Tracking AI-Related Job Loss Using Unemployment Insurance Claims Data in California

BEN HYMAN, TILL VON WACHTER, DIANA HERRERA, ROOZBEH MOGHADAM, JACOB MORRIS, SWAPNIL MOTGHARE, AND KARA SEGAL

JUNE 2026

CALIFORNIA POLICY LAB

TABLE OF CONTENTS

UNIVERSITY OF CALIFORNIA

The California Policy Lab generates research insights for government impact. We are an independent, nonpartisan research institute at the University of California with sites in Berkeley, Los Angeles, and Sacramento.

This research publication reflects the views of the authors and not necessarily the views of our funders, our staff, our advisory boards, the California Employment Development Department, or the Regents of the University of California.

Executive Summary

Artificial intelligence (AI) has advanced at an extraordinary pace, generating both concern and optimism about its implications for the future of work. While generative AI tools only began reaching the public in late 2022, they have been integrated into the workforce with remarkable speed. For example, by 2025, hundreds of millions of people were using generative AI tools and more than a third of U.S. workers reported using generative AI to enhance productivity in their jobs.[1] While these technologies may deliver large productivity gains and complement human labor, they may also result in workers being displaced. Given the speed and potential scale of this transformation, policymakers need timely data to identify emerging trends of AI-related job losses and to inform safety-net and workforce development policies to support these workers.

The California Policy Lab (CPL), in partnership with the California Employment Development Department (EDD), is studying AI-related job loss by analyzing Unemployment Insurance (UI) claims combined with measures of occupational AI exposure. These measures estimate the share of job tasks that can be performed by modern AI and are an indirect proxy for actual AI use. They are widely used as the best available indicator to assess where AI is impacting the labor market. UI claims reflect laid-off workers who do not quickly transition into new employment and instead rely on temporary income support. UI claims data is used as a leading indicator of labor market strength or weakness by state and federal agencies, economic forecasters, and the general public.

This research combining UI claims data with AI exposure measures is synthesized in the newly released California AI-Unemployment Tracker (CAIT), which provides near-real-time data on AI-exposed UI claims in California, including detailed breakdowns by age, gender, race/ethnicity, education, industry and region. California is the first state to proactively track UI claims by AI exposure, and the CPL-EDD partnership enables timely monitoring to detect early signs of possible AI-induced job loss. The findings below draw on CAIT, updated through May 2026, and provide crucial input for potential policy solutions. CAIT will be updated monthly along with key takeaways. Its underlying tabulated data are available for public use.

1 "How People Use ChatGPT," (Chatterji et al., 2025)

Key Findings

- **Finding 1: No Evidence of Rising Statewide UI Claims in AI-Exposed Occupations**
  Since the release of ChatGPT-3.5 in 2022, statewide UI claims through May 2026 show no evidence of a surge in AI-related layoffs. The proportion of UI claims coming from AI-exposed workers has also not seen a statistically significant increase relative to prior to the pandemic. This is consistent with existing Current Population Survey estimates that show no nationwide link between AI exposure and unemployment.

- **Finding 2: UI Claims Increased for College-Educated Workers under AI Exposure**
  While AI-related job loss is not detected across all UI claims statewide, claims from college-educated workers in occupations with high AI exposure increased after ChatGPT-3.5's release in 2022, and remained elevated through May 2026, compared to low-AI-exposure workers who did not experience a noticeable change in trend, relative to their pre-pandemic level in 2019 and early 2020.

- **Finding 3: UI Claims Rose in the SF Bay Area and Technology Sectors**
  Over this same period (Nov 2022 - May 2026), claims from workers in occupations with high AI exposure in the San Francisco Bay Area also experienced a sharp and sustained increase. AI-exposed UI claims were also particularly elevated statewide in technology sectors such as *Information* and *Professional Services* over the same period. Claims levels from many of these subgroups remained elevated through May 2026, even relative to their pre-pandemic level in 2019 and early 2020.

Measurement validation: Contrary to the belief that AI-exposed workers are unlikely to file for UI, we find the UI program is widely used by workers whose prior occupations were exposed to AI. From 2023–25, 30.3% of initial California UI claims were filed by workers in occupations that had high AI exposure prior to their job loss. This suggests UI claims can be a useful data source for detecting the potential impact of AI on job loss.

Important context and ongoing monitoring: CAIT enables near real-time monitoring of AI-related UI claims, serving as a leading indicator of broader layoffs should they emerge in the future. Currently, the tracker shows no association between AI exposure and statewide UI claims. While the tracker reveals elevated UI claims among subgroups of tech-related, AI-exposed workers — levels that exceed their pre-pandemic baseline — these patterns are descriptive and cannot be conclusively attributed to AI alone. During the period studied, the technology sector experienced substantial post-pandemic shifts, including the unwinding of rapid hiring during the remote-work boom. While it is difficult to identify alternative explanations other than AI for patterns that stretch into 2025 and 2026, there may be other factors at play. As a result, the findings should be interpreted as an early, descriptive signal of potential AI-related disruption, rather than a full accounting of all AI-related job loss or causal evidence of the impact of AI on jobs. Although statewide UI claims do not yet indicate widespread AI-related job loss, persistently higher claims among AI-exposed subgroups through May 2026 suggest impacts may be starting to surface.

I. Measuring AI Exposure in UI Claims

This research introduces a new data tool, the California AI-Unemployment Tracker (CAIT) to monitor potential links between AI exposure and job loss in near real time. We connect information from UI claims in California to recently developed measures of "AI exposure," which capture the extent to which the tasks typically performed in an occupation can be carried out by modern AI. While these AI-exposure measures cannot directly determine an employer's motivation for layoffs, they are widely used and represent the current industry standard, providing the best available indication of where AI may affect jobs. Tracking AI exposure within UI claims in California provides a level of granularity and timeliness not available in existing data sources.

This analysis focuses on AI-related job loss that leads to UI claims. Eligible workers with sufficient prior earnings and who were laid off through no fault of their own, can receive weekly UI benefits up to 50% of their prior weekly earnings (up to $450 per week) for up to 26 weeks. Typically, many people who are laid off find new jobs without filing for UI (Bell et al., 2022). By focusing on job losses that subsequently result in UI claims, the tracker is designed to capture the subset of displaced workers who require social insurance or job search assistance — providing a timely signal of when AI-related job loss generates measurable labor market disruption.

Using UI Claims Data to Track AI-Related Layoffs

There is considerable discussion about how AI may lead to job loss, but tools to track these impacts in real time are limited. When a worker files for unemployment insurance, they select an occupation code from a drop-down menu to describe their pre-layoff job. We use this information to measure a UI claimant's AI exposure, defined as the extent to which tasks performed in that occupation can be performed by large language models (LLMs) like ChatGPT or Claude. We use two standard approaches to measuring AI exposure — one based on *potential* AI use and the other on *observed* AI use:

1. **Potential exposure measure:** This measure assesses whether LLMs are capable of reducing the time required to accomplish an occupation's task by at least 50%. This approach was developed by a team of OpenAI and academic researchers in the journal *Science* (Eloundou et al., 2024), and is widely used by AI researchers and considered an industry standard.

2. **Observed exposure measure:** This measure assesses how often workers use Anthropic's Claude platform to complete occupational tasks. This measure was developed as part of the Anthropic Economic Index (Handa et al., 2025), is intended to capture actual or *observed* AI use among Claude users, and can be updated as the capabilities of Claude evolve.

Each measure has strengths and weaknesses, but can be thought of as roughly representing the share of an occupation's tasks exposed to AI. The potential measure covers all occupations, but cannot definitively indicate whether AI is actually being used to complete tasks in those occupations. The observed measure, on the other hand, can describe whether Claude was actually used to accomplish an occupation's tasks. It is also regularly updated, capturing ongoing developments of LLMs' capacities. However, Claude's user base is skewed toward occupations in the technology sector (like software developers), reflecting AI use in selected industries among the overall market. We use both measures to build a more comprehensive picture. While we use the measure developed from Claude usage data released in Feb 2025, we confirm in the technical appendix that our results do not depend on this choice: our results remain unchanged when using a newer AI exposure measure released in March 2026 (Massenkoff and McCrory, 2026) that captures more recent advancements in Claude's capabilities.

We divide UI claimants into three groups based on the degree of AI exposure in their prior occupations. Those with "high" AI exposure are defined as having exposure scores in the top 25% of all scores. These workers were employed in occupations most susceptible to AI-related disruption. Those with "moderate" AI exposure (labeled "mid" in figures) have scores in the middle 25% to 75%. Individuals with "low" AI exposure have scores in the bottom 25%. Some examples of high-AI-exposure occupations are *customer service representatives* and *software developers*. Examples of low-AI-exposure occupations are *heavy truck drivers* or *nursing assistants*. We provide full methodological details behind these AI exposure measures and discuss robustness to using alternative percentile cutoffs to define "high" exposure in the technical appendix.

II. Main Findings from the California AI-Unemployment Tracker

In partnership with the California Employment Development Department (EDD), CPL developed the California AI-Unemployment Tracker (CAIT), an interactive tool that tracks initial UI claims[2] and measures the level of AI exposure in these claims. The tracker provides a regularly updated snapshot of potential AI-related labor market disruptions and is designed as a real-time monitoring tool.

UI Claims are a Relevant Data Source for Examining AI-Related Job Loss

Our analysis of UI claims data from 2023–25 shows that workers in AI-exposed occupations account for a large share of claims. This stands in contrast to the view that AI-exposed workers — who may have higher education and stronger reemployment prospects — would be less likely to take up UI benefits.

Most UI claims come from occupations with moderate to high AI exposure, especially among workers with bachelors, masters, or PhD degrees. Using the potential AI exposure measure, Panel A of Figure 1 shows that about 70% of all UI claims filed between 2023 and 2025 came from occupations with moderate or high exposure to AI. This compares to 68% from 2017–2019 which we report as a baseline pre-pandemic benchmark. Specifically, 30.3% of claims were in occupations with high AI exposure, while another 39.6% came from occupations with moderate AI exposure. In contrast, 30.2% of claimants were in occupations with low AI exposure, reflecting roles that are least susceptible to AI. Panel B shows that UI claimants with higher levels of education are also concentrated in occupations with greater AI exposure.[3]

2 Initial UI claims include "new" initial claims and "additional" claims. An additional claim occurs when an earlier claim was filed by the same individual, at least one week of certification was skipped because the claimant returned to work, and then the claim was subsequently re-opened before the benefit year expired.

3 See the tracker for data disaggregated by education level from 2017-2019.

FIGURE 1: Share of Initial UI Claims in California by potential AI exposure and education level

PANEL A: Share of total UI claims by ***potential AI exposure***

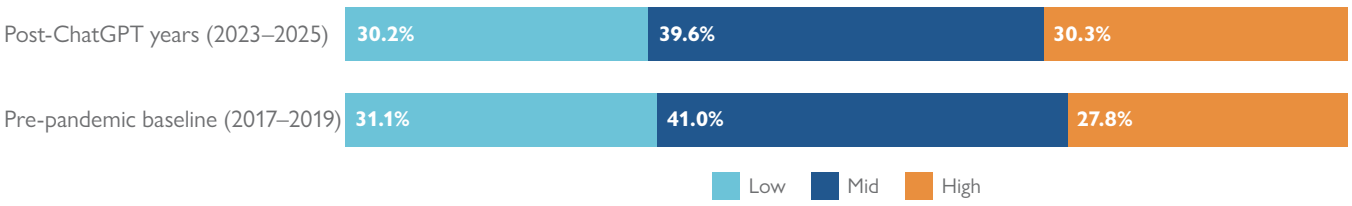

PANEL B: Share of new California UI claims from 2023 to 2025 by education level and ***potential AI exposure***

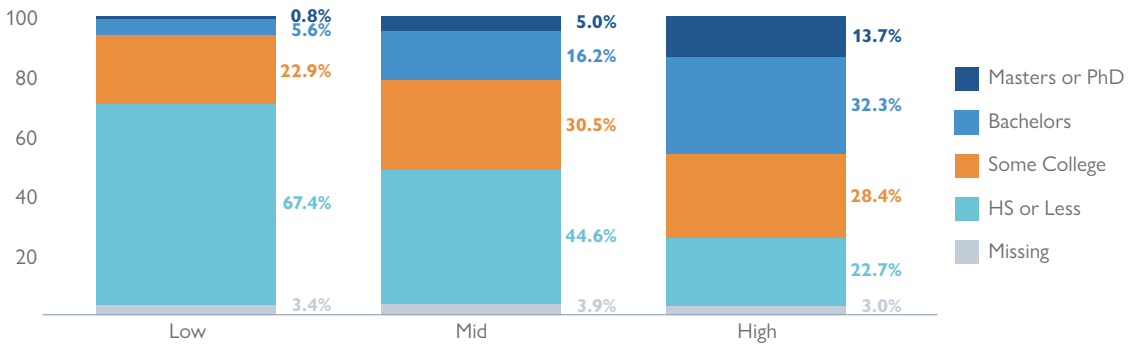

Notes: Claims are grouped by their potential AI exposure scores using the "beta" method from Eloundou et al. (2024). High-AI-exposure claims have AI exposure scores of 0.49 or higher (top 25%), moderate AI exposure claims have scores between 0.48 and 0.12 (middle 25%-75%), and low-AI-exposure claims have scores below 0.12 (bottom 25%).

The substantial share of UI claims from occupations with moderate to high AI exposure indicates that UI is a widely used support for AI-exposed workers, and thus a viable source of information on AI's impact on the workforce. However, the share of claims from AI-exposed occupations has remained broadly stable over time, including relative to the pre-pandemic period.

Finding 1: Since the Release of ChatGPT-3.5 in 2022, Statewide Data Through May 2026 Show No Evidence of Rising UI Claims from AI-Exposed Occupations

Figure 2 plots initial monthly UI claims from Jan 2019 through May 2026 for occupations in the low (bottom 25%), moderate (25–75%), and high (top 25%) AI exposure groups. The vertical dashed line in Nov 2022 marks the release of ChatGPT-3.5, a watershed event that sparked wider adoption of the technology, with some sources reporting overall business adoption rising from under 10% in 2022 to over 50% by Feb 2026. Panel A uses the *potential* exposure measure, while Panel B uses the *observed* exposure measure.

Panel A shows initial UI claims across the three AI exposure groups move roughly in parallel from 2019 through May 2026, with no strong divergence following the ChatGPT-3.5's release.[4] Both before the pandemic and after ChatGPT-3.5's release, there were a higher number of claims from occupations in the moderate exposure group, with fewer monthly claims from the low and high exposure groups. While high-AI-exposure claims exhibited an increase in mid-2022 and mid-2023, these changes fall within normal seasonal patterns in UI claims. Panel B uses Anthropic's usage-based measure, and shows similar findings; the number of claimants in each exposure group remains steady after ChatGPT-3.5's release, though which claims are tagged as "high," "mid," or "low," is different.

FIGURE 2: Overall Number of Initial UI Claims in California by AI Exposure Group

PANEL A: By ***potential AI exposure*** of claimants' prior occupation

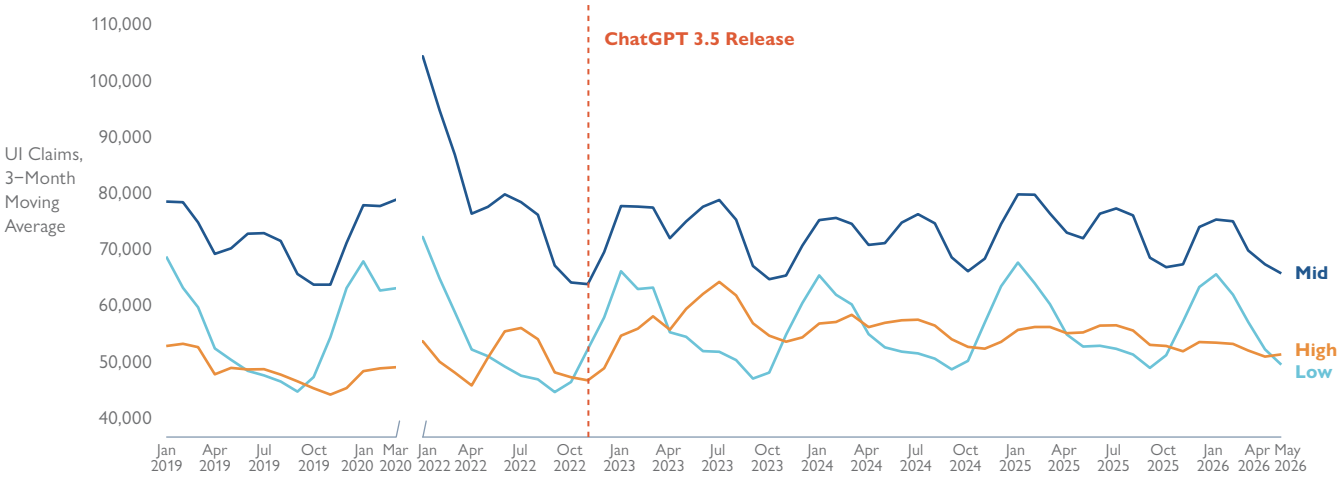

Notes: March 2020 to Jan 2022 is suppressed due to the unprecedented surge in UI claims during the pandemic, which would otherwise dominate the scale and obscure important comparisons between pre-pandemic and post-ChatGPT-3.5 trends. Dashed vertical line indicates the release date of ChatGPT-3.5. In Panel A, claims are grouped by potential AI exposure scores using the "beta" method from Eloundou et al. (2024). High-AI-exposure claims have AI exposure scores of 0.49 or higher (top 25%), moderate AI exposure claims have scores between 0.48 and 0.12 (middle 25%-75%), and low-AI-exposure claims have scores below 0.12 (bottom 25%). In Panel B, claims are grouped by observed AI exposure scores from Anthropic Claude use (Handa et al., 2025). Claims in the top quartile have Claude-based AI exposure scores of 0.107 or higher. Claims in the middle 50% have scores between 0.106 and 0.011. Claims in the bottom quartile have AI exposure scores below 0011. Data show the 3-month moving average (mean over the last three months) of UI claims. Claims data are through May 30, 2026.

4 The pandemic-era surge (not shown in the figures due to its outsized scale) appears as a sharp but temporary disruption.

PANEL B: By ***Observed AI Exposure*** of claimant's prior occupation

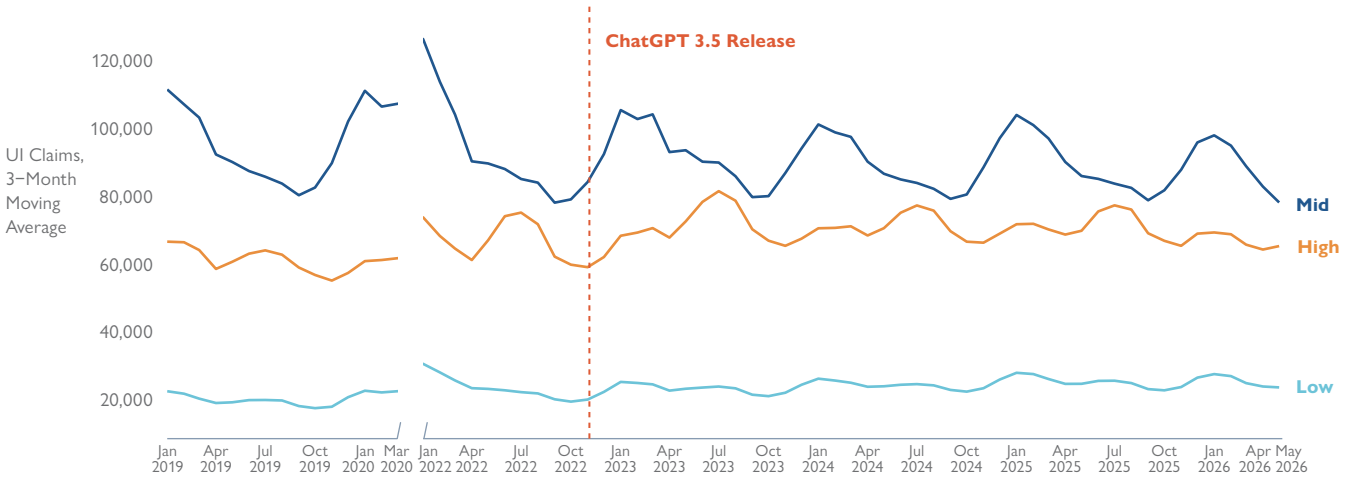

Notes: March 2020 to Jan 2022 is suppressed due to the unprecedented surge in UI claims during the pandemic, which would otherwise dominate the scale and obscure important comparisons between pre-pandemic and post-ChatGPT-3.5 trends. Dashed vertical line indicates the release date of ChatGPT-3.5. In Panel A, claims are grouped by potential AI exposure scores using the "beta" method from Eloundou et al. (2024). High-AI-exposure claims have AI exposure scores of 0.49 or higher (top 25%), moderate AI exposure claims have scores between 0.48 and 0.12 (middle 25%-75%), and low-AI-exposure claims have scores below 0.12 (bottom 25%).In Panel B, claims are grouped by observed AI exposure scores from Anthropic Claude use (Handa et al., 2025). Claims in the top quartile have Claude-based AI exposure scores of 0.107 or higher. Claims in the middle 50% have scores between 0.106 and 0.011. Claims in the bottom quartile have AI exposure scores below 0.011. Data show the 3-month moving average (mean over the last three months) of UI claims.

This absence of strong aggregate AI-related displacement effects is consistent with a recent Yale Budget Lab study that finds no significant trends when examining AI exposure among unemployed Current Population Survey (CPS) respondents (Gimble et al., 2025), as well as administrative-data evidence from Denmark showing that firms adopting AI do not significantly change employment patterns (Humlum and Vestergaard, 2025).

While CAIT at present does not show rapid increases in UI claims among occupations with high AI exposure, in the remainder of this report we provide evidence suggesting our measurement approach is able to pick up some AI-related job loss. First, we show that the slight statewide rise in UI claims among occupations with high AI exposure reflects some larger and more sustained increases in UI claims among salient subgroups after the release of ChatGPT-3.5.

Second, as an additional check that our instrument is detecting AI-related job loss, we examine major layoff announcements most plausibly linked to AI adoption and find clear spikes in UI claims among high-AI-exposure individuals following these events. This indicates that the tracker is sufficiently sensitive to detect AI-driven job loss when it occurs. However, such effects have not reached a scale large enough to be detected in aggregate, statewide trends.

CAIT also allows users to see how AI exposure among UI claimants varies across worker characteristics and industries. These include age, gender, race/ethnicity, education, industry and region in California.

This section describes AI-related UI patterns across selected subgroups. We begin by focusing on education cohorts, building on prior work showing that entry-level workers in more AI-exposed, "white-collar" occupations experienced disproportionate headcount reductions following the introduction of ChatGPT-3.5 (e.g., Brynjolfsson et al., 2025; Lichtinger & Hosseini Maasoum, 2025). While this earlier evidence identifies occupations where impacts are concentrated, prior studies either lack direct measures of worker education or do not observe longitudinal UI outcomes. By contrast, UI administrative records include worker education information and allow us to examine how AI-related displacement patterns vary across education groups. While earlier measures of changes in company headcounts combine layoffs, attrition, and slower hiring, our analysis of UI claims also focuses solely on layoffs, as measured by people who lose their jobs and file for benefits. Because it tracks new job loss rather than overall staffing levels, CAIT is likely to detect even temporary increases in unemployment. The full set of results — using both observed and potential exposure measures — are reported in the technical appendix.

Finding 2: UI Claims Increased for College Graduates in Occupations with High Potential AI Exposure

Figure 3 shows trends in initial UI claims by education level and AI exposure group, relative to the month of ChatGPT-3.5's release (Nov 2022). Expressing the data indexed to Nov 2022 allows us to interpret the percent growth in UI claims relative to the baseline of UI claims prior to the ChatGPT-3.5 release as a benchmark. Looking backward, the index also helps better assess to what extent UI claims in different groups were stable (or not) before ChatGPT-3.5's release.

Panel A shows that after ChatGPT-3.5's release, high-AI-exposure claims increased among more educated workers, peaked in July 2023, and remained elevated thereafter. Claims among workers with a bachelor's degree rose by over 50% from Nov 2022 to July 2023 (an index value exceeding 150), increasing in levels from roughly 13,000 to more than 22,000 claims per month. Compared to pre-pandemic trends, high-AI-exposure claims exceed previous claim rates well into 2025 and 2026. In May 2026, this group had an elevated claims level of about 16,000 per month. While higher educated groups also experienced increases in UI claims among occupations with low AI exposure, these were smaller in magnitude. Moreover, Panel B illustrates a noisier pattern among UI claims with low AI exposure, in part due to higher seasonality among those types of jobs.

Overall, these results suggest that job losses among highly educated workers have been more concentrated in high-AI-exposure jobs since ChatGPT-3.5's release. In the technical appendix, we provide event study estimates that quantify the relative differences in claims between high- and low-AI-exposure workers for different education groups. This direct comparison approach confirms that there was a persistent increase in UI claims from more educated workers with high AI exposure after ChatGPT-3.5's release in Nov 2022, compared to similarly educated workers in low-AI-exposure groups, and compared to lower educated workers.

As stated at the outset, while the claims are persistently higher, the extent to which they are driven by AI remains unclear, since other factors might be causing this specific pattern; this could include reduced hiring in Information Technology-related sectors after the pandemic, restrictive monetary policy, or fluctuating economic conditions.

FIGURE 3: Initial UI Claims in California by Education and ***Potential AI Exposure***

PANEL A: High-AI-Exposure UI Claims Indexed to ChatGPT-3.5 Release

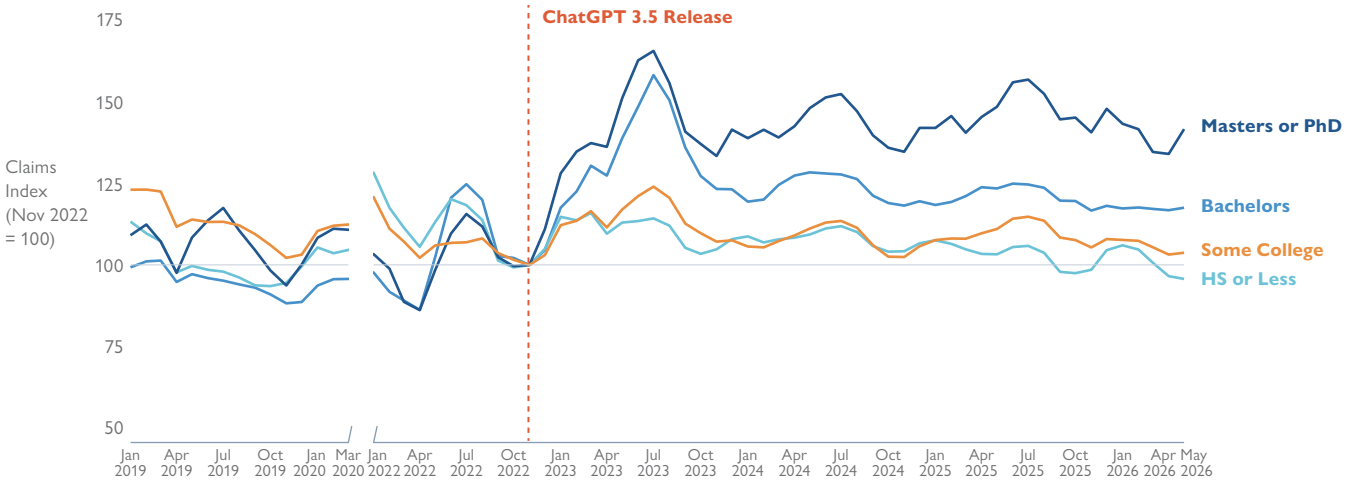

PANEL B: Low-AI-Exposure UI Claims Indexed to ChatGPT-3.5 Release

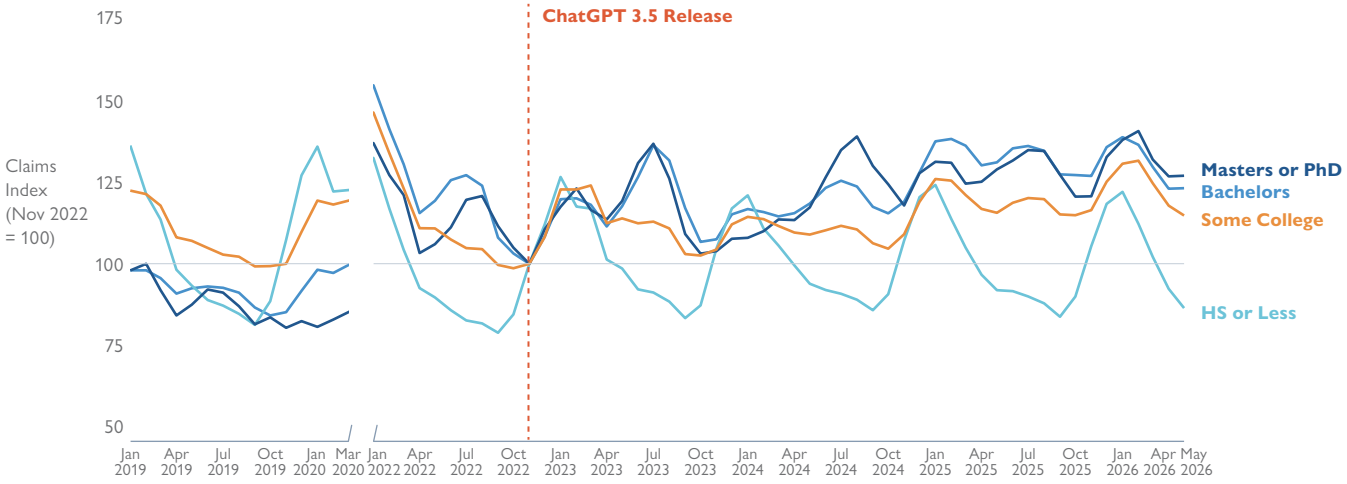

Notes: March 2020 to Jan 2022 is suppressed due to the unprecedented surge in UI claims during the pandemic, which would otherwise dominate the scale and obscure important comparisons between pre-pandemic and post-ChatGPT-3.5 trends. Dashed vertical line indicates the release date of ChatGPT-3.5. Claims are grouped by potential AI exposure scores using the "beta" method from Eloundou et al. (2024). High-AI-exposure claims have AI exposure scores of 0.49 or higher (top 25%). Data show the 3-month moving average (mean over the last three months) of UI claims.

Finding 3: UI Claims Among Occupations with High AI Exposure Increased, and Remained Elevated, in the San Francisco Bay Area and in Prominent Technology Sectors

The tracker also enables comparisons across California's 14 economic regions (regional planning units). Figure 4 contrasts trends in the San Francisco Bay Area ("Bay-Peninsula Region") with the rest of the state, which we highlight given the Bay Area's role as a center of AI innovation and early adoption — making it a natural place to detect emerging effects. While trends in UI claims by degree of AI exposure are similar across our two measures, Anthropic's observed AI exposure metric provides clearer signals of regional differences in overall exposure. Results for all regions are available in the tracker.

High-AI-exposure UI claims increased in the Bay Area by over 50% after ChatGPT-3.5's release, and remained elevated relative to the rest of the state (Panel A). Panel B shows that the AI "content" of these claims — the average share of occupational tasks that can be completed by AI — rose steadily beginning several months prior to ChatGPT-3.5's release, with a peak in mid-2023, and remained higher compared to the mean in 2019. While the early increase suggests that the tracker is picking up technology trends that may be less related to AI prior to ChatGPT-3.5's release, the fact that AI exposure of UI claims remained higher throughout 2025 and into May 2026 indicates AI-related developments have played a role in elevated UI claims. Overall, these patterns point to the specific industry mix of the Bay Area being important for understanding which segments of the California economy may be impacted by AI layoffs first.

FIGURE 4: Initial UI Claims in San Francisco Bay Area vs. Rest of California by ***Observed AI Exposure***

PANEL A: Initial UI Claims with High ***Observed AI Exposure*** Indexed to ChatGPT-3.5 Release

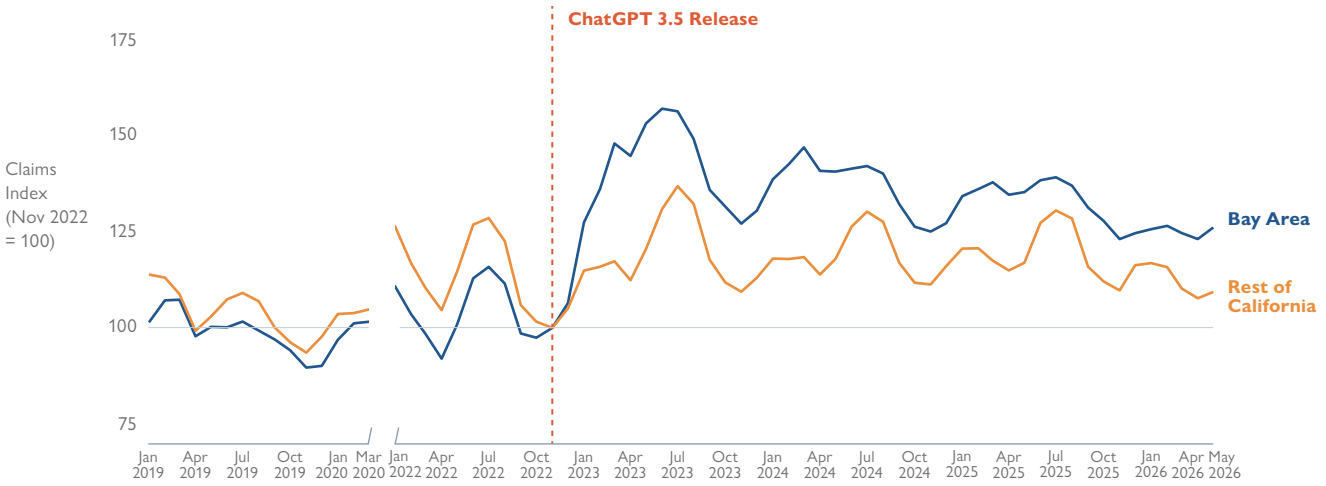

PANEL B: Mean of ***Observed AI Exposure*** Score Among UI Claims

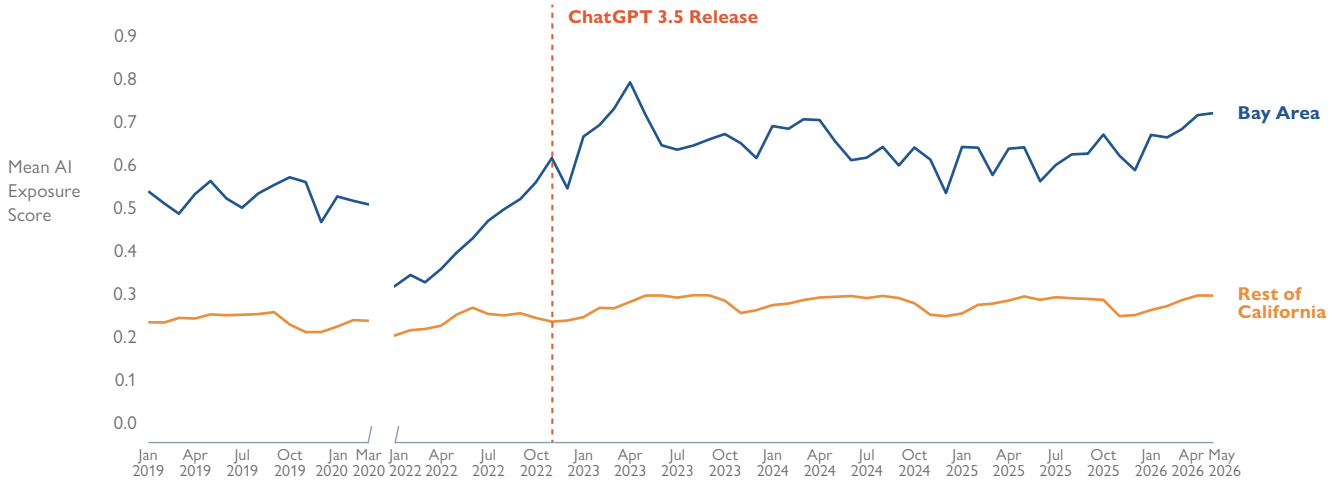

Notes: March 2020 to Jan 2022 is suppressed due to the unprecedented surge in UI claims during the pandemic, which would otherwise dominate the scale and obscure important comparisons between pre-pandemic and post-ChatGPT-3.5 trends. The dashed vertical line indicates the release date of ChatGPT-3.5. AI exposure scores are based on the Anthropic Economic Index (Handa et al., 2025) and capture observed use of Claude.ai across occupations. The mean AI exposure score reflects the average exposure among UI claimants, weighted by claims; higher values indicate greater exposure. Data show the 3-month moving average (mean over the last three months) of UI claims.

Consistent with elevated AI-exposed claims in the Bay Area, Figure 5 shows a rise in initial UI claims with high potential AI exposure in prominent tech sector industries across the state. It also presents the share of UI claims originating from occupations with high AI exposure across the full set of California NAICS 2-digit industries. Panel A focuses on two industries associated with the tech sector — *Information* and *Professional Services* — relative to remaining industries in California. Both sectors show a large increase in 2023, with high-AI-exposure claims in the *Professional Services* industry (which often contains occupations such as software developers) remaining elevated relative to the state as a whole. While this is not shown in the figure, high-AI-exposure claim rates in 2019 were also similar to levels reported in 2017 and 2018.[5]

Panel B shows the share of all claims that came from occupations with high-AI-exposure by calendar year. The *Finance and Insurance* industry consistently has the largest share of high-exposure claims, while sectors such as *Agriculture* which are more manual in nature, are less susceptible to AI. Importantly, this breakdown shows that while high-AI-exposure UI claims increased from 2022 to 2023 in a handful of industries, they have for the most part, remained relatively stable over time. In fact, the share of high-AI-exposure claims in 2022 was similar to rates reported in the earlier 2017–2019 pre-pandemic period.[6]

FIGURE 5: Initial UI Claims in California from Occupations with High ***Potential AI Exposure*** by Industry

PANEL A: UI Claims from High-AI-Exposure Occupations Information and Professional Services vs. Rest of California Industries

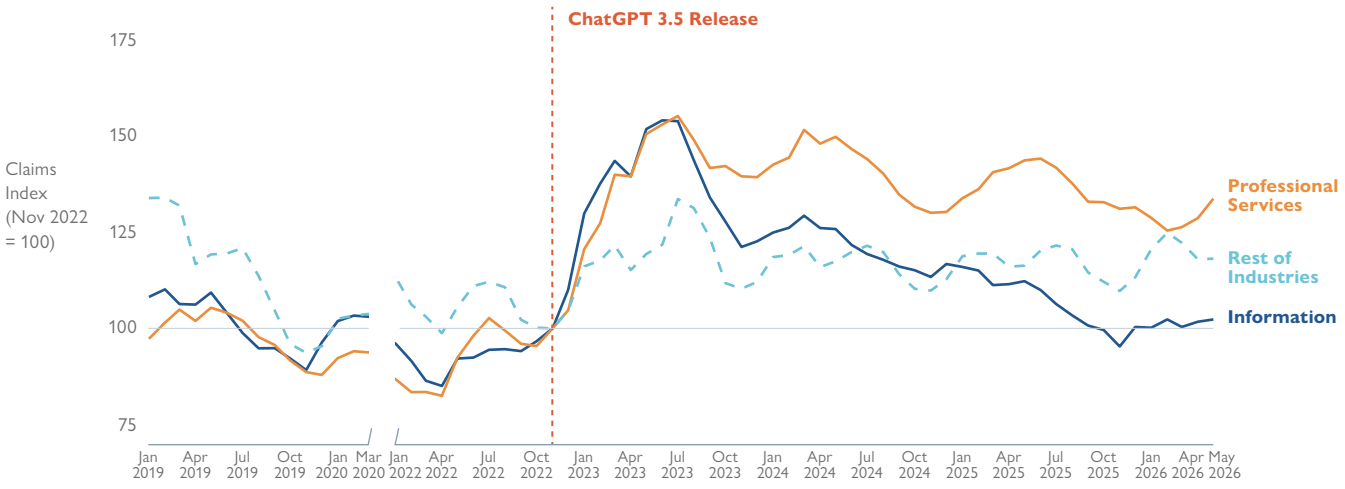

5 See the tracker for data from 2017–2018.

6 See the tracker for data from 2017–2018.

PANEL B: Share of Initial UI Claims from High-AI-Exposure Occupations by Industry from 2022–2025

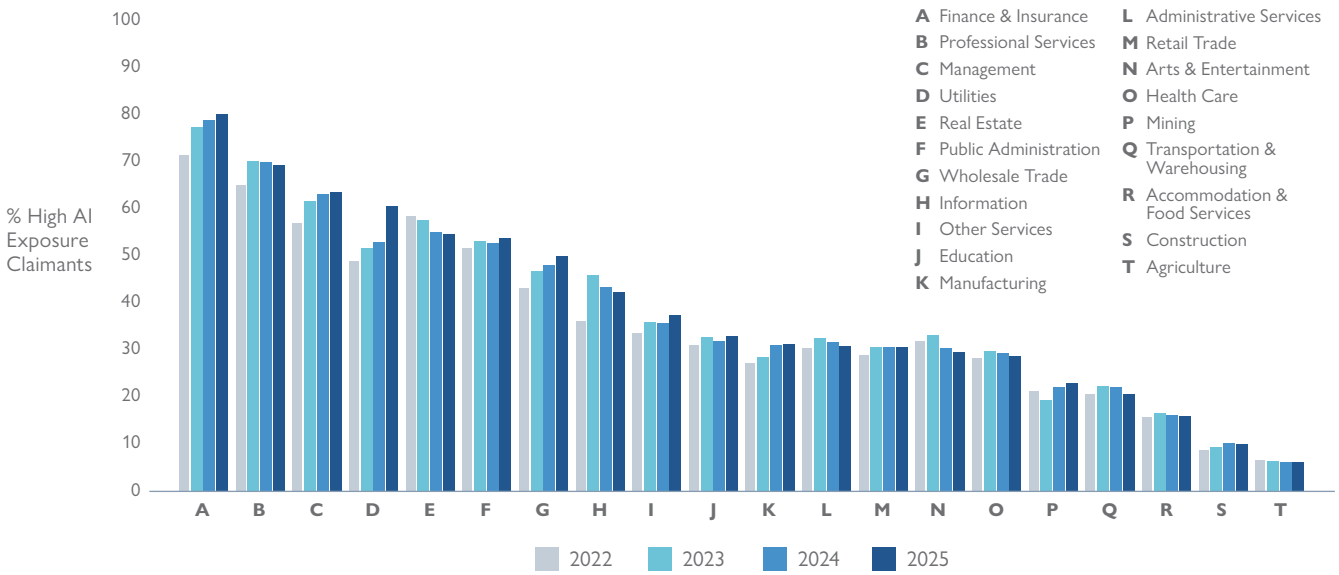

Notes: Claims are grouped by their potential AI exposure scores using the "beta" method from Eloundou et al. (2024), with only the "high" group (top 25%) shown here. The dashed vertical line indicates the release date of ChatGPT-3.5. Data by industry are available through 2025Q2.

We do not find large disproportionate increases by race, ethnicity, gender, or age, in the number of high-AI-exposure UI claimants (see the Technical Appendix for more details).[7] This may change with increasing diffusion of AI, and future iterations of CAIT will seek to measure if there are potential disparate impacts of AI. Future CAIT analyses may also seek to identify differences within education, region, or industry groups.

7. In contrast to prior research on reduced employment among entry-level workers in AI-exposed occupations (e.g., Brynjolfsson et al., 2025; Lichtinger & Hosseini Maasoum, 2025), we do not see disproportionate increases among younger UI claimants from occupations with high AI exposure. This is not surprising, because younger workers are typically less likely to claim UI benefits, which could be due to lack of awareness or lack of eligibility.

III. Validation Analysis

Because CAIT is a new policy tool, a key question is whether it can pick up *known* AI-related layoffs and thus function as an early detection system. To this end, we conducted a validation analysis. We identified six large employers that publicly announced "AI-driven" mass layoffs (with workforce reductions ranging from March 2024 to April 2025), where downsizing was plausibly linked to AI adoption rather than broader cost-cutting measures or other factors. Employers were identified from a list of technology layoff announcements and their stated rationale for downsizing, drawing on company statements and concurrent media coverage. For each employer, we determined the planned date of layoff using legally mandated public Worker Adjustment and Retraining Notification (WARN) filings. We then used UI claims data to examine the number of initial UI claims associated with these employers around those dates.

FIGURE 6: Initial UI Claims in California from Six Publicly Announced AI-Driven Mass Layoffs

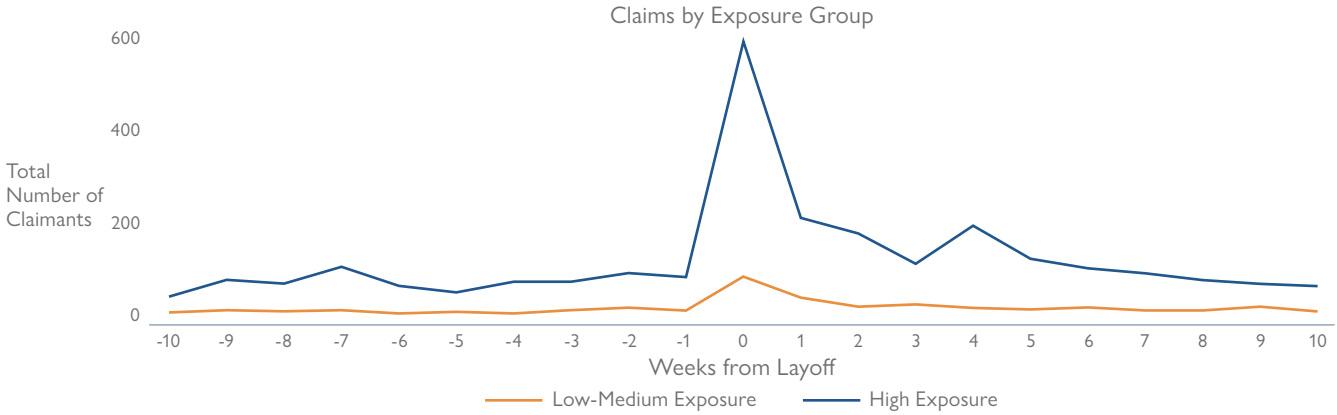

Notes: While suppression rules to safeguard the anonymity of these layoffs preclude reporting any specific totals for announced layoffs relative to UI claims for each of the six identified events, our estimates suggest that among these announced layoffs, roughly 58% of impacted employees subsequently filed for UI benefits within 10 weeks of layoff after adjusting for normal turnover. High-AI-exposure claims have observed AI exposure scores of 0.107 or higher (top 25%); Low-Medium exposure claims have scores below 0.107.

These publicly announced layoffs are clearly reflected in the UI claims data. Figure 6 shows the number of UI claims before and after each employer's planned layoff date. Prior to the layoff date, claims are relatively stable — hovering around 100 per week — consistent with typical levels of worker turnover. However, in the week of the layoffs, claims spike sharply and remain elevated in the weeks that follow. This pattern confirms that UI claims increase when AI-related mass layoffs occur. The figure also disaggregates claims by AI exposure group: "High" includes observations at or above the 75th percentile of exposure, while "Low-Medium" includes all observations below the 75th percentile using the observed AI exposure measure. This decomposition confirms that among likely AI-related layoffs, we disproportionately detect claims from workers with high AI exposure, further validating AI exposure as a meaningful tool for measuring AI-related job loss.

IV. Limitations

There are important limitations to this analysis. As stated at the outset, our measures of AI exposure capture the extent to which job tasks could be or have been performed by AI, not necessarily whether AI is being adopted or causing displacement at a specific workplace. It is also well known that UI claims capture only workers who claim UI benefits after losing their jobs, rather than all people who lose their jobs. In particular, UI claims do not capture workers who do not file for UI, either because they are unaware of it, because they quickly find new employment, they exit the labor force altogether, or they are not eligible for UI due to their legal status, or because they are self-employed (including gig workers), or in the case of younger workers, because they haven't worked long enough to qualify for UI benefits. Finally, multiple factors could explain the increases in high-AI-exposure UI claims we observed. While the initial rise in claims in 2023 likely reflects broader post-pandemic labor market dynamics, the more persistent elevation thereafter, especially in college-educated groups and in tech sectors well into 2025 and May 2026 — is less easily explained by COVID-era factors. Yet, this pattern is descriptive, and cannot be definitively attributed to AI without information on firm-specific AI adoption. Going forward, continued tracking of technology-related layoffs provides a way to calibrate and contextualize claims about AI-driven job loss in the public discourse and to help inform the policy response.

V. Conclusion

This report introduces the California AI-Unemployment Tracker (CAIT), the first tool in the nation to track AI-related layoffs using unemployment insurance claims data in near-real-time. Developed in partnership with the California Employment Development Department (EDD), the tracker links detailed UI claims with measures of both potential and observed AI exposure across occupations to provide new descriptive insights into the scope of AI- and technology-related labor market trends and to identify potential impacts.

UI claims data reveal that a large share of initial UI claimants from 2023–25 were from occupations with high and moderate AI exposure, similar to pre-pandemic shares. When tracking initial UI claims after the release of ChatGPT-3.5 in late 2022, we find a sustained increase in UI claims among workers from occupations with high AI exposure for key subgroups that are plausibly most exposed to the effects of AI at their workplace, including workers with a college degree or more, workers living in the Bay Area, and workers from technology-heavy industries.

These findings are broadly consistent with widely cited research showing that AI does not seem to have affected national unemployment rates (Gimbel et al. 2025), but has affected the headcounts of certain occupations with more exposure to AI (e.g., Brynjolfsson et al., 2025; Lichtinger & Hosseini Maasoum, 2025). This analysis extends that research by showing that lower firm headcounts capture not only slower hiring but also layoffs of existing workers, especially among higher-educated workers, and workers in the San Francisco Bay Area and in IT-related sectors. However, the analysis is descriptive and cannot causally disentangle tech sector layoffs after the COVID-19 remote-work boom from emerging AI-related labor market pressures.

By capturing the subset of displaced workers most likely to require adjustment support, the California AI-Unemployment Tracker provides policymakers with an early-warning system for detecting AI-driven layoffs. These insights can inform more rapid, targeted responses, spanning social insurance and workforce services. It could also inform longer-term strategies to help workers adjust, including workforce development and training programs. As AI continues to reshape the labor market, timely and detailed evidence will be critical. Through regular updates to CAIT, California is positioned to monitor these changes in real time, respond more effectively to emerging risks, and serve as a model for other states seeking to better understand and manage AI-driven labor market transitions.

VI. Acknowledgements

The California Employment Development Department partnered with the California Policy Lab to conduct research to measure AI-related job loss trends in California, using Unemployment Insurance claims data, combined with AI exposure measures. The California Policy Lab developed the California AI-Unemployment Tracker (CAIT), which will be updated regularly. The underlying data are available for public use.

Please see the associated technical appendix for more in-depth information about the data, measurements of AI exposure, and additional results supporting the analysis in the report: "Technical Appendix: Tracking the Impact of AI on Job Loss Using Unemployment Insurance Claims Data in California."

This research is produced through a partnership between the Labor Market Information Division of the California Employment Development Department and the California Policy Lab, a nonpartisan research center at the University of California, with sites in Berkeley, Los Angeles, and Sacramento. Any statements should only be attributed to the California Policy Lab, and do not reflect the views of the Labor Market Information Division of the California Employment Development Department. Any errors or omissions are the responsibility of the California Policy Lab, not of the Labor Market Information Division of the California Employment Development Department

For media or other inquiries about the definitions, methodology, and findings of this report, please contact reach out to Dr. Ben Hyman or Professor Till von Wachter.

To obtain further information about the data underlying the tabulations in the report and technical appendix, please contact: Juan Barrios, Chief, Labor Market Information Division, California Employment Development Department.

This work has been supported (in part) by grant #2510-51111 from the Russell Sage Foundation and with support from The Alfred P. Sloan Foundation (Grant #: G-2025-25158). We also gratefully acknowledge supporters of the California Policy Lab, including The James Irvine Foundation and the Woven Foundation for their generous support. Any opinions expressed are those of the author alone and should not be construed as representing the opinions of any funders. All errors should be attributed entirely to the authors.

Suggested citation: Hyman, B., von Wachter, T., Herrera, D., Moghadam, R., Morris J., Motghare, S., Segal, K. (2026). Tracking AI-Related Job Loss Using Unemployment Insurance Claims Data in California. California Policy Lab, University of California. https://capolicylab.org/wp-content/uploads/2026/06/Tracking-AI-Related-Job-Loss-Using-Unemployment-Insurance-Claims-Data-in-California.pdf

References

Bell, A., Hedin, T., Mannino, P., Moghadam, R., Romer, C., Schnoor, G., & Von Wachter, T. (2022). *Increasing Equity and Improving Measurement in the U.S. Unemployment System: 10 Key Insights from the COVID-19 Pandemic.* California Policy Lab. https://www.capolicylab.org/publications/increasing-equity-and-improving-measurement-in-the-us-unemployment-system-10-key-insights/

Brynjolfsson, E., Chandar, B., & Chen, R. (2025). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.* Stanford Digital Economy Lab Working Paper. https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/

Chatterji, A., Cunningham, T., Deming, D. J., Hitzig, Z., Ong, C., Shan, C. Y., & Wadman, K. (2025). *How People Use ChatGPT.* NBER Working Paper No. 34255. National Bureau of Economic Research. https://doi.org/10.3386/w34255

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). *GPTs Are GPTs: Labor Market Impact Potential of LLMs. Science*, 384(6702), 1306–1308. https://doi.org/10.1126/science.adj0998

Frank, M. R., Ahn, Y.-Y., & Moro, E. (2025). *AI Exposure Predicts Unemployment Risk: A New Approach to Technology-Driven Job Loss.* PNAS Nexus, 4(4), pgaf107. https://doi.org/10.1093/pnasnexus/pgaf107

Gimbel, M., Kinder, M., Kendall, J., & Lee, M. (2025). *Evaluating the Impact of AI on the Labor Market: Current State of Affairs.* The Budget Lab at Yale. https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs

Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., Troy, K. K., Amodei, D., Kaplan, J., Clark, J., & Ganguli, D. (2025). *Which Economic Tasks Are Performed with AI? Evidence from Millions of Claude Conversations.* arXiv:2503.04761. https://arxiv.org/abs/2503.04761

Hosseini Maasoum, S. M., & Lichtinger, G. (2025). *Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data.* SSRN Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5425555

Humlum, A., & Vestergaard, E. (2025). *Large Language Models, Small Labor Market Effects*. NBER Working Paper No. 33777. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5250742

Massenkoff, M., & McCrory, P. (2026, March 5). *Labor Market Impacts of AI: A New Measure and Early Evidence.* Anthropic. https://www.anthropic.com/research/labor-market-impacts

Rainey, K. (2026, March 10). *AI-Related Layoffs Test New York›s Ability to Track Job Losses.* Bloomberg Law. https://news.bloomberglaw.com/payroll/ai-related-layoffs-test-new-yorks-ability-to-track-job-losses

Yotzov, I., Barrero, J. M., Bloom, N., Bunn, P., Davis, S. J., Foster, K. M., Jalca, A., Meyer, B. H., Mizen, P., Navarrete, M. A., Smietanka, P., Thwaites, G., & Wang, B. Z. (2026). *Firm Data on AI.* NBER Working Paper No. 34836. National Bureau of Economic Research. https://doi.org/10.3386/w34836