



CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

SEPTEMBER 2026

The 2026 Carnegie California AI Survey

Ian Klaus
Mark Baldassare
Rachel George
Marissa Jordan
Abigail Manalese

Carnegie California

The 2026 Carnegie California AI Survey

Ian Klaus
Mark Baldassare
Rachel George
Marissa Jordan
Abigail Manalese

September 2026

About the Carnegie Endowment for International Peace

In a complex, changing, and increasingly contested world, the Carnegie Endowment generates strategic ideas, supports diplomacy, and trains the next generation of international scholar-practitioners to help countries and institutions take on the most difficult global problems and advance peace. With a global network of more than 200 scholars across twenty countries, Carnegie is renowned for its independent analysis of major global problems and understanding of regional contexts.

© 2026 Carnegie Endowment for International Peace.
All rights reserved.

Carnegie does not take institutional positions on public policy issues; the views represented herein are those of the author(s) and do not necessarily reflect the views of Carnegie, its staff, or its trustees.

No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from the Carnegie Endowment for International Peace. Please direct inquiries to:

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, DC 20036
P: + 1 202 483 7600
F: + 1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost at
CarnegieEndowment.org.

TABLE OF CONTENTS

Summary	1
Carnegie California AI Survey 2026	3
Economy, Work, and the Labor Market	4
Trust and Safety	10
Democracy and Government	15
Looking Forward	21
Advisers	21
Acknowledgments	22
Notes	23



Carnegie California

Carnegie California links developments in California and the West Coast with national and global conversations around technology, subnational affairs, and trans-Pacific relationships. At a distance from national capitals, and located in one of the world's great experiments in pluralist democracy, Carnegie California engages a wide array of stakeholders as partners in its research and policy engagement.

About the Authors

Ian Klaus is the founding director of Carnegie California, the West Coast center of the Carnegie Endowment for International Peace.

Mark Baldassare is a nonresident scholar at Carnegie California. He is the survey director at the Public Policy Institute of California, where he holds the Arjay and Frances Fearing Miller Chair in Public Policy. He is also a nonresident senior scholar at the Leonard D. Schaeffer Institute for Public Policy and Government Service at the University of Southern California.

Rachel George is a nonresident scholar at Carnegie California and a lecturer in international relations at Stanford University.

Marissa Jordan is the program manager of Carnegie California at the Carnegie Endowment for International Peace.

Abigail Manalese was an intern and research assistant with Carnegie California at the Carnegie Endowment for International Peace. She is a JD candidate at Loyola Law School and a graduate of the University of California, Los Angeles.

Summary

Launched in 2025, the Carnegie California AI Survey is now in its second year. Core questions around the economy, work, trust and safety, and democracy and governance remain benchmarks in the survey, offering longitudinal perspectives on the priorities of Californians. However, given the rapid pace of AI development, the survey has been adapted to reflect current discussions around AI, including its potential to impact the 2026 California election and broader workforce trends.

Overall, the 2026 survey revealed a growing expectation that AI will significantly alter work and the U.S. political economy more generally. As a result, candidates, policymakers, and scholars are increasing their attention to, and providing more detailed ideas around, significant changes to the social safety net. These proposals include centralized wealth pools funded by private companies, specialized taxes on top-earning companies, government-supported job loss trackers, and universal basic income.

Key survey findings point to increased uptake of AI technology and uncertainty around its impacts and the policies needed to address them. While the results reveal notable demographic, economic, and geographic divides, they also show, in the state and national context, important areas of substantial alignment between Republicans and Democrats.

- Most Californians believe AI is of strategic importance.** A majority says that AI is important to the nation's economic growth and competitiveness, and a similar share believes that AI is important to the state's economic growth and competitiveness. These results are similar to the 2025 results, and there is political consensus. Many Californians also express concerns about the negative impacts of AI on jobs and the gap between the rich and the poor. Most think that large businesses will benefit from AI and, of note, an increased share of respondents says that small businesses will also benefit from AI.
- Uptake of AI at work is increasing rapidly, but productivity gains remain limited.** About six in ten Californians say that they use AI daily or occasionally at work, which is an increase over the past year. And the share of respondents who say they rarely or never use AI for work has dropped by fifteen points. The use of AI at work varies by age group and level of education, while it is similar across genders and partisan groups. About half say that AI use is extensive or moderate at their work, which is also an increase from last year. Most say they have experienced minor or no improvement in productivity as a result of AI use.

- **Workforce training and support for AI skills building is a popular policy response to AI-driven changes in the workplace.** A majority of Californians supports redistributing profits from large companies by reinvesting in the workforce including support for increased AI skills training, and a similar majority reports being eager to receive such training. About half of Californians say that they expect to use AI tools at work more in the future, which is a growing share. Most believe that AI is important to a worker's success in today's economy. Californians are increasingly concerned about losing their jobs due to AI, and few expect there to be more job opportunities resulting from the technology. In this context, there is bipartisan support for a state policy that would protect workers from job displacement.
- **AI model safety is not just the concern of policymakers or experts.** Nearly three quarters of Californians agree that the government should require AI companies to test their most advanced systems for safety and provide a detailed plan for how they will prevent harm.
- **Californians are most concerned about privacy, cybersecurity, spam or fraud, misinformation, and deepfakes.** Notably, while the share of Californians concerned about cybersecurity harms is the same this survey year, the issue has risen in salience. Californians now rank this AI-related concern second, up from fourth in 2025.
- **The public is skeptical of AI use at any government level.** Under 10 percent of Californians support federal, state, or local government agencies using AI to make decisions that affect them, although Californians are more open to government use of AI in cases where use definitively improves decisionmaking.
- **Confidence in spotting AI-generated content is growing among California residents, even as concerns about its role in politics remain high.** About half of Californians now say they are at least somewhat confident they can tell real from AI-generated information, a notable increase from last year. Concern about AI-fueled election deepfakes and political polarization remains widespread, though it has eased somewhat since 2025.

Carnegie California AI Survey 2026

The findings in this paper are based on a survey of 1,500 adult residents of California, conducted via YouGov between April 21 and May 11 in English and Spanish according to respondents' preferences. Prior to the survey, in March, the Carnegie California survey team invited input, comments, and suggestions from policy experts and its own advisory group—including advisers from state and local government, California universities and think tanks, industry, and civil society. The team then determined the final survey methods and designed and categorized the questions into four topic areas.

YouGov initially surveyed 1,627 adult residents. The respondents were then matched down to a sample of 1,500 to produce the final data set. The respondents were matched to a sampling frame based on gender, age, race, and education. The sampling frame was constructed by stratified sampling from the full 2023 American Community Survey (ACS) one-year sample with selection within strata by weighted sampling with replacements (using the person weights on the public use file). The matched cases were weighted to the sampling frame using propensity scores. The matched cases and the frame were combined, and a logistic regression was estimated for inclusion in the frame. The propensity score function included age, gender, race/ethnicity, years of education, and home ownership. The propensity scores were grouped into deciles of the estimated propensity score in the frame and post-stratified according to these deciles. The weights were then post-stratified on 2024 presidential vote choice and on gender, age (four categories), race (four categories), and education to produce the final weight.

The YouGov panel includes information about each respondent's demographic and political profile. We present results for four racial/ethnic groups: Asian, Black, Latino, and White. Residents of other racial and ethnic groups are included in the results reported for all adults, but sample sizes for these less populous groups are not large enough to report separately. We present results for five geographic regions, representing approximately 90 percent of the state population. The Central Valley includes the counties of Butte, Colusa, El Dorado, Fresno, Glenn, Kern, Kings, Madera, Merced, Placer, Sacramento, San Joaquin, Shasta, Stanislaus, Sutter, Tehama, Tulare, Yolo, and Yuba. The San Francisco Bay Area includes the counties of Alameda, Contra Costa, Marin, Napa, San Francisco, San Mateo, Santa Clara, Solano, and Sonoma. Los Angeles includes Los Angeles County. Inland Empire includes the Riverside and San Bernardino counties. And Orange/San Diego includes the Orange and San Diego counties. Residents of other geographic areas are included in the results reported for all adults, but sample sizes for these less populous areas are not large enough to report separately. We also report the results for those who identify as Democrat, Republican,

independent, or other, but sample sizes for other voters are not large enough to report separately. Lastly, and important to this survey with its focus on jobs and the economy, we disaggregate the results by union and nonunion members, as well as age group, level of education, annual household income, and employment status.

The overall margin of error is +/- 3.5 percent. The margin of error is calculated at the 95 percent confidence interval. When applicable, we compare the findings to the 2025 Carnegie California AI Survey results and the 2023, 2024, and 2025 Carnegie California Global Affairs Survey results, which were conducted with the same methodology and used national surveys with similar questions, such as Gallup, Ipsos, Politico, Pew, and YouGov.¹

Economy, Work, and the Labor Market


In May 2026, California Governor Gavin Newsom signed a first-of-its-kind executive order (EO), Executive Order N-6-26, which directs California to prepare workers, small businesses, and communities for workforce disruptions due to artificial intelligence. The EO is meant to balance tracking workforce disruption caused by AI with enabling workers to benefit from “AI-driven productivity.”² But while policymakers are evidently aware of AI’s potential growing impact on California’s economy, work, and labor market, how are Californians viewing the changing conditions as AI use continues to increase?

Economy

Most Californians are keenly aware of AI’s economic significance. Seven in ten adult residents say that AI is important (22 percent “very,” 48 percent “somewhat”) to both the nation’s economic growth and competitiveness and to California’s economic growth and competitiveness (23 percent “very,” 44 percent “somewhat”). The belief that AI is important to the nation’s and California’s economic growth and competitiveness has not significantly changed in the past year. These views are widely held across demographic groups and state regions. Notably, majorities across partisan groups think that AI is important for the nation’s economic growth and competitiveness (73 percent Democrats, 81 percent Republicans, 64 percent independents). There is similar consensus among Democrats and Republicans when they are asked to rate the importance of AI for the state’s economic growth and competitiveness (see figure 1).

Figure 1. AI and California's Economic Growth

How important do you think AI is to California's economic growth and competitiveness?



Source: 2026 Carnegie California AI Survey
N: 1,500

Californians continue to have mixed views when asked if AI will make California's economy "better" in the next three years (27 percent in 2025, 31 percent in 2026), "worse" (31 percent in 2025, 31 percent in 2026), or "stay the same" (20 percent in 2025, 19 percent in 2026). Today, fewer than half across partisan and demographic groups and state regions expect that AI will make the California economy better in the next three years. The belief that AI will make California's economy better is most widely held by those under thirty years old (44 percent). About half of Californians continue to think that the increased use of AI will make the job market worse in their part of California (50 percent in 2025, 47 percent in 2026). And about half of Californians continue to think that the increased use of AI will make the gap between the rich and the poor worse in the next three years (52 percent in 2025, 48 percent in 2026). Today, pluralities of Californians across partisan, age, education, income groups, and regions think the increased use of AI will make the job market and the gap between the rich and the poor worse in their part of the state in the next three years.

A large but shrinking majority of adults believes that large companies will benefit from AI (68 percent in 2025, 61 percent in 2026), while a smaller and increasing share of adults think that small companies will benefit from AI (35 percent in 2025, 43 percent in 2026). Today, majorities of Democrats (65 percent), Republicans (70 percent), and adults across

demographic groups and state regions think that large companies will benefit from AI. Majorities of Republicans (60 percent), those under thirty years old (58 percent), those earning \$100,000 or more (53 percent), and full-time workers (51 percent) believe that small companies will benefit from AI.

The belief that AI is important to the nation's and California's economic growth and competitiveness has not significantly changed in the past year.

How would Californians prefer to spend the profitability gains from corporate implementation of AI? Half or more would opt for redistributing AI profit by reinvesting it in the workforce (54 percent) and passing it on to customers in the form of lower prices (50 percent). Fewer want the AI profit to be reinvested in innovation, research, and development (36 percent); data safety and security efforts (36 percent); severance and job search assistance to laid off workers (36 percent); or environmental sustainability (29 percent). Only one in four want to redistribute AI profit by delivering returns to shareholders (26 percent).


Work

AI is playing an increasingly prominent role in the California workplace, and this is resulting in changes in the nature of work itself and growing concerns among employees about their futures. About half of employed residents, and a growing share, believe that AI is being used in their workplace or industry “extensively” or “moderately” (45 percent in 2025, 51 percent in 2026), while fewer are saying that it is being used “minimally” or “not at all” compared to last year (48 percent in 2025, 36 percent in 2026). Today, six in ten or more say that AI is being used “extensively” or “moderately” among those under thirty years old (60 percent), college graduates (60 percent), and those earning over \$100,000 (62 percent).

Sixty percent of employed adults say they use AI for work either “daily” or “occasionally,” compared to 44 percent in last year’s survey. The share who say they “rarely” or “never use” AI for work has dropped by fifteen points (55 percent in 2025, 40 percent in 2026) (see figure 2). Today, the main reasons given for not using AI more at work include a lack of trust in the security of personal data (25 percent), a lack of confidence in the accuracy and reliability of the AI output (25 percent), and a lack of training and education in how to use AI (22 percent). Two in three or more workers are using AI daily or occasionally among those living in the San Francisco Bay Area (73 percent), under thirty years old (70 percent), college graduates (66 percent), and those earning over \$100,000 (66 percent). Similar shares across partisan groups (59 percent Democrats, 63 percent Republicans, 62 percent independents) are using AI daily or occasionally at work.

Figure 2. Current Use of AI at Work

How often do you use AI for work?



Source: 2026 Carnegie California AI Survey
N: 508

About one in four workers continues to say that AI has led to “great” improvement in their work or productivity at their job. A shrinking share of workers reports that AI has led to a “minor” improvement in their productivity at their job (46 percent in 2025, 36 percent in 2026), while a growing proportion says there has been “no noticeable” improvement (26 percent in 2025, 35 percent in 2026). Today, employed adults under thirty years old are the most likely to say that AI has led to “great” improvement of their work or productivity at their job (43 percent). In the broader context, 18 percent say that AI is speeding up work “a lot” in their industry overall. There are fairly pronounced racial differences in reported productivity gains from AI at work. Black Californians are the most likely to report a “great improvement” in their work or productivity (44 percent), compared with 32 percent of Whites, 26 percent of Hispanics, and 11 percent of Asian Californians.

While Californians report mixed gains in the workplace, relatively few California workers say they have taken any classes or received any training on AI tools in the past twelve months (17 percent in 2025, 13 percent in 2026). Today, those living in the San Francisco Bay Area (20 percent) and those with incomes over \$100,000 (21 percent) are the most likely to have had AI classes or training. Nonetheless, six in ten employed residents continue to say they are “very” or “somewhat” interested in being offered training and courses on AI uses

in the future at work (58 percent in 2025, 61 percent in 2026). Today, about half or more California workers across partisan and demographic groups and state regions have at least some interest in being offered training and courses on AI uses in the future at work.

Relatively few California workers say they have taken any classes or received any training on AI tools in the past twelve months.

Lastly, majorities of employed adults continue to say that AI has not changed their career plans (57 percent in 2025, 58 percent in 2026), and few say that AI has changed their career plans (5 percent in 2025, 4 percent in 2026) or shifted their career plans (10 percent in 2025, 12 percent in 2026). Today, those under thirty years old are the least likely to report that AI has not changed their career plans (38 percent) and most likely to say that AI has impacted their career path or career plans.

Labor Market


Many Californians believe AI will play a larger role at work in the future. They have increasing concerns about the negative impacts of AI on their current employment and future prospects in the labor market. In this context, Californians want to see the government make policies and employers take actions to protect employees from the effects of job displacements caused by AI.

A growing share of California's employed adults says that they expect to be using AI more in the future at work (48 percent in 2025, 53 percent in 2026) (see figure 3). While about half or more adults across California's regions hold this view today, there have been noteworthy increases in this perception in the Inland Empire (44 percent in 2025, 60 percent in 2026) and Orange/San Diego counties (41 percent in 2025, 53 percent in 2026) in the past year. Interestingly, Democrats (58 percent) and Republicans (61 percent) have similar expectations about using AI tools more in the future at work. The expectations for using AI declines with age and increases with education and income.

An overwhelming majority believe that the ability to understand and use AI is important for a worker to be successful in today's economy (77 percent in 2025, 75 percent in 2026). And this perception is widely held across partisan and demographic groups and state regions. Half of Californians also believe that AI affects early career employment either "extensively" or "moderately" in their part of California. This view is most often held among those under thirty years old (69 percent), those earning \$100,000 or more (64 percent), college graduates (62 percent), and San Francisco Bay Area residents (62 percent).

Figure 3. Future of AI Use at Work

Do you personally expect to be using AI tools more in the future at work?



Source: 2026 Carnegie California AI Survey
N: 698

These perceptions of AI's profound impacts on the labor market are creating anxiety about current jobs and future employment prospects. An increasing share of California's employed adults is concerned about losing their job due to being replaced by AI in the next three years (31 percent in 2025, 37 percent in 2026). Today, San Francisco Bay Area residents report the highest and growing concern for losing their job due to AI (35 percent in 2025, 47 percent in 2026), followed closely by Orange/San Diego (25 percent in 2025, 41 percent in 2026). Moreover, relatively few Californians think that they will have more job opportunities because of the use of AI in the workplace (8 percent in 2025, 11 percent in 2026), while most believe they will have fewer job opportunities (49 percent in 2025, 44 percent in 2026) or it will not make much difference for their job prospects (26 percent in 2025, 27 percent in 2026). Pluralities across partisan and demographic groups and regions expect fewer job opportunities.

Do Californians favor a role for state government policy and specific initiatives to deal with the changing labor market due to AI? Two in three adult residents support a state policy that would protect workers from job displacement caused by AI (69 percent in 2025, 66 percent in 2026). Majorities across partisan groups (79 percent Democrat, 61 percent Republican, 69 percent independent) and in all demographic groups and state regions support this state policy on AI job displacement. When asked about several company safeguards for

job displacement due to AI, about half of Californians are in favor of requiring advanced notice to employees and offering other employment within the company (49 percent), and about half are in favor of ensuring that high-risk decisions are made with human oversight (47 percent). When asked to choose among several new initiatives that they would be most likely to support, about half of Californians name either mandatory employer-funded retraining for employees displaced by AI (22 percent), a wealth tax on the highest-earning AI companies to fund support for workers displaced by AI (14 percent), or specialized taxes on companies implementing AI in ways that impact jobs (11 percent). In sum, many Californians are in favor of state government policy, company safeguards, and new initiatives to deal with concerns about job displacement due to AI.

Trust and Safety

In January 2026, California’s vanguard AI safety law, SB 53, went into effect. The law offers an approach to balancing innovation and safety that has been adopted elsewhere, but essential implementation gap questions remain around incident reporting, frontier AI frameworks and transparency reports, whistleblower protections, independent and third-party evaluations, internal deployments, scoping of SB 53, the state and federal relationship, and “trust but verify” principles both within and beyond California. Meanwhile, in the first half of 2026, questions on trust in AI also coalesced into conversations around child safety and data centers. Important legislative debates early in the year saw calls for the regulation of chatbot usage by minors and high-profile legal disputes involving the regulation of “addictive” social media.³

AI development in California has led to some of the earliest and most dynamic policy debates between state and federal concerns about trust and safety. The nation saw tensions over California-based AI company Anthropic and the federal government in the *Anthropic v. Department of Defense* legal dispute, the release of Mythos and the White House executive order preempting state AI-regulation, and the public release of Fable 5 restricted by the following U.S. Export-Control Order.⁴ At the center of these debates were often questions surrounding infrastructure concerns, such as potential cybersecurity harms and supply chain disruption. When asked about mandated testing and harm reduction in AI safety, a majority of Californians across all regions (over 50 percent) prefer national standards over state or regional standards

The 2026 iteration of the AI survey included a greater focus on questions regarding trust in AI companies, social companionship, autonomous decisionmaking, data centers, and overall risks. For the most part, Californians’ trust and safety concerns have remained consistent over the past year. The most drastic changes—those involving changes of ten or more percentage points—show milder sentiments increasing neutrality toward AI. Percentages of residents who are “very concerned” have decreased since last year, with some now saying


“somewhat concerned” or “don’t know.” The data suggest a potentially less polarized landscape, but also perhaps increased uncertainty due to a surge of AI-related information in the past twelve months.

Trust

Californians remain broadly concerned about a wide swath of AI-related risks, and the priority issues are also largely the same: privacy harms, such as misuse of personal information (63 percent); cybersecurity harms, such as helping malicious actors identify and exploit vulnerabilities in computer systems (61 percent); use for spam or fraud (59 percent); use to manipulate others or spread misinformation (56 percent); and use to alter or manufacture false content, such as audio or video deepfakes (53 percent) (see figure 4). Democrats are most concerned with privacy harms (71 percent), Republicans with cybersecurity harms (64 percent), and independents equally with privacy harms (58 percent) and use for spam or fraud (58 percent).

Figure 4. AI Risk Concerns

Which categories of risks are you most concerned with (select all that apply):



Source: 2026 Carnegie California AI Survey
2025 N: 1,601; 2026 N: 1,500

Trust in AI companies remains unchanged from 2025: Most Californians continue to distrust AI companies to adequately protect consumers. Fifty-three percent of Californians reported disagreement (32 percent “strongly,” 21 percent “somewhat”) when asked if they agreed or disagreed with the notion that companies will build and use artificial intelligence systems in a way that protects their personal data. On top of this distrust, Californians remain in the dark about what sorts of data companies collect and share about them (65 percent 2026, 67 percent 2025). They are not comfortable with AI systems accessing any personal content (72 percent) regardless of whether it is published publicly (69 percent in 2025) or privately (79 percent in 2025). They also have mixed attitudes about companies utilizing facial recognition (34 percent “not comfortable,” 27 percent “somewhat comfortable,” 23 percent it “depends on the purpose,” 9 percent “very comfortable,” 7 percent “don’t know”). Sentiments over the past year are clear: California residents still do not trust AI companies to act in their best interests.


**Trust in AI companies remains unchanged from 2025:
Most Californians continue to distrust AI companies
to adequately protect consumers.**

The tradeoffs between transparency and privacy on the one hand and consumer cost on the other hand are notable. There is less support for transparency when it comes with cost. Last year, we asked Californians, “How strongly do you agree or disagree with the following statement: Companies should be required to tell you the types and sources of your personal data that an automated decisionmaking system analyzes when helping to make important decisions about you?” This year, the survey team added a caveat—that this transparency condition would add extra costs to company services. Responses were more mixed compared to last year. Most respondents across gender, age, party affiliation, and residence all shifted ten or more percentage points away from “strongly agree.” For example, 59 percent of Los Angeles residents reported “strongly agree” last year, and less than half (48 percent) reported the same this year (see figure 5).

Trust in AI models themselves is a slightly different story. Californians trust AI models to provide useful information, but do not trust them to make independent decisions. A plurality (44 percent) of Californians trusts that AI systems are at least somewhat able to provide useful information. Majorities of those eighteen to twenty-nine years old (55 percent), college graduates (51 percent), those earning over \$100,000 (54 percent), those with a child in the home (55 percent), nonunion members (54 percent), part-time workers (63 percent), and big city residents (56 percent) trust that AI can provide useful information. But there is a disconnect between information and decisionmaking. Many California residents do not trust AI systems to usefully make decisions and/or take actions; 42 percent of Californians report not trusting AI systems (21 percent “do not trust much,” 21 percent “do not trust at all”).

Figure 5. Transparency and Personal Data Use for AI Models

How strongly do you agree or disagree with the following statement: Companies should be required to tell you the types and sources of your personal data that an automated decisionmaking system analyzes when helping to make important decisions about you (such as determining your access to lending, insurance, housing, education, or employment opportunities) even if it increases the cost of their services?



Source: 2026 Carnegie California AI Survey
N: 1,500

Safety

Social Companions

A handful of bills in the California 2026 legislative session, as well as a prominent referendum, consider the issue of child safety in interactions with AI models. This section of the survey asked Californians about their perceptions—both positive and negative—on social companionship with AI models.

Nearly half (46 percent) of Californians disagree with the idea that AI-powered social companions can effectively alleviate feelings of loneliness in their communities. On child protections, Californians with a child in the home are evenly split on whether AI social companions can effectively alleviate loneliness (41 percent “agree,” 41 percent “disagree”) and 19 percent report “don’t know.” Younger Californians tend to hold a more positive view of potential for social companionship, with 43 percent of eighteen- to twenty-nine-year-olds at least in somewhat agreement versus only 27 percent of those ages sixty-five and older in somewhat agreement.

Overall, Californians disagree on perceived social companionship benefits and simultaneously indicate similar levels of concern. We asked Californians, “How concerned are you about your family members relying on AI for social companionship?” Their responses were mixed: 43 percent reported at least some level of concern (17 percent “very,” 26 percent “somewhat”), 45 percent reported no concern, and 12 percent were unsure. Similarly, we also asked respondents to rank confidants as sources of personal advice, including an AI chatbot, a faith or community leader, a professional counselor or therapist, a family member or a friend, and a mentor or professional colleague. Most Californians ranked a family member or a friend first (59 percent) and AI chatbots last (60 percent). Within regions, 12 percent of San Francisco Bay Area residents ranked AI chatbots as their first choice, compared to 11 percent of the Central Valley, 8 percent of the Inland Empire, 6 percent of Los Angeles, and 4 percent of Orange and San Diego residents.

Data Centers

Data center development has been a polarizing issue for Californians and other Americans across the country. For instance, the Nevada-based company that supplies California residents in South Lake Tahoe with electricity recently ended an agreement with its California counterpart to focus on building data centers in Nevada⁵—the long-term impacts of which can only be met by further expansion of energy infrastructure in both states.

A national survey conducted by the Pew Research Center in January 2026 asked Americans about their attitudes on data centers and related policy issues.⁶ Pew defined the centers as “buildings filled with computer equipment that powers streaming services, AI, online banking, and other things.” Carnegie California asked the same questions in April and May of 2026 to compare how Californians felt in comparison to the nation.

On hearing about data centers, California residents’ awareness is on par with the national average (45 percent “a little,” 24 percent “a lot”). But Californians consistently hold a more pessimistic view of data centers than the average American regarding their negative impacts on the environment (47 percent of Californians versus 39 percent of Americans), local quality of life (44 percent versus 30 percent), and local jobs (30 percent versus 15 percent)—though they are slightly more optimistic about benefits to local tax revenue (31 percent versus 23 percent).

Californians consistently hold a more pessimistic view of data centers than the average American.

In sum, Californians’ main trust and safety concerns have remained the same over the past two years, with a slight shift in priorities toward cybersecurity. Aside from general harms, Californians are overall more pessimistic than their national counterparts on newer issues concerning social companionship and data centers.

Democracy and Government

California voters head to the polls in November 2026 to elect a new governor, a race in which candidates' positions on AI are at least "somewhat important" to a majority of Californians (52 percent) or "very important" (18 percent). Against this backdrop, the Carnegie California survey team asked Californians how they view AI's growing presence in public services, its potential to shape elections and democratic processes, its role in schools, and its place in the future of government and public safety.

Similar to last year, this survey found that the public is generally cautious about AI's impact on government and politics. Californians remain broadly skeptical of governments using AI in decisionmaking, even as a growing share reports that AI has made accessing public services somewhat easier. Concerns about AI-generated content influencing politics remain widespread, though are somewhat less acute than a year ago. Support for a larger state role in AI, whether in schools or through a public computing option, is divided, with some schisms falling along partisan and demographic lines.

Public Services

As government agencies increasingly use AI, including for citizen-facing AI chatbots and public-facing assistants as well as for internal productivity tools, few Californians report that AI has dramatically changed their ability to access government services—although the share reporting that they see some improvement is growing.⁷ In 2026, nearly a quarter of Californians (24 percent) say that AI has "significantly" or "slightly" improved their ability to access public or government services, up from 17 percent in 2025; while a similar share (18 percent in 2026, 18 percent in 2025) reports that it has hindered their access. The share saying they "don't know" fell from 37 percent in 2025 to 32 percent in 2026, suggesting Californians are forming more settled views on how AI impacts their engagement with government services, even if those views remain mixed.

Californians' reported personal use of digital tools to access government and public services has remained relatively unchanged since 2025. Among those who report using digital tools for engaging with government services, health services remain the most common use case (21 percent in 2026, 23 percent in 2025), while use of digital tools to access legal advice rose (13 percent in 2026, 10 percent in 2025). As in 2025, a gender gap persists in the use of digital tools for civic purposes: Men are more likely than women to report using digital tools to access election information (16 percent versus 10 percent) and legal advice (14 percent versus 12 percent), a pattern nearly identical to last year's survey. Younger Californians also continue to use digital tools at higher rates than older residents, particularly for election information (22 percent of those ages eighteen to twenty-nine, compared with 9 percent of those ages forty-five to sixty-four).


Geography also reflects differences in Californians’ engagement with digital government services. Residents of the Inland Empire (73 percent) and Orange/San Diego (68 percent) are considerably more likely than those in Los Angeles (53 percent), the San Francisco Bay Area (56 percent), and the Central Valley (56 percent) to say they have not used any digital tool to access government or public services in the past six months. The pattern varies somewhat by service: Central Valley residents report the highest use of digital tools for health services (29 percent), Bay Area residents the highest use for election information (18 percent), and Los Angeles residents the highest use for transportation services (20 percent), while Inland Empire and Orange/San Diego residents report the lowest use across most categories measured.

Mis/Disinformation

Californians remain deeply concerned about the impact of AI on mis- and disinformation in the political sphere.⁸ Yet their confidence in their ability to distinguish real from AI-generated information has grown, even as concern about the technology’s effects on politics remains high. About half of Californians (50 percent) say they are “very” or “somewhat” confident that they can tell the difference between real and AI-generated information, up from 40 percent in 2025, while the share saying they are “not confident” dropped (50 percent in 2025, 39 percent in 2026) (see figure 6). This growing confidence is most

Figure 6. Confidence in Identifying AI-Generated Information

How confident are you that you can tell the difference between real and AI-generated information?



Source: 2025 and 2026 Carnegie California AI Survey
2025 N: 1,601; 2026 N: 1,500

pronounced among younger Californians: 68 percent of those ages eighteen to twenty-nine report at least some confidence, compared with 36 percent of those age sixty-five and older. Republicans (57 percent) and Democrats (54 percent) report similar levels of confidence, while independents lag behind (47 percent) and are also the most likely to say they “don’t know” (20 percent).

Californians’ confidence in detecting AI-generated content somewhat outpaces the national picture. A Pew Research Center survey conducted in June 2025 found that 53 percent of American adults said they were “not too” or “not at all” confident they could tell AI-generated content from human-made content, while only 39 percent of Californians reported the same in 2026 (though the two surveys slightly differed in their response scales).⁹

Despite this growing confidence among Californians in detecting AI content, concern about the influence of AI-generated content on elections and political life remains high. Half of Californians (51 percent) say they are “very concerned” about the role of deepfakes and other AI-generated content in elections, down from 57 percent in 2025, though the combined share who are at least “somewhat concerned” has fallen only modestly (80 percent in 2026, 84 percent in 2025). A similar pattern holds for concern about AI-generated content heightening political violence and polarization (48 percent “very concerned” in 2026, versus 55 percent in 2025). As in the survey’s trust and safety findings, this softening appears to reflect a shift from “very concerned” into “somewhat concerned” and “don’t know,” rather than a broad decline in underlying concern. Democrats remain the most likely to be very concerned about deepfakes in elections (60 percent), compared with 52 percent of Republicans and 44 percent of independents. Older Californians are considerably more concerned than younger ones (66 percent of those ages sixty-five and older versus 45 percent of those ages eighteen to twenty-nine).

Concern about AI’s role in political life also varies by region. Los Angeles residents are considerably less likely than residents elsewhere in the state to say they are “very concerned” about AI-generated content heightening political violence and polarization (42 percent), compared with 52 percent in Orange/San Diego, 51 percent in the Inland Empire, 50 percent in the San Francisco Bay Area, and 48 percent in the Central Valley. A similar, if less pronounced, regional pattern holds for concern about deepfakes’ influence on elections, where Los Angeles (49 percent) and the Central Valley (49 percent) report somewhat lower concern than the Inland Empire (55 percent) and Orange/San Diego (52 percent).

Californians remain divided, and often unsure, about whether AI can play a constructive role in democratic life. Asked directly whether the use of AI in democratic processes, such as summarizing public policies or facilitating citizen input into government decisionmaking, would be helpful or harmful, more Californians say “helpful” (38 percent) than “harmful” (26 percent), though more than a third (36 percent) say they “don’t know.” Views on whether AI tools can help Californians become more informed and engaged voters and citizens have changed little since 2025 (33 percent “yes” in 2025 and 2026), as has the share who think emerging technologies can facilitate productive civic dialogue (28 percent “yes”


in 2026, 26 percent in 2025). When asked what would be the best way for AI to improve citizens’ engagement with government, no single answer commands a majority; the top response, improving access to information about one’s rights and applicable policies, draws only 20 percent, while more than a third (34 percent) say they “don’t know,” underscoring how unsettled Californians’ expectations remain on this question.

Government Use of AI

Californians remain cautious about government use of AI in decisionmaking, though this year’s survey offered new insight into the conditions under which citizens are more open to government using AI. Both the 2025 and 2026 surveys asked whether federal, and separately state and local, government agencies should use AI to make decisions that directly affect respondents and their communities. In 2026, however, the survey added a new middle option, “only if it improves the quality of the decisions,” which had not been offered in 2025. Because adding a new option necessarily draws some respondents away from the preexisting options, the shares choosing outright “yes” or “no” both declined for federal government use of AI (yes: 12 percent in 2025, 7 percent in 2026; no: 65 percent in 2025, 45 percent in 2026), with 37 percent of Californians choosing the new conditional option in 2026. A similar pattern holds for state and local government use of AI (yes: 11 percent in 2025, 7 percent in 2026; no: 67 percent in 2025, 45 percent in 2026), with 36 percent choosing the conditional option in 2026 (see figure 7).

Figure 7. State and Local Agency Use of AI

Do you believe that state and local government agencies should use AI to make decisions that directly affect you or your community?



Source: 2025 and 2026 Carnegie California AI Survey
2025 N: 1,601; 2026 N: 1,500
Note: The answer “Only if it improves the quality of the decisions” was not offered in the 2025 survey.

A similar pattern is evident in views of AI use by local police, fire departments, and emergency services: Outright opposition fell from 57 percent in 2025 to 40 percent in 2026, outright support fell from 19 percent to 9 percent, and 39 percent of Californians chose the new conditional option in 2026. Republicans are somewhat more supportive of unconditional AI use by emergency services than Democrats (17 percent “yes” versus 10 percent “yes”), a pattern consistent with 2025’s finding that men were more supportive than women of AI use in this context.

This measured, conditional stance toward AI in policing and emergency services echoes broader national findings. A 2026 study of 1,500 American adults found that public support for government use of AI depends less on potential cost savings or efficiency gains than on whether systems are accurate, explainable, and subject to meaningful human oversight. Concerns about fairness and privacy become especially salient when AI is used in criminal justice decisions, such as sentencing, rather than more routine government services like benefits administration.¹⁰ Notably, these priorities were remarkably consistent across party lines, a bipartisan pattern that mirrors this survey’s findings about relative cross-party agreement on questions of AI’s use in government.

New questions in the 2026 survey asked Californians which government uses of AI they find most acceptable and what concerns them most. Asked which uses of AI in the federal government they would support most, Californians most often point to national security (18 percent) and government surveillance (18 percent), though a plurality (40 percent) says they “don’t support AI use by the federal government for any of the listed purposes.” Republicans are considerably more open than Democrats to AI use for immigration and border enforcement (15 percent versus 3 percent). When asked what concerns them most about federal or local government use of AI, Californians most often cite government access to and use of private citizens’ data (22 percent for both federal and local government), well ahead of concerns about bias, transparency, or accountability.

When asked what concerns them most about federal or local government use of AI, Californians most often cite government access to and use of private citizens’ data.

Californians hold mixed, and largely unchanged, views about AI’s potential to improve public safety outcomes. Today, about as many Californians are optimistic as in 2025 that AI-assisted tools can improve prediction of and response to wildfires, floods, and other extreme events (47 percent in 2026, 46 percent in 2025). Optimism about AI’s potential to reduce and combat crime is similar (36 percent in 2026, 35 percent in 2025). This year, Republicans are considerably more optimistic than Democrats about AI’s crime-fighting potential (51 percent versus 36 percent), while views on wildfire and disaster response show a similar, if smaller, partisan gap.

Broader agreement that AI safety requires a shared, multistakeholder effort has softened somewhat since 2025, echoing a pattern seen elsewhere in this year's survey of "very concerned" or "strongly agree" responses shifting toward more moderate or uncertain ones. The share who "strongly" or "somewhat" agree that making AI safe and trustworthy for public use needs to be a nationwide effort across civil society, industry, government, and universities fell from 67 percent in 2025 to 63 percent in 2026. Agreement that organizations investing in AI should help establish guidelines for its safe and responsible use fell from 72 percent to 67 percent, and agreement that government should require AI companies to test advanced systems for safety fell from 77 percent to 73 percent. Even with this softening, majorities across parties continue to agree with each of these statements; Democrats are considerably more likely than Republicans or independents to strongly agree that AI companies should be required to test their systems for safety (65 percent versus 48 percent and 41 percent, respectively).

Education

Support for a state policy to advance AI literacy in California schools has held roughly steady, if slightly softer than last year: 43 percent of Californians support such a policy in 2026, compared with 46 percent in 2025, while opposition is little changed (27 percent in 2026, 25 percent in 2025). This modest decline appears across parties, with support among Democrats slipping from 54 percent to 52 percent and support among Republicans falling from 48 percent to 45 percent. As in 2025, a substantial gender gap persists, with men (49 percent) far more supportive than women (37 percent). Regional patterns have shifted somewhat. In 2025, support for AI literacy instruction was highest in the San Francisco Bay Area (56 percent); in 2026, Los Angeles residents report the highest support (48 percent), with the Bay Area close behind (45 percent) and the Inland Empire lowest (32 percent).

New to the 2026 survey, a question asked Californians whether the emergence of AI tools will affect the integrity and value of academic degrees for the better or worse. Californians are considerably more likely to say AI will affect the integrity and value of academic degrees for the worse (48 percent) than for the better (19 percent), with a third (33 percent) unsure. Younger Californians are notably more optimistic than older ones: About a quarter of those ages eighteen to twenty-nine (24 percent) and thirty to forty-four (25 percent) say AI will affect degree integrity for the better, compared with 13 percent each of those ages forty-five to sixty-four and sixty-five and older. Republicans (30 percent) are also considerably more likely than Democrats (20 percent) or independents (11 percent) to see a positive effect. Because this question is new to the survey, there is no 2025 baseline for comparison.

Public Option

Californians remain closely divided over whether the state should build a public AI option, with overall views largely unchanged from 2025 but a notable narrowing of the partisan gap. Forty-four percent of Californians support the State of California creating a publicly available, shared virtual platform, or public cloud computing cluster, to advance AI research and development, compared with 45 percent in 2025. Opposition, however, fell more substantially, from 27 percent in 2025 to 19 percent in 2026, with the share saying they “don’t know” rising from 28 percent to 37 percent. This shift is concentrated among Republicans: Republican support for the initiative rose from 38 percent in 2025 to 44 percent in 2026, while Republican opposition fell from 42 percent to 27 percent over the same period. Democratic support, by contrast, was essentially unchanged (56 percent in 2025, 53 percent in 2026).

Support for state government prioritizing AI safety research and development has softened somewhat. About half of Californians (53 percent) “strongly” or “somewhat” agree that state government should focus more time and funding on AI safety research and development, down from 57 percent in 2025. Democrats remain considerably more likely than Republicans or independents to “strongly” agree that state government should prioritize AI safety funding (24 percent versus 17 percent and 16 percent, respectively).

Looking Forward

The 2026 Carnegie California AI Survey speaks to the importance of the next two years in AI policymaking and politics. State residents are using the technology more and recognize its importance to economic competitiveness, their work, their government, and their rights. They are therefore simultaneously eager for workforce training and anxious about harms.

AI may be a transformational technology, but the policy solutions Californians are mostly supportive of—enabling workers to adapt to its impacts and advancing safety—are hardly radical and are often bipartisan. This is a moment of opportunity in AI policymaking and politics. We will see where we are in 2027.

Advisers

Stephen Caines Esq., former chief innovation officer and budget director, City of San José

Neil Chilson, Head of AI Policy, Abundance Institute

Elena Cryst, Director of Policy and Society, Stanford Institute for Human-Centered Artificial Intelligence

Jasmine Dhaliwal, Policy Team, METR

Samantha Gordon, Chief Advocacy Officer, TechEquity

Courtney Russell, Chief Executive Officer, Invest Long Beach

Scott Singer, Fellow, Technology and International Affairs and Co-Director, China AI Initiative, Carnegie Endowment for International Peace

Micah Weinberg, Nonresident Scholar, Carnegie California, Carnegie Endowment for International Peace

Acknowledgments

We would like to thank Alexander Marsolais, Kara Eytcheson, and Sarah Higdon of YouGov for conducting the survey with immense care and rigor. We'd also like to thank Anamitra Deb, Anmol Chaddha, and in particular David Nahmias for their collaboration. A special thank you also goes to Gustavo Romero, former junior fellow of the Global Order and Institutions Program, for his Spanish translation. And for their amazing work on editing this paper and creating the incredible graphics to go along with it, we thank Carnegie's communications team: Alana Brase, managing editor; Haley Clasen, senior editor; Helena Jordheim, assistant editor; Amy Mellon, senior graphic designer; and Jocelyn Soly, creative director. This work is supported by a grant from Omidyar Network.

Notes

- 1 Andy Kemp, “Rising AI Adoption Spurs Workforce Changes,” Gallup, April 12, 2026, <https://www.gallup.com/workplace/704225/rising-adoption-spurs-workforce-changes.aspx>; Annaleise Azevedo Lohr and Alec Tyson, “Google / Ipsos Multi-Country AI Survey 2026,” Ipsos, January 15, 2026, <https://www.ipsos.com/en-us/google-ipsos-multi-country-ai-survey-2026>; Katherine Long, “Poll: Trump Voters Skeptical of Artificial Intelligence’s Impact on Jobs,” Politico, May 2, 2026, <https://www.politico.com/news/2026/05/02/poll-trump-voters-skeptical-artificial-intelligence-jobs-00902754>; Luona Lin, “About 1 in 5 U.S. Workers Now Use AI in Their Job, Up Since Last Year,” Pew Research Center, October 6, 2025, <https://www.pewresearch.org/short-reads/2025/10/06/about-1-in-5-us-workers-now-use-ai-in-their-job-up-since-last-year/>; and Alexander Rossell Hayes, “More Expect AI to Affect the Economy Negatively than Positively, but Americans are Split on How AI Will Impact their Own Lives,” YouGov, June 2, 2026, <https://yougov.com/en-us/articles/54888-more-expect-artificial-intelligence-ai-affect-economy-negatively-than-positively-americans-split-how-ai-will-impact-their-lives-may-29-june-1-2026-economist-yougov-poll>.
- 2 “Governor Newsom Signs First-of-Its-Kind Executive Order to Prepare Workers and Businesses for Potential AI Disruption,” Office of Governor Gavin Newsom, May 21, 2026, <https://www.gov.ca.gov/2026/05/21/governor-newsom-signs-first-of-its-kind-executive-order-to-prepare-workers-and-businesses-for-potential-ai-disruption/>.
- 3 “Assembly Bill 2023: Companion Chatbots: Children’s Safety,” California State Legislature, 2026, https://leginfo.ca.gov/faces/billNavClient.xhtml?bill_id=202520260AB2023; and Aaron Mak, “The Lawsuits That Could Give AI Its ‘Big Tobacco’ Moment,” Politico, June 6, 2026, <https://www.politico.com/news/2026/06/06/the-lawsuits-that-could-give-ai-its-big-tobacco-moment-00951764>.
- 4 Dario Amodei, “Where Things Stand With the Department of War,” Anthropic, March 5, 2026, <https://www.anthropic.com/news/where-stand-department-war>; Frontier Red Team, “Assessing Claude Mythos Preview’s Cybersecurity Capabilities,” Anthropic, April 7, 2026, <https://www.anthropic.com/research/mythos-preview>; “Promoting Advanced Artificial Intelligence Innovation and Security,” White House, June 2, 2026, <https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/>; Beatrice Nolan, “Anthropic Releases Its First Mythos-Class Model to the Public,” Fortune, June 9, 2026, <https://fortune.com/2026/06/09/anthropic-releases-its-first-mythos-model-to-the-public/>; and Sophia Cai and Cheyenne Haslett, “Trump Administration Partially Lifts Anthropic AI’s Export Ban,” June 26, 2026, <https://www.politico.com/news/2026/06/26/white-house-makes-peace-with-anthropic-for-now-00965675>.
- 5 Kate Gordon and Noah Gordon, “Are Data Centers the Villains in the Battle Over Electricity?,” Carnegie Endowment for International Peace, June 3, 2026, <https://carnegieendowment.org/emissary/2026/06/data-centers-electricity-energy-climate-villains>.
- 6 John Gramlich, Brian Kennedy, Colleen McClain, and Galen Stocking, “How Americans View Data Centers’ Impact in Key Areas, from the Environment to Jobs,” Pew Research Center, March 12, 2026, <https://www.pewresearch.org/short-reads/2026/03/12/how-americans-view-data-centers-impact-in-key-areas-from-the-environment-to-jobs/>.
- 7 Amanda Yeo, “Anthropic to Provide California Government With AI Tools,” Yahoo News, June 30, 2026, <https://www.yahoo.com/news/politics/articles/anthropic-california-government-ai-tools-073501540.html>; and <https://statescoop.com/california-poppy-state-government-ai-assistant/>.
- 8 Sophia Fox-Sowell, “Only California State Employees Can Use Poppy, a New AI Assistant Designed for Security,” StateScoop, February 20, 2026, <https://www.nytimes.com/2026/06/14/us/ai-deepfake-hany-farid.html>.
- 9 The surveys used different response scales. Pew asked respondents whether they were “extremely,” “very,” “somewhat,” “not too,” or “not at all” confident, while the California survey asked whether respondents were “very,” “somewhat,” or “not confident.” Brian Kennedy, Eileen Yam, Emma Kikuchi, Isabelle Pula, and Javier Fuentes, “How Americans View AI and Its Impact on Human Abilities, Society,” Pew Research Center, September 17, 2025, <https://www.pewresearch.org/science/2025/09/17/how-americans-view-ai-and-its-impact-on-people-and-society/>.
- 10 Guimin Zheng, “When Governments Use AI, Citizens Want It to Be Accurate, Accountable and Fair, Especially When Police and the Courts Are Involved,” *LSE United States Politics and Policy* (blog), July 3, 2026, <https://blogs.lse.ac.uk/usappblog/2026/07/03/government-ai-accuracy-accountability-fairness/>.



CARNEGIE
ENDOWMENT FOR
INTERNATIONAL PEACE

1779 Massachusetts Avenue NW
Washington, DC 20036

CarnegieEndowment.org