April 2026

CALIFORNIA'S AI FUTURE

A Leadership Agenda for a People-First Economy



TECH EQUITY

Contents

Introduction	**4**
The Political Dynamics of AI in California: How Public Opinion and Corporate Lobbying are Shaping the Debate About Regulating AI	**6**
Where is the public on AI?	7
How AI attitudes split across California	10
Those who've faced bias see it coming again	12
The landscape: how industry shapes the debate	14
The narrative environment: shaping how people think	17
AI & California's Labor Market: Reckoning with AI's Impact on Workers	**22**
Harms and costs: What's at stake?	23
AI tools are being used by employers to demean and exploit workers	25
The importance of prioritizing worker voice	30
Solutions to support California's workers and economy	31
The future is ours to build	34
Appendix: Protecting Workers Amid Potential AI-Driven Labor Displacement	**36**
A Policy Agenda for AI-Driven Labor Displacement	37
Conclusion	41
AI & Affordability in California: Land of Opportunity or Hardship?	**42**
Harms and costs: What's at stake?	43
How California can lead	47
Building a California for all	51
Beyond Big Tech: An Alternative Vision for California's Tech Economy	**52**
Harms and costs: What's at stake?	54
The opportunity ahead	62

INTRODUCTION

The next generation of California leaders will take office at the dawn of a new technological age. AI is transforming society and, whether or not it achieves the promise its boosters claim, will have profound consequences for us all. California, home to most of the biggest AI developers, is where the future of this technology will be decided. And this state's leaders are uniquely positioned to determine whether it will be a force for exploitation and extraction or one that enhances human dignity and advances shared prosperity.

This moment requires courageous leadership to chart our course. TechEquity's hope with this project is to provide California's leaders and the public with a guide.

The outlook is murky. Despite both the hype and the skepticism, no one knows exactly where the technology is headed. It could be as epochal as the advent of fire or as "normal" (but still revolutionary) as the birth of the internet. Mapping out a policy agenda to influence the development of AI towards its most beneficial outcomes will require new ideas and different ways of thinking.

On top of that, the political dynamics are unique. AI does not map neatly onto existing political mental models. Support for regulation cuts across a wide swath of the public who do not normally stand together: Republicans and Democrats, working and upper middle classes, and large proportions of all racial and gender demographic groups. Overall, 70% of Californians say they want state leaders to act to create strong guardrails for AI.

This fluidity presents California's leadership with immense challenges, but also opportunities. They will be called upon to make sense of this moment and chart a path toward future shared prosperity, while decision-makers in Washington remain committed to an irresponsibly unregulated approach. To help inform these leaders and the public as we face these awesome challenges, TechEquity has prepared a set of memos that together paint a picture of what it would mean for AI to work for humans rather than against us. These memos represent a set of policy solutions that can form the basis for a people-first AI future. This package includes:

- Original survey and focus group research that shows how Californians think about AI and what they want from their leaders.
- A look inside the Big Tech influence-peddling playbook and detailed insights into how narrative and framing are shaping the AI debate.
- An alternative vision for an AI future that prioritizes competition and fairness, invests in a publicly funded tech ecosystem, and reforms capital markets to catalyze entrepreneurship.
- A study of AI's contribution to the affordability crisis facing California families, including algorithmic rent-fixing, data centers' effect on energy costs, and surveillance pricing's ability to squeeze every last dollar out of consumers—and policy solutions on each front.
- An examination of AI's impact on workers and the labor market, from mass job displacement to AI-enabled exploitation and surveillance in the workplace to global supply chain issues that echo the sweatshop labor scandals of the past, together with ideas for how to build a pro-worker future.
- A deep dive into the critical issue of job displacement, including how to separate truth from fiction on AI-driven job loss predictions and what California leadership can do to prepare the state and its working families for what is ahead.

These memos focus on economic issues, which is TechEquity's core area of expertise. We recognize, of course, that AI raises urgent concerns beyond economics—including child safety, environmental concerns, energy use, and threats to democratic institutions. Many of the recommendations here are relevant in those domains as well, and we are grateful for the important work that our partners and other organizations are doing to shape the broader policy landscape in these areas.

We stand ready to inform the conversation about the important work of making these policy ideas a reality.

> As Tim Wu notes in *The Age of Extraction:* "During its best and most forward-thinking years, the United States actually sought to balance economic power, whether through the broad distribution of land and other productive assets, the breaking up of monopolies, protecting worker organizing, or building a social safety net… What technology will do for us depends on how we design it and where economic power resides. That's why the real challenge for humanity is to design an economic future that works for more people, rather than assuming that future will just arrive someday."

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.

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.

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

"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 *"f 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.

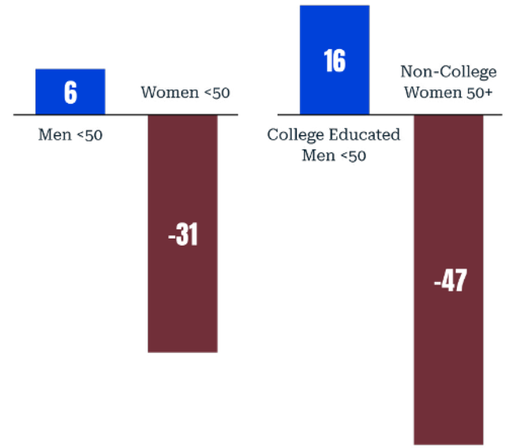
The data:

- 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.

AI & CALIFORNIA'S LABOR MARKET

Reckoning with AI's Impact on Workers

California is a tale of two states: home to both the country's highest concentration of billionaires and the highest poverty rate. This dichotomy has been exacerbated in recent years by the growth of the AI industry. AI-powered wealth consolidation, paired with low union density and reduced worker power overall, is resulting in historic inequality. In this context, productivity gains from AI are increasing profits for shareholders and executives at the expense of worsening working conditions or even loss of jobs for workers.

California's leadership has to confront at least three distinct issues that AI presents for workers.

First, there are real risks that the proliferation of AI systems across the economy will displace workers, driving up the unemployment rate and causing economic chaos. Given the urgency many policymakers feel to address this potential threat, we provide a more detailed framework for solutions in the appendix.

Second, AI tools are being used by employers to demean and exploit workers across industries. These systems set wages artificially low, surveil workers in their private moments, and make hiring and firing decisions that are often biased and without recourse.

Third, the global supply chains that AI companies rely on to build, train, and secure their systems consist of jobs that are often underpaid and underprotected, contributing to a race to the bottom in global labor standards.

These risks are heightened by decades of policy that has undermined worker agency and consolidated power in the hands of employers.

But just as policy choices have brought us to the current moment, different policy choices by California's leaders can put us on a better path. The trajectory of AI's growth and its use by employers is not inevitable. The societal impacts and outcomes of technology are not predetermined. The future is up to our leadership and up to us.

JUST AS POLICY CHOICES HAVE BROUGHT US TO THE CURRENT MOMENT, DIFFERENT POLICY CHOICES BY CALIFORNIA'S LEADERS CAN PUT US ON A BETTER PATH.

Given widespread plans for AI adoption across industries, it has never been more urgent for elected officials to act to prevent harms and ensure the economic potential of AI does not rely on exploitation to succeed. As Liz Shuler, president of the AFL-CIO, stated this February, "The decisions that we are making right now on AI are going to shape the next 50 years. We're demanding our politicians do what 80% of Americans want you to do: pass common-sense guardrails. Protect working people. Bring us to the table now, not after the fact. The labor movement and working people across this country are watching, and we demand real leadership."

This memo articulates the harms workers are experiencing from AI and outlines solutions that elected officials can adopt to improve the lives of working people, while also better serving businesses and consumers.

Harms and costs: What's at stake?

AI creates risks for job displacement, worker exploitation, and erosions in job quality and worker power across the labor market

AI and digital technologies impact nearly every aspect of people's livelihoods, from their workplace rights to whether they can find and keep a job. Below are some examples of how employers' use of AI is determining the life chances of workers in California now.

AI may result in job displacement and increased precarity

The same CEOs who design and sell AI technology are simultaneously raising alarm bells about the potential for massive job displacement—stoking widespread fears about AI and automation.

It is very difficult to fully grasp the impact of AI on job displacement at this stage. Some research suggests that "the broader labor market has not experienced a discernible disruption since ChatGPT's release." Other research finds evidence that entry-level employment declined in occupations that are most exposed to AI and in which AI tools are used to *automate* work rather than *augment* human workers.

But while the picture across the economy is still unclear, within specific industries and workplaces, AI technologies have already posed a direct challenge to workers. Actors and writers have recognized that AI and automation have become an existential threat to their livelihoods. AI tools have contributed to replacing workers in these occupations and also used their creative output to generate new content for which workers are not remunerated. Similarly, computer programmers are employed to train AI models, which are likely to deskill or possibly displace their current roles.

Workers have begun to push back against such incursions. The Screen Actors Guild-American Federation of Television and Radio Artists (SAG-AFTRA) and the Writers Guild of America (WGA) carried out strikes in 2023 and negotiated contracts that addressed, among other things, job security, compensation for the use of their creative output or likeness by AI systems, and notice and consent rules informing workers and their unions about the use of AI. Similarly, the Culinary Workers Union negotiated a new contract in 2023 that included a severance package for employees whose roles were replaced by newly introduced technology, as well as training for new jobs created by such technology.

Instead of cutting jobs en masse, employers may also introduce technology that worsens jobs. Automation of tasks can lead to deskilling, as workers at the start of their careers are prevented from performing tasks that are relatively easy to automate, but that serve as a critical foundation for building their experience and networks. The opportunity to learn on the job and develop the repertoire needed to succeed is denied to workers when entry-level jobs become automated.

Contrary to utopian claims by tech executives of widespread abundance or one-day work weeks, history suggests increased productivity achieved through AI will result in fewer workers and remaining workers doing more for the same or reduced wages. Employers might also cut working hours or resort to zero-hour contracts. Our recent history already shows this to be true: since the late 1970s, our policy choices have led to a stark divergence between productivity and worker pay.

A lack of good data makes it harder for us to understand how AI is changing the labor market. Most companies are not releasing data that could independently validate claims of AI-driven displacement and help California's policymakers protect our workforce. Lack of access to data on AI's enterprise use and outcomes makes it more difficult for public interest researchers and policymakers to verify whether statements by CEOs about mass layoffs due to AI are "AI washing" or empirically valid, cautionary signals that require immediate action.

Beyond job loss alone, researchers have begun to identify an overall sense of dismay and cynicism that CEOs' outsized pronouncements about AI's takeover of the workforce are creating. As Abigail Marks, Professor of the Future of Work at Newcastle University, has said, "the greatest risk may not be that AI replaces everyone overnight, but that the fear of replacement becomes widespread and corrosive—damaging wellbeing, undermining dignity and building resentment."

Whether it happens through deteriorating job quality, lower wages, a potentially widespread loss of employment, or eroding belief in the social contract, technological advancements in automation and AI can be catastrophic for workers across industries and occupations if left unaddressed.

AI tools are being used by employers to demean and exploit workers

AI used in hiring raises the risk of automated discrimination

AI systems are now widely used by recruiters at all stages of the hiring process, from sourcing applicants to final selection. According to a study by Upturn, applicants can be exposed to a range of biases during hiring that harm their prospects. For instance, predictive tools use past data patterns to determine which job postings are visible to applicants—if women are less likely to click on openings for managerial roles, a predictive tool might determine that such roles should not be visible to a female job applicant. Separate from publicizing postings and sourcing applicants, AI-enabled "automated-decision systems"

(ADS) are being used to vet candidates and move them through the hiring process. These systems have the potential to "exacerbate existing biases and contribute to discriminatory outcomes," as noted by California's Civil Rights Council. Allegations of AI-related hiring discrimination are already the basis of a number of ongoing legal disputes.

Although companies often claim to use human review to avoid such biases, there is evidence to suggest that reviewers are highly credulous of AI model outputs. In an experiment in which human reviewers worked with an AI model that exhibited racial bias, researchers found that the reviewers accepted the model's biased decisions 90% of the time. The same reviewers were able to make selections without racial bias when working without AI or when working with an AI model without racial bias.

These tools are making highly important decisions in the lives of California's workers, with few regulations and little to no transparency from employers. As a result, job applicants are unlikely to know that they were screened by predictive technology or to have insight into how they may have been impacted. This lack of transparency undercuts workers' ability to enforce their legal right to nondiscrimination.

Algorithmic management undercuts worker dignity

Employers across sectors make widespread use of algorithms to determine work scheduling, track and measure worker productivity, and make consequential decisions regarding promotions, pay, disciplinary action, and dismissal. This can wreak havoc for workers in a variety of ways.

App-based workers, for instance, continue to struggle to make ends meet in the face of algorithmic wage discrimination, in which black-box algorithms are used to calculate differing wages for broadly similar work based not only on productivity but also on factors beyond the workers' control—and even personalized data. In a 2025 survey of 2,500 Uber drivers conducted by PowerSwitch Action and Gig Workers Rising, 73% reported that declining low-fare rides caused their income to drop because the app either gave them fewer rides or only low-fare rides. 33% of drivers also observed that rides on the app slowed as they came closer to achieving a bonus or reaching a higher level in a rewards program. A majority of the workers surveyed reported psychological distress due to a lack of control over their earnings on the Uber app.

Gig workers are far from the only category of workers subjected to algorithmic management. A study conducted by the UC Berkeley Labor Center found that algorithmic management is used in call centers, warehouse and distribution centers, home care, retail and grocery stores, janitorial and security services, transportation, restaurants, hotels, healthcare, construction, and the public sector.

In call centers, AI tools monitor calls, prompt workers in real time to adjust how they speak to consumers, and measure their productivity. Workers who take longer than the pre-determined time for a call are penalized and receive lower ratings, which could, in turn, affect their pay and career advancement.

Similarly, the use of AI and automated decision-making in healthcare can lead to overworked healthcare workers and underserved patients as algorithmically determined nurse-to-patient ratios are imposed. AI tools are also used for remote patient monitoring, medical condition prediction, and automation of care plans. This technology leaves both workers and patients vulnerable to imperfect technologies and strips out the personal judgment, expertise, and touch that make difficult life experiences humane.

Productivity-at-all-costs results in work intensification and occupational health & safety hazards

Productivity tracking tools place workers under immense physical and mental pressure. Warehouse workers at Amazon report suffering from injuries and working under dangerous and traumatic conditions in order to achieve the productivity benchmarks determined by tools that map their movement through the warehouse. Workers commonly report stress caused by being observed and evaluated nonstop in the workplace.

A study of Amazon delivery drivers found that drivers are required to connect to at least nine different apps during their workday, subjecting them to continuous and close observation. In the same study, drivers also reported that they may be penalized for safety behaviors that slow their productivity, such as hard braking to avoid obstacles.

Among gig workers, the previously mentioned survey by PowerSwitch Action and Gig Workers Rising reported that over 40% of Uber drivers continue to drive when tired and even when in pain in order to earn enough to make ends meet under the uncertain conditions created by algorithmic management tools.

It's not just warehouse and platform-based gig workers who experience work intensification due to the productivity pressures of being measured by technology: even hospice chaplains have reported, "metrics prevented them from fulfilling their calling."

Surveillance technologies mean extractive "bossware"

Digital technologies enable employers to continuously monitor and collect data about workers. Through "limitless worker surveillance," employers leverage data collected from workers at in-person workplaces and at home, often without their knowledge, infringing on their autonomy and privacy. Employers are able to collect data on workers' physical well-being, including fitness levels, pregnancy, and other health measures. As a visceral example of how this manifests at a work site, Burger King will soon use AI to check if employees say "please," "thank you," and "welcome to Burger King." The data collected by employers and vendors is then leveraged in management decisions that extract more from workers.

Workers can be subject to monitoring even while working remotely. Workers' digital actions are closely recorded through the use of technology that tracks and logs keystrokes, idle time, website and application usage, bandwidth utilization, screenshots, and webcam photos, including hundreds of tools organizations like Coworker refer to as "bossware."

The National Employment Law Project (NELP) found that the use of bossware "intensified a range of existing job quality problems, including harmful disciplinary practices, job precarity, lack of autonomy, poor working conditions, work fissuring, independent contractor misclassification, exploitative pay, unfair scheduling practices, barriers to accessing benefits, discrimination, inequity, and suppression of worker collective action."

Surveillance technology can also be used by employers as a tool to monitor conversations about workplace conditions and to facilitate violations of workers' rights to protected concerted activity. The AFL-CIO notes that "surveillance systems can alert management when people are talking in a bathroom, or if the word 'union' is said, tracking 'troublemakers' and manipulating scheduling to prevent organizing or even retaliating."

Global AI supply chains contribute to a race to the bottom on labor standards

AI tools exist because of a vast global supply chain that encompasses everything from mineral extraction to manufacturing, shipping, data storage, and handling e-waste. Starting from Big Tech companies in Silicon Valley, chains of subcontracting relationships stretch across the world, connecting workers and communities across continents. Global subcontracting has long been a common business strategy for companies seeking to increase margins by cutting labor and compliance costs, but the ease with which AI companies can offshore this work—and the critical role these workers play in the production of AI systems—has made the practice explode. Companies and workers who contribute different knowledge and resources to AI systems are part of a global value chain in which they co-create these powerful tools.

The segment of this supply chain that enhances datasets and trains the algorithms that form the basis of AI—referred to as "data work"—provides a helpful case study of how work is subcontracted globally to fuel the current AI boom. Data work tends to be low-wage and precarious, with reports of workers in the Global South being paid as little as $1.50 per hour, and workers like content moderator Daniel Motaung in Kenya have been dismissed in alleged retaliation after raising questions about poor pay and working conditions. Data workers struggle to make ends meet even as companies raise billions of dollars in private equity and venture capital. Workers also report severe mental health impacts from certain aspects of this work. As part of a lawsuit filed against Meta and outsourcing firm Samasource in Kenya, it was revealed that more than 140 Facebook content moderators were reportedly diagnosed with severe post-traumatic stress disorder, as the legal filing alleges that "images and videos including necrophilia, bestiality and self-harm caused some moderators to faint, vomit, scream and run away from their desks."

These companies frequently shift work around the world, leading to loss of livelihoods with little or no notice. Workers on digital platforms have found themselves suddenly laid off, locked out of their accounts, or even denied payment for work already completed. With little recourse to legal protection and no clear channels of communication with these companies, workers are left to fend for themselves.

This is not just a problem for workers overseas. U.S.-based data workers surveyed and interviewed by TechEquity, Alphabet Workers Union, and Communications Workers of America shared similar challenges, including struggling to make ends meet, inadequate training, pressure to meet unrealistic work quotas, and limited access to benefits, among other issues.

The global and fragmented nature of data work, which takes place both in physical offices and on platforms in workers' homes, and the constant shifting of companies' contracts pose challenges for workers and unions seeking to bring data workers together to fight for improved conditions.

As fears mount in the U.S. about the possibility of AI-related labor displacement, tech leaders have discussed the potential for this technology to create new kinds

of work opportunities. LinkedIn has identified data annotator as one of the fastest-growing jobs in the U.S., with data annotation among the fastest-growing skills in demand. If data work represents one of the jobs of the future, elected officials in California should ensure that these roles offer dignity to workers.

The importance of prioritizing worker voice

Enacting policies that support workers also result in positive outcomes for businesses and consumers

Engaging workers throughout the process of developing technology that impacts them at work is good for workers and for businesses. It allows business cases for the technology to emerge from the ground up and increases trust, ultimately driving sustainable and widespread adoption. Drawing on workers' direct knowledge of their workflows increases the likelihood of successful adoption, which in turn can improve worker satisfaction as well as business outcomes.

Despite these clear benefits, workers are largely not involved in the adoption of new technologies. Gallup found that "relatively few employees currently have a meaningful say in tech decisions—just over a third say they have 'a lot of' (12%) or 'some' (25%) influence."

An MIT Project NANDA study found that 95% of enterprise generative AI pilots failed, with major contributing factors including a lack of trust in the products and the failure to align with employee workflows. Meanwhile, a survey of 1,500 U.S. workers and AI experts by the Stanford Institute for Human-Centered AI and the Digital Economy Lab found that workers viewed AI integration as unwanted in over 40% of the AI-enabled tasks reviewed by researchers.

In contrast, studies have shown that workplace technologies that take into account workers' needs, existing work processes, and general welfare are more likely to succeed than systems imposed on workers. When Kaiser Permanente launched a project to bring technology-enabled ambulatory care clinics to their patients, workers and their representatives across the organization were involved at every stage of the process, resulting in an environment of trust in which workers could benefit from new technologies. This led to the successful rollout of the clinics and improved care for patients.

Excluding workers from decision-making does not just lead to low adoption—it can result in significant harm to consumers. For example, in a survey conducted by National Nurses United (NNU), 69% of nurses who tried a tool that used algorithms to analyze electronic health records reported disagreeing with the computer-generated estimate of a patient's illness severity and the number of hours of nursing care they needed. These measures are used to make staffing decisions, so incorrect estimates can lead to understaffing and critical gaps in patient care.

In another example, employees at Brookdale Senior Living, one of the nation's largest operators of senior homes, reported that the company used an algorithm-based system for determining staffing that "failed to capture the nuances of caring for vulnerable seniors" and "underestimated the amount of labor they needed to meet resident needs." The company has even faced patient wrongful death lawsuits in which family members and employees allege that staffing levels may have played a contributing role.

Solutions to support California's workers and economy

Policymakers must take a holistic approach, with a focus on enhancing worker agency, to address the risks of AI in the labor market

In order to improve conditions for workers and society at large, it is necessary not only to regulate current technological development but also to address the underlying economic and political conditions that make it easier to extract from and exploit workers. With the right policies in place, technological innovation in California can create shared prosperity, rather than inequality.

Priorities for elected officials in California should include:

Shift the balance of power in the workplace by expanding collective bargaining and access to unions: Workplace technology functions best when it is co-designed with workers rather than imposed on them. To enable workers to have a voice in technology decisions, the power imbalance between workers and employers must be addressed directly. Elected officials in California can create an enabling environment for workers to develop their collective voice. Unions have proven to be the most effective tool for channeling worker voice

and have a strong track record of negotiating collective bargaining agreements with employers that shape technology. These contracts serve as models for what is possible when employers, employees, and their representatives develop solutions together.

However, the vast majority of working people are not currently covered by union contracts. Policies that expand access to collective bargaining and protect workers' right to unionize are essential to ensure that the adoption of technology is fair and mutually beneficial. As California continues to develop its governance approach to technology, workers and their representatives should be included on any boards, commissions, or other multistakeholder entities established by the state related to AI and other digital technologies.

Make a plan to address the potential for AI-related job loss: While the true impacts of AI on job automation and displacement are still unclear, California should take proactive steps to prevent widespread unemployment. One early step would be to require employers and AI developers to provide transparent data around the use and impact of AI products within enterprises. This data can help guide policy interventions to ensure that AI is used to enhance job quality and prosperity rather than fuel mass unemployment. This data could also inform research that helps identify workers most likely to be affected, design effective workforce training programs, and develop new models of safety-net transition assistance that can aid affected workers.

Establish and enforce workers' technology rights: The state government has an important role to play in creating baseline, common-sense protections that will benefit all workers in California. As we've described, emerging technology can act like a black box used to obscure workers' rights to minimum wages, health and safety protections, and non-discrimination. As Lorena Gonzalez, President of the California Federation of Labor Unions, AFL-CIO, says, "We need stronger guardrails to make sure there is human review and oversight of any decision made by a machine that impacts a worker's job and paycheck."

Specifically, workers need transparency around how tech affects them at work—from hiring and firing decisions, to how their wages are set, to whether and how they are surveilled. This includes informed consent, meaningful human oversight, and clear guardrails around when and how tech can be used for critical decisions, as well as the right to access, correct, and control the use of their data. California has been a leader in establishing workers' foundational data rights through the California Privacy Rights Act. These rights can be expanded to ensure that protections evolve to be responsive to emerging technologies.

However, expanding transparency and data rights for workers alone is not enough. Merely providing information to individual workers still places an enormous burden on them to act on and enforce their rights. And providing notification only *after* workers have been subject to AI and algorithmic decision-making systems denies workers the collective opportunity to actively shape these systems in a way that enhances their work. That is why workers need additional collective rights such as pre-use notification, prohibitions on the use of AI for disciplinary or termination decisions, and the ability for workers to contest AI-driven decisions to ensure that workers retain their agency over decisions made by technology that affect their livelihoods.

In recent years, hundreds of state-based legislative proposals have advanced across the nation to put in place sensible guardrails related to electronic monitoring, algorithmic management, workers' data privacy, human-in-the-loop mandates, automation, and other key considerations at the intersection of technology and labor. California should be at the leading edge of innovating new protections and models for collaborative governance when it comes to AI and the future of work.

Invest in state regulatory and enforcement infrastructure: While there are existing laws on the books that protect workers from many workplace harms, when those harms are perpetrated by technology, enforcement can be challenging. California's Labor and Workforce Development Agency needs a robust set of resources to evaluate the claims, outcomes, and interactions of this technology within California's workforce. In order to ensure California can continue to thrive as a state with strong labor standards, California's leaders must ensure that the people vested with the authority to investigate and enforce our labor laws have the time, money, and talent they need to do so.

In addition, we can institute policies that hold employers accountable for replacing workers with AI. Nobel laureates Daron Acemoglu and Simon Johnson have proposed new tax structures that can incentivize employers to retain workers. California could pursue these and other policy proposals to ensure humans remain firmly in control in critical situations, update layoff protections for the modern workplace, and implement other mechanisms that disincentivize the use of AI that is there merely to displace workers.

Proactively prevent the most harmful impacts of technology: California should ensure that no one is subject to the most egregious applications of workplace technology. The state should align its regulations with other jurisdictions that have prohibited the use of technologies like facial recognition, predictive behavioral tools, monitoring, and quotas that create health and safety risks;

the use of AI tools for discipline or termination; and the utilization of intrusive surveillance technology.

Ensure that California-based AI companies are held accountable for improving labor conditions throughout their supply chains: Tech companies based in California are at the top of global AI supply chains, giving them both enormous leverage and responsibility to ensure that the workers who contribute to their profits are treated fairly. Companies often claim to have sustainable labor practices, but this rarely extends across their global network of contractors and subcontractors. Elected officials in California should exercise their authority to improve transparency and accountability in this industry.

When the state procures AI technology, it has an obligation to ensure that public funds do not contribute to the exploitation of workers, but rather support good jobs. Supply chain labor abuses are not new or unique to the AI industry, and policymakers can learn from improvements in the garment industry and other sectors, for example, by expanding California's Public Contract Code to apply Sweatfree standards to AI systems as well. If AI technology opens up new job categories—like data annotation—California should ensure that the future of work is built on good-quality jobs for Californians and for workers around the world.

The future is ours to build

California's leaders hold the future of the state's workers in their hands. As AI spreads throughout a labor market in which employers already have enormous power over workers, the health of the state's economy hinges on our leaders making bold policies that protect and empower workers.

As SEIU President April Verrett has written, "Right now, AI and algorithmic management are concentrating unprecedented control at the top while systematically disempowering everyone else. But this outcome isn't inevitable. We can build an alternative future where technology enhances rather than diminishes human work. Where working people participate as decision-makers, not mere data points. Where care, connection, and community take priority over profit margins and efficiency metrics. This better future won't arrive on its own."

Adopting the policy priorities outlined in this brief—proposals that strengthen workers' rights to negotiate collectively, secure new technology rights for workers, and protect them from harm—will create a labor market for the twenty-first century where technology can be used to enhance workers' livelihoods and the economy overall.

Acknowledgements

Thank you to the workers, unions, organizations, and researchers who informed this memo, both publicly and confidentially, including Annette Bernhardt (UC Berkeley Labor Center) and Aiha Nguyen (Data & Society Research Institute).

APPENDIX

Protecting Workers Amid Potential AI-Driven Labor Displacement

AI CEOs have made expansive predictions about the widescale disruption that their products could create across the labor market. However, the development, use, and impact of technology in the workplace are always shaped by political choices: "Technologies do not evolve deterministically. They are shaped by the people who use them and the organizations that guide how they are used."

The responsibility to determine whether AI does, in fact, cause an employment shock that undermines the social fabric, or whether it can be used to improve job quality and access to opportunity, will fall largely on California's next generation of leaders.

But charting this path will not be straightforward. The approaches we have taken to soften unemployment shocks of the past are insufficient. In previous waves of labor market transformation, policymakers and corporate leaders relied almost entirely on workforce development and retraining programs that, on their own, were not enough to successfully transition workers into a new economy.

Most recently, the "learn to code" movement made promises that reskilling workers for high-tech jobs would build wealth for communities that had been left behind in the aftermath of the Great Recession. But over the last fifteen years, the success rates of these programs have been minuscule. The ranks of the tech workforce continue to lack diversity, and inequality has increased—driven in large part by the concentration of wealth fueled by the tech industry. And now, those very same jobs that desperate workers were told to pursue are at risk of being automated by AI.

The labor market transformations for which we have precedent have also typically been contained to specific industries, whereas AI is threatening to transform everything from screenwriting to long-haul trucking. In this context, it will be impossible to design worker retraining programs that meet the vast and varied potential need.

What we most hope decision-makers take from this memo is that workforce training commitments alone are an insufficient response to the current and unfolding impact of emerging technology on workers. The state will need a bold, multifaceted agenda to ensure workers are empowered rather than exploited or displaced en masse by AI. These include policies to address the increasing imbalance of power between workers and corporations, and the fair distribution of the gains AI is creating for the economy.

A Policy Agenda for AI-Driven Labor Displacement

While the threat of AI-driven labor displacement has gotten much of the headlines, it is important to understand how these worries are connected to other aspects of AI's impact on workers.

Often, policymakers will be inclined to jump to solutions that are narrowly tailored to job automation or displacement. In doing so, leaders will have enacted weak scaffolding to support the broader workforce. It is critical that workers have a robust set of rights that acknowledge the realities of a technology-driven, 21st-century workplace. These rights serve as a foundation on which specific solutions to potential automation can be added. The policy ideas we outline in the main body of this memo are key to addressing the threats workers face from AI. They set a strong foundation to ensure that the introduction of technology is beneficial to workers, not exploitative.

Below, we outline four additional principles that can inform California policymakers as they proactively plan for the potential for mass workforce disruptions related to AI—whether real or a form of "AI-washing."

Fund research that uncovers what is real and what is hype in AI job loss

"A white-collar bloodbath." "AI Is Wrecking an Already Fragile Job Market for College Graduates." "AI is coming for entry-level jobs. Everybody needs to get ready." We've all come across recent headlines like these. They shape the public's and policymakers' beliefs about how AI is impacting the workforce. But the studies behind these headlines reflect a more nuanced reality—and point to serious gaps in the data that must be closed if we are to make sound policy.

A review of the literature shows that many of the studies that make claims about AI automation focus primarily on "task exposure," those specific pieces of a job that an AI tool may automate. The more tasks that can be automated by AI tools that make up a specific job role, the logic goes, the more likely that job is to be automated.

But task exposure only allows us to infer from existing AI capabilities. It doesn't provide any context for whether a particular job is *actually* being automated away. Most importantly, task-exposure studies say nothing about the company-level dynamics that determine AI adoption and use. We will need much richer datasets in order to create a clearer picture of AI's impact on labor displacement.

Some proposals for how to structure this data regime have already been made. But we will need a more robust conversation that includes researchers, companies, organized labor, and lawmakers to define what we really need to know to protect workers from AI and offer proposals for how to establish the datasets that will help us answer those questions.

California can lead the way. State leadership can convene these conversations and fund data collection and research that can create a clearer picture of the AI job displacement landscape. These research approaches must capture both rich qualitative data from the workers who are at the frontlines of AI usage as well as quantitative data from the employers who are making decisions on how, when, and which tools to adopt. The state could also take steps to require not only AI developers but also employers who purchase and use AI tools and systems to provide transparent data on the use and impact of AI products within enterprises.

Empower workers to co-design solutions

Gallup found that "relatively few employees currently have a meaningful say in tech decisions—just over a third say they have 'a lot of' (12%) or 'some' (25%) influence." *How* AI enters workplaces has an enormous impact on whether employers use technology as a way to cut labor costs or to improve jobs. We know that when clear guardrails are in place, and workers are given a voice, "input from workers can increase the likelihood that organizations use generative AI tools effectively, *and* workers' job quality improves." Employers and technology companies should authentically engage workers and unions, who are closest to the technology and best suited to speak to their own needs more broadly.

Worker voice mechanisms can enable workers to ensure that technology actually benefits them and their communities, not just in this wave of technological adoption but in future ones as well.

Engaging workers in the technology adoption process can also help to stem the tide of job loss. Strengthening workers' bargaining power "increase[s] pressure on companies to pursue alternatives to layoffs and provide greater support for those who do lose their jobs, as well as offer protections against other worker abuses that new technology such as AI may facilitate." For example, unions have already negotiated collective bargaining agreements that include job security protections, pay protection and severance programs, comprehensive technology transition agreements, and training and workforce development for technological change.

Finally, as California continues to develop its governance approach to technology, workers and their representatives should be included on any boards, commissions, or other multistakeholder entities established by the state related to AI and other digital technologies.

Take a trust-based approach to worker training

Programs to ensure that the skills of the labor force match the job market are only one piece of the picture when policymakers are considering how to ensure that working people are protected as technology develops. As Serena Oduro and Anuli Akanegbu of Data and Society write, "Investments in workers' skills can be most valuable when they are connected with efforts to address the broader socioeconomic conditions that shape career outcomes."

Oduro and Akanegbu recommend that workforce development efforts focused on AI "clearly define 'AI literacy,' provide paid work, ensure AI literacy pipelines actually lead to quality jobs, and expand the stakeholders involved in designing and implementing AI literacy efforts," most critically, workers themselves and labor unions. Paternalistic solutions that are designed around the needs of employers or policymakers are much less likely to succeed. California can also "require and fully fund rigorous evaluations of training programs."

In addition to training programs, the state should invest in designing new programs that allow workers to self-determine their future. One of the most promising economic opportunities that AI may create is to lower the cost of starting and running a business. In addition to creating training programs that allow workers to transition into new jobs, they can also provide grants and loans directly to workers to start their own businesses using AI.

Ensure the gains created through AI technology are fairly distributed

AI tools are made possible by the vast amounts of data they consume, which are created by the public. However, Big Tech companies are using their power to entrench the prosperity gains AI creates in the hands of a few individuals. In order for workers and the general public to realize the economic opportunity that AI promises, policymakers must ensure a fairer distribution of the benefits.

Economists have proposed many promising ideas that would more broadly share the gains created by AI, including:

- Nobel laureates Daron Acemoglu and Simon Johnson have proposed new tax structures that balance out our current tax-preferencing for capital over workers and thus incentivize employers to retain employees rather than replace them with technology. There are other tax proposals that would capture value from AI gains to pay for more robust unemployment insurance programs and other social safety nets that meet the scale of the potential threat from AI.
- Acemoglu and Johnson, alongside David Autor, also recommend "discouraging 'expertise theft'" from workers by "building legal frameworks that support workers' ownership of their capabilities and creative output [to] give workers greater control over how their expertise is used, reduce the commodification of human knowledge, and preserve incentives to invest in skill development and innovation."
- Scholars such as Liam Epstein have proposed creating sovereign wealth funds that are funded by AI companies and which pay out a dividend to members of the public, much like Alaska's oil fund. Other proposals of this type suggest the government could take equity stakes in AI companies and distribute the gains from those stakes via tax refunds or other payments to members of the public.
- There is some evidence that wage insurance policies can smooth the transition for some workers into new jobs. These programs cover some difference in lost pay for workers who are forcibly transitioned into lower-paying jobs. One such program, the Obama-Era Reemployment Trade Adjustment Assistance Program, was found to help workers transition to new jobs faster and increase their wages over the following four years. We can consider expanding and strengthening these programs for the AI era.

Beyond workforce development

Workforce development programs on their own will not meet the potential threat posed by AI to the labor market. Taken together, the principles outlined above provide a foundation on which we can build an economy where AI is creating widespread value and opportunity rather than facilitating even greater inequality, worker exploitation, and extraction.

AI & AFFORDABILITY IN CALIFORNIA

Land of Opportunity or Hardship?

In partnership with Economic Security California

Affordability has come to dominate the national conversation and emerged as a decisive issue in several high-salience races last year in Virginia, New Jersey, and New York City. In contests shaped by national polarization and economic anxiety, successful campaigns framed rising costs not as isolated consumer problems but as the result of concentrated corporate power shaping prices, access, and quality of life. The undertold part of that story is tech, and specifically AI's role in driving our affordability crisis.

Today, working families feel the squeeze everywhere—from rent and utility costs to groceries. In each instance, the tech sector is a behind-the-scenes player in driving up the cost of these goods. Utility rates are creeping higher because of how AI saps energy supply. Corporate actors from Delta Airlines to Kroger grocery stores are using pricing technology to set "personalized" prices and to coordinate with each other to constrain consumer choices. Landlords are using rent-setting technologies to coordinate higher rents while skirting and undermining traditional anti-collusion laws.

Businesses that once competed against each other are working together to limit competition and drive up prices—and AI is making it faster, more widespread, and easier to hide. Instead of creating better products or lowering prices, corporations now compete with you, the customer, to find what you're willing to pay based on how desperate or vulnerable you are.

The existing guardrails against corporate overreach have not held, and California's working families are literally paying the price. With bold leadership, we can reclaim our reputation as a home for dreamers; to do so, our leaders must make California a state that works for everyone, not just corporate interests.

This memo identifies AI's role in California's affordability crisis, from algorithmic collusion on rents to data centers' inflationary impact on utility rates to surveillance pricing that drives up the costs of everyday essentials. It then identifies solutions. The policy prescriptions here can increase competition to drive down costs, stop AI companies from passing energy costs on to households, and give regular people control over their data. Together, they are the building blocks of a plan for a just economy in the AI age.

Harms and costs: What's at stake?

Unchecked use of AI drives inflation, harming California households

Far from its initial promise of connecting more people and democratizing information, the technology industry is increasingly fueling its growth by exploiting every opportunity to squeeze more out of Californians. While the sector advertises innovation, time-saving products, and services that create convenience, users are seeing their own data wielded against them to set discriminatory prices and increase costs across the economy.

Algorithmic price fixing is helping landlords drive up housing costs

Housing is typically a household's largest monthly expense. In California, housing costs relative to other expenses have long been out of balance. The state has the dubious distinction of having the second-lowest homeownership rate, and the second-highest rate of housing cost burden (renters spending more than 30% of their income on rent). Now, landlords are turning to technology to stack the deck even further.

Companies like RealPage, which sell rental pricing algorithms, facilitate landlords setting their rents together to maximize profit. Large property management companies and "mom-and-pop" rental owners alike—who would otherwise compete on price and quality—collude to keep rents high. By sharing information with RealPage that the public does not have, landlords were able to

increase their revenue by an additional 2-7% each year on top of regular market increases. In internal communications, these property owners even expressed a preference that their competitors use the software, as widespread use constrains renters' alternative options. Research shows that landlords who use the technology added $3.8 billion in costs for renters in 2023 alone. In California, pricing coordination added $110 a month to market rent for Riverside renters, $99 in San Diego, $62 in San Francisco, and $24 in Los Angeles.

ONE LANDLORD EVEN REMARKED, "I ALWAYS LIKED THIS PRODUCT BECAUSE YOUR ALGORITHM USES PROPRIETARY DATA FROM OTHER SUBSCRIBERS TO SUGGEST RENTS AND TERMS. THAT'S CLASSIC PRICE FIXING..."

Pricing middlemen like RealPage acknowledge that their products are designed to extract the most through anticompetitive practices. RealPage has touted that its software "driv[es] every possible opportunity to increase price," while suggesting that among landlords, "there is greater good in everybody succeeding versus essentially trying to compete against one another in a way that actually keeps the entire industry down." One landlord even remarked, "I always liked this product because your algorithm uses proprietary data from other subscribers to suggest rents and terms. That's classic price fixing..."

Renters, meanwhile, are shouldering the cost created by this unaccountable technology while outdated laws struggle to keep up. Enforcement agencies like the FTC and state attorneys general are responsible for catching price fixing, but the introduction of AI-driven middlemen has made enforcement far more difficult. Traditional investigations look for communication between competitors, but algorithmic price fixing allows companies to coordinate through shared software without ever exchanging a word. Regulators only catch the most egregious instances of price fixing, and even then, enforcement is weak: The Department of Justice settled a lawsuit with RealPage in late 2025, requiring it to stop using competitors' data to set prices and accept a court-appointed compliance monitor, but notably did not exact any financial penalty or prevent the company from using competitor data as long as it is more than a year old.

AI data centers sap the utility grid as rates increase for consumers

Alongside housing and food, utilities account for one of the most inelastic elements of a family's budget. AI-driven data center expansion is now increasing these costs independent of consumer behavior. A recent Bloomberg News analysis found that, across the US, utility bills increased 267% over 5 years in areas near data centers, and California is already experiencing similar data center-fueled rate hikes.

The most comprehensive analysis of data center location and usage available to the public reveals that California has the second-highest electricity demand related to data centers of any state after Virginia. Pacific Gas and Electric Company (PG&E), California's largest electricity utility, saw a 40% increase in data centers requesting to hook up to the power supply in 2025. PG&E recently estimated that so much data center development is planned in San Jose that the city could increase its energy use nearly threefold above its current peak.

Santa Clara is a bellwether for how a few companies demanding outsized electrical power can affect a community. The city has the highest concentration of data centers in California, in large part because the rates Silicon Valley Power offers are consistently 40% lower than those of PG&E. As nearby AI companies looked to the town to meet their demands, the utility company increased its rates to fund the necessary infrastructure buildout. City residents, previously used to a consistent 2-3% rate increase each year, had to absorb an 8% increase in early 2023, another 5% increase in June of that year, and a 10% increase in 2024. Today, 60% of the city's energy is used to power data centers.

California data centers nearly doubled their electricity consumption between 2019 and 2023, from 5.5 to 10.8 terawatt-hours, and could reach the equivalent power use of 2.4 million homes by 2028. The California Public Utilities Commission's consumer watchdog, the Public Advocates Office, has warned that customers could end up paying for billions of dollars in grid upgrades if projects never materialize or use far less power than promised. With data center electricity demand projected to double or triple over the next three years, some California cities face the same pricing pressures that drove utility costs up by as much as 267% in other data center hubs.

AI-enabled data surveillance raises our prices and threatens privacy, autonomy, and economic stability

AI has made it easier for businesses to collude to raise prices, but even companies acting alone can use it to extract as much as possible from consumers.

The data broker industry is one of the most influential and least regulated sectors of the economy. Companies using data from data brokers can decouple the prices of goods and services from traditional supply-and-demand factors, such as the availability of goods or competitive forces, by instead leveraging immense troves of personal data against us. This widespread tracking is known as "surveillance capitalism," and the rates it sets are "surveillance pricing." Surveillance pricing pits customers against their data, diminishing their purchasing power in the process.

Companies now build profiles of potential buyers using data from apps, websites, and loyalty programs, and then use those profiles to charge each customer the most they are predicted to pay. Using a host of factors—demographics, location, purchase history—algorithms calculate individualized prices designed to extract maximum revenue from each customer. San Diego County's District Attorney fined Target $5 million for increasing the price of goods once a customer's location data showed they were near a store. Pricing algorithms showed Bay Area residents hotel rates $500 higher than those shown to people from Phoenix or Kansas City looking at the exact same room and dates. After recent public scrutiny, Instacart ended AI pricing experiments that charged different customers up to 23% more for identical items.

While California's Delete Act gives consumers the right to delete their data from data brokers, it does not prevent companies from setting discriminatory prices based on data they collect directly through apps, websites, and loyalty programs.

The push to extract more from consumers is driving companies to go beyond traditional sources of personal data. In 2025, advocates found that Wegmans supermarket was collecting biometric data on customers entering its retail locations. A notice posted outside the stores informed customers that their facial images, eye scans, and voice prints would be collected and used to identify them while they shopped. In 2023, Rite Aid agreed to a five-year ban on the use of facial recognition technology after the FTC found that the technology discriminated against customers of color.

How California can lead

The state can achieve wins on affordability by restricting surveillance pricing and enforcing competition laws

The current guardrails that could keep AI tools from raising prices and protect against concentrated economic power and anticompetitive corporate behavior are inadequate. Intense tech industry lobbying efforts to stop policymaking and enforcement—coupled with underresourced regulators and the inherently opaque systems that power these tools—have resulted in an anemic governmental response to these cost-driving technologies. Antitrust law, rent stabilization, privacy regulations, and antidiscrimination protections are all failing to hold the line against predatory uses of technology to raise prices for Californians.

California can course correct. The state can and must act swiftly to rein in AI's impact on distorted markets and runaway prices. California's leaders must take bold action to ensure competition, fair markets, and consumer choice, ultimately making the state an affordable place to live.

Drive competition to strengthen the economy

Consumer prices are creeping up in part because big businesses are working together, with a major assist from AI, to exploit their customers. As discussed in more detail in our other memo, *Beyond Big Tech: An alternative vision for California's tech economy*, to restore market equilibrium, California needs real competition that drives choice and enables countervailing consumer power. California's leaders must:

- **Update California's Cartwright Act to address "single-firm conduct" and anticompetitive conduct:** Modernize state antitrust law to clearly prohibit conduct that allows a single company to use its size and power to illegally crush competition, also known as "single-firm conduct." Most states already have antitrust laws on the books that regulate anticompetitive behavior by a single entity. California, which only regulates antitrust violations between two or more entities, does not. A change of this kind would include holding dominant AI platforms accountable when they block competitors, self-preference their own services, or leverage power in one market to squeeze out rivals.

- **Build on recent anti-collusion legislation:** Expand on AB 325, which passed in 2025, which prohibits algorithmic collusion, to strengthen antitrust enforcement against pricing power abuse, corporate practices that undermine competition, and companies that control production, distribution, and sales.
- **Implement structural solutions that promote competition and innovation:** Require data portability so consumers are not locked into dominant platforms, mandate interoperability standards so smaller competitors can enter the market, and consider structural separation to prevent tech giants from leveraging dominance in one market to push out competition in others.
- **Invest in CalCompute and public AI infrastructure:** To reduce dependence on Big Tech infrastructure, driving up costs across the economy, California should fully fund CalCompute, established by SB 53 (2025), to provide affordable AI compute resources for small businesses, the public sector, researchers, and startups.

Stop AI from passing energy costs onto households

Across the country, the buildout of data centers and associated energy infrastructure has driven up energy bills for households and small businesses. Despite committing to pay their fair share, the industry has fought prohibitions on passing its costs along to ratepayers.

To prevent ratepayers from shouldering the cost burden for the AI data center buildout. California must:

- **Establish strong ratepayer protections.** To safeguard households and small businesses from absorbing any stranded costs of building data center infrastructure, the state should require that data centers pre-pay or frontload 100% of the costs to serve them (including transmission, generation, and capacity). Additionally, California can establish separate tariff structures for large-load customers and require protective contract terms, including a minimum contract length, minimum demand charges, early termination penalties, and exit fees.
- **Ensure the costs of data center expansion are not coming from the public purse.** The state should prohibit the use of tax incentives—whether that's through use tax exemptions, corporate income tax credits, property tax abatements, utility tax exemptions, or sales tax exemptions. In 2019, Illinois established a substantial tax incentive for data centers. Since then, 27 sites have been constructed, with an average of 22 workers per site. This extremely low job creation threshold qualified the companies to receive $983 million in tax breaks. Ultimately, costing Illinois taxpayers $1.7 million dollars for every single job created (not including temporary, one-time hires of construction workers). At $1.7 million per job, this is not economic development—it is a direct wealth transfer from taxpayers to some of the world's most profitable corporations. To put that figure in perspective, $1.7 million per job would be enough to hire a teacher, a firefighter, and a nurse—and still have money left over. In early 2026, responding to public backlash over the economic and environmental impact of data centers, Gov. Pritzker of Illinois proposed a two-year break on any tax incentives for data centers. The experience is not isolated to Illinois, as many other states have offered similarly expensive tax breaks for very few jobs in return. California should learn from these cautionary tales rather than repeat them.
- **Require transparency from data center operators.** California should mandate public disclosure of energy consumption, water use, and emissions data, and ensure that data centers fund independent cost studies. This would build on previous efforts such as the policies in SB 57 (2025), which directed CPUC to study the impacts of data centers on households, or AB 93 (vetoed by Gov. Newsom in 2025), which sought to require data center operators to report information about their water use and energy consumption.

Communities across the country have begun to stand up forcefully against the unfettered expansion of data centers and their impact on the public. AI Now Institute developed a comprehensive policy report and toolkit for state and local policymakers that includes a variety of policy proposals, ranging from data center moratoriums to specific policy proposals to ensure economic burdens are not placed on our communities, which California should evaluate and consider implementing.

Modernize transparency and discrimination protections for the age of Big Data

With unequal access to data underpinning so much of how AI works, the prices Californians pay, and why it's so hard to hold bad actors accountable, public officials must modernize existing protections for how information flows today. California led when it enacted the California Consumer Privacy Act (CCPA) in 2018, but efforts by the tech industry and its business partners have intentionally slowed its implementation and effectiveness. Meanwhile, the tech industry has accelerated its development of AI and emerging technology as it uses its resources to stall regulation and oversight.

To ensure Californians do not pay the price for this toxic dynamic, our leaders must:

- **Expand civil rights laws for algorithmic discrimination.** The state can extend existing anti-discrimination protections to cover AI-driven decisions in housing, employment, lending, and insurance by requiring that individuals have some legal recourse, such as notice and the right to an explanation, the right to correct, and the right to alternative human review.
- **Strengthen enforcement and programs of CalPrivacy.** California's new Automated Decision-Making Technology (ADMT) regulations established limited consumer protections around AI decision-making in housing, lending, and employment, but will not take effect until 2027, and do not explicitly cover surveillance pricing. The state should ensure that our privacy regulations do not allow for the collection, use, or sale of our data to power harmful uses of AI technology. Additionally, California must invest in programs such as the DROP data deletion tool to ensure proactive enforcement of existing data rights.
- **Build up individual rights to control the flow of personal data.** Californians need the right to control how their information is used. In addition to expanding existing protections, new data rights must include prohibitions on the collection and use of location and other sensitive data, stopping the flow of data to entities that did not obtain express consent to access that information, and the right to remedy any violation of individual data rights.
- **Remedy the information asymmetries that allow privately-held data to be weaponized against consumers.** The RealPage case was exposed by investigative reporting, but with better public data on rental housing, the exploitative actions of landlords could have been caught sooner. The state should pursue the creation of a public repository of rental price and ownership data that would allow regulators and the public to see the same information that landlords can see about how prices are moving. This transparency would ensure accountability, allow renters to enforce their rights as consumers, and provide policymakers with better intelligence about the housing market. Legislation, including AB 1018 (2025), could also protect notice, consent, and appeal requirements in instances when a consequential decision is made using personal data. California must ensure that data and notice can also be used for the public benefit, not just for private profit.

Building a California for all

The public's mistrust of Big Tech, data centers, and AI is unmistakable—nearly 80% of voters across partisan lines say they have little or no faith that tech leaders will act in the best interest of the people of California, and 70% of Californians think we need strong laws to regulate AI.

These sentiments are driven by a sense that Big Tech is not on our side—new technology is being used to exploit and extract rather than to improve our quality of life and increase shared prosperity. Californians can see that, in the current economy, Silicon Valley and business executives profit while everyday households struggle.

This memo articulates a way to reverse these trends. Working households need a fair shot that includes prices based on actual market competition, ratepayer protections for utilities, civil rights shields against algorithmic discrimination, and, most importantly, courageous leaders invested in holding corporations accountable and making the state work for the many, not just the few. Any official who wants to show they are fighting for economic security for Californians must have a clear plan to address AI and tech's role in growing the affordability crisis.

BEYOND BIG TECH

An Alternative Vision for California's Tech Economy

In partnership with AI Now

Californians have always been known for turning big dreams into reality. It's how people from every background have come together to build the ideas, tools, and industries that have reshaped our world. The current enthusiasm around AI has sparked a meaningful opening to extend this vision, creating opportunities for California to incubate best-in-class technologies and simultaneously reckon with what it can do to ensure that the development of this industry meets the needs and aspirations of everyday Californians.

Right now, we're at a critical juncture. The AI hype of the past few years has distracted from the reality that the current trajectory of the tech industry is consolidating—not expanding—California's innovation ecosystem and actively working against the interests of everyday people.

California's next generation of leaders can and must create the economic conditions for a pro-competition and people-first version of AI and the tech sector overall. Doing so is key to unleashing a healthier and more distributed innovation economy. This memo holds recommendations for how to do that.

Historically, California has created the conditions where two people with a great idea in someone's garage can compete with behemoth, legacy corporations. This environment of competition, disruption, and technological innovation challenged the old order and birthed an industry that is the envy of the world, creating entirely new generations of products, companies, and ways of thinking.

This spirit of entrepreneurship now faces an existential threat. Over the past two decades, our regulatory environment has allowed a few large technology companies to gain unaccountable power over the tech ecosystem, preventing new entrants from gaining a foothold in the market. The recent AI boom is only deepening that power. Google, Microsoft, Amazon, and Meta control access to key inputs necessary for building AI. This means that instead of competing with Big Tech firms, AI start-ups are *dependent* on them for resources, growth, and even survival—hampering the sprawling, innovative tech economy core to California's identity. At the same time, California's overreliance on gains from a handful of companies jeopardizes the stability of our state budget and the pocketbooks of millions of Californians who never agreed to this gamble.

Californians are eager to see technology solve pressing social challenges and deliver real benefits for people's lives. When asked, a significant number of Californians think that AI is likely to find solutions to problems that we cannot solve (43%) and improve creativity and innovation (42%). When asked in small-group settings what people mean when they say "innovation," participants shared their hopes for cheaper groceries, lower energy bills, and easier access to doctors.

But instead of tackling the real problems that Californians want to see solved, AI is largely making these problems worse. In fact, AI is being deployed in ways that raise our grocery prices, drive up our energy bills, depreciate the quality of our healthcare systems, and materially reduce our access to life-saving services. These tools are largely being deployed in ways that concentrate power upstream, creating less and more expensive access for everyday Californians to critical needs like healthcare, housing, and employment.

A product ecosystem driven by a single vision—set by a handful of firms chasing the same profit motives—is costing our state a truly diverse horizon of possibility, filled with innovations that could tackle people's real-life needs.

This memo names pro-competition solutions that will build an innovative tech ecosystem that allows all companies to compete on the merits and build products that deliver widespread social benefits. That goal will require California's leaders to have a clear vision for the future—one that is centered on improving everyday people's lives, solving intractable public problems, making public investments that can catalyze growth and innovation, and spurring a new generation of builders and dreamers.

Harms and costs: What's at stake?

The current AI industry is defined by concentration and dependence, precluding a truly diverse, innovative, and competitive ecosystem from flourishing.

Big Tech companies are exploiting their monopoly positions to further entrench power and prosperity in the hands of a few

The AI market is structured in a way that entrenches dominance in a few firms that maintain control over critical inputs and funding resources. Rather than creating broad-based economic gains, this market reality is consolidating gains among the existing winners.

- **Big Tech controls the resources necessary to enter the AI market, breeding concentration and dependence.** The largest Big Tech firms in the U.S., including Meta, Google, Microsoft, and Amazon, have used their dominant market positions to gain control over key inputs at all levels of the AI stack. These firms control access to resources (data, chips, compute, data center infrastructure, foundational models), the talent to make the most of these resources, and the ability to monetize this incredibly expensive business through their control over distribution channels and their unparalleled access to the consumer. Even where new entrants are able to enter the market, they are dependent on the resources controlled by dominant firms in order to succeed, creating an ecosystem of dependence rather than competition.
- **Big Tech firms do not just control access to inputs in the AI market: they control access to consumers.** These firms, which already benefit from outsized access to key inputs, can capture value downstream by pushing their own AI products onto customers or by controlling the platforms on which new AI technologies will roll out. For example, Google is able to push its generative AI products into existing products like Search. Microsoft is able to leverage its control over the enterprise market by pushing its AI assistant product, Copilot, across its work apps. Amazon can push generative AI products across its shopping platforms. A new entrant may find it all but impossible to compete with entrenched players who can easily integrate AI into existing products—whether customers ask for them or not.
- **Capital structures financing tech incentivize monopoly creation as a feature rather than a bug.** A major driver for consolidation in the tech industry is venture capital (VC). The VC funding system is driven by the idea of the "power law," meaning that one hugely successful company could hedge against the losses generated by the rest of a failing portfolio. To ensure one's company becomes the next multi-billion dollar unicorn, massive growth (and eventually monopoly consolidation) is the goal. This creates a market dynamic where VC companies are choosing winners and losers, rather than allowing good business ideas to compete on the merits. VC firms push startups to continue raising more money and taking on more risk, rather than settle on a sustainable model, because VC firms prove their value to institutional investors by pointing to the size of subsequent funding rounds their startups raise. In baseball terms, VC ignores potential doubles and triples in an all-out pursuit of grand slams. The incentives of VC also push hard for companies to be acquired, putting massive pressure on startups to be swallowed up and consolidated into larger companies—and thus fueling the problem of market concentration.

Corporate consolidation and lack of diversity in the capital ecosystem have led to a crisis of innovation

Lack of competitive pressure does not fuel innovation: it stalls it. The "AI race" dynamics have already stalled the publication of new technical AI research, and the lack of industry-wide competitive pressure is generating a narrow set of similar products all driving to the same result.

- **Unparalleled spending on AI infrastructure without a pathway to profitability is a road to crisis rather than innovation.** A recent report from Bain & Company estimates that $2 trillion in annual revenue will be required by 2030 to make the AI infrastructure build-out profitable. Given the need to monetize as quickly as possible in order to justify the massive investments in build-out, companies are defaulting to known, extractive business models and repackaging existing technologies. For example, OpenAI's recent decision to start placing ads on the low-cost versions of ChatGPT points to the pressure investors are putting on the AI model builders.
- **Depletion of public research creates opportunity costs for innovation and solving public problems.** Over the past decade, AI's research trajectory has been captured by corporate influence, with more and more critical AI research taking place in the research labs of private companies. Private industry has siphoned researchers from their independent academic

institutions at an alarming rate, with 70% of PhD researchers in the field taking jobs in corporate labs as opposed to 20% just two decades ago. Big tech firms are able to provide lucrative compensation packages and unrestricted access to the financial and technical resources that AI companies own and control in-house, limiting the incentives for researchers to conduct research at public institutions that lack similar benefits and resources.

> **"HOW TECHNOLOGY IS USED IS ALWAYS INTERTWINED WITH THE VISION AND INTERESTS OF THOSE WHO HOLD POWER."**
>
> – Daron Acemoglu & Simon Johnson in Power and Progress

- **Scaling methods adopted by private industry are extractive and resource-intensive.** The development of AI has centered around the pursuit of one hypothesis at the expense of all others: that scale is the only path to progress. As a result, all of the hyperscale model builders are pursuing AGI by consuming an ever-increasing and unsustainable amount of resources. This approach requires consuming training data without regard to ownership and natural resource extraction, which exacerbates the climate crisis. This "bigger is better" paradigm has also crowded out opportunities to pursue other hypotheses that may be less resource-intensive, or to pursue smaller, more discrete models built for specific purposes.

Concentration threatens California's fiscal and economic stability, putting everyday people at risk

The growing concentration of capital within a few firms in the AI market poses a significant risk to California's economy. In 2024, it was widely reported that a single company—Nvidia—was suspected to have played a major role in closing California's budget deficit. The gains attributed from AI firms' soaring—but unstable—valuations are masking an otherwise bleak economic picture: California is banking on personal income tax windfall while failing to fund gaps in critical access to services, such as Medicaid, impacting millions of low-income Californians' access to healthcare. In the event of a market correction, California's budget crisis could turn catastrophic.

Such a market correction would also hit the pocketbooks of everyday Californians. As our economy is increasingly concentrated in AI gains, everyday people are increasingly exposed to financial risk. Tech companies comprised over 80% of U.S. stock market gains in 2025, and stock market ownership represents nearly 30% of the net worth of American households.

The perils of tech concentration also threaten our economy's broader security and stability. This structure creates single points of failure for security breaches and devastating effects during outages, such as the July 2024 CrowdStrike software outage that crashed over 8.5 million Windows devices, including those belonging to hospitals, airlines, banks, and government services.

Californians' future, and our state's budget, depend on our ability to build a more resilient tech economy in which opportunity, innovation, and prosperity are more widely distributed.

A roadmap for prosperity in the age of AI

California thrives when its people have the opportunity to turn their vision and ambition into reality. Achieving what everyday people want from technology—scientific breakthroughs, solutions for climate change, better health care, and a more affordable life—requires boldness from the state's leadership, who will take the helm at the dawn of a new age. California's elected leaders must adopt an agenda that:

- reshapes our markets towards competition and fairness,
- advances a bold vision for a publicly funded tech ecosystem, and
- reform capital markets to catalyze entrepreneurship.

Foster competition to unleash real innovation

Competition policy to build California's future

Competition policy has long played an important role in shaping California's tech economy. A series of regulatory actions against large monopolies in the 1970s-1990s—interventions at such companies as AT&T, IBM, and Microsoft—broke open the innovation ecosystem and enabled a range of new technologies to thrive. In fact, the collective actions against Microsoft's monopoly in the late 1990s created space for new players to enter into the search and browser markets, namely a small startup that was barely on the radar screen when the federal antitrust case was initially decided in 1998: Google.

Google and the slew of tech companies that were birthed in its wake would likely not have been possible without antitrust enforcement, creating a more competitive environment for internet search. Twenty years later, history is poised to repeat itself as Google, Meta, and Amazon—companies now known as tech hyperscalers—face their own set of antitrust lawsuits and investigations for dominating various markets.

Innovation thrives in competitive markets. When no single player dominates, new ideas and opportunities have room to emerge—and the economic benefits are shared beyond the hands of a few powerful companies. To harness the full potential of AI and emerging technology—both in what these tools can do and in who gets to prosper from them—California must establish brightline rules that preserve a competitive ecosystem for its businesses, workers, and consumers.

- **Update California's Cartwright Act to address single-firm conduct:** In 2022, the California legislature requested that the California Law Review Commission (CLRC) conduct a study on California's antitrust law to better understand if California's existing laws were sufficient to the challenge of ensuring competition amongst a growing set of unprecedented technology titans. After three years of extensive work, the CLRC voted in January 2026 to recommend that California's legislature modernize state antitrust law to clearly ban conduct by a single entity that restrains competition or has a harmful effect on its overall and connected markets. These changes would include holding dominant AI platforms accountable when they block competitors, self-preference their own services, or leverage power in one market to squeeze out rivals. This recommendation was supported by a wide range of consumer advocacy groups, economic justice organizations, proponents of tech start-ups such as Y-combinator, and AI safety organizations. In addition, the CLRC is slated to provide recommendations on additional antitrust issues, including mergers and acquisitions.
- **Investigate and enforce robust merger review** to scrutinize mergers, acquisitions, and partnerships, including those involving Big Tech hyperscalers. A comprehensive review of these types of deals would include restrictions or pre-merger review of Big Tech companies that enter into partnerships with AI startups—including those that may look unconventional on their face like credit arrangements where an investment includes a commitment to exclusive use of that investor's services (e.g. a $12B Microsoft's investment into OpenAI that includes "cloud credits" and exclusivity agreement to use Microsoft Azure as its cloud provider), or takes the form of acqui-hires where a company can hire competitors' talent and bypass them as a competitor through a licensing agreement that avoids triggering a mergers and acquisitions review. Reviews should also take a look at potential conflicts of interest through cross-board representation across Big Tech firms and AI startups and dataset acquisitions that could further market concentration.
- **Pass bright-line rules to prohibit anti-competitive market behavior before it can occur.** To prevent the concentration of power in a small number of dominant technology companies from distorting the market, California must investigate potential competitive harms, including AI partnerships and investments, self-preferencing, and blocking rivals' access, and establish clear rules governing how cloud providers and AI developers treat their customers and competitors. Cloud providers should be prohibited from self-preferencing or offering more favorable terms to similarly-situated customers. They should be blocked from using data obtained through hosting customer infrastructure to develop competing products and from penalizing customers across their ecosystems for building competitive products and services. Foundation-model companies should be subject to the same self-preferencing prohibitions and additionally required to ensure their models are interoperable across cloud infrastructure providers, so that no single platform becomes an inescapable bottleneck for the businesses and developers who depend on them.
- **Build a genuinely pro-competitive economy in California using a full toolkit—not a single lever.** We must take a multi-faceted approach to encouraging competition and innovation. That means ensuring that workers, consumers, and businesses each have meaningful, accessible rights to bring complaints, raise concerns, and pursue legal action when market power is being abused. It also means robustly funding the California Department of Justice with the talent, resources, and clear mandate to pursue an affirmative competition agenda.

Build anew

Set a bold vision for public innovation

Our tech ecosystem is in need of a paradigm shift that allows for the emergence of new approaches to innovation. We are at a time when students, researchers, and entrepreneurs across California are feeling more energized than ever by the immense potential of technology to tackle our state's most pressing challenges, from revolutionizing healthcare and education to transforming our transportation systems and tackling climate change.

This prompts the question: what could AI look like if it were guided by public interest objectives and non-market imperatives, freed from private industry's dominance over the AI stack?

There is a strong appetite across civil society, government, and academia—supported by the public writ large—for public investment in the core infrastructure and research pathways that can fuel tech innovation outside of private AI companies. In 2025, California passed a law mandating the development of CalCompute, a publicly owned-and-operated cloud computing cluster to be established within the University of California system that fosters "research and innovation that benefits the public." CalCompute presents California with an opportunity to open up space and funding for alternative research trajectories and invest in technological development that is focused on solving important, public problems in ways that are not currently incentivized by the market. CalCompute's ultimate success will require leadership from the next Governor to see it through.

Now is the time for California to reclaim ownership over the direction of technological innovation in our state and ˙beyond, inspiring the next generation of big ideas and divesting California's tech future from the incentives of private firms.

- **Invest in public AI infrastructure and alternative research trajectories that** reduce dependence on Big Tech infrastructure. California should fully fund CalCompute to provide affordable AI compute resources for small businesses, researchers, and startups. The program should incentivize AI research trajectories outside the large-scale, "bigger is better" scaling paradigm, including smaller-scale models, alternative ownership structures (such as cooperative AI and data coops), federated AI, and more.
- **Launch a California "Grand Challenges" program that directs AI towards public missions,** defining a series of public benefit objectives that are set by the public, mobilizing public resources to drive technological breakthroughs that tackle the toughest challenges facing our state.
- **Place enforceable, binding conditions on all public AI investment** to hold public AI investment accountable to public-oriented goals. These include enforceable resource limits to curb the harmful environmental effects of AI and processes to ensure structurally disadvantaged groups have deliberate authority in the development of public infrastructure.

Catalyze innovation

Overhaul capital systems to enable more companies to thrive

To fund a more robust and competitive innovation ecosystem, the state can catalyze different methods of allocating capital to startups and innovators.

California already has a vehicle for public financing of infrastructure and private development in alignment with the state's values and priorities—the California Infrastructure and Economic Development Bank (IBank), which was created in 1994. As part of the American Rescue Plan in 2021, the state utilized new federal funding to establish the Expanding Venture Capital Access program within the IBank. The VC Access program, which will invest an initial $250m in VC funds and startups, presents a unique opportunity for the state to not only fund a wider range of businesses but to encourage private investors to move away from their steadfast commitment to producing "grand slam" results and instead build portfolios around doubles and triples. By doing so, California can shape capital markets to encourage investments in the kinds of businesses that are building AI to address the use cases the public desires—while at the same time creating more dynamism and sustainability in the economy.

While the VC Access Program is a promising effort, it can be leveraged to catalyze more private capital to seed a robust innovation ecosystem. California leadership can direct the IBank to:

- **Expand the VC Access Program to include later-stage companies.** Many startups with mid-sized potential have a relatively easy time finding early-stage capital, but then struggle to find investors who can provide growth capital. Elected leaders should consider increasing the capital allocated to the VC Access Program so that it can fund startups beyond their earliest stages.
- **Evangelize its success to attract private investors.** The biggest challenge with developing new modes of venture capital investing is lack of awareness among the startup community. Venture capital funds are notoriously opaque, making it hard to propagate data that demonstrates other approaches can work. The next Governor should direct GO Biz to develop partnerships that can amplify the learnings and successes of the VC Access Program in order to socialize a broader set of methodologies for investing in innovation.
    - These partnerships can and should include research partnerships that encourage private fund managers to share data with researchers who can provide wide-scale analysis into the performance of different fund structures.

- **Leverage the state's assets to fund more sustainable growth.** In addition to being a fund manager itself through the VC Access Program, the State of California is also an institutional investor in many private venture capital funds. The next Governor and legislative leadership should commission an analysis of the State's various institutional investments in venture capital and make recommendations for how California could lead institutional investors in moving capital to fund managers who take a "doubles and triples" approach to investing. While many of these institutional funds are not managed directly or indirectly by elected leaders (the funds are, for example, held in pension plans or university endowments), elected leaders can use their leadership to encourage these institutional fund managers to adjust their thinking.

The opportunity ahead

Throughout history, our government has driven technological breakthroughs by setting ambitious targets for our most pressing challenges—from the federal government's formation of the Advanced Research Projects Agency in 1958, which became the backbone of the early internet, to the National Institute of Health's sustained funding for exploratory medical research that delivered the mRNA vaccine when we needed it most. These public investments for the public good have produced some of the last century's most impactful scientific advances.

Today, the growth of artificial intelligence offers a once-in-a-generation opportunity to transform our work, economy, and society. Unfortunately, the infrastructure, research, and resources driving this change are held almost entirely in private hands with incentives aimed at maximizing profit rather than public benefit. California's major tech companies—Google, Meta, Apple, Nvidia—have grown to eye-popping valuations enabled by the infrastructure, resources, and culture provided by the state of California, from our world-class public university system to the heavily, publicly-subsidized medical device and military research facilities that line our coasts. These companies and their products exist because of the public's century-long investment in innovation, people, and technology.

If we want AI to unlock the same broad-based dynamism the internet once did, we need policy that puts society's needs first—not the balance sheets of a few tech companies.

California's next leader has a once-in-a-generation opportunity: to ensure this state isn't just remembered as where the AI companies were born, but one where we proved advancements in technology could benefit everyone. This means imagining a new and different trajectory—one that seeds opportunities broadly, rather than concentrating them into fewer and fewer hands, and that treats our public challenges as problems all Californians are both invested in and equipped to solve together. The question isn't whether AI will reshape our future; it's what kind of future do Californians want?



Find us at techequity.us

Reach out at info@techequity.us