THE POLITICAL DYNAMICS OF AI IN CALIFORNIA

How Public Opinion and Corporate Lobbying are Shaping the Debate About Regulating AI

In partnership with Diffusion.Au

AI has become a tier-one economic and safety concern for California voters, and a potential fault line in the 2026 elections. This memo presents original research on how Californians think about AI, what they want from their leaders, and the forces shaping the debate.

The Electorate: While Californians express qualified optimism about technology itself, they are deeply anxious about AI and, in particular, the industry that is building it. They see AI as a contributor to economic instability through job displacement and consumer manipulation, and as a longer-term structural risk to employment and opportunity. They are furious at the tech elite, whom they see as stealthily rewriting society's rules.

The Coalition Math: Demand for strong AI regulation is broadly bipartisan (81% Democrats, 73% Republicans) and holds across regions (72%+ in six of seven regions; Sacramento is the exception at 63%). However, notable divides exist along gender, age, and education lines. Women, especially older and non-college-educated women, are more concerned and represent a key intensity bloc.

What Voters Are Looking For: Californians want leaders who will act as an independent check on the technology industry, ensuring AI serves working families, not just billionaires in Silicon Valley. Credibility on AI comes from

linking it to everyday consumer and labor concerns and offering tangible policy actions, not abstract principles.

The Trust Gap: Voters do not trust the federal government on AI; they are wary of the state legislature; and they actively dislike tech executives. They're looking for leadership that is visibly independent and willing to enforce rules that protect jobs, privacy, and freedoms without deferring to industry.

KEY NUMBERS

71%
want strong laws to regulate AI

54%
concerned about AI's future direction

48%
think AI is moving too fast

60%
believe only ultra-elites will benefit

9%
more trust California over feds on AI

77%
favor AI privacy protections

Where is the public on AI?

Everyday Californians see AI more as a threat than an opportunity

TechEquity, in partnership with Diffusion.Au, conducted a representative statewide survey of Californians and a secondary wave of focus groups to better understand Californians' attitudes toward AI. Below are key lessons for the state's leaders.

"Privacy" means protection from harm

When voters say privacy, they mean deepfakes, predatory pricing, and discrimination

When describing their concerns about AI, voters often refer to "privacy violations" or wanting enhanced "privacy rights." In focus groups, we found that "privacy" functions as policy shorthand. When we asked people to unpack these violations, they consistently translated them into specific material harms: companies using information about them for predatory pricing; personalized deepfakes; workplace surveillance; bias and discrimination in housing, credit, and job recruitment tools; and other endpoints in their lives affected by AI and automated decision-making systems.

The data:

- Support for protecting personal privacy rights polls at nearly 90%.
- Overall, 64% are concerned about AI enabling deepfakes, including 40% who are extremely concerned.
    - This is especially acute among women (46% are extremely concerned vs. 34% of men) and among women over 50 (49% are extremely concerned).
- 60% are concerned that AI will enable the spread of disinformation.

What the data suggests: Messaging that connects AI policy to personal security (rather than abstract or conceptual data rights) resonates most strongly with all voters, especially suburban women and parents, demographics that index highest on fears about stolen likenesses and AI-enabled scams targeting family members.

AI is now a cost-of-living and job-security issue

Kitchen-table economics drives voter behavior more than sci-fi catastrophes

While the catastrophic scenarios (bio-weapons, nuclear control) generate fear, the voting issue is economic displacement.

The data: There is deep concern about replacing low-paying jobs (51% concerned) and reducing wages (54% concerned).

What the data suggests:

- Alignment with labor is a credibility signal on AI. Voters trust unions (+42 net favorability) far more than tech executives (-2 net) or figures like Elon Musk (-30 net).
- The research also reveals a messaging asymmetry. Promises to "retrain workers for the AI future" tend to backfire, as voters hear them as confirmation that job losses are on the way. Messages that emphasise corporate accountability ("companies should pay a price if they automate without contributing to the community") perform significantly better. This pattern is especially pronounced among Latino voters, of whom 59% are concerned AI will replace low-paying jobs, and 59% believe it will reduce people's wages.

“Built for them, used on us” is the dominant mood

Bipartisan class resentment about who controls AI and who benefits is greater than any concern about the technology itself

One of the most politically potent findings is the class resentment over who benefits from AI.

The data:

- 59% of Californian voters say AI will mostly benefit wealthy households and corporations.
- 54% of voters think AI will increase inequality and shift wealth upwards.
- Elites associated with the industry, like Elon Musk, have a net favorability of -30 (26% Fav / 56% Unfav). Musk, in particular, is underwater with every single demographic group except Republicans.
- In focus groups, even wealthy voters shared the view, articulating the idea that *“if you’re not one of the few people who own the corporation, you’re going to get screwed.”*

What the data suggests:

- Class resentment around AI is bipartisan. The sentiment is as strong among Inland Empire Republicans as among LA progressives. A “working families vs. Silicon Valley elite” frame has crossover appeal because it speaks to a shared perception that the system is rigged.
- On language, split testing shows that “Silicon Valley executives” elicits more negative sentiment than “tech industry executives.” The underlying driver isn’t identifiable from this test, but the difference doesn’t travel. In polling outside California, concern often focuses on the “tech industry”, which reads like a sector; inside California, “Silicon Valley” reads like a specific power centre with a locally charged reputation.

Voters want leaders who aren’t “bought”

Voters believe legislators can protect them from AI-related harms, but worry they won’t act independently under industry pressure

Voters doubt both the state and federal legislature’s ability to regulate AI; however, this is not because they think lawmakers lack technical expertise, as

industry sometimes claims. Instead, they fear politicians are "bought out" or compromised by their relationships with industry. Voters are asking: Can I trust you to act in our interest?

The data:

- Only 35% trust the CA government to control AI.
- Focus groups abounded with quotes like, *"Are their campaigns funded by these companies?", "You can't serve two masters", "[I don't trust them], they'll just do whatever the companies want."*
- The "speed" mandate: 48% say AI development is moving "too fast." Only 3% say "too slow." Californians want to see leaders tap the brakes rather than hit the accelerator.

What the data suggests:

- **A note of caution:** Proximity to the tech sector carries reputational risk, and association with industry inherits its negatives.
- **The polling points to an unoccupied position:** Voters are looking for leaders who can credibly demonstrate independence from tech money and who propose concrete measures to ensure regulatory independence. No candidate is currently filling this space.
- When asked what would build confidence in AI governance, voters prioritize:
  - Expert consultation over lobbyist access: Voters want legislators to develop their response by "speak[ing] with experts, not lobbyists."
  - Transparency on conflicts: with the most commonly suggested action being bans on policymakers holding AI-affected stocks
  - Enforcement with teeth: Voters are skeptical of vague principles and are looking for "meaningful penalties" for harmful behavior.
  - Accountability mechanisms: independent and publicly disclosed tracking and reporting on AI's environmental and economic impacts

How AI attitudes split across California

Young men and older women see different futures

Young educated men see opportunity; older non-college women see danger

When reviewing the data to identify correlations that move together, the most significant indicator is gender.

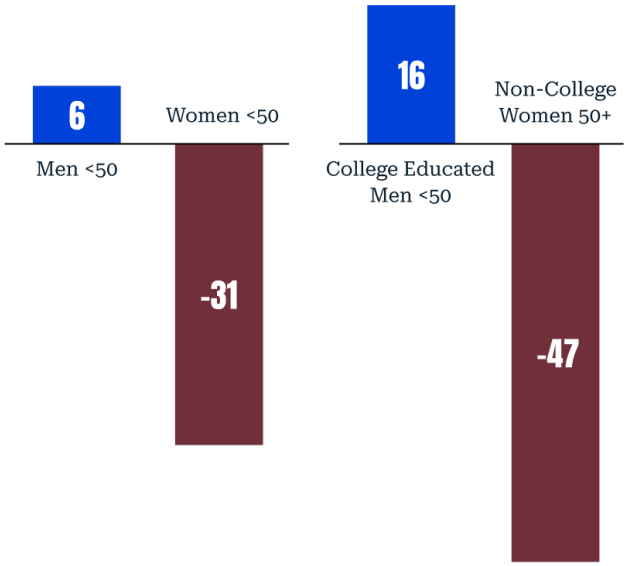
The data:

- Among voters under 50, men are net 44 points more positive than women about AI's future advances.
- Stratifying further by education and age widens the gap massively, with college-educated men under 50 being +24 net excited, while non-college women aged 50+ are –61 net concerned, an 85-point spread.

What the data suggests:

- Taken together, the gender, age, and education split signals two fundamentally different realities. Younger, college-educated men remain wary of elite influence but, on balance, see AI as an exciting career opportunity. For older, non-college-educated women, it reads as an uncontrolled threat to themselves, their families, and our broader society.
- **The implication:** These groups also respond to different frames—younger adults engage with economic opportunity messaging, while older voters, especially women, respond more to reassurance and protection. Older women are the most concerned about AI in our data; that combination of high concern and high turnout makes them a powerful constituency shaping the political dynamics on this issue. As a result, the community contains two parallel conversations.

Difference in Sentiment by Age, Gender, and Education



Those who've faced bias see it coming again

Black and Latino Californians recognize the patterns of algorithmic discrimination

Although gender, age, and education are by far the most significant demographic factors, there are also distinctive racial concerns related to AI technologies. In simple terms, those who have experienced structural discrimination are more sensitive to the way these tools are developed and implemented.

The data:

- **Black voters:** Are cautious, but more tolerant than other people of color. 29% of Black voters believe AI will increase discrimination (vs 26% overall). For them, it's a matter of principle and process rather than a personal crisis, with some of the strongest support for civil rights protections in AI: 82% favor new regulation (vs 75% overall). They're also the most confident in the government's ability to control AI, especially at the state level, with 42% saying they trust Californian legislators, compared with 36% overall.
- **Latino voters:** Meanwhile, Latino voters express a greater sense of urgency and economic risk. They are most concerned about the impact on jobs, with 36% convinced AI will erode lower-paying jobs (vs 29% overall), 47% believing it will reduce workers' rights (vs 43% overall) and 29% convinced it will reduce people's wages (vs 23% overall) and that AI will make discrimination worse (32% vs 29% overall). They are also the most convinced that these tools will exacerbate economic inequality, with 63% believing that only elites will benefit from these tools (vs 60% overall), and are therefore strongly supportive of job protection policies.
- **Asian American voters:** AAPI voters are the most AI-familiar demographic (70% have used LLMs) and many work in the tech sector (17% compared to a median of 7% for other non-white groups). They are favorable towards regulation (43% favor new regulations "a great deal", 75% favor new regulations), but not as acutely as other non-white groups (for example, 49% of Latino voters favor new regulations "a great deal" while black communities are 82% favorable for new regulations as a whole).

What the data suggests:

- The variation here suggests that race is a weaker predictor of AI attitudes than economic position or industry proximity.

The Bay Area knows tech; the Inland Empire fears it

Proximity breeds skepticism of industry; distance breeds anxiety about displacement

California is not a monolith. Regional variation in AI attitudes reflects both economic exposure and lived experience with the tech industry.

The Data:

- **Bay Area:** Highest AI literacy (77% familiar), some of the heaviest LLM users (67%), but sharpest home-turf skepticism, with only 23% favorable toward Silicon Valley executives and 75% support strong laws (while 25% do not trust industry to self-regulate). This is lived experience and local grievance from people who've watched tech's promises collide with its reality.
- **LA, Inland Empire, Central Valley:** 40%+ are extremely concerned about job displacement. AI is seen less as an opportunity and more as a direct threat to livelihoods for everyone, from Hollywood and the creative industries through to service work and logistics. They do not trust the tools or the people who are building them.
- **Sacramento:** The outlier at 63% support for strong regulation (vs 72%+ elsewhere).

What the data suggests:

- In the **Bay Area**, messaging that acknowledges the value of innovation while demanding accountability resonates: "guardrails so innovation can thrive without the public paying the costs again." Leaders who can credibly criticize tech excess while championing innovation may find an unusual coalition here.
- In **LA, the Inland Empire, and the Central Valley**, protection framing lands: a firewall between livelihoods and algorithms. "Clear rules, real penalties, worker-first safeguards." These audiences want tangible action and enforcement with teeth—principles without consequ
- In **Sacramento**, lower support for "strong regulation" may reflect cynicism about legislative politics and doubts that new laws change outcomes. Competence framing should land: "clear standards, real oversight, and enforcement that actually sticks."

The landscape: how industry shapes the debate

The previous sections of this memo describe what Californians want from AI governance. But public preferences don't translate automatically into policy. Between voter sentiment and legislative outcomes sits a well-resourced effort to shape what regulation is seen as reasonable and what becomes politically viable. Understanding this landscape is essential context for anyone seeking to act on what the polling reveals.

The technology industry is running a sustained, well-resourced campaign to control the terms of the AI debate in California—how the technology is understood, what regulation is seen as "reasonable," and what legislation becomes viable. It operates through legislative lobbying, narrative framing, and shaping the mechanics of the policy-making process itself. If you want AI policy to reflect Californian interests, you need to understand how this machine works.

The scale of investment

Industry spending to influence legislation and elections in California has skyrocketed since 2023

A small number of firms are underwriting a sustained and large-scale AI lobbying push in California.

Between Q1 2023 and Q3 2025, major technology companies and trade associations reported $57.6 million in lobbying expenditures in California's official lobbying filings. That spending is heavily concentrated: Google ($14.3m) and the CA Chamber of Commerce ($12.5m) together account for nearly half of the total, and the top four actors (Google, California Chamber of Commerce (CalChamber), Computer & Communications Industry Association (CCIA), Waymo) account for 71%.

The volume here matters less than the pattern. Spending spiked dramatically in Q3 2024—the quarter California's proposed Journalism Preservation Act (AB 886) and Frontier AI Safety Bill (SB 1047) were being contested—and remained elevated through 2025, more than double the early-2023 baseline. Google alone spent $10.7 million in that single quarter, nearly 90 times its spending a year earlier, and more than it had spent in the previous 20 years combined.

The most striking feature, however, is how spending flows through intermediaries. Google paid $7 million to CCIA in Q3 2024—95% of CCIA's California lobbying budget Meta paid $3.1 million to CalChamber in Q3 2025. The coordination between major AI companies isn't visible in their direct spending trajectories; it's visible in their parallel funding of trade association infrastructure that lobbies on their behalf.

The disclosed figures capture only part of the investment. They exclude campaign contributions, non-lobbying PR and narrative work, and other spending structured to sit outside reporting requirements. In addition to direct lobbying, companies invest in large internal corporate policy teams and create lobbying subsidiaries that shield expenditures from public view.

Industry is also moving into direct electoral intervention, modelling its efforts on Fairshake, the crypto-backed super PAC that spent over $10 million to defeat Katie Porter's California Senate bid and more than $40 million to unseat Ohio Sen. Sherrod Brown, the Senate Banking Committee chair, in 2024.

Leading the Future, a new super PAC network launched in August 2025 with over $125 million in funding, is explicitly built on this template: run by the same political operatives who led Fairshake, and backed by Andreessen Horowitz, OpenAI's Greg Brockman, and other major AI investors. California is one of its four initial target states. Meta has separately launched two super PACs—one California-specific, one national—pledging "tens of millions," and has since put at least $65 million behind the effort to back candidates, regardless of party, who will resist new restrictions on the industry.

The scale of this electoral spending dwarfs anything available to consumer advocates, labor organizations, or civil society groups working on AI governance in California.

The policy process: bad faith at scale

The AI industry's playbook is to bog down legislative negotiations in order to drag out lawmaking and essentially institutionalize the permissionless status quo

Interviews conducted with legislators, advocates, and former industry insiders at the end of the 2024 legislative session describe a consistent pattern in how tech industry lobbyists engage: collaborative in tone on the surface, but obstructive in practice. The method is less about winning a single argument than about controlling the terms, tempo, and venues of the debate.

It begins with **scale and message discipline.** A large number of lobbyists deliver consistent talking points across multiple forums: committee testimony, op-eds, stakeholder meetings, and private conversations with legislators. The same arguments recur in slightly different forms, creating the impression of broad agreement where, in fact, there is coordinated positioning by a small number of decision-makers. Further, lobbyists have also begun strategically deploying sympathetic non-tech household-name companies in their efforts, using familiar brands to amplify pressure on legislators even when those companies are not directly affected by particular bills.

It is then escalated with **high-profile pressure.** Senior executives, such as Sam Altman and Elon Musk, engage directly with the most senior legislators, exerting private, discreet pressure for specific outcomes. The dynamic is double-edged: legislators report feeling simultaneously flattered that someone so prominent is calling them directly and threatened by the intensity of the lobbying. The implied message is that crossing these figures carries consequences.

That message discipline is reinforced by **strategic timing**. Companies and trade associations often participate constructively throughout much of the process, then introduce amendments or "concerns" at the last possible moment, when scrutiny is at its hardest and legislative staff capacity is stretched. In interviews, legislative insiders describe lobbyists who "pretend to play ball," stating they have only a few issues with the language, only to backtrack right at the last minute, when there is no time for staff or legislators to address the issues they claim are non-negotiable. The effect is predictable: fundamental objections arrive with no time to resolve them, and legislation stalls or collapses without industry having to appear to oppose reform outright.

The arguments themselves are often designed to be unfalsifiable. The most common is the **"too broad / too specific" contradiction**: proposals are criticized as simultaneously too broad (inviting overreach) and too specific (unable to anticipate future developments). This rhetorical pincer creates paralysis. Any concrete measure can be attacked as either over- or under-inclusive, shifting the burden onto advocates to produce "perfect" legislation while industry bears no equivalent standard.

Alongside this sits a moral framing that **separates "good actors" from "bad actors."** Companies position themselves as responsible partners willing to address harms, while locating the real problem elsewhere: rogue individuals, foreign adversaries, or small operators without safety teams. In the SB 1047 debate, opponents argued that the bill would burden responsible developers

while doing nothing about malicious actors, reframing the question from "should powerful AI systems carry safety requirements?" to "why punish the good guys?" The result is a quiet redefinition of the problem: regulation becomes a conversation about policing outliers and individual misuse rather than about the design choices and business models of the firms shaping the market.

When a legislator hesitates, **companies often play the economic-anxiety card.** Lobbyists argue that proposed legislation will cost jobs, reduce revenue, and drive companies out of California. Given the state's budget pressures, this lands hard—no legislator wants to be blamed for economic decline. The argument that strong regulation risks California "falling behind" taps into real anxieties about the innovation economy, even when the underlying claims deserve scrutiny. The threat of relocating out of state is the sword of Damocles invoked often by industry lobbyists and leaders, but California's concentration of talent, venture capital, research institutions, and existing infrastructure makes large-scale departure logistically and economically implausible for most firms. The ecosystem that makes these companies valuable is not easily replicated elsewhere—a fact that rarely features in the lobbying conversation.

Finally, there is an endgame that reaches beyond Sacramento: **federal preemption**. This part of the playbook isn't new: it mirrors Big Oil and Big Tobacco's decades-long resistance to state-level regulation, following a now well-established three-step sequence. First, industry groups promote multiple, conflicting state bills, often via trade associations that provide cover for individual companies, generating fragmentation. Second, they denounce the resulting "patchwork" of laws as unworkable, emphasizing compliance costs (especially for small businesses) even though industry itself created the complexity. Third, they lobby Congress for weak, carve-out-filled federal legislation that, crucially, preempts states such as California from adopting stronger standards. The manufactured "patchwork" becomes the justification for locking in minimal protections nationally.

The narrative environment: shaping how people think

Industry not only lobbies legislators; it also deliberately shapes how AI is conceptualized. With a topic as new and abstract as AI, research consistently shows that the metaphors people encounter become the mental models they use to decide what is possible, what is risky, and who should be in charge. Much of

this framing work is achieved through a disciplined, intentional set of narrative structures that normalize industry-friendly assumptions.

A 2024 Diffusion.Au/University of Cambridge research project analyzed more than 300,000 words of AI policy discourse across the EU, UK, and US. It finds that industry representatives draw on a narrower, more consistent set of frames than academics, journalists, or civil-society voices; a consistency that suggests intentional strategy rather than coincidence.

One frame dominates industry discourse: the "Journey/Race." Technology companies repeatedly describe AI development through racing metaphors: a competition with a prize, opponents to outpace, and milestones along the way. The language is familiar: *"companies race to be the first"; "we're leading the pack"; "we're on the fast track."*

Industry's messaging is powerful

The AI Industry's disciplined use of the "journey/race" narrative frame helps create the logic for their preferred policy outcomes

The race metaphor does more than add color. It supplies a ready-made story about what is happening, what matters, and what counts as responsible behavior; and it does so without having to argue for any of it explicitly.

- **It turns a political choice into a natural event.** A "race" is not something you vote on; it's something that is already underway. The frame quietly shifts AI development from a decision we could shape to a process we can only keep up with.
- **It manufactures urgency by making delay feel like defeat.** If the defining risk is "falling behind," then the default answer to uncertainty is speed. Caution becomes a luxury; deliberation becomes dithering; regulation becomes friction. The frame doesn't need to specify what we lose; the emotional logic of competition does the work.
- **It smuggles in a presumed prize.** Races are run for something worth having. The benefits of AI can remain vague because the story's structure treats the destination as self-evidently desirable and skepticism as negativity rather than as discernment.
- **It narrows the meaning of "good governance."** In a race, rules exist to stop cheating and legitimate the winner, not to ask whether the race is safe, who set the route, or what happens to bystanders along the track. Regulation is reframed as officiating: make it "fair," not slow it down, re-route it, or pause it.

- **It defines who gets agency, and who doesn't.** Races have competitors and spectators. The public is positioned as an audience that should back "our" side, not as a co-designer of the system or a setter of constraints. Those who question the premise of the race can then be cast as naïve, anti-progress, or simply "not on the team."
- **It makes harms look like collateral, not design.** Racing stories normalize trade-offs: crashes, mistakes, and rule-bending become predictable side-effects of going fast. That subtly lowers the moral temperature around preventable harms and pushes accountability into the background.

In other words, the race frame does not simply motivate speed; it crowds out accountability. And you can see that in what industry discourse leaves unsaid.

If instead, one uses a "building" metaphor, the conversation shifts to blueprints, standards, inspectors, and structural failure, and to who is responsible when corners are cut. Call it a machine or critical infrastructure, and it naturally raises safety regimes, oversight, insurance, warranties, and liability, because most people cannot independently verify the system's integrity. Those frames make governance feel ordinary and necessary. Their absence from the public narrative the tech industry is driving is not incidental: it keeps the debate oriented around competition, not responsibility.

The invisible actor problem

The personification of AI shields human actors from accountability

Beyond specific metaphors, there is a subtler pattern in how AI is discussed: the technology is routinely written about as if it has personhood and agency. Headlines announce that *"AI is coming for your job,"* that *"AI will revolutionize healthcare,"* and that *"AI is changing everything."* In this construction, AI does things—it acts, it decides, it transforms.

This linguistic habit is not neutral. It is strategically cultivated by the people who build and sell these systems, and it has consequences. When AI is the subject of the sentence, the people making decisions disappear: the executives choosing what to build, the engineers deciding how to train systems, the companies determining how tools are deployed, the investors shaping incentives. The chain of reasoning from decision to outcome becomes invisible, and if something isn't people-caused, the public won't believe that it can be solved by people.

This is why AI is such a convenient abstraction for industry. It is not "AI" that decides to automate customer service jobs; it is a company choosing to deploy a

system in a particular way. It is not "AI" that discriminates in hiring; it is a vendor that built a flawed tool and an employer that purchased it without adequate testing. When actors are named, accountability becomes possible. When they are hidden behind the abstraction of "AI," the technology appears autonomous and inevitable, and the political question of who decides never gets asked.

Plotting a path forward

While the AI industry's spending and narrative dominance make safeguards harder to achieve, there is a path available

Industry's investment is large, and its tactics are sophisticated. That can feel overwhelming. But the fundamentals still favor determined leadership for three reasons.

- **Public opinion is a headwind, not a tailwind, for industry.** Californians are skeptical of technology companies, their executives, and their motives. Trust in self-regulation is thin; support for strong oversight is broad and bipartisan. Industry is spending to change that reality, not to ride it.
- **Its core narrative is more fragile than it looks.** The "race" frame only works if people accept the contest and desire the prize. But 59% of Californians believe the benefits of AI will mainly flow to the wealthy and corporations. Once voters see AI less as a prize and more as a tool being wielded against them, the race narrative's promise collapses. Reframing AI as something to be built, with standards, guardrails, and liability when it malfunctions, changes what "responsible" action looks like.
- **And the procedural games depend on obscurity.** The industry's complaint about "patchwork" legislation at the state level loses force when it is understood as a manufactured route to federal preemption. The "too broad / too specific" maneuver weakens when it is named for what it is. Visibility itself is a countermeasure.

This briefing has set out what Californians want on AI, and the forces organized to deny it. The path forward is not about matching industry dollar-for-dollar. It is about clarifying the stakes, naming the actors, and offering an AI agenda that serves working families rather than Silicon Valley shareholders.

About the survey

This research combined (1) a representative survey of n=1,400 California adults (n=1,000 statewide plus oversamples of Black n=100, AAPI n=100, and tech-savvy adults n=200), fielded 29 April–8 May 2025 (±3.1% MoE), and post-stratified by gender, age, race, region, and education. All findings in this memo are based on the 1,248 respondents who are registered voters; (2) two online focus groups (10 participants each) on 10 July 2025, screened to reflect key survey-identified segments; and (3) in-survey real-time dial testing, where respondents rated randomly ordered messages on a 0–100 slider (starting at 50) updated at least every two seconds, producing weighted, smoothed approval traces and segment lifts tested at 95% confidence.

Full toplines and cross tabs can be accessed at diffusion.au/aicompass.

Acknowledgements

This memo was informed in part by research conducted by the CITRIS Policy Lab at UC Berkeley. The authors gratefully acknowledge the Lab's contributions to the analysis and findings presented here.