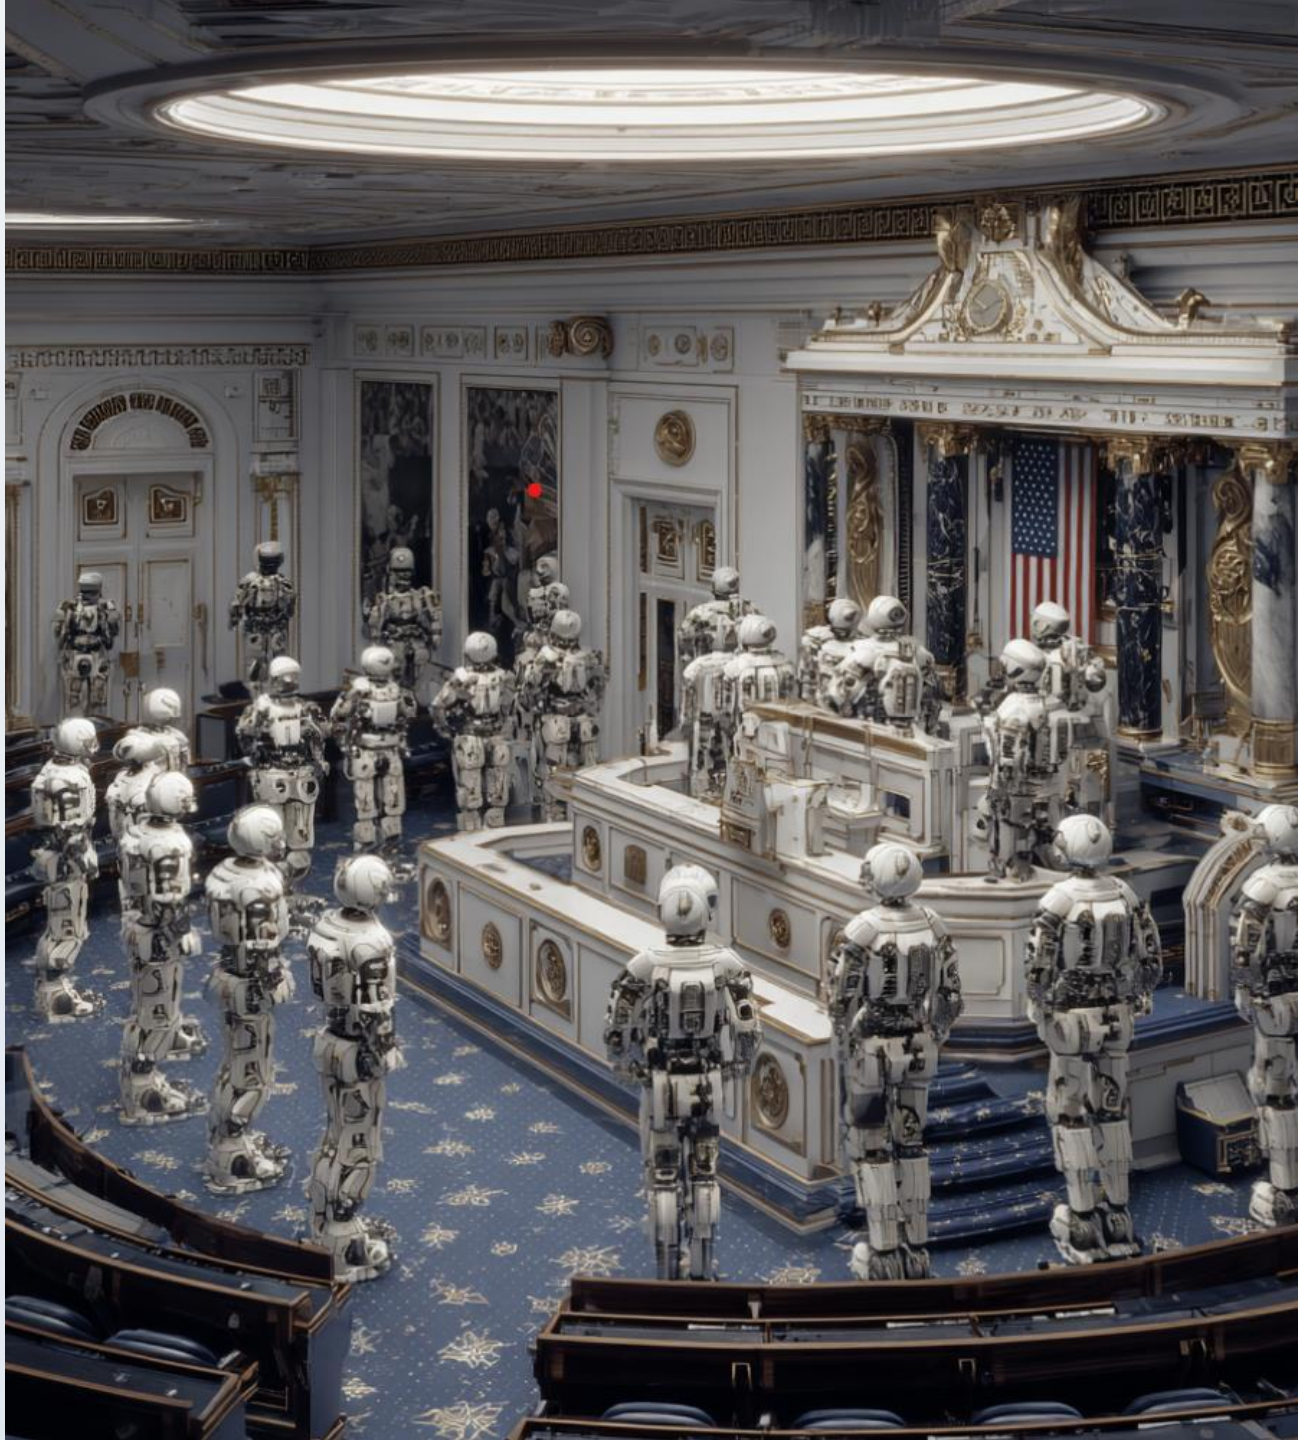


AI-Created Attack Jingle Against Nate Persily

Say no to Nate, he’s not for us,
Caught in theories, lost in fuss.
We need roots, we need truth,
Not just chatter, that’s aloof.
Say no to Nate, yes to the real,
Yes to work, to a deal we can feel.
Elites detached, they don’t understand,
Say no to Nate, let’s take a stand!



Law and Policy



Policy Options

Bans or disclosure rules

On use of AI in advertising or organic posts

Disclosure, transparency, and verification

Know how AI tools are being used and by whom (including APIs and open-source tools); encourage watermarking and disclosure — technically difficult, and uncertain whether it works at the speed necessary

Auditing

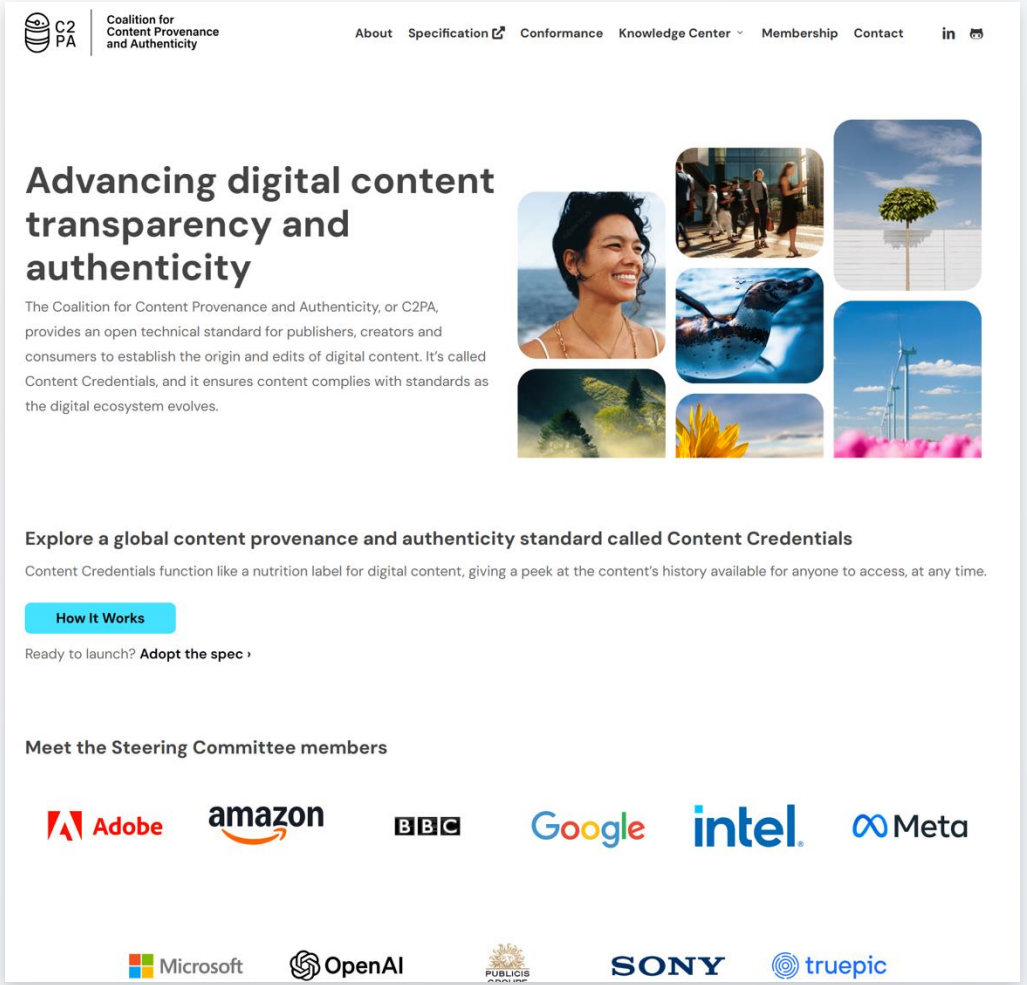
For political/racial bias and accuracy

Curb chatbots for critical election information

E.g., polling place location — but banning use of tools for politics is a fool's errand

Watermarking

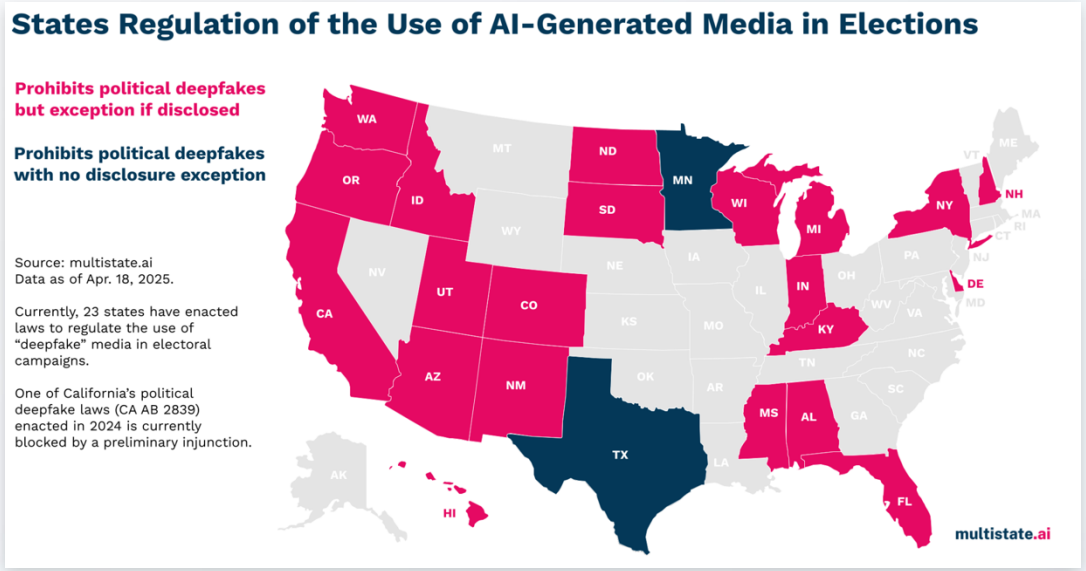
- Just AI-generated imagery or all imagery?
- How much AI in an image makes it synthetic?
- Will people trust watermarks? (Will they distrust all unmarked content?)
- What happens when they fail?
- Fact-checking labels provide a cautionary tale



Coalition for Content Provenance and Authenticity (C2PA)

State Regulation of Election Deepfakes

- 31 states now require labels on AI-generated political ads
  - Most: disclosure safe harbor rather than outright bans
  - Map (Apr. 2025) shows the earlier wave
- No federal statute
  - FEC (Sept. 2024): fraudulent-misrepresentation rule applies regardless of technology; no AI-specific rule
  - FCC (Feb. 2024): AI-generated voices in robocalls are "artificial" under the TCPA



Map: MultiState.ai (Apr. 2025); Al Jazeera, Oct. 9, 2026; NCSL

Election Deepfake Laws in Court

Kohls v. Bonta (E.D. Cal. 2025)

AB 2655 (platform takedown) struck Aug. 5, 2025; AB 2839 (deceptive election media) struck on First Amendment grounds (Aug. 2025). Court suggests narrower laws on the time, place, and manner of voting could survive.

Kohls v. Ellison (8th Cir. Feb. 9, 2026)

Preliminary injunction denied on standing and delay grounds — not the merits; Minn. Stat. § 609.771 remains in force

Minnesota Voters Alliance v. Ellison (8th Cir. Aug. 20, 2026)

Pre-election injunction provision is a prior restraint; ban on “materially false” election speech likely fails strict scrutiny (preliminary)

Election Law Blog (Aug. 30, 2025); Bloomberg Law (Aug. 6, 2025); RecordingLaw (2026)

Beyond Elections



AI in Deliberative Democracy

Analyzing public comments

From town halls or surveys, to identify shared values and points of consensus

More human-centered understanding of public opinion

Examples

AI helped analyze 4,000+ responses in Fort Collins, CO, and 8,000+ voices in Bowling Green, KY



See Sarabi, "How AI Can Unlock Public Wisdom and Revitalize Democratic Governance," Carnegie Endowment (July 22, 2025)

Enhancing Government Responsiveness

Voter information

AI helps draft social media content and powers chatbots that answer basic questions about polling locations and hours

Streamlining operations

Optimizes resource allocation, like mail-in ballot processing equipment

Examples

DonorAtlas created 20 million donor profiles; Legislaide reduced legislation drafting from 3–6 months to days



See Sarabi, Carnegie Endowment for International Peace (July 22, 2025)

Future AI Democracy Tools: The Habermas Machine (2024)

1. **Gather opinions**
   Human participants submit their initial written views on a topic

2. **AI synthesizes**
   A language model analyzes all views and generates “group statements” that bridge divides

3. **AI selects**
   A separate model predicts and selects the statement most likely to earn group support

4. **Refine with feedback**
   Humans critique the winning statement; the AI uses the feedback to generate a better version. Cycle repeats.

5. **Results**
   AI-mediated groups more likely to find consensus; participants rated AI statements clearer and less biased than human mediators’

Tessler et al., “AI Can Help Humans Find Common Ground in Democratic Deliberation,” Science 386, eadq2852 (2024)

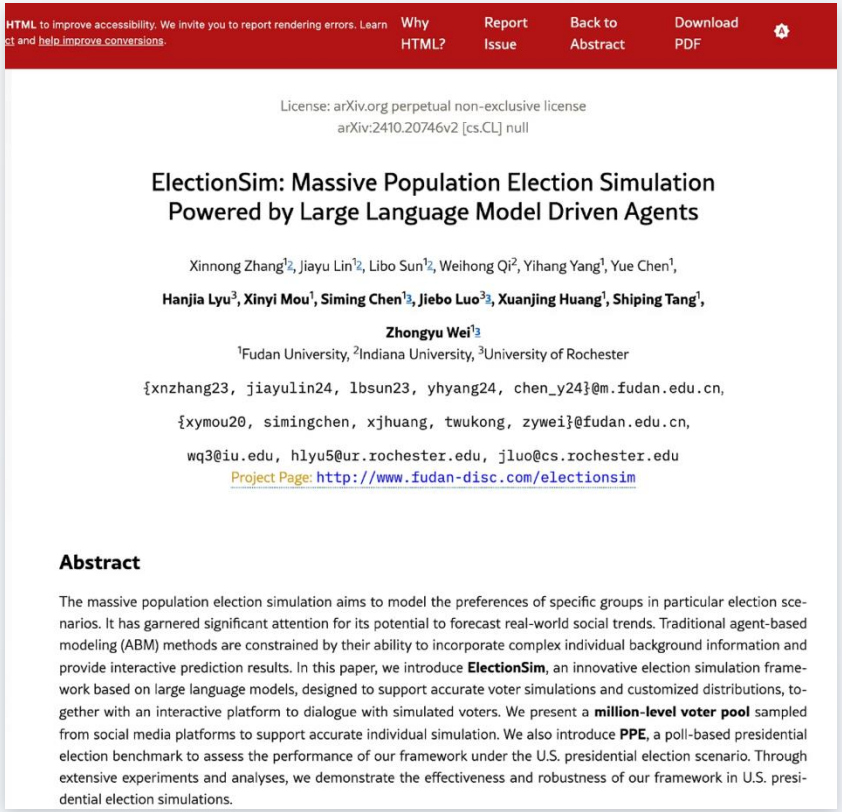
ElectionSim: A “Digital Twin” of the Electorate (2024)

1. **Construct digital twin**
   From real-world demographics and social media data

2. **Inject scenarios**
   Political messages, policies, or campaign strategies

3. **Simulate reactions**
   AI voter agents mimic public opinion and voting behavior

4. **Analyze outcomes**
   Correctly predicted winners in 47 of 51 states in 2020 — but could be misused for perfecting manipulation



HTML to improve accessibility. We invite you to report rendering errors. Learn ct and help improve conversions. Why HTML? Report Issue Back to Abstract Download PDF

License: arXiv.org perpetual non-exclusive license
arXiv:2410.20746v2 [cs.CL] null

ElectionSim: Massive Population Election Simulation Powered by Large Language Model Driven Agents

Xinnong Zhang[1,2], Jiayu Lin[1,2], Libo Sun[1,2], Weihong Qi[2], Yihang Yang[1], Yue Chen[1],
Hanjia Lyu[3], Xinyi Mou[1], Siming Chen[1,3], Jiebo Luo[3,3], Xuanjing Huang[1], Shiping Tang[1],
Zhongyu Wei[1,3]
[1]Fudan University, [2]Indiana University, [3]University of Rochester
{xnzhang23, jiayulin24, lbsun23, yhyang24, chen_y24}@m.fudan.edu.cn,
{xymou20, simingchen, xjhuang, twukong, zywei}@fudan.edu.cn,
wq3@iu.edu, hlyu5@ur.rochester.edu, jluo@cs.rochester.edu
Project Page: http://www.fudan-disc.com/electionsim

Abstract

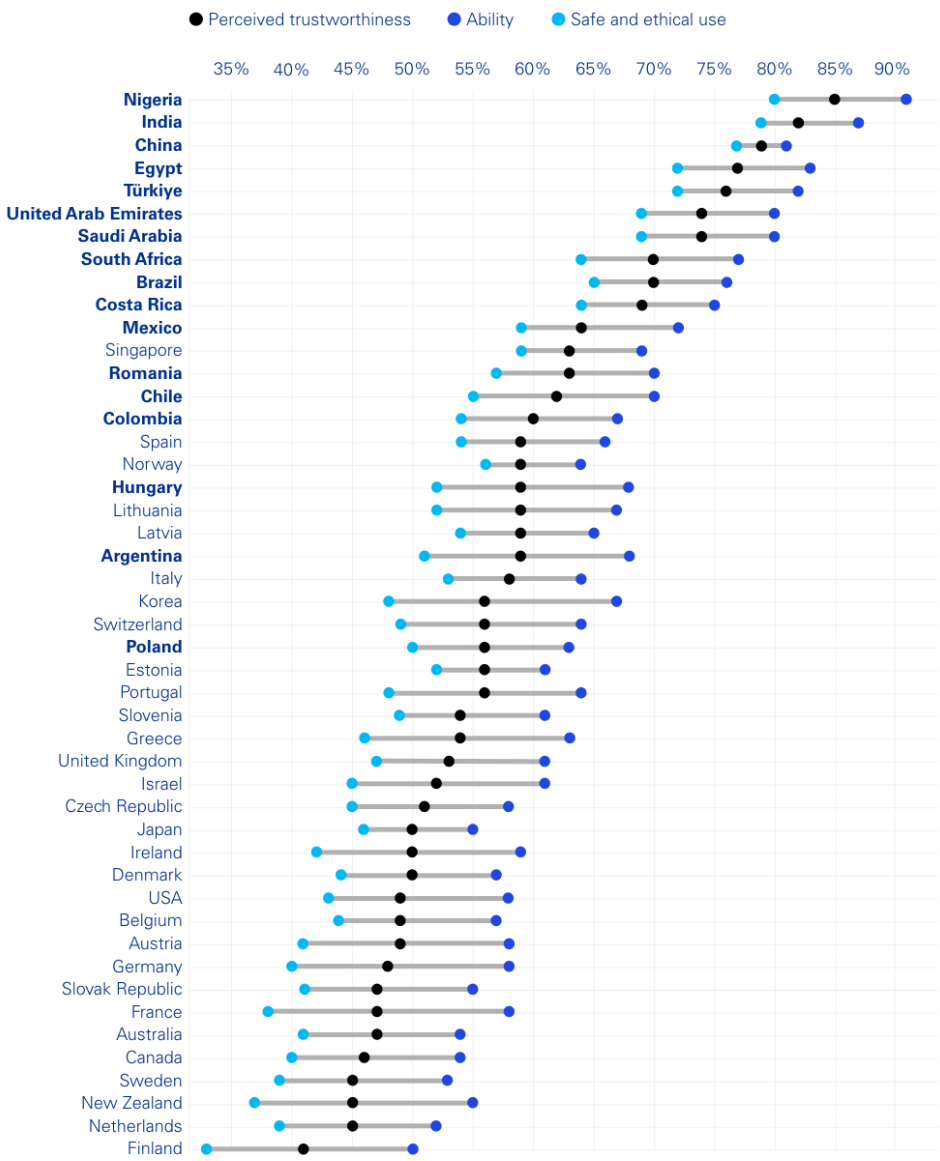
The massive population election simulation aims to model the preferences of specific groups in particular election scenarios. It has garnered significant attention for its potential to forecast real-world social trends. Traditional agent-based modeling (ABM) methods are constrained by their ability to incorporate complex individual background information and provide interactive prediction results. In this paper, we introduce **ElectionSim**, an innovative election simulation framework based on large language models, designed to support accurate voter simulations and customized distributions, together with an interactive platform to dialogue with simulated voters. We present a **million-level voter pool** sampled from social media platforms to support accurate individual simulation. We also introduce **PPE**, a poll-based presidential election benchmark to assess the performance of our framework under the U.S. presidential election scenario. Through extensive experiments and analyses, we demonstrate the effectiveness and robustness of our framework in U.S. presidential election simulations.

Zhang et al., “ElectionSim: Massive Population Election Simulation Powered by LLM-Driven Agents” (2024 preprint)

Public Opinion on AI

Cross-national comparisons · Structure of U.S. opinion · Data centers

Figure 9: Perceptions of the trustworthiness of AI systems

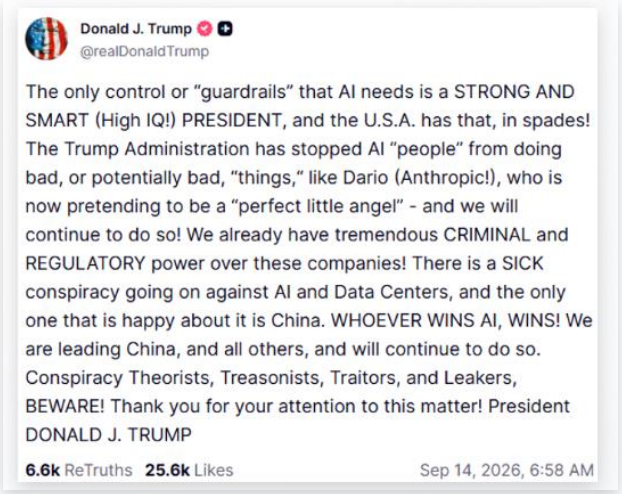


% Agree = 'Somewhat agree', 'Agree', 'Strongly agree'. Ordered by perceived trustworthiness.
Bolding indicates countries with emerging economies.

© 2025 The University of Melbourne.

© 2025 Copyright owned by one or more of the KPMG International entities.
KPMG International entities provide no services to clients. All rights reserved.

AI Is Having Its Political Moment







"Senators from both parties question OpenAI on breach of AI startup Hugging Face" (AP, Sept. 11, 2026)





"Sanders, Casar Want to Outlaw AI 'Superintelligence' and Pause Advanced AI" (Sept. 2026)



"The data center boom meets resistance in Maine as lawmakers pass a yearlong moratorium" (AP, Apr. 15, 2026)

PBS NewsHour/AP, Sept. 11, 2026; TechRepublic, Sept. 2026; WBUR/AP, Apr. 16, 2026. Photos: U.S. Senate / U.S. House; Mike Stewart/AP

The Global Paradox

AI Optimism Outweighs Pessimism in Most Countries

Top five and bottom five countries, by expected impact

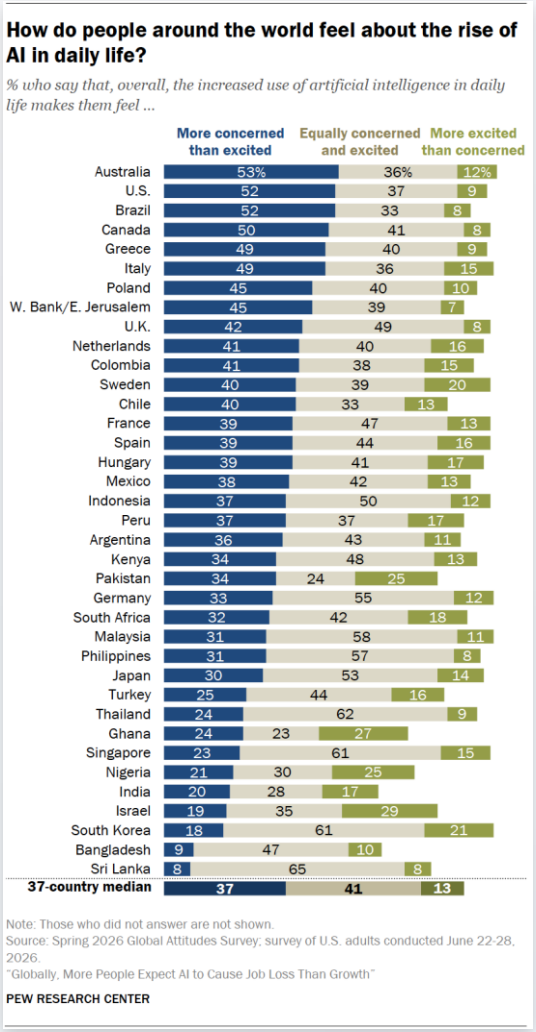
AI will mostly help or mostly harm people in this country		AI will make your daily life mostly better or mostly worse	
Country/Area	**% Mostly help**	**Country/Area**	**% Mostly better**
Top five			
China	93	China	93
Vietnam	91	Vietnam	91
Singapore	77	Singapore	83
Israel	69	Malta	71
Nigeria	68	Israel	69
Bottom five			
Afghanistan	37	State of Palestine	41
United States	36	Malawi	38
Malawi	35	Afghanistan	36
Egypt	35	Egypt	35
Bangladesh	34	Bangladesh	32

Microsoft-Gallup AI Adoption Survey, 2026

Get the data • Download image

GALLUP

- U.S. builds the most capable models, yet its public is among the most apprehensive
- 37-country median: 37% more concerned
  - In half the countries, modal response is “equally concerned and excited” (median 41%)



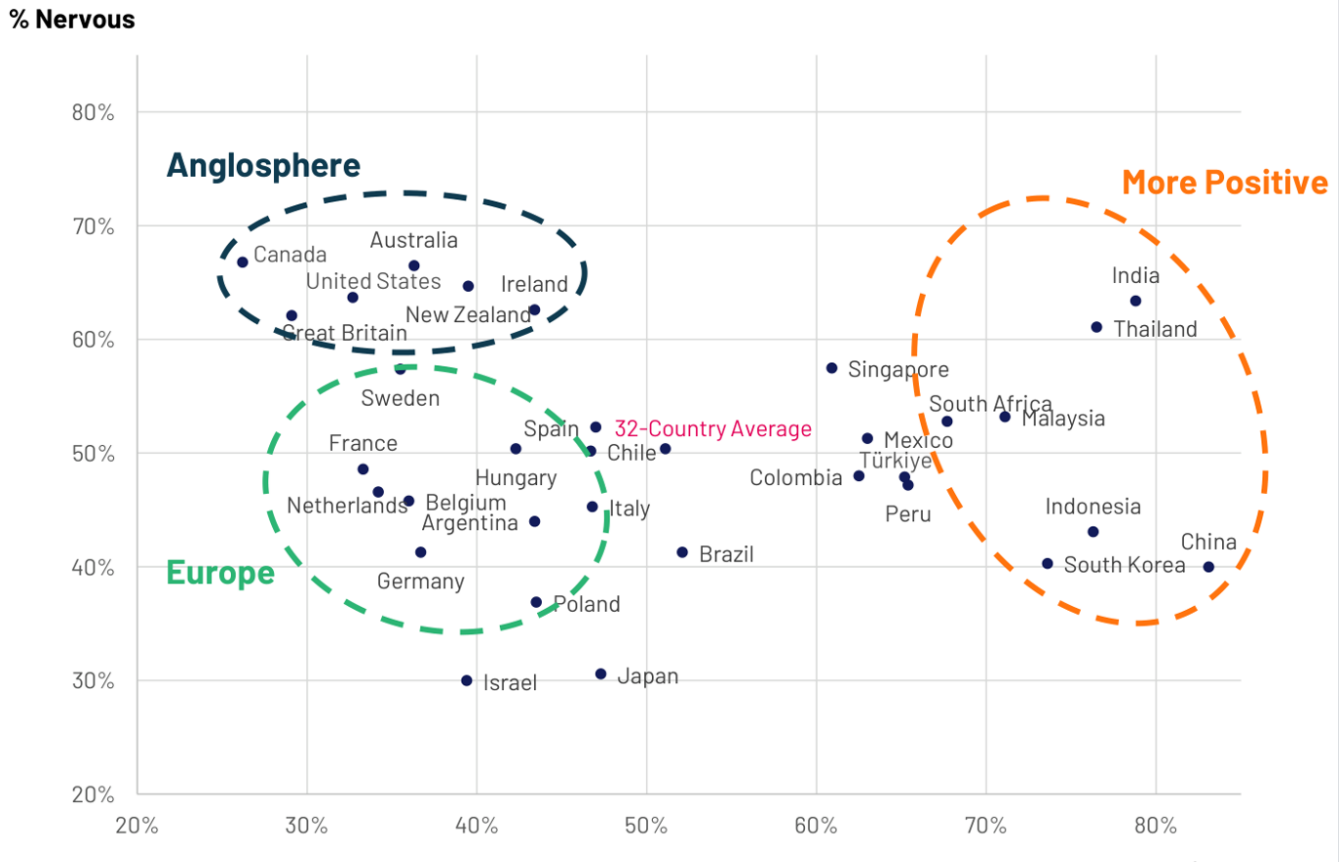
Pew Research Center, Sept. 17, 2026; Gallup–Microsoft (2026)

Excitement v. Nervousness Across 32 Countries

How much do you agree or disagree with the following?

- **Products and services using artificial intelligence make me nervous**
- **Products and services using artificial intelligence make me excited**

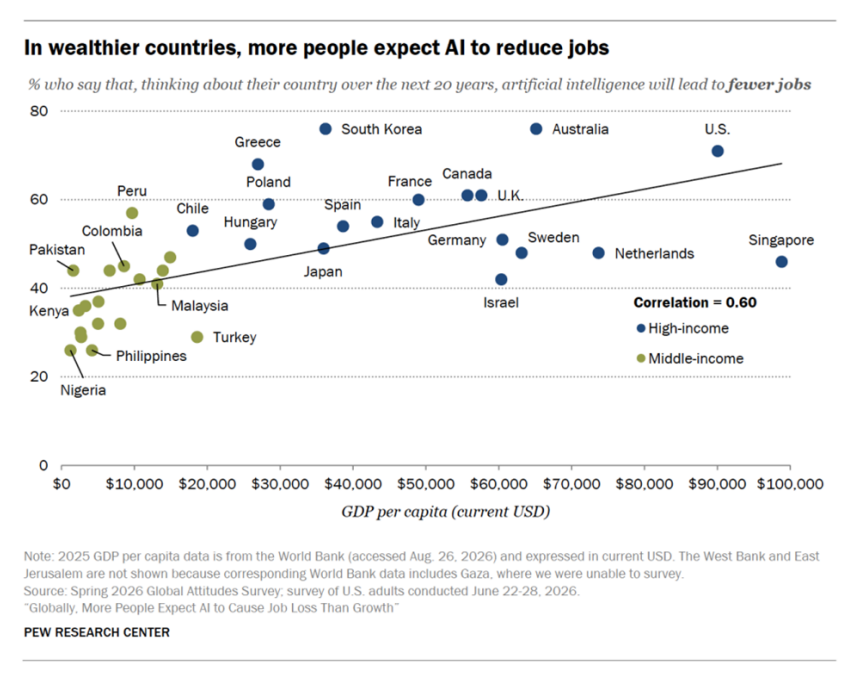
Base: 23,532 adults under the age of 75 across 32 countries between March 20, 2026 and April 3, 2026.



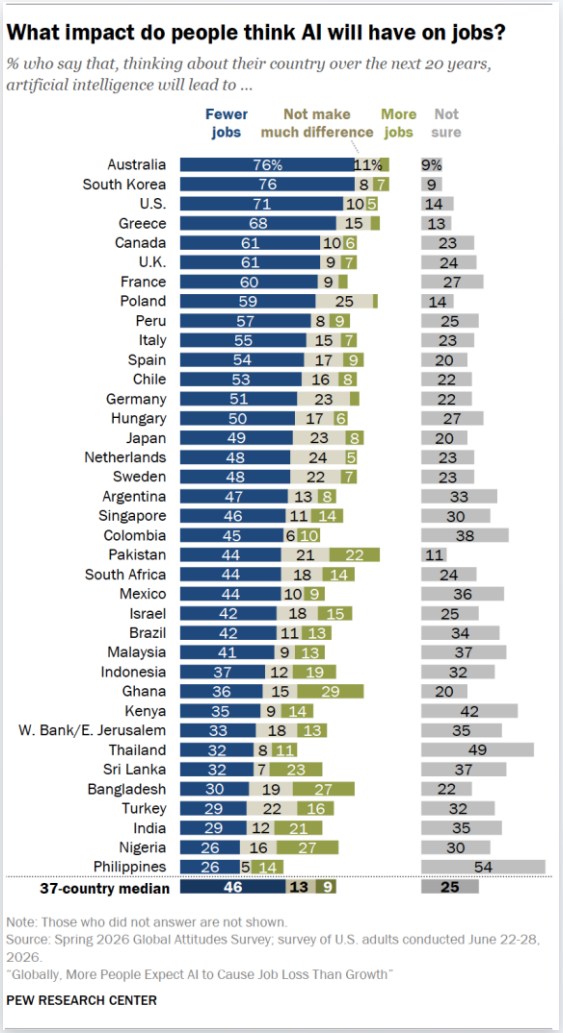
© Ipsos AI Monitor

14

Expected Effects on Jobs

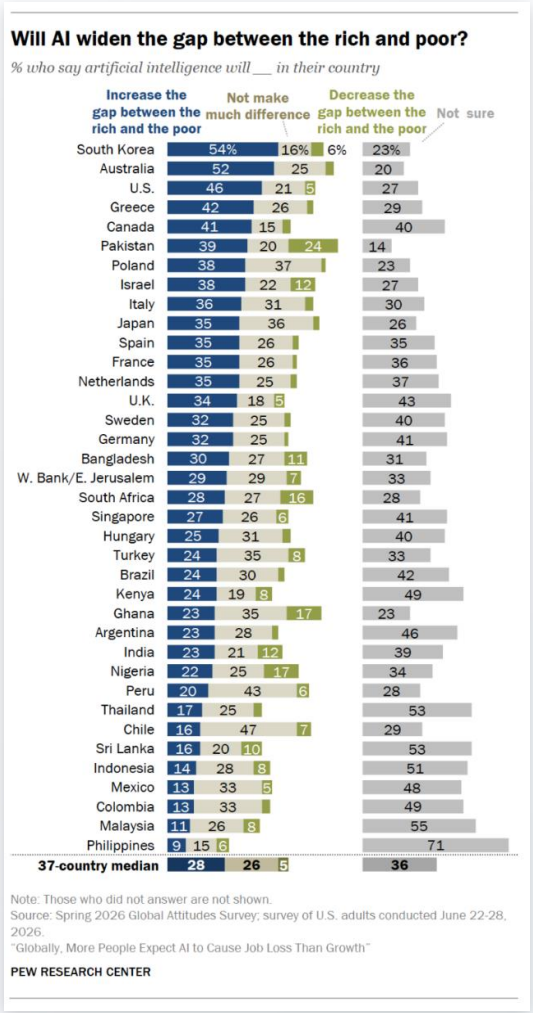


- Expect fewer jobs over next 20 years
  - High-income median: 55%
  - Middle-income median: 36% (about a third unsure)
- Tracks GDP per capita
- U.S. share expecting fewer jobs: 64% (2024) → 71% (June 2026)
- Widening rich–poor gap: U.S. 46% v. Colombia 13%
  - High-income median 35%; middle-income 22%



Pew Research Center, Sept. 17, 2026; Pew, Aug. 18, 2026

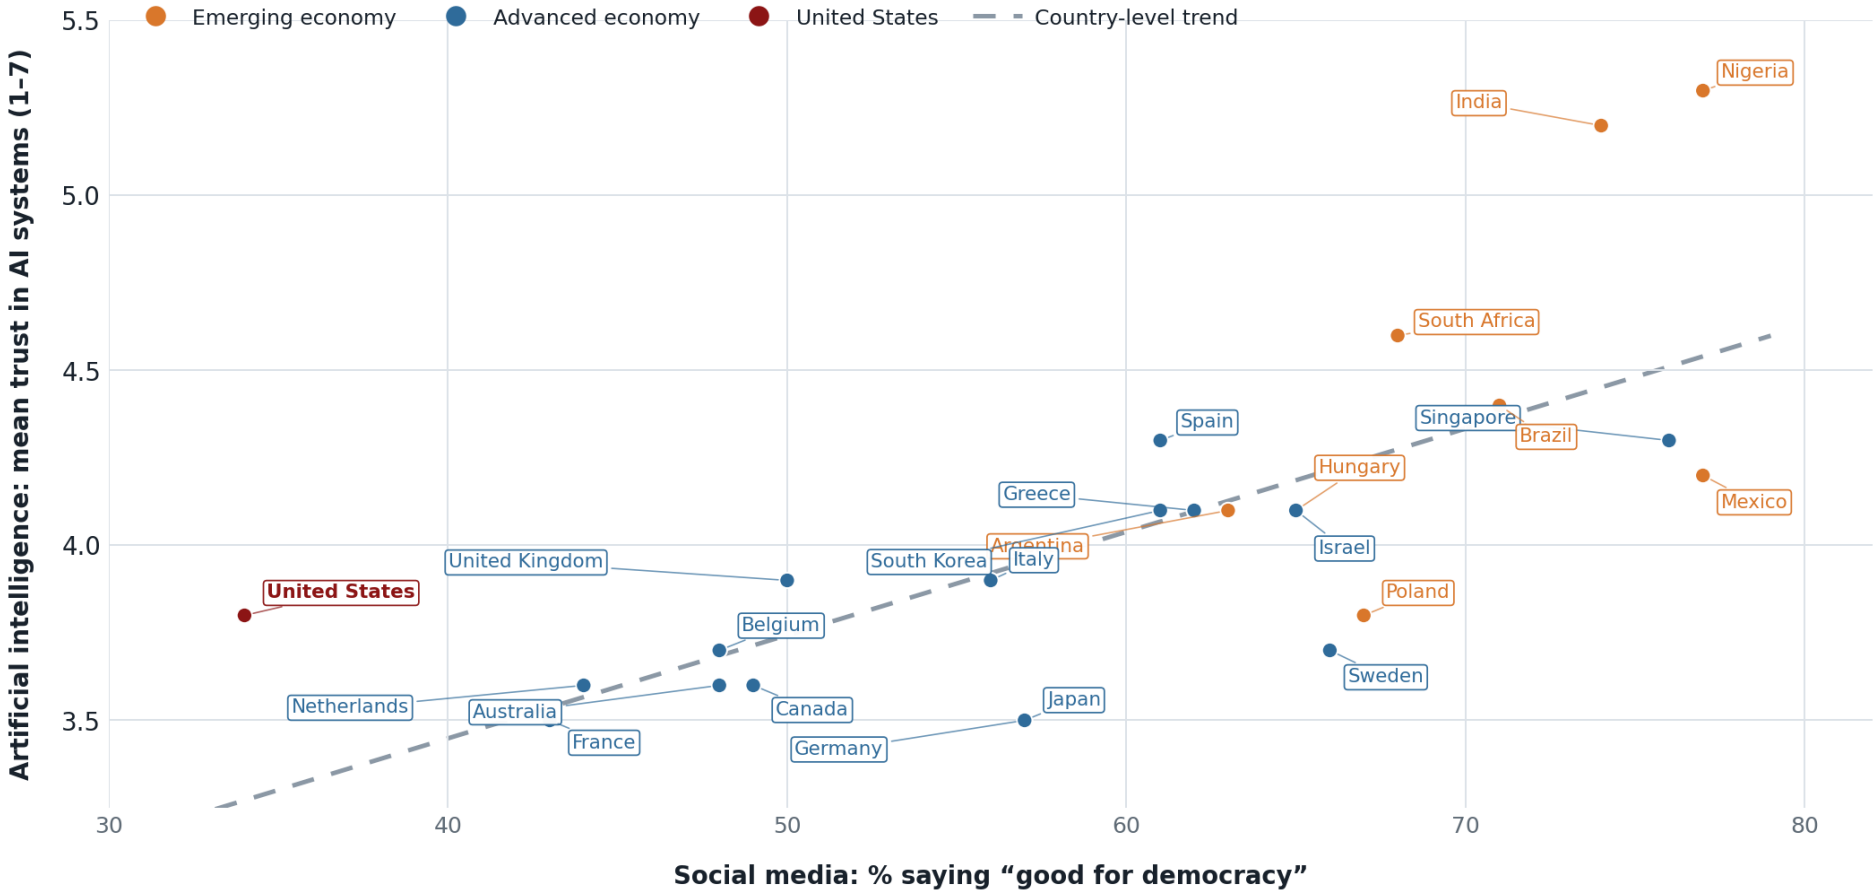
Expected Effects on Inequality



Pew Research Center, Sept. 17, 2026

Social-media optimism is associated with greater trust in AI

Country-level comparison across 24 countries; each label marks one country (descriptive r = 0.71)

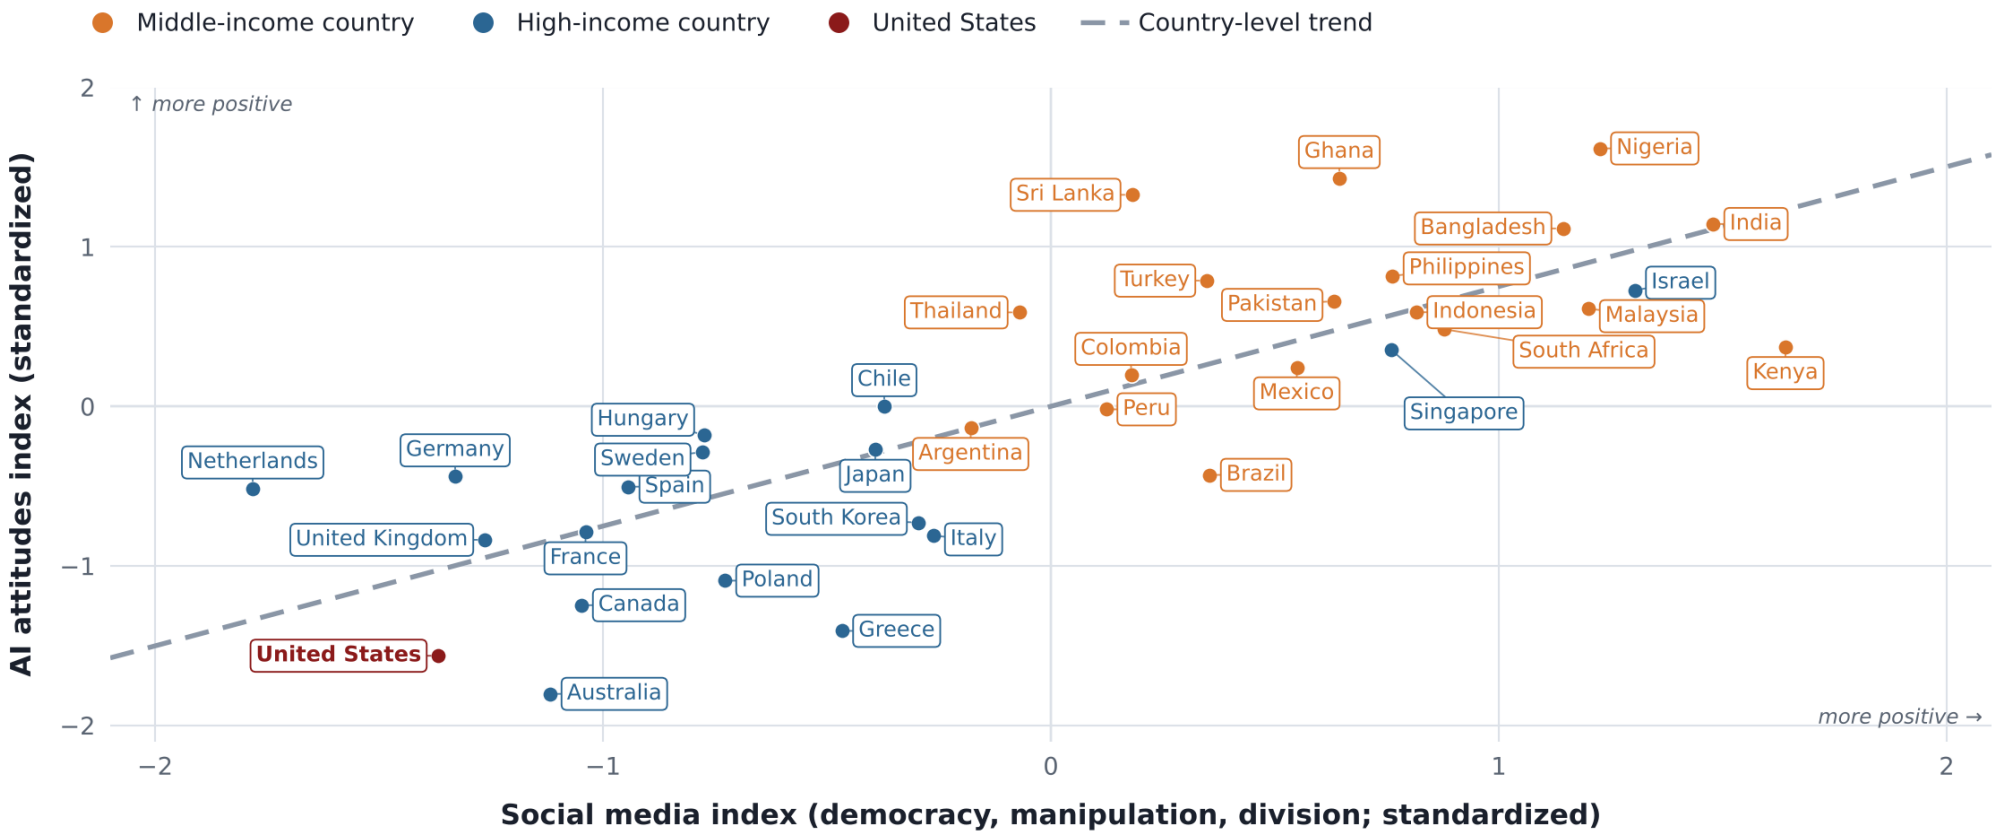


The measures are not identical and were fielded at different times. The correlation is descriptive at the country level and should not be interpreted as an individual-level relationship or a causal effect.

Sources: Pew Research Center Global Attitudes Surveys (2022–23), Q28/Q54; University of Melbourne & KPMG, Trust, Attitudes and Use of AI (fielded Nov. 2024–mid-Jan. 2025), Appendix 3. Economy groups follow KPMG/IMF classification.

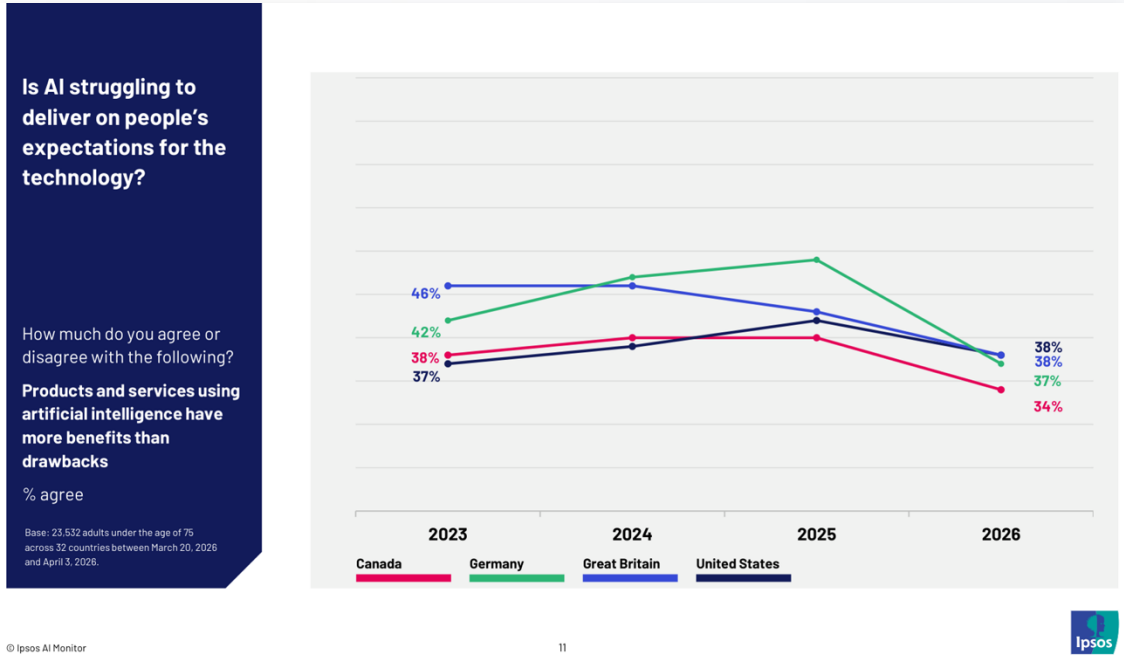
Countries more positive about social media are more positive about AI

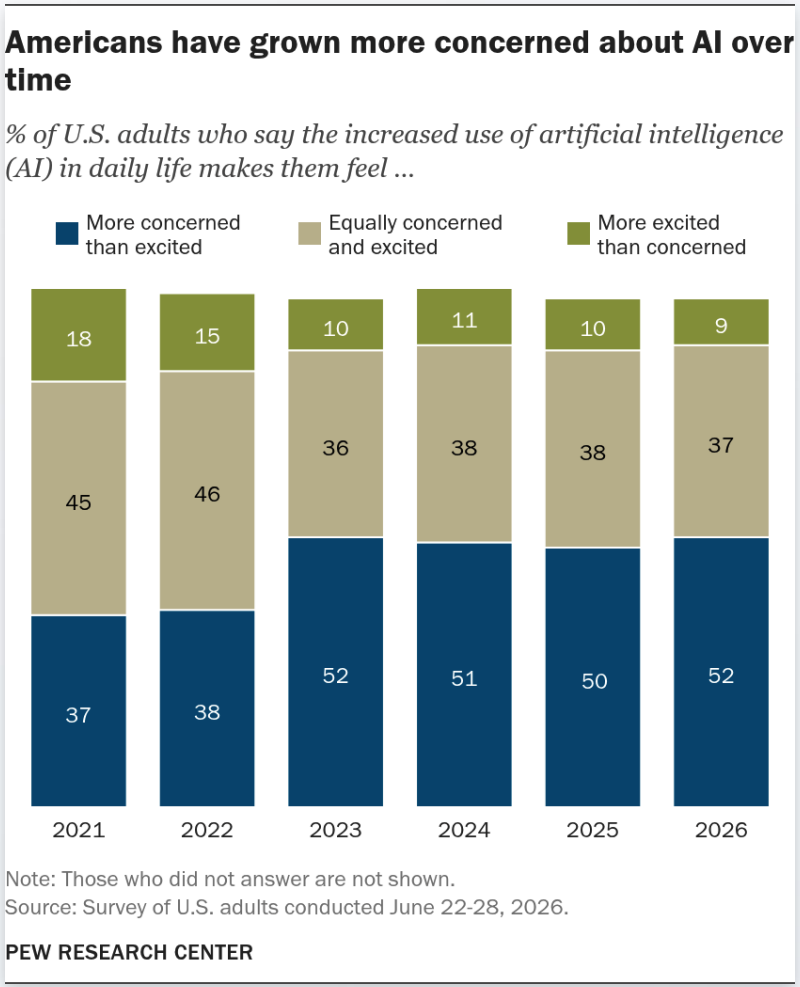
Composite indexes from one Pew survey, 36 countries, spring 2026 (descriptive r = 0.79)



Social media index: mean of standardized country scores on (1) social media good minus bad for democracy, (2) % saying internet/social media made people easier to manipulate (reversed), (3) % saying it made people more politically divided (reversed); α = 0.90. AI index: mean of standardized net excitement (excited minus concerned), net jobs (more minus fewer), and net inequality (decrease minus increase); α = 0.83. 0 = 36-country average; units are country-level standard deviations. Correlation is descriptive, not individual-level or causal.
Sources: Pew Research Center, Spring 2026 Global Attitudes Survey (Feb. 8–May 13, 2026); U.S. from separate Pew panel surveys. Reports of Oct. 1, 2026 (social media) and Sept. 17, 2026 (AI, topline Q55–Q57). Income groups per Pew's World Bank classification. West Bank/East Jerusalem omitted.

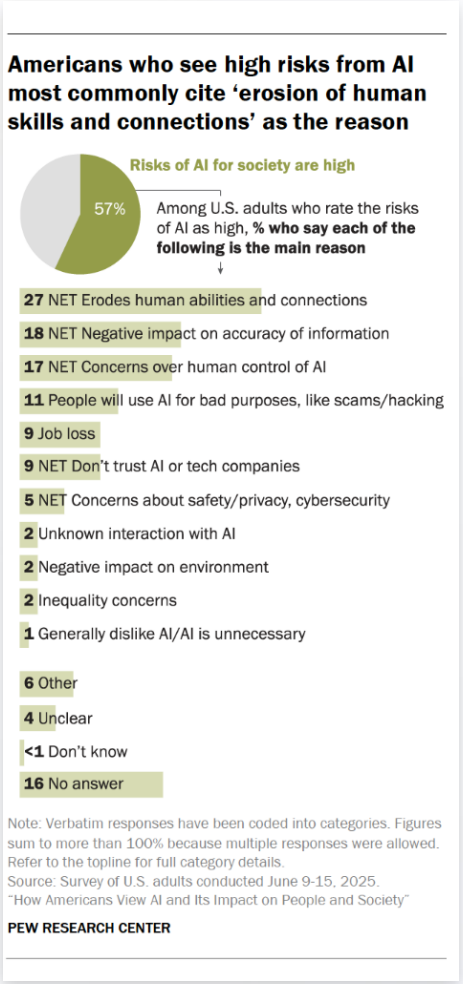
Benefits v. Drawbacks, 2023–2026

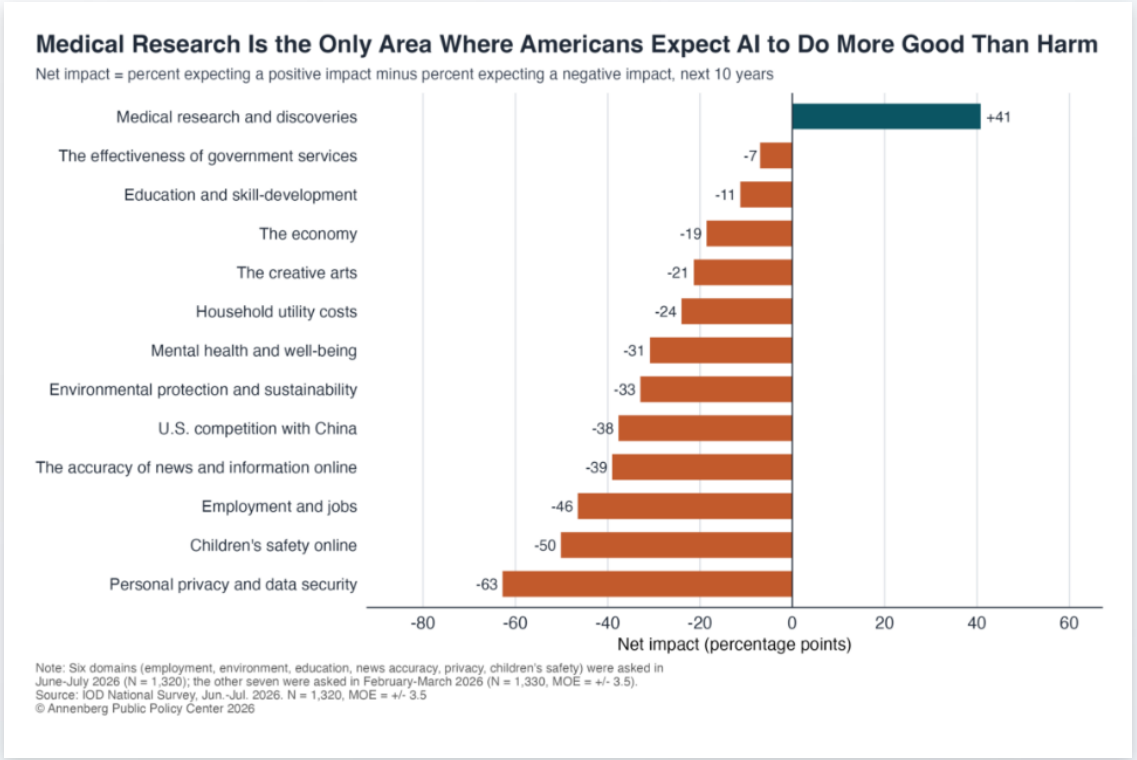




Ipsos AI Monitor 2026; Pew Research Center (June 2026)

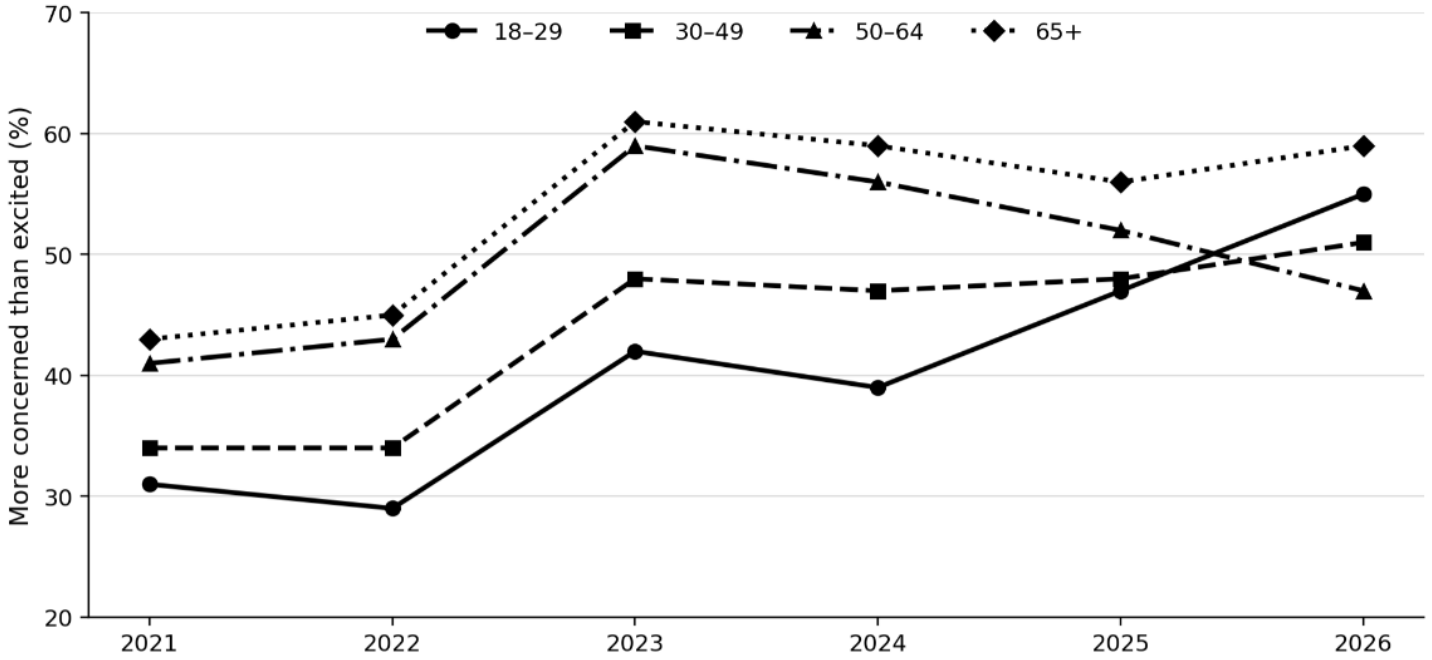
Americans' Pessimism About AI Is Pervasive





Pew Research Center

Concern about artificial intelligence, by age, 2021–2026



Question: “How does the increased use of artificial intelligence (AI) in daily life make you feel?”
Shown: Percentage responding “More concerned than excited.”

Source: Colleen McClain & Eugenie Park, “Young Adults in the U.S. Are Increasingly Wary of AI, Concerned It Will Take Jobs,” Pew Research Center (Aug. 18, 2026); Pew Research Center, “Young Adults’ Worry Over AI Continues to Climb, While Their Excitement Has Dropped Off” (chart).

The Structure of U.S. Opinion on AI: Demography

Do you think that the impact of artificial intelligence (AI) on society over the next 20 years will be...

U.S. adults by race and ethnicity, age, and gender | February 17-23, 2026

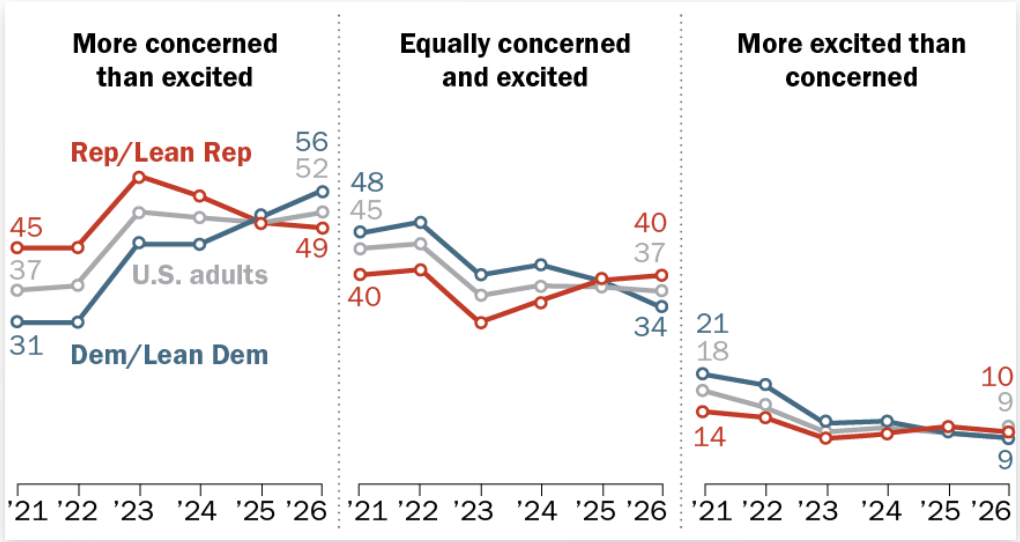
Demographic group	Negative	Positive	Equally positive and negative	Not sure
ALL U.S. ADULTS	40%	16%	31%	13%
RACE AND ETHNICITY				
White	43%	15%	30%	12%
Hispanic	36%	18%	29%	18%
Black	35%	15%	32%	18%
Asian*	24%	30%	38%	8%
GENDER				
Men	36%	22%	30%	12%
Women	43%	11%	31%	15%
AGE				
18-29	48%	14%	—	8%
30-49	39%	17%	—	12%
50-64	38%	18%	—	12%
65+	35%	15%	—	21%

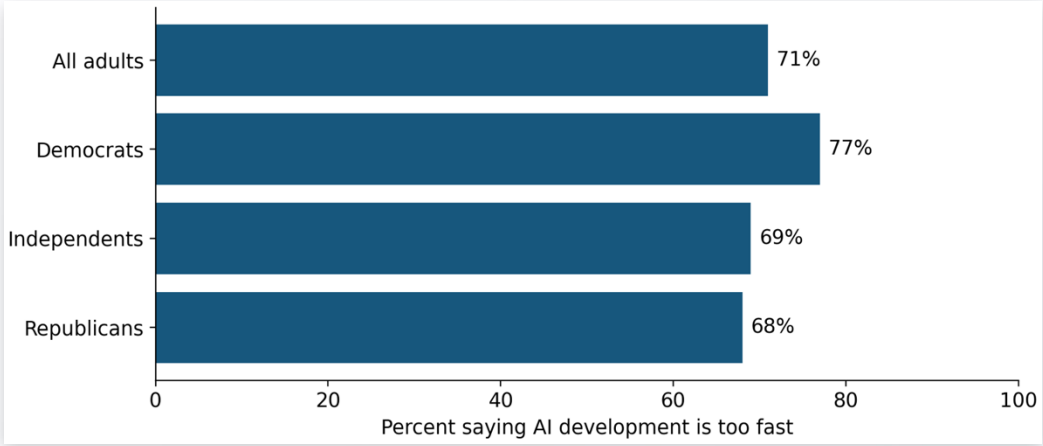
Note: Negative and positive combine very/somewhat responses. A dash means the precise figure was not published in the reviewed Pew age chart; it is not zero. Percentages may not total 100 because of rounding or nonresponse. *Asian estimates represent English-speaking adults only. White, Black and Asian groups are non-Hispanic single-race; Hispanic adults may be of any race.

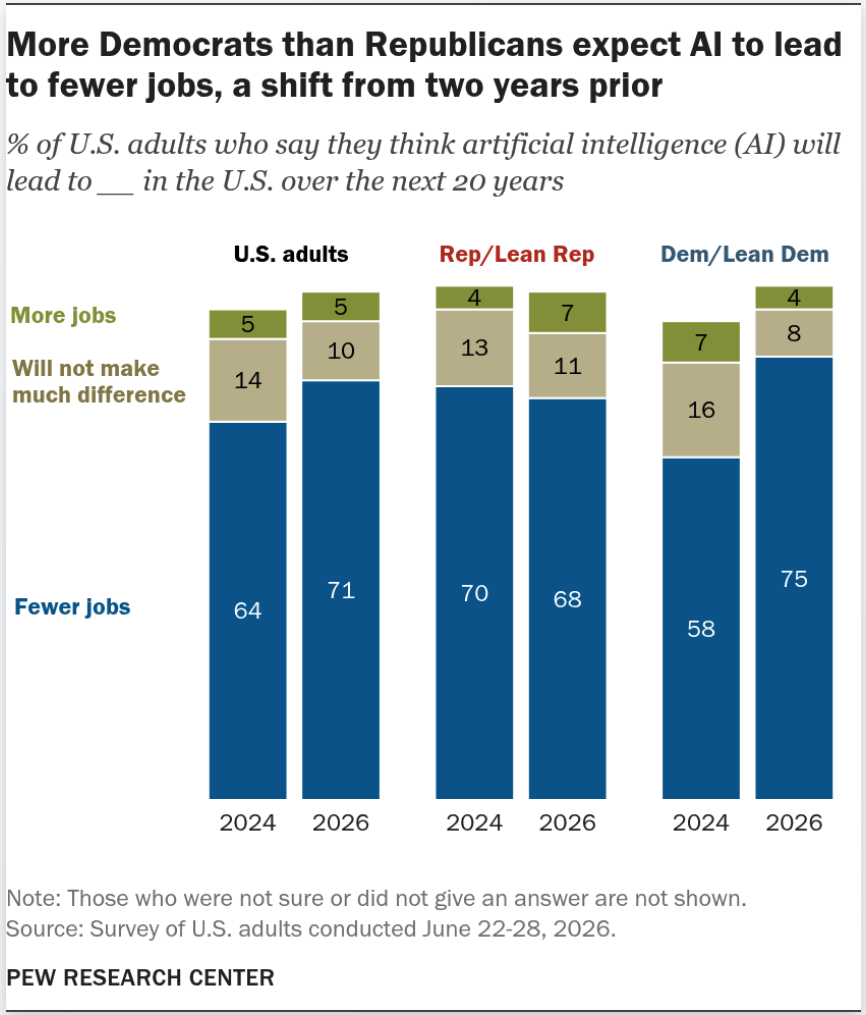
Source: Pew Research Center, *Americans and AI 2026: Chatbots, Smart Devices and Views on Impact* (June 17, 2026), survey of 5,119 U.S. adults, Feb. 17-23, 2026. pewresearch.org/internet/2026/06/17/americans-and-ai-2026-chatbots-smart-devices-and-views-on-impact/

Pew Research Center, Sept. 17, 2025

Partisanship: Minor Differences



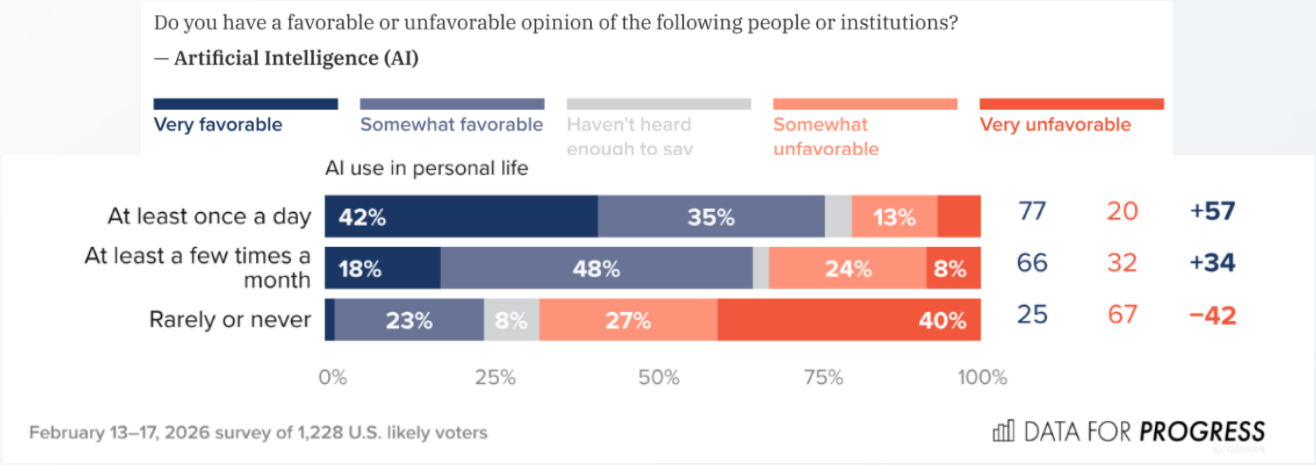




Economist/YouGov, May 2026; Pew Research Center, Nov. 6, 2025 and Sept. 16, 2026

Attitudes by Frequency of Use: To Know AI Is to Love AI… (or at Least Hate It a Bit Less…)

Survey · measure	Frequent users	Occasional users	Non-users
Gallup–Microsoft · trust AI for accurate info	45% (daily)	26% weekly; 15% rarely	8%
Gallup–Microsoft · worried about AI	68% (daily)	80% (less often)	74%
SCSP/Gallup · trust AI to be fair and unbiased	43%	—	12%
Annenberg · expect negative societal impact	29% (heavy)	40% (light)	54%
Data for Progress · net favorability	+57 (daily)	—	−42
Annenberg · oppose local data center	60% (heavy)	60% (light)	64%



Gallup–Microsoft, 2026; SCSP & Gallup, Sept. 2026; Annenberg Public Policy Center, 2026; Data for Progress, Feb. 2026

Research Agenda

Information ecosystem

- Scale — share of consumed media that is synthetic (average and tails)
- Strength — is synthetic media more persuasive/powerful, and when?
- Character — who propagates synthetic media and why?

Election ecosystem

- Elite use (parties, candidates, campaign consultants, media)
- Mass utilization (election information, decision making)
- Election administration

Government

- Simulated legislatures/referendums
- Delivery of services
- Surveillance/propaganda

Thank You

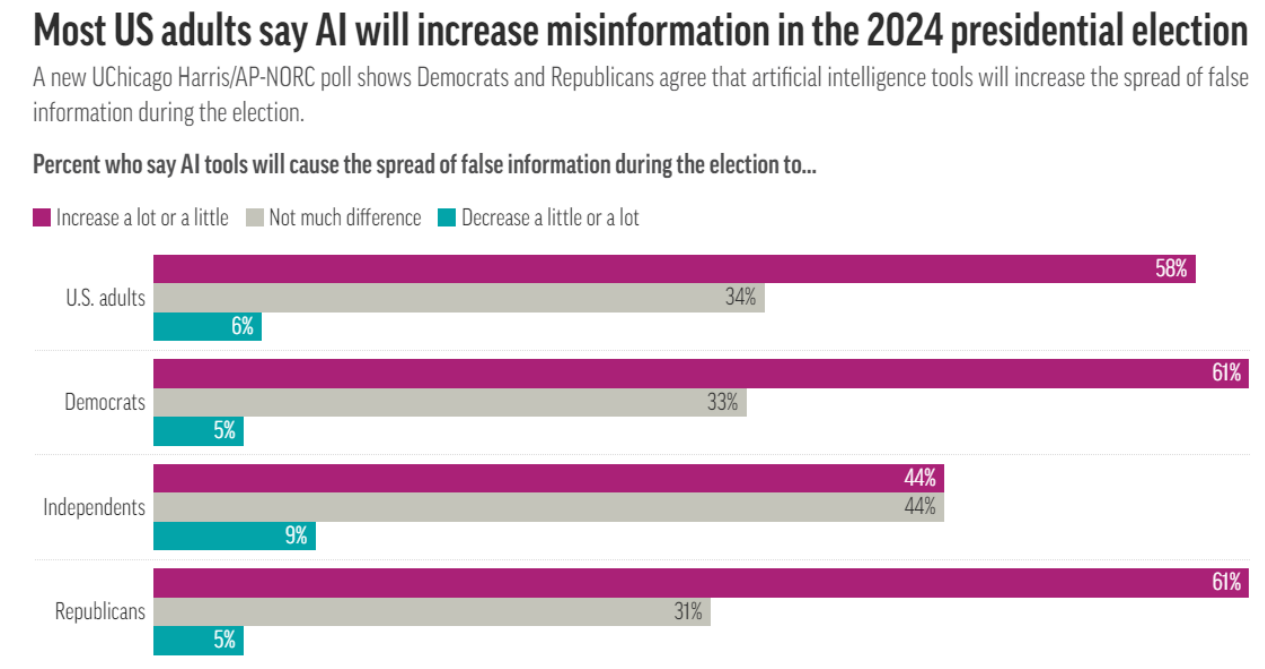
npersily@law.stanford.edu





Additional Slides

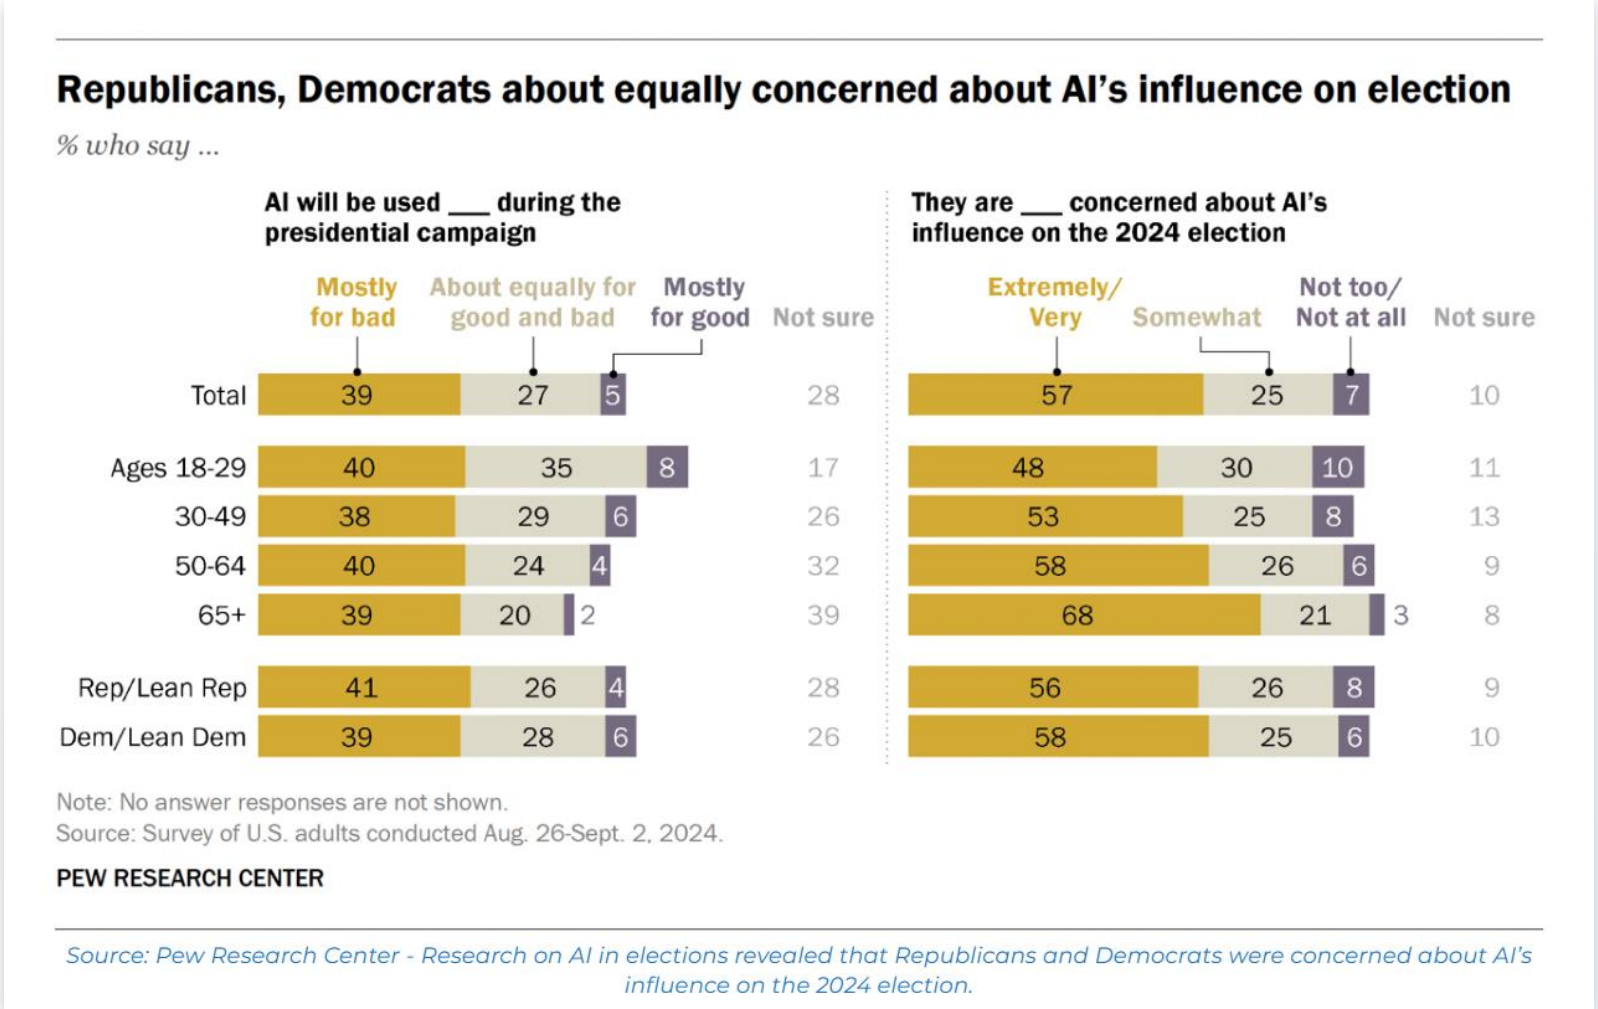
What the Public Thinks About AI and Elections (2023)



Do you expect artificial intelligence (AI) to increase or decrease your trust in the following, or do you think it will have no impact? Advertisements I see for candidates in upcoming U.S. elections		
Increase trust	320	15%
No impact	1114	51%
Decrease trust	768	35%
Do you expect artificial intelligence (AI) to increase or decrease your trust in the following, or do you think it will have no impact? The outcome of U.S. elections		
Increase trust	332	15%
No impact	1098	50%
Decrease trust	773	35%
Do you think misinformation spread by artificial intelligence (AI) will have an impact on who wins the upcoming 2024 U.S. presidential election?		
Yes, definitely	510	23%
Yes, probably	657	30%
No, probably not	344	16%
No, definitely not	135	6%
Don't know / No opinion	557	25%

Morning Consult/Axios, Aug. 10–13, 2023; UChicago Harris/AP-NORC (2023)

AI and the 2024 Election, by Party



Source: Pew Research Center - Research on AI in elections revealed that Republicans and Democrats were concerned about AI's influence on the 2024 election.

Pew Research Center, Aug. 26–Sept. 2, 2024



AI, Democracy, and Authoritarianism

Overview of AI and Democracy?

A broader frame: AI and its impact on different political regimes
- Democracy and Authoritarianism

Does AI favor Authoritarianism?
- First order effects
- Second order effects

Starting claim: democracy > authoritarianism

Democracies perform better over time

- Growth: about 20% higher income per person in the long run
- Stability: fewer economic disasters, less volatility
- No major famines with elections and a free press
- Peace: democracies rarely fight each other
- Error correction: mistakes get exposed and fixed

Liberty and Equality matter

- Self-direction and dignity
- Equal voice in collective decisions
- Freedom is both the means and the end
- Extending liberty to people on an equal basis undercuts caste and second-class citizenship

Today's question: could AI erode these advantages?

Democracies dominate the world economy

9 of 10

largest economies are democracies

World Bank, 2021

17 of 20

G20 economies are democracies

IMF

19 of 20

richest countries per capita are democracies

United Nations, 2023

“

No one pretends that democracy is perfect or all-wise. Indeed it has been said that democracy is the worst form of government except for all those other forms that have been tried from time to time.

Winston Churchill, House of Commons, November 11, 1947



Source: Imperial War Museums (NYP 45063)



Does AI Favor Autocracy?

AI, power, and the future of democracy

Two hypotheses

1

Architecture

AI is centralizing: it rewards concentrating data and decisions

Harari

2

Practical effects

AI makes control cheaper, faster, and scalable

Feldstein & Joshi

Three levels of impact



Epistemic

What citizens know and believe



Material

How elections and government work



Foundational

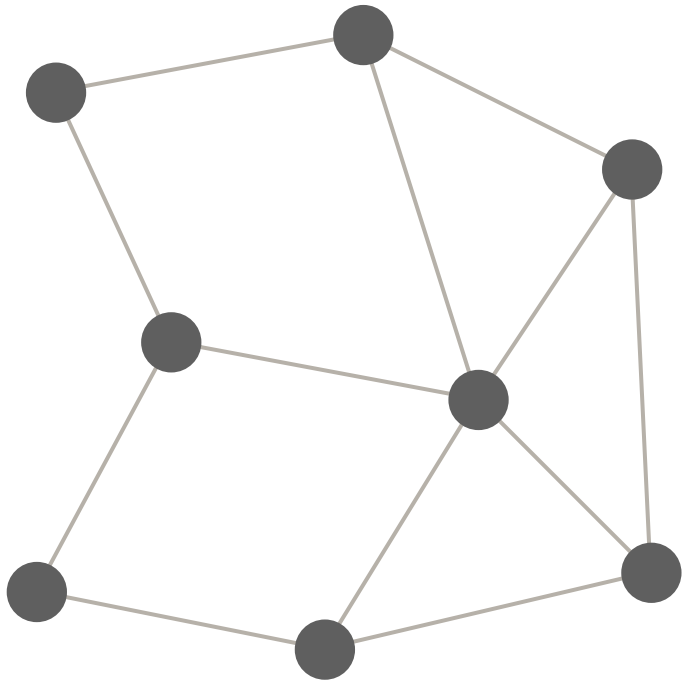
Who holds power and is accountable

Framework: Summerfield et al., 2024



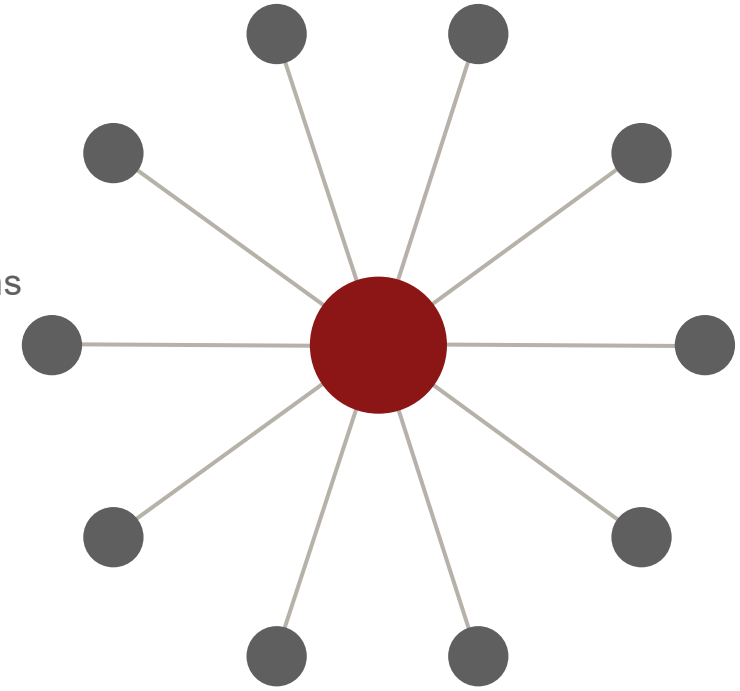
Hypothesis 1: Is AI a centralizing technology?

Harari: regimes as data-processing systems



Democracy: distributed

20th century: distributed wins
AI era: centralization becomes efficient?



Dictatorship: centralized

Harari: more data, better AI?

> If you disregard all privacy concerns and concentrate all the information relating to a billion people in one database, you can train much better algorithms than if you respect individual privacy and have in your database only partial information on a million people.

Harari, 21 Lessons for the 21st Century (2018), p. 66

Holds for surveillance AI: data on behavior (faces, movements) helped China's facial-recognition firms

Stress test

$286B vs $12B
US vs China private AI investment, 2025

2.7%
US lead at the frontier, March 2026: near parity

Compute, talent, open web
the inputs that drive frontier AI, not citizens' private data

Verdict: right about surveillance AI, wrong so far about frontier AI

Is Harari Right?

For

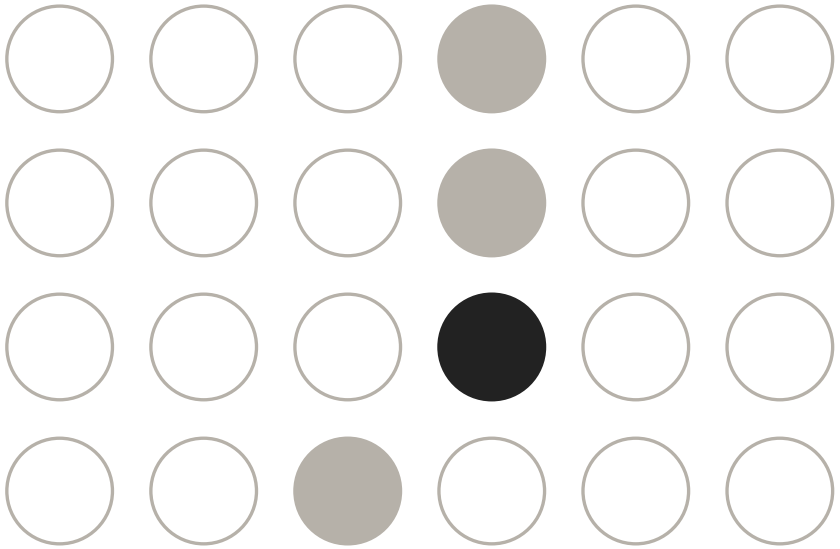
- Unrest → more AI surveillance → less unrest
- Censorship scales: 1.05B users online
- State-built networks are easier to control

Against

- Repression distorts data about what people think
- No feedback → spirals of delusion
- Still needs humans and huge budgets

AI can watch what people do; also what they think?

Bostrom: the vulnerable world hypothesis



Most inventions are white. One black ball could devastate civilization.

$140

per person per year for a “High-tech Panopticon”

< 1% of world GDP

AI makes ubiquitous surveillance affordable

A utilitarian disaster for open societies?

Bostrom's case

- Survive a black-ball technology
- Less crime, more trust
- Standby beats a crisis build

The realistic risk

- Turnkey totalitarianism
- Captured systems; no error correction
- Crisis builds are never repealed



Hypothesis 2: AI makes control cheaper

Digital repression is spreading

86 → 91

countries above the mean on digital repression

0.12 → 0.16

global mean repression score

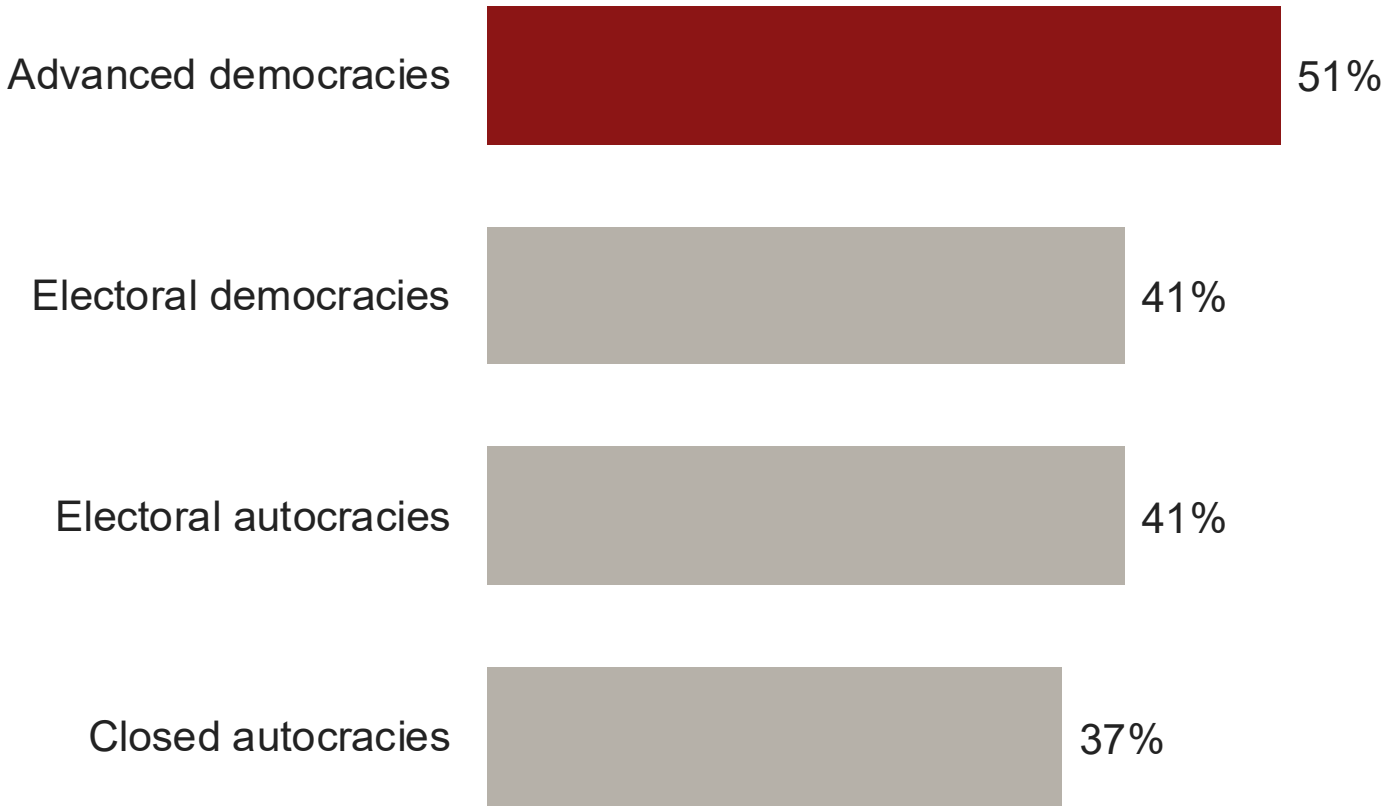
2021 → 2025 · Rising even in the US, Argentina, and Italy

Source: Feldstein & Joshi, Carnegie, 2026

The repression stack

War	50 AI targets a day vs. ~100 a year
Transnational	Thai CCTV used against foreign exiles
Cognition	LLMs that recite the state's line
Flooding	Generative fakes in the Iran–Israel war
Content	WeChat silently drops messages
Surveillance	75 of 176 countries by 2019
Access	Iran's “white SIMs” for loyalists
Infrastructure	Junta-linked ISPs in Myanmar

Democracies use AI surveillance too



Governance, not technology, best predicts abuse

Source: Feldstein, Carnegie, 2019 (176 countries)



Inside democracies: threats and promise

AI persuasion: real but modest

2–4%

shift from AI-written policy messages

≈ 0

extra effect from microtargeting

–20%

conspiracy belief after AI dialogue

The bigger change: persuasion is now nearly free to produce

Can AI strengthen democracy?

Promise

106,000 children enrolled in food aid, no forms (NJ)

50,000 public comments sorted in ~2 hours (UK)

16,000 residents deliberating online (St. Louis)

Peril

20,000+

AI-generated emails, some in real people's names, helped sink an air-quality rule

“Democracy requires friction to function.”



Who governs the stack?