Executive Summary

California stands at the epicenter of the artificial intelligence revolution, home to both the companies building frontier systems and the workers most exposed to disruption. In Sacramento, a new wave of executive orders and bills signals that state leaders now see AI-driven job displacement not as a distant threat but as an active policy problem demanding a coherent response.

This report argues that framing Sacramento's dilemma as a binary choice between "protecting jobs" and "softening the blow" is a false one. The real question is how to construct a social contract in which the state sets hard guardrails on the pace and form of displacement, while aggressively investing in worker power, income supports, and pathways into AI-complementary work. That means combining ex-ante protections (like notice requirements, limits on automated management, and collective bargaining over AI deployment) with ex-post cushioning (like retraining, portable benefits, and sectoral transition funds).

The Scale of the AI Displacement Threat

Global and National Projections

Major studies suggest that AI and broader automation could displace a non-trivial share of the global workforce by 2030, though estimates vary widely. The McKinsey Global Institute has estimated that up to 15 percent of workers worldwide—around 400 million people—could be displaced by automation between 2016 and 2030 in a midpoint scenario, with faster adoption scenarios pushing displacement as high as 30 percent. In the United States alone, McKinsey has suggested that tens of millions of jobs could be lost or transformed, with

midpoint scenarios showing 39 million U.S. workers affected and rapid automation scenarios up to 73 million.

The pattern is not uniform across occupations. Routine, predictable tasks—whether manual or cognitive—are significantly more automatable, while jobs requiring high interpersonal skills, creativity, or physical dexterity in unstructured environments remain more resilient. This asymmetry means that office support, some customer service roles, and parts of transportation and logistics are acutely exposed, while care work, many trades, and complex professional roles may be reshaped rather than fully displaced.

Why California Is Especially Exposed

California is particularly vulnerable because it combines a large concentration of tech firms building and deploying AI with a labor market heavily populated by both knowledge workers and low-wage service workers. The state's 2021 Future of Work Commission report highlighted how rising automation and digitalization intersect with inequality, warning that low-income, Black, Latino, and immigrant workers are more likely to be in jobs at risk of automation while having less access to training and safety nets.

Recent events have made this abstract risk concrete. California has already seen AI-linked layoffs from major firms, including Meta, Oracle, and UPS, as companies cut tens of thousands of jobs while simultaneously investing billions in AI infrastructure and tools. These layoffs are often justified as "efficiency gains" or shifts toward AI-enabled business models, but for displaced workers the immediate reality is income loss and uncertain prospects for reemployment.

Sacramento’s Emerging AI Policy Architecture

Executive Orders and Strategic Studies

In 2023, Governor Gavin Newsom signed Executive Order N-12-23 directing state agencies to study the development, use, and risks of generative AI across state government, including its impact on the public-sector workforce. The order called for inventories of AI tools used by state agencies, development of evaluation frameworks, and training programs to equip state workers with skills to use AI responsibly and equitably.

By 2026, as AI-linked layoffs mounted, Newsom went further, issuing a first-of-its-kind executive order instructing agencies to “anticipate, measure, and mitigate” AI-driven job losses statewide. This newer order calls for a comprehensive assessment of which industries and occupations are most vulnerable, recommendations for training programs, and ethical guidelines for AI use, effectively making California a testing ground for proactive AI labor regulation.

Legislative Front: From Data to Guardrails

On the legislative side, Sacramento lawmakers are now moving beyond abstract hearings to concrete bills targeting AI’s workplace impacts. Assembly Bill 2545, the AI Deployment & Workforce Displacement Assessment bill authored by Assemblywoman Pilar Schiavo, would create an advisory panel combining workforce organizations, workers, industry, and experts to systematically assess how AI deployment is affecting job loss, industry disruption, and what policies and funding are needed for displaced workers. The bill, which has passed the Assembly and is pending in the Senate, is explicitly a data-gathering and agenda-setting

mechanism rather than an immediate brake on AI deployment.

In parallel, Senator Eloise Gómez Reyes has introduced SB 951, modeled on the state WARN Act, to require employers to provide at least 90 days' notice to workers and the Employment Development Department before conducting mass layoffs attributable to AI. Backed by the California Labor Federation as the "AI Job Killer Notice Act," SB 951 is emblematic of a strategy that does not ban AI-driven layoffs but aims to slow and surface them, giving time for intervention and support.

Other bills, such as versions of the "No Robo Bosses Act" (SB 947), seek to restrict employers from using AI systems to discipline or fire workers without human oversight and to bar certain predictive or discriminatory uses of AI in employment. These efforts focus less on displacement per se and more on preventing AI from becoming an unaccountable manager, recognizing that how AI is used in the workplace can be as consequential as whether it eliminates jobs.

The Legacy of AB 5 and Proposition 22

California's recent history with gig work—via Assembly Bill 5 and Proposition 22—casts a long shadow over the AI debate. AB 5 sought to reclassify many independent contractors, especially in platform work, as employees entitled to full labor protections, positioning California as a pioneer in tackling "false self-employment" in the digital economy. Gig platforms like Uber, Lyft, and DoorDash mounted a massive backlash, culminating in Proposition 22, which carved app-based drivers out of AB 5's framework and preserved their contractor status with a more limited set of protections.

This episode illustrates both California's willingness to push aggressive worker protections and the power of tech-driven business models to fight back through lobbying and ballot initiatives. It also foreshadows how AI policy fights may play out: organized labor pushing for

strong guardrails on displacement and algorithmic management, while tech firms and some investors warn that stringent rules will weaken California's innovation edge and push companies to friendlier jurisdictions.

Stakeholders and Their Competing Narratives

Labor Unions and Worker Advocates

California labor unions have moved rapidly to treat AI as a core bargaining issue rather than a distant technological curiosity. At a landmark conference in Sacramento, more than 200 union members and technologists gathered to strategize how to confront AI and other technologies that threaten workers, emphasizing workers' rights to negotiate over AI deployment, create joint tech committees, and use AI to strengthen organizing.

Unions and allied groups are backing bills like SB 951 and SB 947, arguing that workers must have notice, transparency, and bargaining power before AI reshapes or eliminates their jobs. Their narrative frames AI policy as a continuation of long-standing fights over workplace surveillance, just-cause protections, and offshoring—only this time the "offshore" is a model in a data center.

Tech Firms and Industry Advocates

Technology companies, for their part, emphasize AI's potential to raise productivity, create new kinds of jobs, and keep California at the forefront of global innovation. They often acknowledge displacement risks but argue that overly rigid regulation, especially ex-ante bans on certain uses, will push development out of California, jeopardizing tax revenue and long-term job growth.

The industry narrative leans heavily on historical analogies in which technology destroyed some occupations but generated new ones and raised living standards overall, from mechanized agriculture to the computer revolution. However, critics point out that those transitions were often brutal for specific communities and that the pace and scope of contemporary AI adoption may outstrip past adjustment capacities, especially without robust safety nets.

State Leaders: Walking a Tightrope

California's political leadership is pulled between these poles. On one hand, state leaders want to cement California's role as the global hub of generative AI, which promises high-value investment and prestige. On the other hand, they face a state where inequality is already extreme and housing, healthcare, and education costs leave many workers one layoff away from crisis.

The executive orders and bills emerging from Sacramento reflect this tension. They rarely prohibit AI outright but focus on studying impacts, increasing notice and transparency, and experimenting with training and safeguards. The question is whether this strategy will be enough to shape AI's labor impact, or whether it amounts to documenting damage while hoping that market-driven job creation eventually catches up.

The False Binary: "Protect Jobs" vs. "Soften the Blow"

Why Pure Job Protection Is Unrealistic

A pure "job protection" strategy—trying to freeze the current occupational structure by

prohibiting AI from displacing workers—runs into both economic and political constraints. Technologically, many AI applications genuinely do certain tasks faster, cheaper, or more accurately than humans, especially in data-heavy back-office work, logistics optimization, and some customer service. Even if California attempted to ban such uses, firms could shift operations or development to other states or countries while still serving California customers, limiting the state's leverage.

Legally, broad bans on AI-enabled efficiency measures could be challenged as interfering with interstate commerce or as arbitrary, pushing courts to scrutinize the state's rationale. Politically, a hard prohibitionist stance risks eroding support among voters who benefit from cheaper, more reliable services and among localities courting AI-related investment.

Why Passive Cushioning Is Inadequate

At the same time, a purely reactive "soften the blow" approach—accepting displacement as inevitable and focusing only on retraining and unemployment insurance—fails on both moral and practical grounds. Evidence from past transitions shows that many displaced workers, especially older workers and those without college degrees, struggle to regain previous earnings levels even years after job loss, and some never fully recover.

In high-cost California, where housing and living expenses are among the highest in the nation, prolonged earnings losses translate quickly into housing insecurity, health stress, and intergenerational harm. A policy posture that simply tracks which communities are hit hardest and offers training vouchers after the fact risks deepening existing racial and regional inequalities, particularly in inland and rural counties that are already economically fragile.

The Case for "Managed Acceleration"

The alternative is what might be called "managed acceleration": allowing AI-driven productivity gains to proceed but subjecting them to democratic planning, strong labor standards, and social insurance robust enough to make transitions survivable and, ideally, upwardly mobile. This model accepts that some jobs will vanish or shrink but insists that the timing, pace, and distribution of those losses are political choices, not neutral outcomes of technology.

Under managed acceleration, Sacramento's role is not to decide whether AI happens, but to decide how its costs and benefits are shared—who gets to say when AI is deployed, who bears the risk, and who captures the gains.

Policy Levers on the "Protect Jobs" Side

Notice, Transparency, and Worker Voice

California is already experimenting with notice-based approaches that give workers and policymakers time to respond to AI-driven changes. SB 951's 90-day notice requirement for AI-related mass layoffs effectively extends the WARN Act into the AI era, buying time for unions, workforce agencies, and local governments to mobilize support, negotiate alternatives, or attract replacement employers.

Beyond notice, laws like the proposed "No Robo Bosses Act" would give workers a say in how AI is used to manage and evaluate them, requiring human oversight for disciplinary decisions and restricting intrusive or discriminatory analytics. Contracts negotiated by unions are also starting to address AI directly, with demands for joint committees and veto power over certain forms of automation—a model Sacramento could encourage through public-sector

bargaining and procurement rules.

Substantive Limits on Harmful Uses

Sacramento can also set substantive limits on particularly harmful applications of AI, even while allowing others. For instance, the state could bar AI from being used to infer protected characteristics (like race or union status) or to make high-stakes employment decisions without meaningful human review, building on provisions in SB 947 and related bills. Strong enforcement and meaningful penalties would be essential to ensure that such limits are not merely symbolic.

Another avenue is linking AI deployment to minimum labor standards. The state could require that firms using AI at scale meet thresholds for wages, benefits, and job quality, effectively preventing companies from using automation as a way to dodge obligations rather than to complement well-paid work. This echoes debates around AB 5, where concerns about misclassification intertwined with fears that technology firms were exploiting regulatory gaps to undercut traditional employment standards.

Public-Sector Leadership by Example

As one of the state's largest employers, California's government is itself a major site of AI experimentation. Executive Order N-12-23 already calls for training state employees to use generative AI and for evaluating its workforce impacts. Sacramento can go further by committing not to use AI to justify net headcount reductions in critical public services without prior worker consultation, impact assessments, and plans for redeployment.

By embedding AI in ways that enhance public workers' productivity and reduce drudgery—rather than simply cutting positions—the state can model a "complementarity

first" approach that it expects private employers to emulate.

Policy Levers on the "Soften the Blow" Side

Robust, Targeted Retraining and Education

No realistic policy mix avoids the need for significant retraining and upskilling as AI reshapes occupations. Newsom's executive orders explicitly call for state agencies to recommend training programs focusing on digital literacy and AI-assisted skills, as well as to prepare workers for roles in the "GenAI economy." However, what matters is not only the existence of training but its quality, targeting, and integration with real employer demand.

California can use insights from the Future of Work Commission and partnerships with institutions like the Stanford Digital Economy Lab to design sector-based training that connects displaced workers to growing fields such as clean energy, healthcare, and AI support roles. This requires sustained funding, wraparound services (like childcare and housing assistance), and metrics that track not just course completions but long-term earnings outcomes across demographic groups.

Income Supports and Portable Benefits

When displacement happens, the short-term question is whether workers can keep a roof over their heads while they retrain or search for new jobs. Existing unemployment insurance and safety-net programs in California provide some cushion but often fall short in duration and generosity relative to high local costs of living. Sacramento could pilot wage insurance programs that temporarily top up earnings for workers who accept lower-paid jobs after displacement, smoothing transitions while incentivizing reemployment.

The state could also revisit the idea of portable benefits that attach to workers across employers and employment types, ensuring that health care, retirement savings, and other protections do not disappear when a job does. In an AI-intensified labor market where more workers may cycle through multiple roles and short-term contracts, such portability becomes even more critical to "soften the blow" without locking workers into obsolete positions.

Regional and Sectoral Transition Strategies

AI-driven job loss will not be evenly distributed. Some regions—such as inland logistics hubs or office-park corridors dominated by back-office work—may see concentrated displacement, while others may gain from AI-driven growth sectors. California's response must therefore include regional transition strategies, potentially funded by state-level adjustment funds supported by taxes on AI-related corporate windfalls.

Sectoral strategies, coordinated with industry, unions, and educational institutions, can proactively map which roles are likely to shrink and which will grow, then plan apprenticeship pipelines and mid-career transitions accordingly. The advisory panel envisioned in AB 2545 could serve as a hub for such planning if given sufficient authority, data, and resources.

Lessons from AB 5: Guardrails Without Backlash

The AB 5 saga offers a cautionary tale about how not to structure ambitious labor protections in the face of powerful tech interests. By designing a sweeping reclassification that many small businesses and freelance workers experienced as rigid and ill-fitted to their needs, lawmakers created fertile ground for a well-funded industry campaign that culminated in Proposition 22, rolling back much of the reform.

For AI policy, this suggests that Sacramento must combine firm guardrails with flexible implementation and clear communication. Policies like notice requirements, human oversight of AI management, and non-discrimination rules are easier to defend than blanket bans because they are framed as due process and fairness measures rather than outright constraints on innovation. Where more aggressive measures are necessary—for example, prohibiting AI use in certain high-risk domains—they must be accompanied by tangible benefits for workers and communities, such as publicly funded alternatives or direct support for affected business models.

Another lesson is to expect and plan for ballot-box battles. Tech companies have already shown their willingness to take fights directly to voters when they dislike legislative outcomes, and AI-related rules are unlikely to be an exception. Building broad coalitions that include not only unions but also community groups, small business associations, and local governments will be essential to sustaining a robust AI labor regime.

A Normative Framework for California's AI Social Contract

Centering Worker Power and Democratic Control

At its core, the question "protect jobs or soften the blow" is about power: who gets to decide how AI is deployed, and whose interests are prioritized when trade-offs arise. A just AI social contract for California requires that workers and affected communities have real power over these decisions, not just a seat in advisory panels.

This implies strengthening collective bargaining rights, supporting unionization drives in AI-heavy sectors, and embedding worker representation in bodies that set AI standards and

oversee public deployments. It also means giving workers rights to data visibility—knowing what systems are used, what data they collect, and how outputs affect their jobs—as a prerequisite for meaningful contestation and consent.

Sharing AI's Productivity Gains

AI's promise is that it can dramatically raise productivity. The danger is that, without deliberate policy, the gains will accrue overwhelmingly to capital owners and a narrow slice of highly skilled workers, while many others see stagnant or falling real earnings. Sacramento can intervene on the distributional margin through tax policy—capturing a portion of AI windfalls for public investment—and through wage and hour laws that ensure AI is used to enhance jobs, not merely to intensify work or cut staff.

More creatively, California could explore models like worker data trusts or public AI infrastructure, where the value generated by large datasets and foundational models feeds into collective goods rather than only private profits. Such approaches would move beyond both job protection and cushioning toward a proactive reshaping of who owns and benefits from AI.

Embracing Innovation With Conditions

The upshot is that Sacramento should neither romanticize existing jobs nor treat displacement as an unavoidable byproduct of progress. Many current roles are low-paying, insecure, and physically or mentally draining; in some cases, AI can and should eliminate the worst work while creating better alternatives. The challenge is to ensure that the workers doing today's worst jobs are first in line for tomorrow's better ones, rather than being discarded as expendable.

That means embracing AI innovation—but only on the condition that it comes with binding commitments: to transparency, to fair treatment, to redistribution of gains, and to genuine worker participation in design and deployment. In this sense, "protecting jobs" and "softening the blow" are not competing goals but intertwined obligations within a broader project of democratic governance over transformative technology.

California's leaders in Sacramento are right to see AI-driven job displacement as both a risk and an opportunity. Early executive orders and bills show a state trying to move from passive observation to active governance, experimenting with notice laws, algorithmic accountability, and large-scale workforce planning. But the choices ahead will determine whether California becomes a model of managed, broadly shared AI prosperity or a cautionary tale of high-tech inequality.

The real policy dilemma is not whether to protect jobs or to soften the blow. It is whether Sacramento will have the courage to insist that AI serve human purposes, including decent work and economic security, even when that means confronting the interests of some of the very companies that define the state's global image. Getting this right will require a politics that treats workers not as collateral damage to be retrained after the fact, but as co-authors of California's AI future.