TECHNICAL APPENDIX JUNE 2026

CALIFORNIA POLICY LAB

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

This technical appendix provides more details on the data, measurement of AI exposure, and additional results supporting the analysis in the report "Tracking AI-Related Job Loss Using Unemployment Insurance Claims Data in California." It describes the construction and linkage of AI-exposure scores to initial Unemployment Insurance (UI) claims [Appendix A], maps the geographic units used in the analysis to counties [Appendix B], and presents trends in high-AI-exposure UI claims by age, gender, and race/ethnicity [Appendix C]. It shows trends in high-AI-exposure claims based on alternate measures of AI usage [Appendix D]. Finally, it fits Difference-in-Differences models that quantify differences in UI claims between more and less AI-exposed groups following the release of ChatGPT-3.5 [Appendix E].

TABLE OF CONTENTS

Appendix A: AI Exposure Measures

The report uses two occupational measures of AI exposure which are linked to Unemployment Insurance (UI) claims: *Potential* and *Observed* AI exposure. To calculate these measures, we take the AI exposure of each task used in an occupation, and combine these to calculate AI exposure scores for occupations. We first use the O*NET task-to-occupation crosswalk to link 6-digit Standard Occupational Classification (SOC) codes to tasks. The crosswalk is based on recurring surveys of job incumbents and occupational experts. Workers report the importance of tasks they perform in their role, and trained analysts review and validate these responses to ensure consistency across occupations.

Table A1 provides an illustrative example of this crosswalk for two distinct occupations — *Online Merchants* and *Gambling Cage Operators*. Executing sales is relatively more important for gambling cage operators, who may, for instance, be exchanging chips for currency. This is reflected in the task's higher importance rating for gambling cage operators.

TABLE A1: Illustrative example of O*NET task-to-occupation crosswalk

SOC-6	OCCUPATION	TASK	TASK IMPORTANCE (1–5)
13-1199	Online Merchants	Record sales	5
		Execute sales	4
43-3041	Gambling Cage Operators	Record sales	5
		Execute sales	5
		Log phone calls	5

Each measure then assigns an AI-exposure score at the task level using its own methodology, and these task-level scores are subsequently aggregated to produce an occupation-level AI-exposure score. The specific procedures used to assign exposure to tasks and to aggregate those scores to the occupation level are described below.

I. *Potential* Exposure Measure: Eloundou et al. (2024)

Eloundou et al. (2024) evaluate all task-occupation pairs — 19,265 tasks across 1,016 occupations — to determine whether a large language model (LLM) can significantly reduce the time taken to complete a given task. Each pair is classified into one of three categories:

1. **AI exposure = 1** — an LLM alone can reduce task time by ≥50%
2. **AI exposure = 0.5** — an LLM + additional software can reduce task time by ≥50%
3. **AI exposure = 0** — the task could not be performed by an LLM

Human annotators evaluate each task–occupation pair by assessing whether (a) an LLM alone could reduce task completion time by at least 50%, or (b) such a reduction could be achieved with additional software built on top of the LLM, taking into account factors such as task complexity, required physical interaction, and the need for contextual judgment. ChatGPT is then used as a secondary means to provide independent assessments of these same criteria. Final task-level exposure labels are determined by aggregating these judgments by majority agreement across annotators; in cases of disagreement, ChatGPT's assessment is used to break ties. The resulting task-level exposure scores (1, 0.5, or 0) are then weighted by O*NET task-importance ratings and aggregated to produce a single AI-exposure score for each occupation, following what Eloundou et al. (2024) call their "β" weighting scheme.

Table A2 uses illustrative data on task-level AI exposure scores combined with our previous task importance weights to produce potential exposure scores for each occupation. The "task importance weight" reflects each task's relative importance within an occupation's full task set. In this example, Online Merchants receive a potential exposure score of 1, whereas Gambling Cage Operators receive a potential exposure score of 0.5. The lower exposure for Gambling Cage Operators reflects the presence of tasks requiring a high degree of physical interaction that could not be done with a machine (exchanging chips, etc.).[1]

TABLE A2: Calculating ***potential exposure*** measure for occupations

SOC-6	OCCUPATION	TASK	TASK IMPORTANCE WEIGHT	TASK AI EXPOSURE	POTENTIAL EXPOSURE
13-1199	Online Merchants	Record sales	5 / (5+4)	1	1
		Execute sales	4 / (5+4)	1	
43-3041	Gambling Cage Operators	Record sales	5 / (5+5+5)	1	0.5
		Execute sales	5 / (5+5+5)	0.5	
		Log phone calls	5 / (5+5+5)	0	

1 The scoring rubric here is motivated by an example developed in Eloundou et al. (2024).

II. *Observed* Exposure Measure: Handa et al. (2025) — Anthropic Economic Index

The observed exposure measure assigns AI-exposure scores to occupations based on how frequently the tasks associated with those occupations are completed using Claude, a large language model developed by Anthropic. Following the Anthropic Economic Index methodology from Handa et al. (2025), we assign an observed AI-exposure score to each 6-digit SOC occupation.

Starting with each task within an occupation, we first import the task share, defined as the share of Claude conversations that involved that task based on conversation logs from Dec. 2024 and Jan. 2025. Unlike the potential exposure measure, task usage is not occupation-specific; it reflects whether the task has been completed with Claude in any context, regardless of the occupation to which the task belongs. To account for tasks that are potentially shared across multiple occupations, we follow Handa et al. (2025) and divide each task share by the number of occupations that use that task, and then sum up "adjusted" task shares within each occupation.[2]

Table A3 illustrates how we calculate observed AI exposure scores for the same occupations shown in Table A2, assuming that Claude was deployed 60% of the time to record sales, 30% of the time to execute sales, and 10% to log phone calls.

TABLE A3: Calculating ***observed exposure measure*** for occupations

SOC-6	OCCUPATION	TASK	SHARE OF CLAUDE USE ON TASK	# OF OCCUPATIONS THE TASK APPEARS IN	ADJUSTED TASK SHARE	OBSERVED EXPOSURE
13-1199	Online Merchants	Record sales	60%	2	30%	0.45
		Execute sales	30%	2	15%	
43-3041	Gambling Cage Operators	Record sales	60%	2	30%	0.55
		Execute sales	30%	2	15%	
		Log phone calls	10%	1	10%	

The illustrative example in Table A3 shows how observed AI exposure can differ from the potential exposure measure. However, such discrepancies may be more limited in practice because 98.6% of tasks are unique to 6-digit SOC occupations. As a result, Anthropic's observed measure largely captures occupation-specific task usage in a way that closely aligns with the underlying task structure used in the potential exposure measure. The primary source of divergence between the two measures therefore arises not from differences in task composition, but from variation in the intensity with which Claude is used across tasks within occupations.

2 In some cases, the Anthropic Economic Index has missing tasks such that adjusted task shares sum to slightly less than 100%. The observed exposure scores by occupation are thus adjusted to account for the share of the total adjusted task percentage.

III. Potential vs Observed Exposure: How do they match up?

Panel A of Figure A1 shows that the two AI-exposure measures are positively correlated. Each dot represents a 6-digit SOC occupation, with the axes corresponding to the potential exposure score and the observed exposure score in logs. Prior to binning, the two measures exhibit a moderate positive correlation (raw correlation = 0.55). This correlation rises to 0.63 when using a variant of the observed measure which computes the share of tasks with any Claude use as shown in Panel B. This pattern suggests that while many occupations have high potential exposure to AI, observed AI usage in those occupations' tasks remains more limited.

FIGURE A1: Comparing potential and observed AI exposure by occupation

PANEL A: Correlation between potential and observed AI exposure

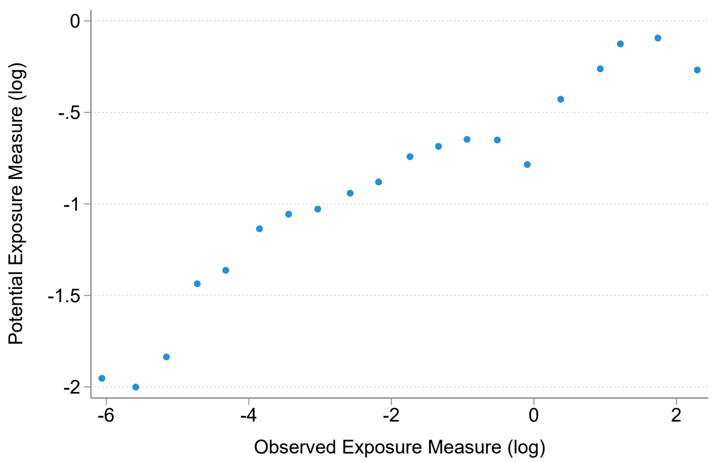

PANEL B: Correlation between potential and alternative formulation of observed exposure based on share of an occupation's tasks with any observed AI use

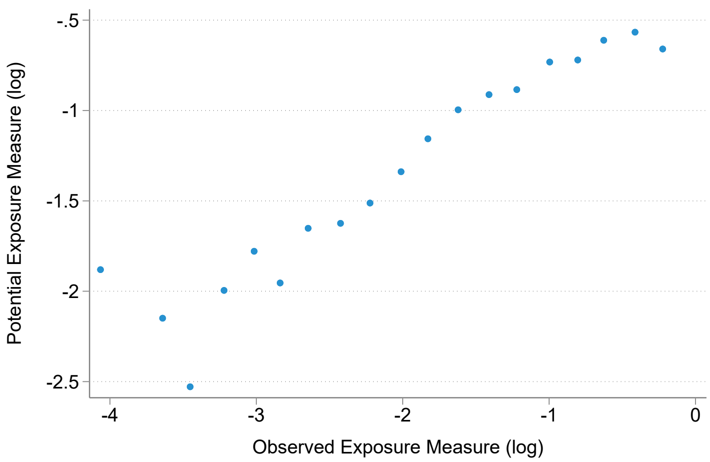

Notes: Each dot represents an occupation according to the 2010 Standard Occupational Classification System at the 6-digit level. In Panel A, the X-axis shows 20 equally-spaced bins of log observed AI exposure scores from Anthropic Claude use (Handa et al., 2025). In Panel B, the X-axis represents 19 equally-spaced bins of log observed AI exposure scores counting the share of tasks with any Claude use, instead of using the Anthropic Economic Index. The Y-axis reports log potential AI exposure scores from Eloundou et al. (2024).

Table A4 lists the ten occupations with the highest potential AI-exposure scores, along with their corresponding observed-exposure values. The lower observed exposure scores suggest that the observed AI usage across occupations remains well below their potential exposure, as has been shown in Massenkoff and McCrory (2026).

TABLE A4: Top Ten 6-Digit SOCs by potential AI exposure and corresponding observed AI exposure score

SOC-6	O*NET-SOC 2010 TITLE	POTENTIAL EXPOSURE	OBSERVED EXPOSURE
152021	Mathematicians	100.00%	0.36%
439081	Proofreaders and Copy Markers	97.50%	0.04%
434021	Correspondence Clerks	96.40%	0.36%
232091	Media and Communication Workers	95.80%	0.08%
151131	Computer Programmers	95.00%	6.16%
151141	Database Administrators	92.30%	1.18%
432021	Telephone Operators	89.50%	0.00%
439021	Data Entry Keyers	89.30%	0.02%
439111	Statistical Assistants	89.10%	3.35%
273091	Interpreters and Translators	88.00%	0.62%

Observed exposure shares simply reflect how much Claude is used on one task relative to another. It is a *relative* measure, and not an *absolute* measure of AI penetration. If Anthropic's Claude users are heavily skewed toward specific industries — like software developers — the measure will simply reflect the company's market concentration.

IV. Assigning AI Exposure Scores to UI Claims data

Both the potential and observed exposure measures assign scores to occupations classified under the Standard Occupational Classification (SOC) system. In contrast, the UI claims data report occupations using the Dictionary of Occupational Titles (DOT). To align these systems, we map SOC codes to DOT codes using standard crosswalks and generate exposure measures for the 3-digit DOT occupations observed in the UI claims data.[3]

Because there are fewer 3-digit DOT occupations than 6-digit SOC occupations, a single 3-digit DOT occupation often maps to several 6-digit SOC occupations, each with its own potential and observed exposure score. To create one potential and observed exposure score for the 3-digit DOT occupation, we take an employment-weighted average of the corresponding SOC-level scores. The intuition behind employment-weighting is straightforward: when a DOT code links to multiple SOC occupations, it is more likely to represent the SOC occupation that employs the most workers. Using employment as the weight embeds this logic directly into the final DOT-level score.

3 The potential measure is originally published on the 2018 SOC taxonomy and is crosswalked to 2010 SOC; the primary observed measure is built on O*NET-SOC 2010 task data. All exposure scores are assigned at the 2010 6-digit SOC level before being mapped to 3-digit DOT.

To implement this, we first use O*NET's 2010 DOT-SOC crosswalk to link each 3-digit DOT occupation to its associated 6-digit SOC occupations. We then use 6-digit SOC employment levels from California's Occupational Wage and Employment Statistics (OEWS) for May 2021 (the year prior to the release of ChatGPT-3.5 in late 2022) from the Bureau of Labor Statistics to create the employment weighted AI exposure scores at the 3-digit DOT level for both the potential and observed measures.

Table A5 demonstrates how the weighting works. Consider the 3-digit DOT code 219 (*Computing and Account-Recording Occupations*), which maps to fifteen different 6-digit SOC codes. For this example — shown here for the potential exposure measure — we calculate the share of total employment contributed by each linked SOC code and then construct a single "DOT-3 AI Exposure" value by summing each SOC-level exposure score weighted by its employment share. The final column reports this employment-weighted AI-exposure score for the 3-digit DOT occupation. We follow the same procedure to generate the corresponding observed exposure score.

TABLE A5: Example of aggregating 6-Digit SOC potential exposure scores to 3-digit DOT

DOT-3	SOC-6	SOC DESCRIPTION	SOC POTENTIAL EXPOSURE	OEWS CA 2021 EMPL. LEVEL	OEWS CA 2021 EMPL. SHARE WEIGHT	DOT-3 POTENTIAL EXPOSURE
219 Computing and Account-Recording Occupations	132082	Tax Preparers	0.625	8,310	0.006	0.533
	272023	Umpires, Referees & Other Sports Officials	0.313	440	0.000	
	431011	First-Line Supervisors of Office & Administrative Support Workers	0.490	153,650	0.116	
	433021	Billing, Cost & Rate Clerks	0.774	55,620	0.042	
	433031	Bookkeeping, Accounting & Auditing Clerks	0.802	170,580	0.129	
	433071	Tellers	0.542	32,010	0.024	
	434011	Brokerage Clerks	0.767	3,140	0.002	
	434131	Loan Interviewers and Clerks	0.594	21,710	0.016	
	435061	Production, Planning & Expediting Clerks	0.596	50,760	0.038	
	435071	Shipping, Receiving & Traffic Clerks	0.500	87,880	0.066	
	435111	Weighers, Measurers, Checkers & Samplers, Recordkeeping	0.293	8,350	0.006	
	436014	Secretaries & Administrative Assistants, Except Legal, Medical & Executive	0.644	172,560	0.130	
	439041	Insurance Policy Processing Clerks	0.833	23,810	0.018	
	439061	Office Clerks, General	0.576	274,700	0.207	
	439111	Statistical Assistants	0.891	300	0.000	
	435081	Stock Clerks and Order Fillers	0.182	261,130	0.197	

V. Sensitivity to Changing Cut-Off for Defining High AI Exposure

To assess the robustness of our results, we check to see if changing the cut-off values that define which claims are classified as high exposure have an impact on our main results. In the main analysis, we use the top 25% of claims to define "high AI exposure." In Figure A2, we use alternative cut-off scores (top 1%, top 5%, top 10%, top 25%) to define high exposure groups, which shifts the number of claims assigned to each group. As seen in the figure, which uses the potential AI exposure score, using these higher percentile cut-offs produces more pronounced spikes but preserves the overall trends in UI claims for the high-exposure group, suggesting that the time trends in high-AI-exposure claims are not sensitive to the specific cut-off chosen.

FIGURE A2: Number of high-AI-exposure claims under alternative percentile cut-offs for defining the high-exposure group: ***potential measure***

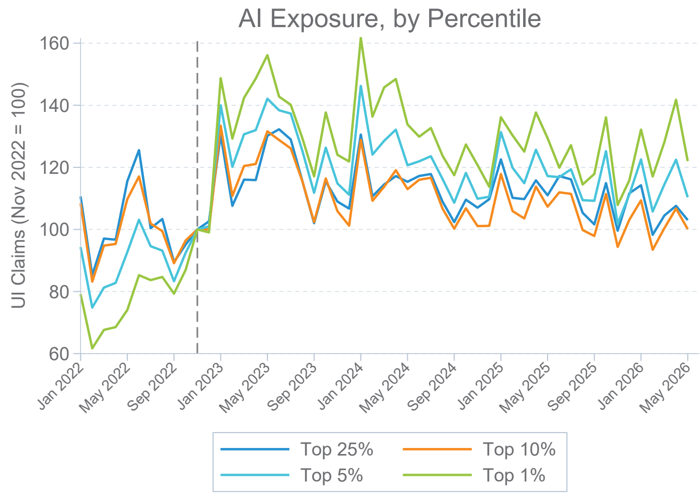

Notes: Dashed vertical line indicates the release date of ChatGPT-3.5. Claims are grouped by potential AI exposure scores using the "beta" method from Eloundou et al. (2024). The top 25% claims have potential AI exposure score of 0.49 or higher, the corresponding thresholds for the top 10%, top 5%, and top 1% are 0.58, 0.67, and 0.89, respectively.

VI. Robustness to Alternative AI-Exposure Measures

The observed exposure measure developed from Claude task-usage data used for our main results was developed as part of a data release in Feb. 2025 (Handa et al., 2025). To ensure that our results are robust to capturing evolving AI capabilities since Feb 2025, we conducted additional analyses using an updated version of the data and methodology that was released in March 2026 (Massenkof & McCrory, 2026).[4] In Figure A3, we show results from our headline finding, which continues to find no evidence of a trend break in any AI exposure group, even using the updated measure.

FIGURE A3: Robustness of main findings to recent technological advances: Initial UI claims using March 2026 update to Anthropic Index

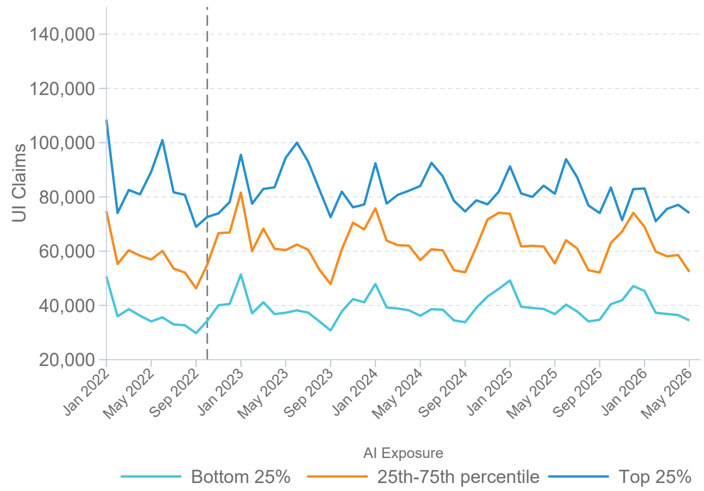

Notes: Dashed vertical line indicates the release date of Chat-GPT-3.5. Claims are grouped by observed AI exposure scores using the March 2026 update to the Anthropic Economic Index (Massenkof & McCrory, 2026). Claims in the top quartile have AI exposure scores of 0.096 or higher. Claims in the middle 50% have scores between 0 and 0.096. Claims in the bottom quartile have AI exposure scores of 0. 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.

4 The March 2026 release represents the fourth update to the Anthropic Economic Index since its initial release in Feb 2025. It differs from prior releases in two key respects: it incorporates Anthropic enterprise usage data and refines the exposure measure by accounting for the share of potential exposure that is realized in practice. In other words, it combines both potential and observed exposure into a single metric. Previous updates were released in March 2025, Sept 2025, and Jan 2026.

Appendix B: California's Regional Planning Units

For the geographic analysis, we use California's 14 Regional Planning Units (RPUs), which correspond to the regions used to organize workforce planning and service delivery under the Workforce Innovation and Opportunity Act (WIOA). The 14 RPUs and the counties they cover are shown in Table B1.

TABLE B1: Counties covered by Regional Planning Units

SR. NO.	REGIONAL PLANNING UNIT	COUNTIES (WITH FIPS CODES)
1	Bay–Peninsula Region	San Benito (069), San Francisco (075), San Mateo (081), Santa Clara (085)
2	Capital Region	Alpine (003), Colusa (011), El Dorado (017), Glenn (021), Placer (061), Sacramento (067), Sutter (101), Yolo (113), Yuba (115)
3	Coastal Region	Monterey (053), San Luis Obispo (079), Santa Barbara (083), Santa Cruz (087)
4	East Bay Region	Alameda (001), Contra Costa (013)
5	Humboldt Region	Humboldt (023)
6	Inland Empire Region	Riverside (065), San Bernardino (071)
7	Los Angeles Basin Region	Los Angeles (037)
8	Middle Sierra Region	Amador (005), Calaveras (009), Mariposa (043), Tuolumne (109)
9	North Bay Region	Lake (033), Marin (041), Mendocino (045), Napa (055), Solano (095), Sonoma (097)
10	North State Region	Butte (007), Del Norte (015), Lassen (035), Modoc (049), Nevada (057), Plumas (063), Shasta (089), Sierra (091), Siskiyou (093), Tehama (103), Trinity (105)
11	Orange Region	Orange (059)
12	San Joaquin Valley Region	Fresno (019), Mariposa (027), Kern (029), Kings (031), Madera (039), Merced (047), Mono (051), San Joaquin (077), Stanislaus (099), Tulare (107)
13	Southern Border Region	Imperial (025), San Diego (073)
14	Ventura Region	Ventura (111)

Appendix C: Results by Age, Gender, and Race/Ethnicity

The CAIT data allow for a rich analysis across a variety of subgroups. While the most striking results are presented in the main text, additional dimensions of heterogeneity are discussed here.

Starting with an analysis of claims by age group, Figure C1 shows that, following the release of Chat-GPT 3.5, high-AI-exposure claims rose proportionally across all age groups. While there is some evidence of peaks in claims around July 2022 and July 2023, claims were highly seasonal from 2019 through May 2026, and show a slight pre-trend prior to the pandemic. Overall, there is no clear divergence in high-AI-exposure claim patterns by age.

FIGURE C1: Initial UI claims by age at claim

PANEL A: High-AI-exposure claims by age indexed to Chat-GPT 3.5 release: ***potential measure***

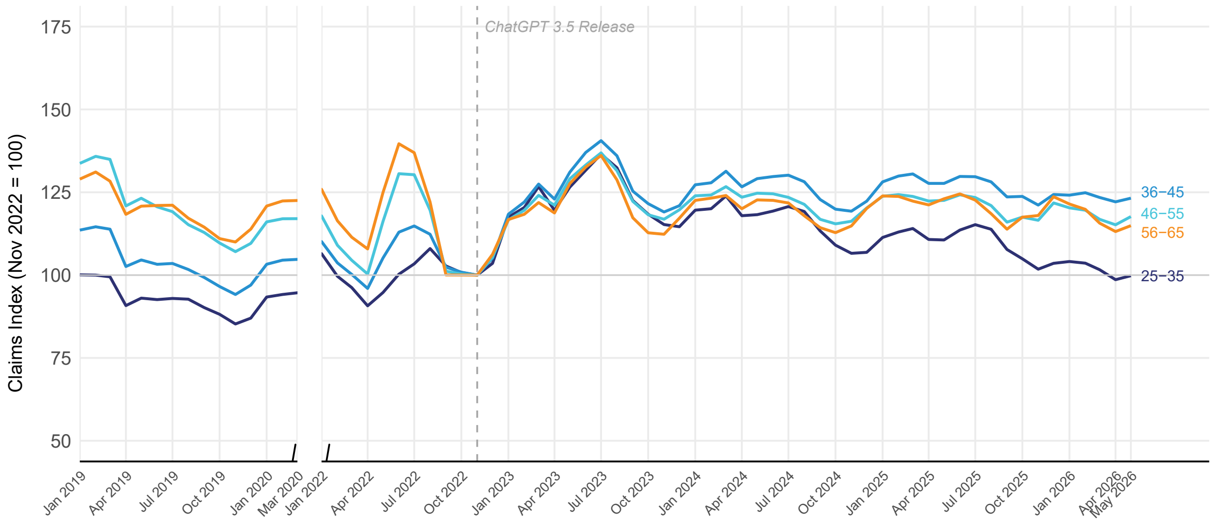

PANEL B: High-AI-exposure claims by age indexed to Chat-GPT 3.5 release: ***observed measure***

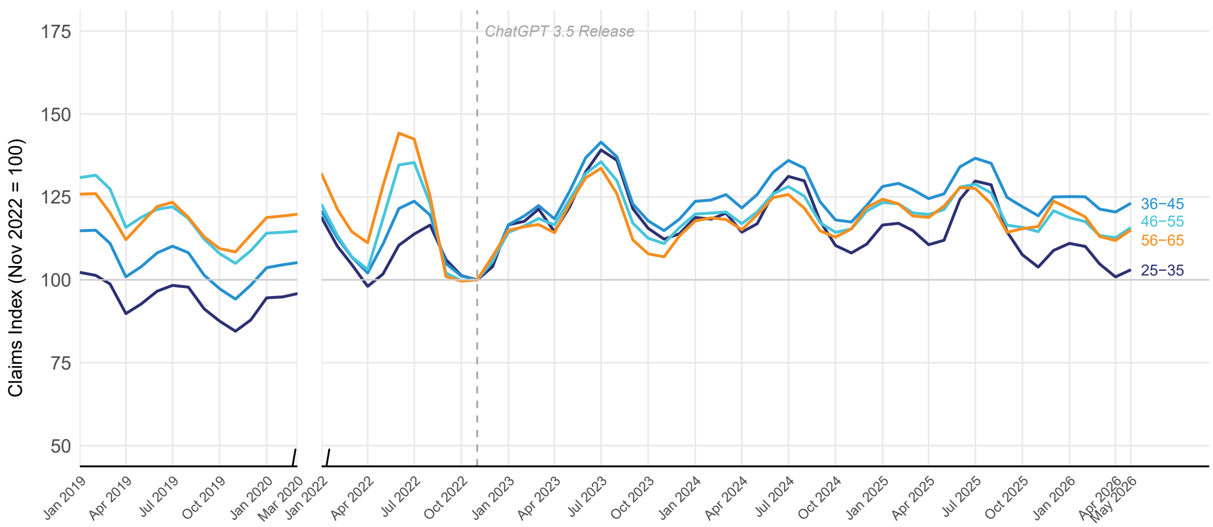

Notes: March 2020 to Jan 2022 is suppressed due to the unprecedented surge in UI claims during the pandemic, which would otherwise dominate the scale and obscure important comparisons between pre-pandemic and post-ChatGPT-3.5 release trends. Dashed vertical line indicates the release date of ChatGPT-3.5. Claims are grouped by potential AI exposure scores using the "beta" method from Eloundou et al. (2024) and observed AI exposure scores from Anthropic Claude use (Handa et al., 2025). High-AI-exposure claims have AI exposure scores of 0.49 or higher for the potential measure and 0.107 or higher for the observed measure. Data show the 3-month moving average (mean over the last three months) of UI claims.

Turning to our analysis of claims by gender in Figure C2, Panel A shows that, following the release of ChatGPT-3.5, high-potential-AI-exposure claims rose proportionally for both male and female workers, peaking in July 2023. Since then, claims levels have reverted back toward their pre-ChatGPT-3.5 mean through May 2026, with slightly larger declines in claims from female workers. Panel B, which compares trends among high observed AI exposure claims, shows a similar overall pattern for males, albeit with more cyclicality.

FIGURE C2: Initial UI claims by gender at claim

PANEL A: High-AI-exposure UI claims by gender indexed to ChatGPT-3.5 release: ***potential measure***

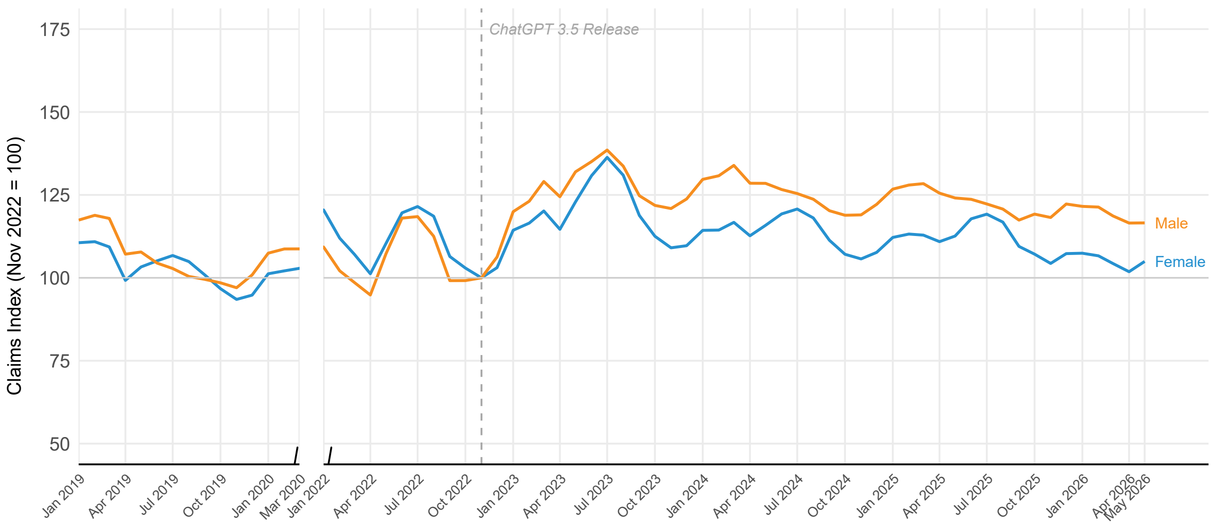

PANEL B: High-AI-exposure UI claims by gender indexed to ChatGPT-3.5 release: ***observed measure***

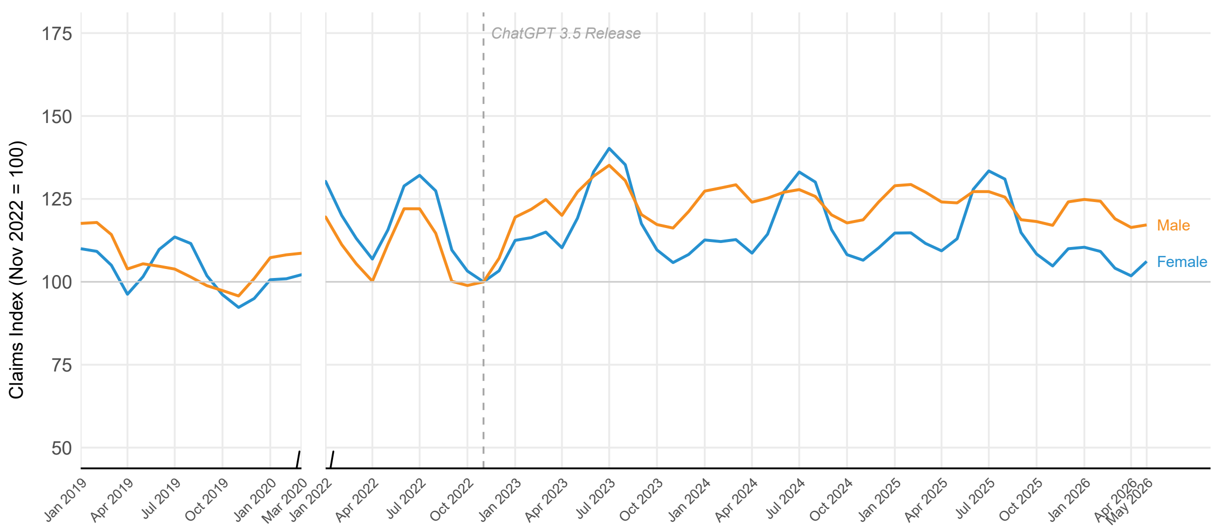

Notes: March 2020 to Jan. 2022 is suppressed due to the unprecedented surge in UI claims during the pandemic, which would otherwise dominate the scale and obscure important comparisons between pre-pandemic and post-ChatGPT 3.5 trends. Dashed vertical line indicates the release date of ChatGPT-3.5. Claims are grouped by potential AI exposure scores using the "beta" method from Eloundou et al. (2024) and observed-AI-exposure scores from Anthropic Claude use (Handa et al., 2025). High-AI-exposure claims have AI exposure scores of 0.49 or higher for the potential measure and 0.107 or higher for the observed measure. Data show the 3-month moving average (mean over the last three months) of UI claims. Date of birth and gender are self-reported by UI claimants. Gender is reported using only two options: Male or Female.

In Figure C3, we turn to an analysis by race and ethnicity. Panel A shows that, since the release of ChatGPT-3.5, high-potential-AI-exposure claims rose proportionally across most racial and ethnic groups, except among Asian claimants, whose claims increased by over 25% through July 2023 and remained elevated through May 2026. This was a reversal of a pre-ChatGPT-3.5 pattern of lower claims from Asian claimants, relative to Nov. 2022. Panel B also shows elevated observed-exposure claims among Asian workers, but the differences are smaller. Furthermore, low-AI-exposure Asian claimants (not reported in Figure C3) also exhibit elevated but smaller claims levels. Overall, we do not find large disproportionate differences in high-AI-exposure claim patterns across groups. The tracker allows monitoring and analyzing these differences on a monthly basis going forward.

FIGURE C3: Initial UI claims by race/ethnicity at claim

PANEL A: High-AI-exposure UI claims by race/ethnicity indexed to ChatGPT-3.5 release: ***potential measure***

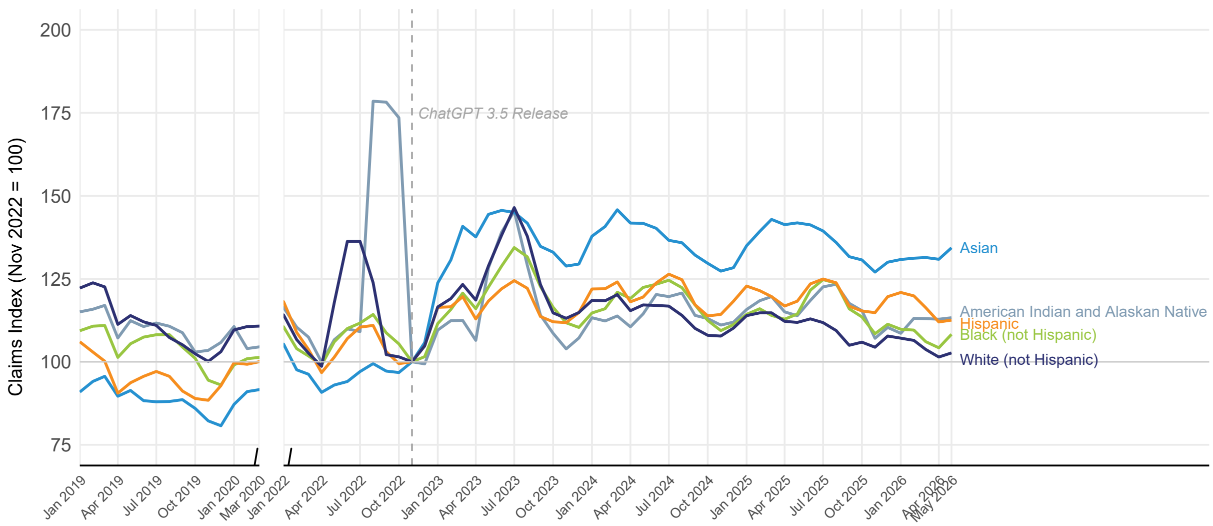

PANEL B: High-AI-exposure UI claims by race/ethnicity indexed to ChatGPT-3.5 release: ***observed measure***

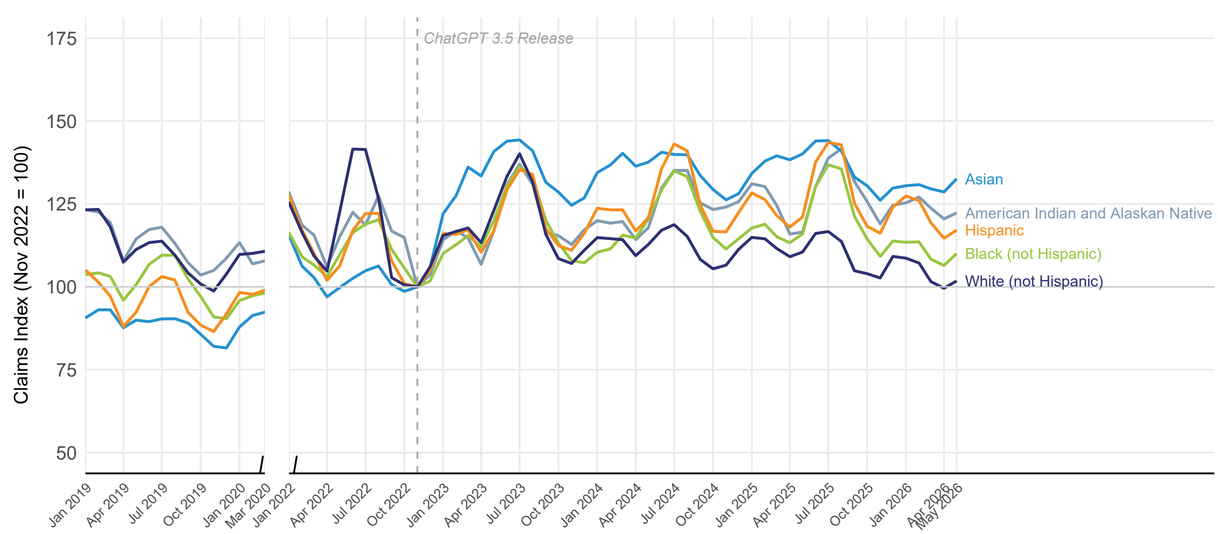

Notes: March 2020 to Jan. 2022 is suppressed due to the unprecedented surge in UI claims during the pandemic, which would otherwise dominate the scale and obscure important comparisons between pre-pandemic and post-ChatGPT trends. Dashed vertical line indicates the release date of ChatGPT-3.5. Claims are grouped by potential AI exposure scores using the "beta" method from Eloundou et al. (2024) and observed AI exposure scores from Anthropic Claude use (Handa et al., 2025). High-AI-exposure claims have AI exposure scores of 0.49 or higher for the potential measure and 0.107 or higher for the observed measure. Data show the 3-month moving average (mean over the last three months) of UI claims. Race/ethnicity is self-reported by UI claimants among five options: White, Black, Hispanic, Alaskan/Native American, and Asian.

Appendix D: Additional Robustness Figures

In the main report, most analyses use the potential exposure measure. In Figures D1 to D4, we report results using the alternate exposure measure to demonstrate that our findings do not depend on this choice. As an example, Figure D1 uses the "observed AI exposure measure," whereas the equivalent figure in the report (Figure 1) uses the potential measure. For the San Francisco Bay Area analysis (Figure D3), we show the results are robust when using the potential exposure measure instead.

FIGURE D1: Share of initial California UI claims from 2023 to 2025 using ***observed AI exposure measure***

PANEL A: Share of total initial UI claims by AI exposure group

PANEL B: Share of initial new UI claims from 2023 to 2025

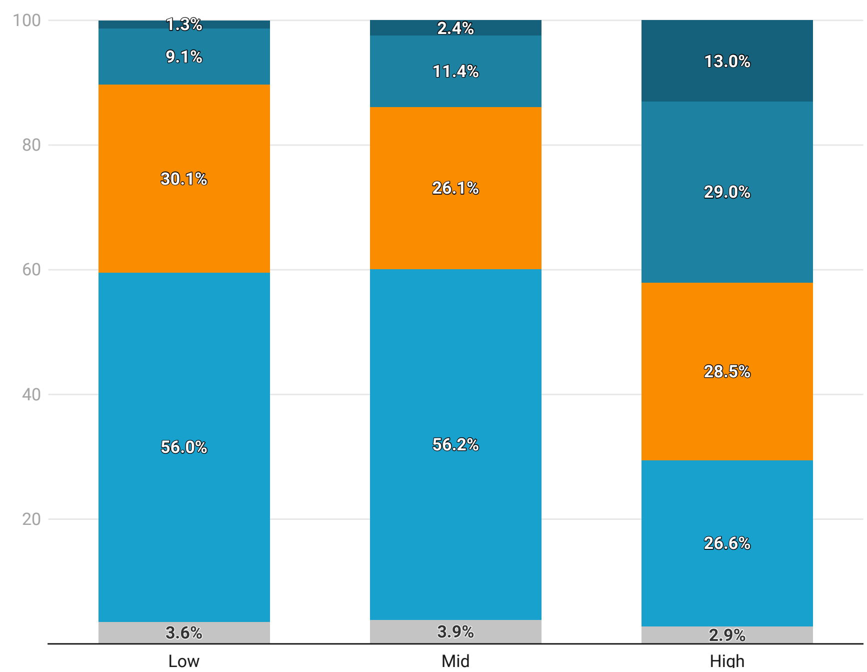

Notes: Claims are grouped by observed AI exposure scores from Anthropic Claude use (Handa et al., 2025). Claims in the top quartile have Claude-based AI exposure scores of 0.107 or higher. Claims in the middle 50% have scores between 0.106 and 0.011. Claims in the bottom quartile have AI exposure scores below 0.011.

FIGURE D2: Initial UI claims by education using ***observed AI exposure***

PANEL A: High-AI-exposure UI claims indexed to ChatGPT-3.5 release

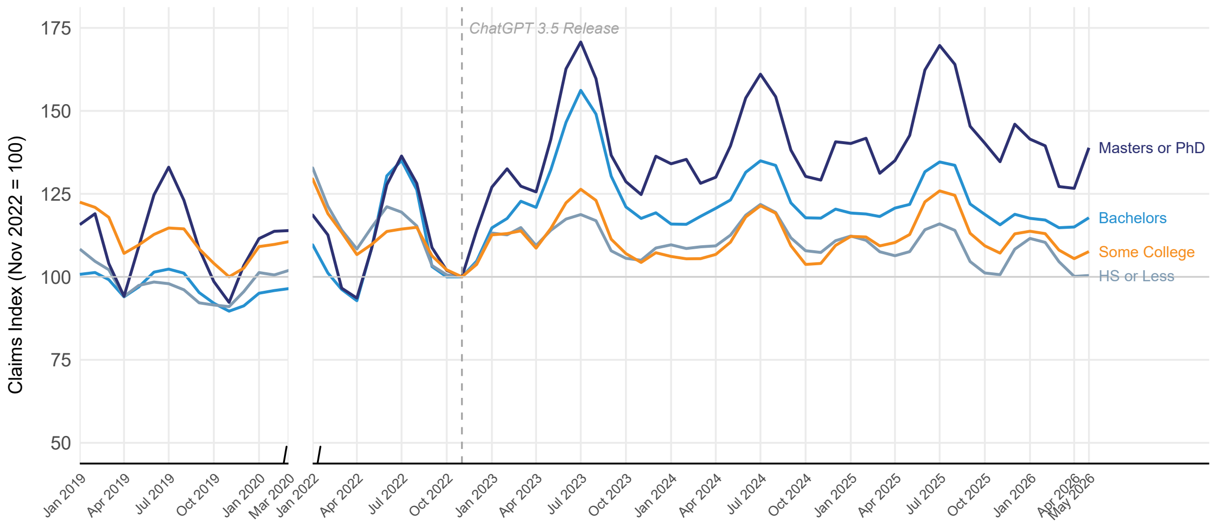

PANEL B: Low-AI-exposure UI claims indexed to ChatGPT-3.5 release

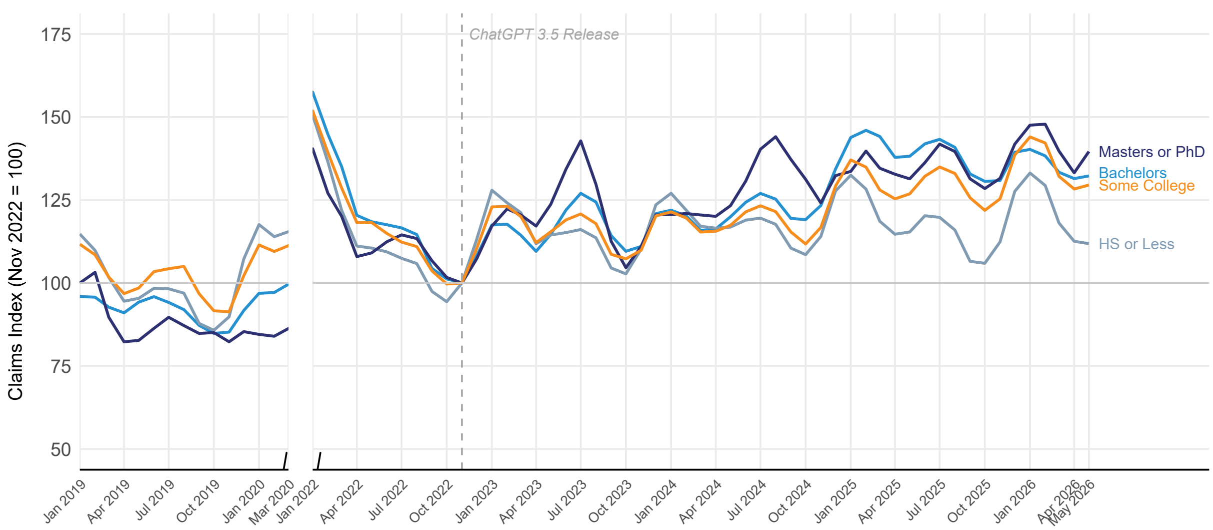

Notes: March 2020 to Jan 2022 is suppressed due to the unprecedented surge in UI claims during the pandemic, which would otherwise dominate the scale and obscure important comparisons between pre-pandemic and post-ChatGPT-3.5 trends. Dashed vertical line indicates the release date of ChatGPT-3.5. Claims are grouped by 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. Data show the 3-month moving average (mean over the last three months) of UI claims.

FIGURE D3: Initial UI claims in San Francisco Bay Area vs. rest of California by ***potential AI exposure***

PANEL A: Initial UI claims with high-potential-AI exposure indexed to ChatGPT-3.5 release

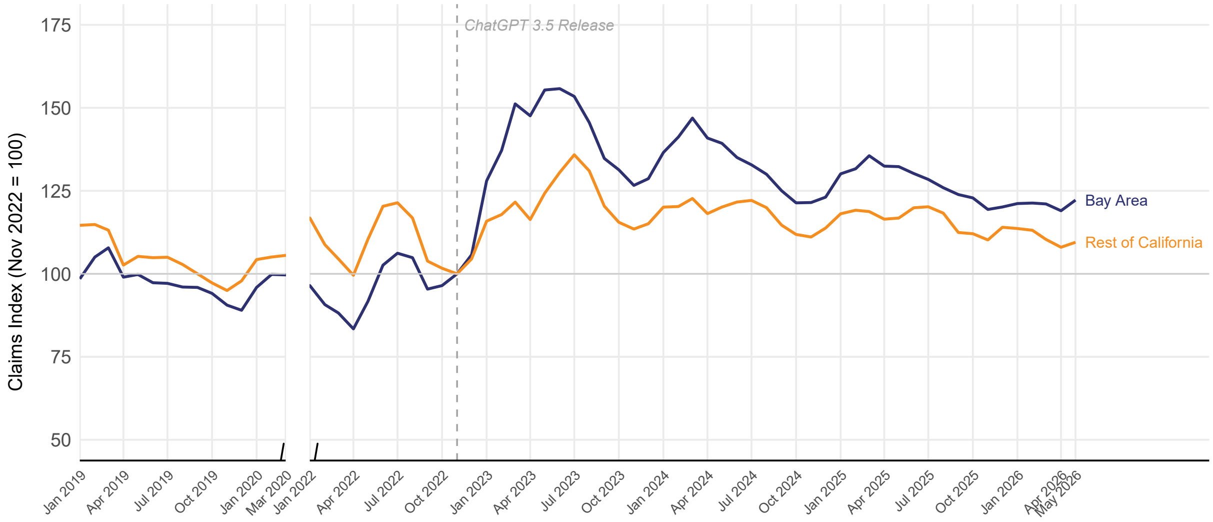

PANEL B: Mean of ***potential AI exposure*** score among UI claims

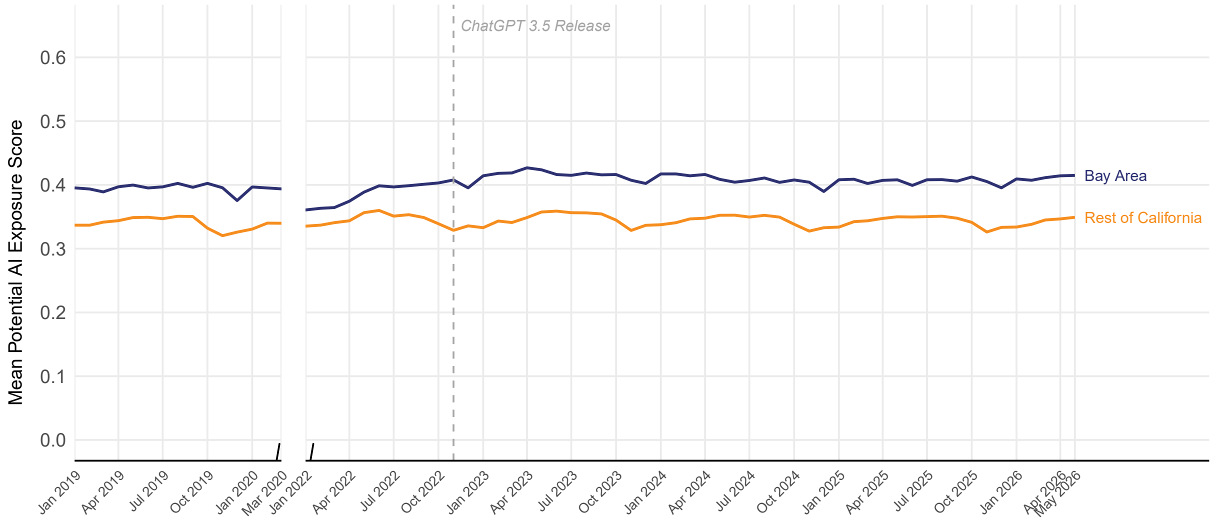

Notes: March 2020 to Jan 2022 is suppressed due to the unprecedented surge in UI claims during the pandemic, which would otherwise dominate the scale and obscure important comparisons between pre-pandemic and post-ChatGPT-3.5 trends. The dashed vertical line indicates the release date of ChatGPT-3.5. 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%). 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.

FIGURE D4: Initial UI claims from occupations with ***high observed AI exposure*** by industry

PANEL A: UI claims from high-AI-exposure occupations Information and Professional Services vs. rest of California industries

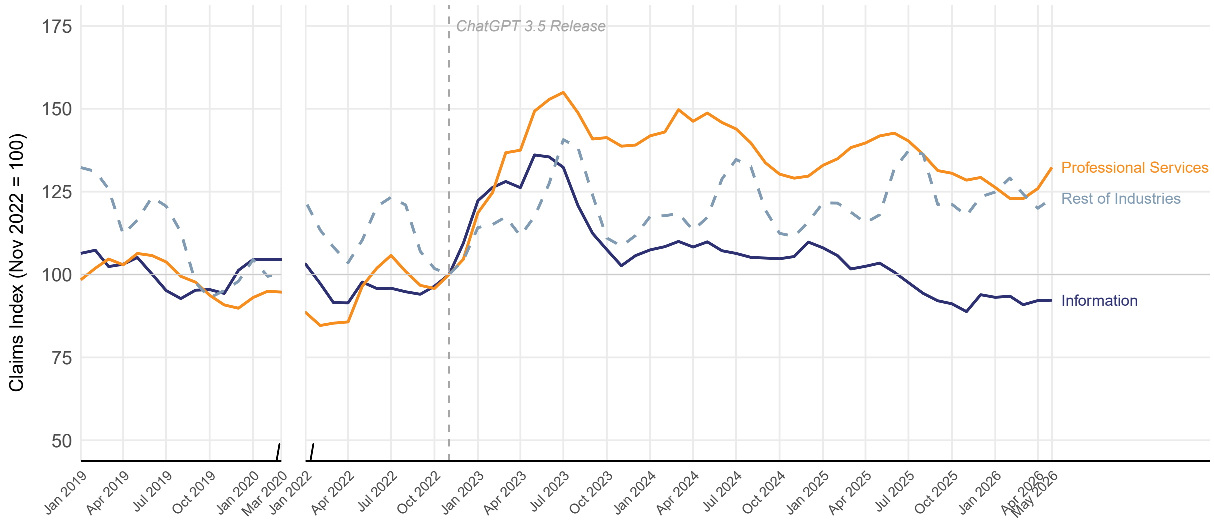

PANEL B: Share of initial UI claims from high-AI-exposure occupations by industry from 2022–2025

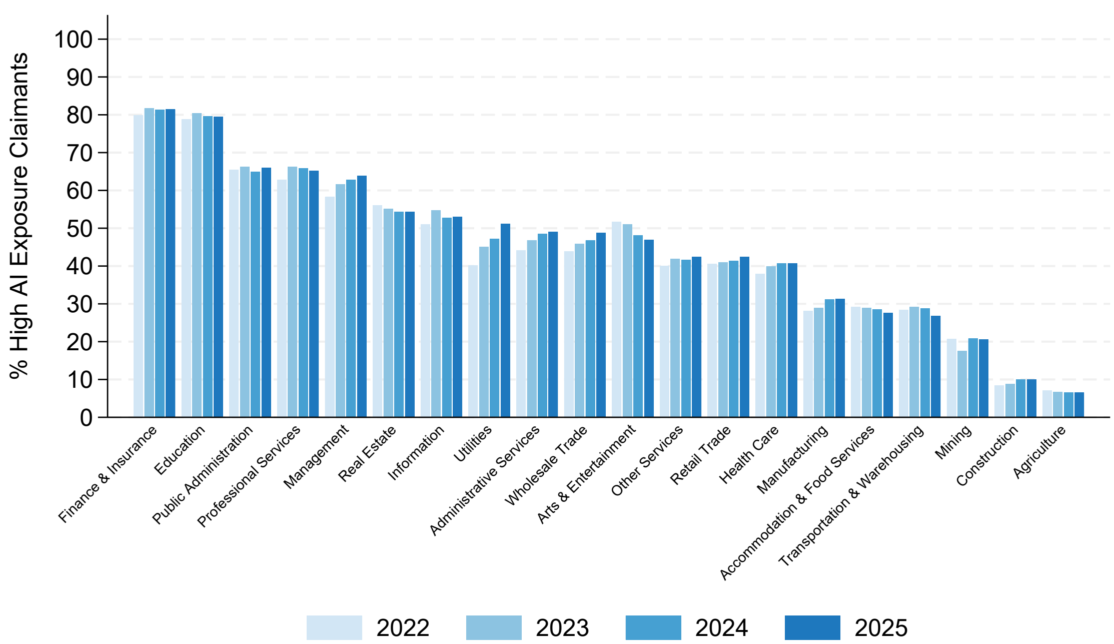

Notes: March 2020 to Jan 2022 is suppressed due to the unprecedented surge in UI claims during the pandemic, which would otherwise dominate the scale and obscure important comparisons between pre-pandemic and post-ChatGPT-3.5 trends. Dashed vertical line indicates the release date of ChatGPT-3.5. Claims are grouped by 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. Data show the 3-month moving average (mean over the last three months) of UI claims.

Appendix E: Difference-in-Differences Accounting Framework

To better quantify differential changes in UI claims around the release ChatGPT-3.5, we fit a series of Difference-in-Differences (DID) event-study models comparing trends across groups with differing levels of AI exposure. The goal of this exercise is purely descriptive: to assess whether changes in UI claims following the emergence of generative AI were disproportionately concentrated among more exposed groups, or instead reflected broader movements shared across groups.

We fit a series of (DID) event-study models of the following form:

$$Y_{it} = \alpha D_i + \delta_t + \sum_{t \neq Nov\ 2022} \beta_t * 1\{t\} * D_i + \epsilon_{it} \quad (1)$$

where,

- Y_{it} : 3-month moving average of number of initial UI claims indexed to Nov 2022 in group i in month t
- D_i : A binary indicator equal to 1 if the claims are for the "treatment" group, which we use as a shorthand to refer to the group of interest in the analysis
- $1\{t\}$: A set of binary indicator variables for each month from Jan 2019 - May 2026 with Nov 2022 as omitted month
- β_τ : estimated impact of treatment for month τ relative to Nov 2022
- δ_t : full set of month fixed effects (coefficients on each indicator variable for calendar months Jan 2019 - May 2026, with Nov 2022 as omitted month) that control for time-varying factors affecting both treatment and comparison groups
- ϵ_{it} : idiosyncratic error term. In practice, estimated residuals are equal to zero in this accounting framework

Because the analysis is conducted on aggregated group-by-month counts, with one observation for the treatment group and one observation for the comparison group in each month, the event-study specification is fully saturated. The group fixed effect, time controls, and treatment-by-month indicators exactly recover the treatment–comparison difference in each month relative to Nov 2022. As a result, there is no residual within-cell variation from which to estimate conventional regression standard errors. We therefore interpret the coefficients descriptively as an accounting exercise of the normalized monthly differences in UI claims, rather than as statistical estimates of the causal effect of the release of ChatGPT-3.5.

To begin, we start with results by geography as one might expect the early labor-market effects of AI to first appear in the San Francisco Bay Area, the center of AI innovation and early adoption. We examine this in two ways: First, we compare the evolution of high AI exposure claims relative to low AI exposure claims *in the Bay Area* and second, we compare the evolution of high AI exposure claims in the Bay Area *relative to those from the rest of California*.

The estimates in Figure E1 indicate that in the Bay Area, the difference between high and low AI exposure claims is persistently larger post-ChatGPT-3.5 relative to the omitted month (Nov 2022), and also higher relative to the 2019 baseline difference. We note the caveat that the large dip from Jan 2022 to Nov 2022 likely reflects relatively higher additional claims among the lower exposure group prior to ChatGPT-3.5's release.

FIGURE E1: Difference between high and low AI exposure claims in the San Francisco Bay Area: ***potential exposure measure***

Notes: The figure plots estimates of β_τ from equation (1) for claimants from the San Francisco Bay Area. The treated group consists of claims with high potential AI exposure, and the control group consists of claims with low potential AI exposure. The dependent variable is a three-month moving average of total UI claims in each group. Estimates from March 2020 to Jan 2022 are suppressed. The dashed vertical line indicates the release date of ChatGPT-3.5.

The estimates in Figure E2 further indicate that within high-AI-exposure claims, the difference in claim levels between the San Francisco Bay Area and rest of California is persistently larger post-ChatGPT-3.5's release relative to the omitted month (Nov 2022), and relative to the 2019 baseline difference. This figure offers some of the strongest evidence, as there is no confounding pre-trend between the San Francisco Bay Area and the rest of California, among high-AI-exposure workers, prior to ChatGPT-3.5's release.

Overall, the evidence shows that there was a sustained, post-ChatGPT-3.5 rise in UI claims from Bay Area workers with high AI exposure, both relative to low-exposure workers in the Bay Area and to similarly exposed workers elsewhere in California.

FIGURE E2: Difference between high-AI-exposure claims from the San Francisco Bay Area and rest of California: ***potential exposure measure***

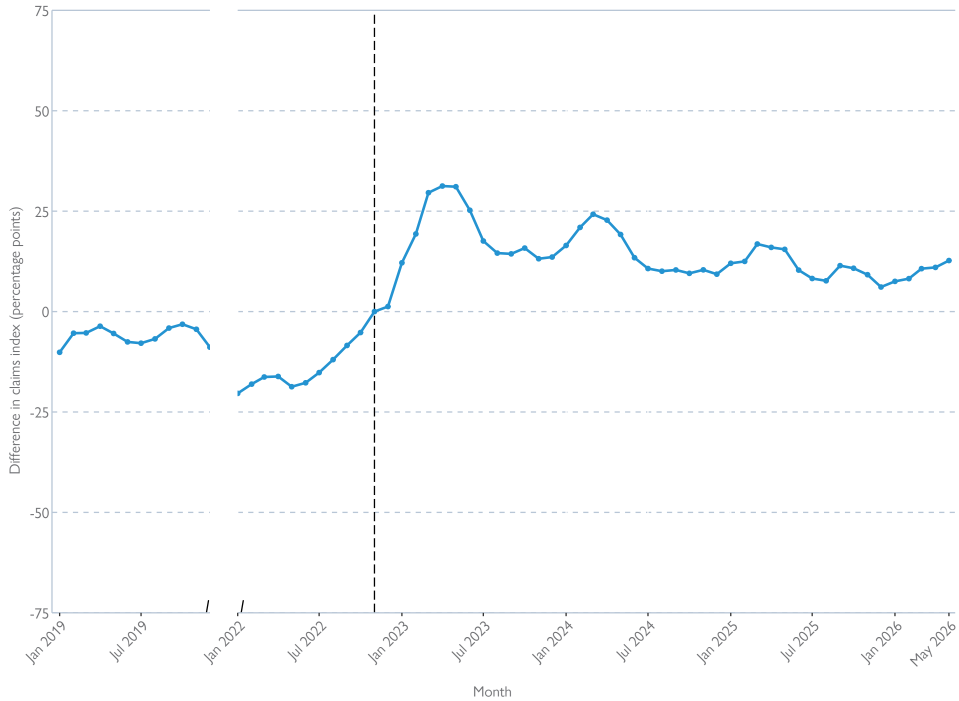

Notes: The figure plots estimates of β_τ from equation (1) for high potential AI exposure claims. The treated group consists of claimants from the San Francisco Bay Area, and the control group consists of claimants from the rest of California. The dependent variable is a three-month moving average of total UI claims in each group. Estimates from March 2020 to Jan 2022 are suppressed. The dashed vertical line indicates the release date of ChatGPT-3.5.

Similar to the Bay Area, we expect the early labor-market effects of AI to appear in the technology sector, which has a larger share of AI-exposed occupations. We examine this in two ways: First, we compare the evolution of high-AI-exposure claims relative to low-AI-exposure claims within two technology sector industries (*Information* and *Professional Services*) and second, we compare the evolution of high-AI-exposure claims in these industries relative to high-AI-exposure claims in all other industries.

The estimates in Figure E3 indicate that in the two technology sector industries, the difference between high- and low-AI-exposure claims is persistently larger in the post-ChatGPT-3.5 period relative to the omitted month (Nov 2022), and also higher relative to the 2019 baseline difference. The estimates in Figure E4 indicate that within high-AI-exposure claims, the difference in claim levels between the technology sector industries and all other industries is also persistently larger in the post-ChatGPT-3.5 period relative to the omitted month (Nov 2022), and relative to the 2019 baseline difference. Together, these patterns indicate that there was a sustained, post-ChatGPT 3.5 rise in UI claims from technology sector workers with high AI exposure, both relative to low-exposure workers in the technology sector and to similarly exposed workers in other industries.

FIGURE E3: Difference between high and low AI exposure claims in information and professional services industries: ***potential exposure measure***

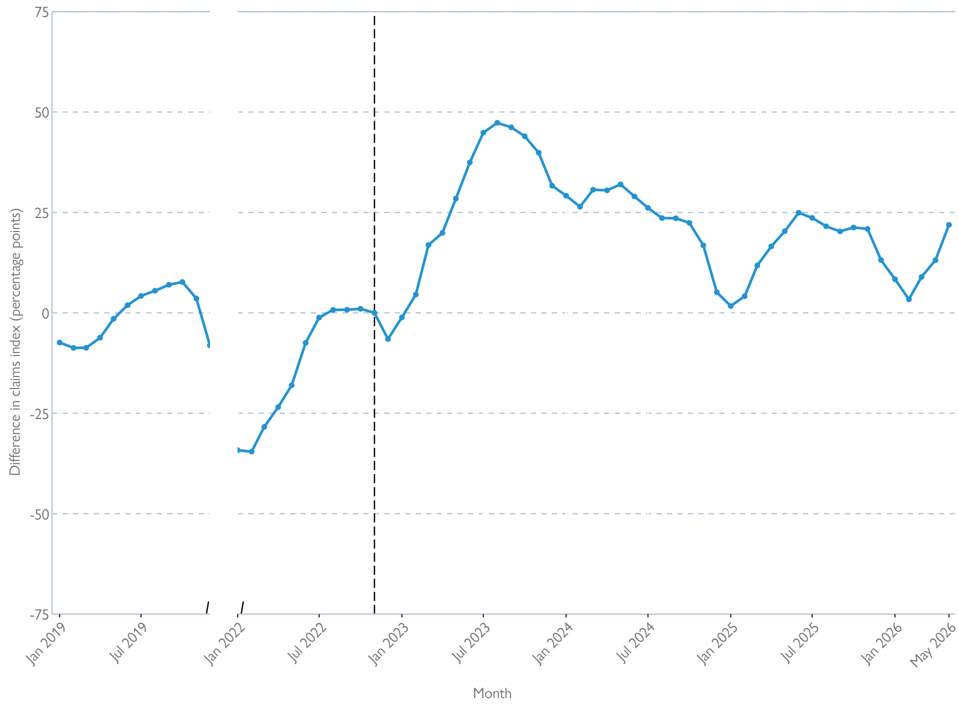

Notes: The figure plots estimates of β_τ from equation (1) for claimants from Information and Professional Services industries. The treated group consists of claims with high potential AI exposure, and the control group consists of claims with low potential AI exposure. The dependent variable is a three-month moving average of total UI claims in each group. Estimates from March 2020 to Jan 2022 are suppressed. The dashed vertical line indicates the release date of ChatGPT-3.5.

FIGURE E4: Difference between High AI exposure claims from Information and Professional Services and all other industries: ***potential exposure measure***

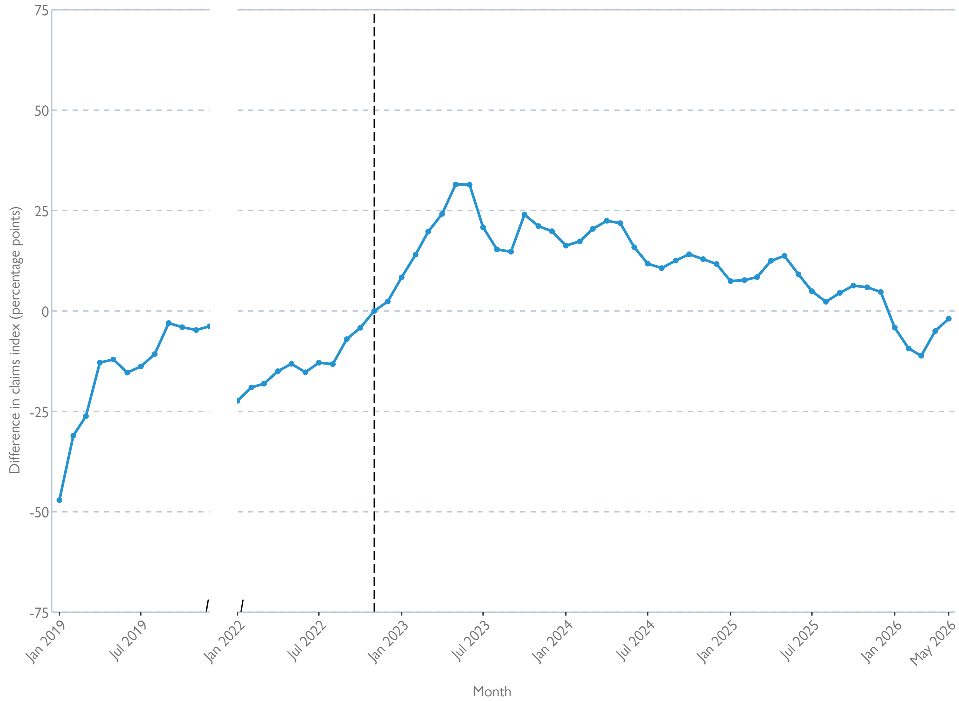

Notes: The figure plots estimates of β_τ from equation (1) for high-potential-AI-exposure claims. The treated group consists of claimants from Information and Professional Services industries, and the control group consists of claimants from all other industries. The dependent variable is a three-month moving average of total UI claims in each group. Estimates from March 2020 to Jan 2022 are suppressed. The dashed vertical line indicates the release date of ChatGPT-3.5.

Next, we compare the evolution of high-AI-exposure claims relative to low-AI-exposure claims among more educated workers, defined as those with a Bachelors, Masters or PhD degree, before and after the introduction of ChatGPT-3.5 in Nov 2022. We focus on this group because high-AI-exposure claims among this group increased and remained elevated after ChatGPT-3.5's release (Finding 3). The estimates in Figure E5 indicate that, within the more educated group, the difference between high and low AI exposure claims is persistently larger in the post-ChatGPT 3.5 period relative to the omitted month (Nov 2022), though there is a pre-trend that makes causal interpretation more difficult. This is not surprising, as high and low AI exposure groups may have been on different paths prior to ChatGPT-3.5's release, even within education groups.

FIGURE E5: Difference between high and low AI exposure claims among highly educated workers: ***potential exposure measure***

Notes: The figure plots estimates of β_τ from equation (1) for claimants with a Bachelors, Masters, or PhD degree. The treated group consists of claims with high potential AI exposure, and the control group consists of claims with low potential AI exposure. The dependent variable is a three-month moving average of total UI claims in each group. Estimates from March 2020 to Jan 2022 are suppressed. The dashed vertical line indicates the release date of ChatGPT-3.5.

Finally, we compare the evolution of high-AI-exposure claims from more educated workers, again defined as those with a Bachelors, Masters or PhD degrees, relative to high-AI-exposure claims from *less* educated workers, defined as those with only a high school diploma or some college education. The estimates in Figure E6 indicate that within high-AI-exposure claims, the difference between claims from more and less educated workers is persistently larger in the post-ChatGPT-3.5 period relative to the omitted month (Nov 2022), and also larger relative to the 2019 baseline difference. The large dip from Jan 2022 to Nov 2022 likely reflects relatively higher additional claims among the lower educated group. Overall, the pattern confirms that more educated workers with high AI exposure experienced a persistent rise in UI claims starting with the release of ChatGPT-3.5, compared to less educated workers with similarly high AI exposure.

FIGURE E6: Difference between high-AI-exposure claims from more educated and less educated workers: ***potential exposure measure***

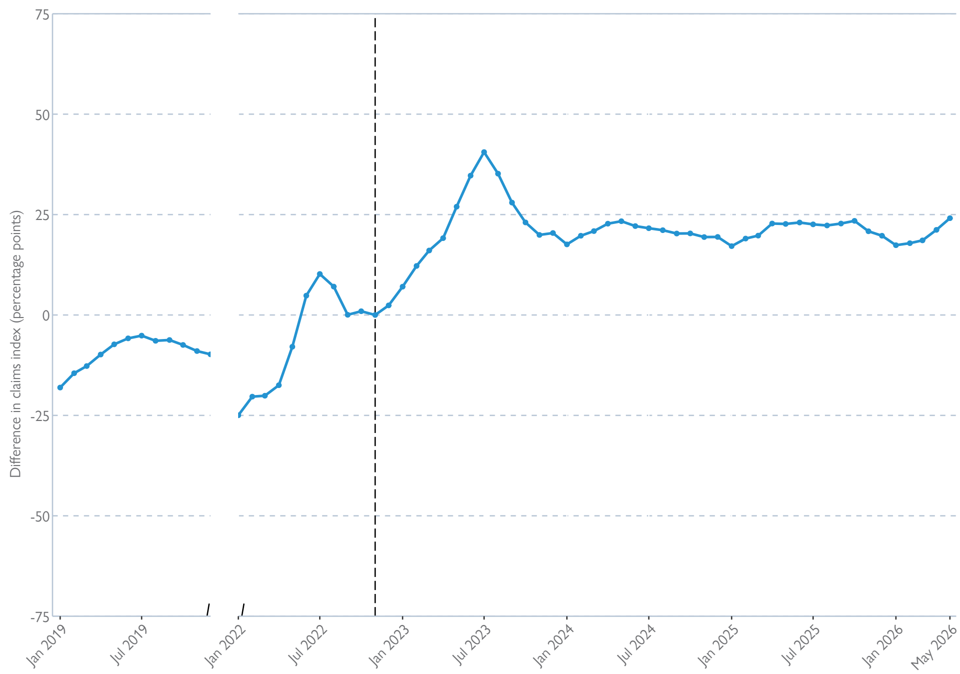

Notes: The figure plots estimates of β_τ from equation (1) for high-potential-AI-exposure claims. The treated group consists of high exposure claimants with Bachelors, Masters, or PhD degrees, and the control group consists of high exposure claimants with only a high school diploma or some college education. The dependent variable is a three-month moving average of total UI claims in each group. Estimates from Mar 2020 to Jan 2022 are suppressed. The dashed vertical line indicates the release date of ChatGPT-3.5.

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 monthly. The underlying data are available for public use.

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-Claims-Data-in-California.pdf

Appendix 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/

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What Can Machines Learn, and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings*, 108, 43–47. https://doi.org/10.1257/pandp.20181019

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). Labor market impacts of AI: A new measure and early evidence. https://www.anthropic.com/research/labor-market-impacts