The Effects of State Unemployment Insurance Benefit Duration Cuts on Labor Market Outcomes

Jesse Linkhorn and Ken Troske*

November 12, 2026

Abstract

This paper examines the effects of state-level reductions in unemployment insurance (UI) benefit duration on labor market outcomes. Between 2016 and 2025, 9 U.S. states reduced their maximum weeks of regular UI benefits below the 26-week standard. I exploit the staggered timing of these policy changes in a difference-in-differences and event study framework. UI recipiency changes by -7.26 percentage points, although the event-study leads reject parallel trends. Comparing posted and reported salaries from Lightcast and Glassdoor data within the same multi-state firm, I find no detectable change in posted wages but find reported starting pay rises, although its leads reject parallel trends. Bounds that permit a differential linear trend of arbitrary slope rule out elasticities of posted wages with respect to UI generosity larger than 0.049 in absolute value. I see no effect on the state unemployment rate either. These reforms were enacted into tight labor markets, which may account for the contrast with earlier evidence.

JEL Codes: J64, J65, J68, H73

Keywords: Unemployment insurance, benefit duration, recipiency rate, starting wages, difference-in-differences, staggered adoption

*Department of Economics, University of Kentucky. Email: jjli230@uky.edu.

1. Introduction

Unemployment insurance is a cornerstone of the American social safety net, providing temporary income support to workers who lose their jobs involuntarily. Since the program's inception under the Social Security Act of 1935, federal law has set no standard for benefit duration (U.S. Department of Labor, 2024), and the standard maximum duration of regular UI benefits was 26 weeks in most states until the Great Recession. Following this downturn, states began to cut benefits to spur labor market recovery. What was once the norm for maximum benefits began to change as states broke away from traditional standards following an extreme period of extended federal UI benefits. These cuts created a precedent following the COVID-19 recession in rising uptake of state UI policies that reduced benefits for unemployed workers. This paper studies the labor market consequences of these cuts. Between 2016 and 2025, 9 states reduced their maximum regular UI benefit duration below 26 weeks, in some cases upwards of 50%, creating substantial variation in the maximum duration of the UI benefits across labor markets. This paper exploits the staggered timing of these policy changes to estimate the effects across three models: the associated effects on short-term UI recipiency rate and unemployment rate, the causal effects on reported and posted wages at multi-state firms, and the causal effects on employment and total paid wages at multi-state firms. The third model is under Census disclosure review and is not reported here yet.

The primary question is whether shorter benefit duration achieves its stated objective of encouraging faster re-employment, and if so at what cost to workers. Standard job search theory predicts that more generous benefits raise reservation wages and extend unemployment spells (McCall, 1970; Mortensen, 1977), so a shorter duration may induce workers to lower their reservation wages and accept jobs more quickly. UI receipt should therefore fall. Unemployment should fall if search effort rises, but it is unchanged if workers leave the labor force instead. The effect on offered pay is ambiguous, since lower reservation wages let employers offer less while a tighter labor market pushes pay up. The welfare implications

then depend on whether shorter spells reduce match quality, that is, on whether workers who exhaust benefits sooner accept lower-paying jobs than a longer search would have found (Chetty, 2008).

The empirical analysis proceeds in multiple parts, each asking one question of one outcome. The first two outcomes and posting volume are estimated on a state-quarter panel, and the two wage outcomes on private firm-level job data. The establishment-level outcomes from administrative data are forthcoming. The first question is whether the cuts bind. One outcome is the UI recipiency rate which is the ratio of average weekly UI claimants (weeks compensated divided by the weeks in the period) to the count of short-term unemployed, those reporting unemployment duration of 26 weeks or less. As a state shortens its maximum duration, fewer weeks of benefits are paid per unemployed worker, which lowers the rate mechanically. I find an associated decrease of this kind, significant at the 5% level under the wild cluster bootstrap, although the event-study leads reject parallel trends. This may induce a behavioral response as well which I do not capture. The second outcome is the seasonally adjusted state unemployment rate, which is not mechanically tied to the maximum duration and therefore captures an aggregate response.

The next question asks whether duration cuts affect the wages at which unemployed workers re-enter employment. The second model uses both reported starting salaries from Glassdoor, restricted to new hires with zero to two years of experience, and posted wages from the Burning Glass Institute/Lightcast job postings data. I use a robust estimator from Borusyak et al. (2024), with firm or employer-job fixed effects alongside state and time fixed effects, which isolate within-firm variation in pay across treated and control states and address the concern that firms in states that cut benefits may differ systematically from firms in states that retained 26-week durations. I look at an additional outcome, the volume of job postings, which looks at the extensive margin and is reported in Section 5.4. The third model (Model 3) examines whether benefit duration cuts affect labor market outcomes at the establishment level: including employment and total wages paid. Using the Longitudinal

Employer-Household Dynamics (LEHD) data, I utilize a similar multi-state firm framework that estimates employment (wages) at establishments in treated states compared to non-treated states under the same firm. The LEHD, which is a dense administrative dataset, provides a near-universe of UI-paying firms in the U.S., which provides the same framework to circumvent issues of selection bias by firms in certain states. Model 3 is estimated in a Census Research Data Center, and its results are under disclosure review, so this chapter does not report them yet.

The key identifying assumption is that, absent the benefit duration cuts, labor market outcomes in treated and control states would have followed similar trends. The models employ both a difference-in-differences and event study framework utilizing additional robust estimators to account for the staggered adoption of policies (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020). The state-level outcomes are estimated with a standard TWFE approach. Robust estimators agree with TWFE in sign but are less precise for recipiency (Table 9 and Appendix A.2). Where the model includes firm or employer-job fixed effects, I use the imputation estimator of Borusyak et al. (2024). All standard errors are clustered by state. The analysis is capped by a bounding exercise which guides the assessment of parallel trends assumptions (Rambachan and Roth, 2023).

I find that short-term recipiency changes by -7.26 percentage points (s.e. = 2.31). The unemployment rate changes by -0.05 percentage points (s.e. = 0.31), and posting volume by -10.61 percent (s.e. = 7.82). Posted wages change by 0.44 percent (s.e. = 0.25). Reported starting pay rises by 1.51 percent (s.e. = 0.32), although its leads reject parallel trends. These cuts reduce the share of the short-term unemployed receiving UI. They do not detectably change unemployment or posted wages, and the posted-wage null excludes the wage response found in the post–Great Recession reforms under parallel trends, but not once the pre-trend is allowed for. I report the other outcomes as imprecise or failing the pre-trend test.

This paper makes three contributions to the literature on UI generosity and labor market effects. First, it provides the first evidence on the wave of state-level duration cuts enacted

since 2016, and primarily after the COVID-19 recession, extending an evidence base built almost entirely on the Great Recession era. This wave of policy has not been examined in the literature, and provides the first look at a set of policies that enact stronger cuts in a tighter labor market environment. This provides evidence of the effects of pre-cyclical policy enactment contrasted with previous literature that examined pro/post cyclical policy strategy. I couple this contribution with the use of comprehensive and more granular administrative data from the LEHD (forthcoming) which provides more precise and dynamic estimation compared to previous literature which uses yearly survey data on employment and is subject to additional assumptions Dahl and Knepper (2022). My results suggest these policies mechanical reduce recipiency among the unemployed but have no significant effect on posted or reported wages in the labor market as a result.

The remainder of the paper is organized as follows. Section 2 provides institutional background on the U.S. UI system and documents the timeline of state-level benefit duration cuts. Section 3 describes the data sources and variable construction. Section 4 presents the empirical strategy. Section 5 reports the estimates. Section 6 reports what those estimates admit once parallel trends is relaxed. Section 7 presents robustness checks, and Section 8 concludes.

2. Institutional Background

2.1. The UI System in the United States

The U.S. unemployment insurance system is a joint federal-state program established by the Social Security Act of 1935. While the federal government sets broad guidelines and provides administrative funding through the Federal Unemployment Tax Act (FUTA), individual states retain substantial discretion over benefit levels, eligibility criteria, and maximum duration of benefits (Whittaker, 2024). State payroll taxes pay regular benefits, and the federal government pays about half the cost of Extended Benefits (U.S. Department of

Labor, 2024). This decentralized structure creates cross-state variation in UI generosity that has been exploited extensively in the empirical literature (Chetty, 2008; Schmieder et al., 2012; Johnston and Mas, 2018).

A worker is eligible for unemployment insurance if they were involuntarily terminated by an employer that pays unemployment insurance taxes, and if their earnings and quarters worked over a preceding base period clear a state-set threshold. An eligible worker receives a weekly benefit amount that replaces a share of prior earnings, subject to a state maximum, for at most the state's maximum number of weeks. This paper exploits changes to the maximum duration of UI benefits at the state level.

The conventional maximum up until the Great Recession was a 26-week maximum benefit duration. While not every claimant earned 26 weeks of benefits (due to whatever formula determined benefits in that state), all claimants could potentially receive 26 weeks. During economic downturns, the federal government extends the duration of benefits through the Extended Benefits program. These extensions are scaled to state unemployment triggers. Previous literature has examined these benefit extensions and their effects on labor market outcomes (Boone et al., 2021). A cut to regular duration lowers the base those extensions build on and therefore reduces overall potential duration during recessions.

2.2. State Variation in Benefit Duration

The Great Recession prompted unprecedented federal UI benefit extensions that challenged subsequent UI policy surrounding benefit duration. Facing depleted trust funds, several states legislated reductions in maximum benefit duration, instead of raising employer taxes. The most prominent early case was North Carolina, which in 2013 reduced its maximum duration from 26 to between 12 and 20 weeks (depending on the state unemployment rate), simultaneously declining to participate in the federal Emergency Unemployment Compensation (EUC) program. Missouri had reduced its maximum duration in 2011 (Johnston and Mas, 2018). Other states followed, with Georgia, Florida, and Kansas implementing cuts

between 2012 and 2014 (U.S. Government Accountability Office, 2015).

Those initial cuts followed the Great Recession and the extended benefit periods that accompanied it. The post-2016 wave cuts have similar motivations but the important difference is the labor markets these cuts were enacted in. The initial reforms were driven by trust fund insolvency amid poor local labor market conditions (Dahl and Knepper, 2022). The second wave of cuts occurred in labor markets that were tighter, on average, than those seen in states part of the initial wave.

The reforms studied here were enacted with an average state unemployment rate of 3.4% (2.7–4.3), 0.4 percentage points below the national rate, against 8.4% (4.8–10.6) and 0.3 percentage points above it for the 7 reforms Dahl and Knepper (2022) study between 2011 and 2014: a gap of 5.0 percentage points. These cuts were enacted into a labor market that had recovered.

2.3. Timeline of Duration Cuts

Table 1 documents the nine states that constitute the treated group in this analysis. These states reduced their maximum regular UI benefit duration below 26 weeks at various points between 2016 and 2025. The reductions range from modest (4 weeks) to substantial (14 weeks), averaging 9.2 weeks, creating variation not only in the timing of treatment but also in the magnitude of the cut.

Table 1: Treated States: UI Benefit Duration Reductions

State	Effective Date	Before (Weeks)	After (Weeks)	Reduction
Idaho	July 2016	26	21	−5
Alabama	January 2020	26	14	−12
Iowa	July 2022	26	16	−10
Kentucky*	January 2023	26	16	−10
Oklahoma	January 2023	26	16	−10
Montana	July 2023	28	24	−4
Tennessee	December 2023	26	12	−14
Arkansas	January 2024	16	12	−4
Louisiana	January 2025	26	12	−14

Notes: "Before" and "After" refer to the maximum number of weeks of regular UI benefits. Several states index duration to the state unemployment rate so these schedules pay fewer weeks when unemployment is low. Idaho's schedule, for example, ranges from 26 weeks at an unemployment rate of 8 percent or above down to 20 at 2.9 percent or below. For the other indexed states the table shows the floor, which binds at their rates. For Idaho it shows 21 weeks, the value at the 3.0 to 3.9 percent rate Idaho held in most quarters. Idaho's schedule rose to the 26-week ceiling in 2020.

* Kentucky's schedule ran 12–24 weeks under HB 4 (2022) from January 2023, and HB 146 (2023) raised the floor to 16 weeks effective July 2023. Its first two treated quarters therefore paid 12 weeks; the dose uses the 16-week floor that binds after.

I utilize two features of this policy environment for this research design. The cuts were enacted through state legislation and typically cycled through several legislative sessions before passing, so the precise timing of enactment is plausibly idiosyncratic. All reform states vary significantly between announcement date and enactment date of their respective policy. Some of them also overlap the COVID-19 pandemic, during which federal programs (PUA, PEUC, FPUC) temporarily expanded generosity. Section 4 examines both of these notions.

3. Data

The analysis draws on five primary data sources covering reforms effective between 2016 and 2025: Department of Labor administrative records, Current Population Survey (CPS)

microdata, Bureau of Labor Statistics (BLS) Local Area Unemployment Statistics, Glassdoor reported salaries, and Lightcast (from the Burning Glass Institute) job posting data. This section describes each source, the construction of key variables, and the resulting panel datasets. Each model is also dependent on data limitations which are detailed as well. All models estimate on the full treated states (Alabama, Arkansas, Iowa, Idaho, Kentucky, Louisiana, Montana, Oklahoma, and Tennessee) and control states (Alaska, California, Colorado, Connecticut, Washington DC, Delaware, Hawaii, Illinois, Indiana, Maryland, Maine, Minnesota, Mississippi, North Dakota, Nebraska, New Hampshire, New Jersey, New Mexico, Nevada, New York, Ohio, Oregon, Pennsylvania, Rhode Island, South Dakota, Texas, Utah, Virginia, Vermont, Washington, Wisconsin, West Virginia, Wyoming) unless stated otherwise.

3.1. Data Sources

The numerator for the recipiency rate is drawn from the Department of Labor's Employment and Training Administration Form 5159, which provides monthly state-level tabulations of UI claims and payment activity. Specifically, this paper uses "Weeks Compensated", which records the total number of weeks for which UI benefit payments were made in each month at the state level. These data are publicly available through the DOL's online reporting system.[1]

The denominator is constructed from the CPS Basic Monthly microdata. The relevant population is the short-term unemployed, those reporting unemployment duration of 26 weeks or less, which aligns the denominator with the population potentially eligible for regular benefits and excludes the long-term unemployed who would have exhausted them under either regime. The microdata are aggregated to state-month counts using the CPS final person weight, and numerator and denominator are then collapsed across months within each quarter. I refer to this measure as the "(short-term) recipiency rate."

[1]ETA 5159 data are available at https://oui.doleta.gov/unemploy/DataDownloads.asp.

I also estimate the effect cuts have on the seasonally adjusted state unemployment rate. I take this measure from the Bureau of Labor Statistics Local Area Unemployment Statistics (LAUS) program, which publishes a seasonally adjusted monthly unemployment rate for every state. The recipiency rate is collapsed to quarters because its CPS denominator is a survey stock that is noisy at monthly frequency so I estimate the unemployment rate at the quarterly level to maintain consistency. The panel covers 49 quarters for the same 42 states used in the recipiency analysis (9 treated, 33 control), yielding 2,058 state-quarter observations with a mean unemployment rate of 4.43 percent.

Glassdoor data provide anonymized, self-reported compensation records at the firm-occupation level. The extract is aggregated at the yearly level but 95.2% of records summarize a single salary report, so the data are treated as worker level reports. I restrict the sample to new hires, defined as workers reporting zero to two years of relevant experience, to focus on starting salaries which proxy for a reservation wage for those exiting unemployment or switching jobs. Employer names are standardized (lowercased, punctuation removed, whitespace normalized) and the sample is restricted to multi-state firms, those with observations in at least one treated and at least one control state. I bound reported salaries below by the annualized federal minimum wage and above at \$500,000. The Glassdoor data, due to access, ends after 2023 which limits the treated pool in this sample to Alabama, Idaho, Iowa, Kentucky, and Oklahoma. The extract has no observations in 2015 or 2016, so Idaho, which cut in 2016, contributes almost nothing to its own pre-period.

The Burning Glass Institute (BGI), which uses Lightcast and Revelio postings data, provides near-universe data on online job postings in the United States, including posted salary information. I restrict postings to those with valid salary information within plausible bounds, from the annualized federal minimum wage (\$7.25 $\times$ 2,080) to a ceiling of \$500,000.

I then collapse to quarterly panels of establishment-jobs, similarly to Hazell and Taska (2025). I define the establishment as an employer in a county, in a given quarter, and a job as a five-digit 2021 SOC occupation at a given pay frequency. The collapse takes 75.7 million

ads to 39.0 million firm-county cells, an average of 1.94 postings per cell. As with Glassdoor, the sample is restricted to multi-state firms that appear in both treated and control states.

3.2. Summary Statistics

Table 2 compares treated and control states before the treatment date. Each treated state is measured over the four quarters preceding its first full treatment quarter and same for control states to show pre-treatment balance. The treated states are the nine listed in Table 1. The 33 control states are those that remain after nine exclusions (South Carolina, Michigan, Missouri, Kansas, Florida, Georgia, North Carolina, Massachusetts, and Arizona). The eight states whose maximum duration moved within the window are dropped rather than recoded, since states that index duration to their own unemployment rate switch repeatedly and without legislation. South Carolina cut its maximum to twenty weeks in 2011, before the window, and is dropped as well.

Treated states were less generous and less attached to the labor market prior to treatment. The pre-reform recipiency rate is 0.190 in treated states against 0.296 in control states. Differences in means are normalized by the pooled standard deviation following Imbens and Rubin (2015). The labor force participation rate is 60.78 percent against 63.96 also confirming treated states have fewer workers participating. These differ for unemployment and employment growth, however. The unemployment rate differs by -0.29 percentage points and four-quarter employment growth by -0.02 percentage points which gives normalized differences of -0.34 and -0.01, respectively, smaller than for the other labor market measures. I discuss concerns with pre-trends in Section 4.4.

Panels B and C measure the wage samples. New hires in treated states report mean base pay of $43,195 compared to $52,931, and posted salaries average $60,052 compared to $71,521 for treated and control states respectively. These gaps are roughly a third of a standard deviation in logs that reflect industry mix and which states cut benefits. Seven of the nine treated states are observed in the year before treatment in the Glassdoor extract.

Table 2: Pre-Treatment Balance: Treated and Control States Before Reform

	Treated		Control		Diff.	Norm. diff.
	Mean	*SD*	*Mean*	*SD*		
Panel A: State × Quarter (public data)						
Recipiency rate	0.190	0.095	0.296	0.148	-0.106	-0.85
Unemployment rate (%)	3.38	0.57	3.67	1.06	-0.29	-0.34
Employment growth, 4-quarter (%)	1.95	1.23	1.97	1.68	-0.02	-0.01
Labor force participation rate (%)	60.78	3.28	63.96	3.94	-3.17	-0.88
Employment-population ratio (%)	58.73	3.31	61.62	4.05	-2.89	-0.78
Panel B: Worker × State × Year (Glassdoor)						
Mean base pay ($)	$43,195	$24,821	$52,931	$31,990	-$9,736	-0.34
log(base pay)	10.559	0.455	10.735	0.513	-0.176	-0.36
Panel C: Firm-Occ × County × Quarter (BGI/Lightcast)						
Mean posted salary ($)	$60,052	$40,366	$71,521	$45,976	-$11,469	-0.27
log(posted salary)	10.841	0.534	11.015	0.545	-0.173	-0.32
Postings per cell	1.6	6.5	2.0	5.2	-0.3	-0.06

Notes: Each treated state is measured over the four quarters before its first full treatment quarter (Panel A and C) or the calendar year before its first full treatment year (Panel B); control states are measured over the same calendar windows, once per treated state, so both groups are observed in identical periods. Panel A rows are state-quarters: 36 treated and 1,188 control (36 and 1,188 for the recipiency rate). Employment growth is the four-quarter log change in the LAUS employment level, in percent. Panel B rows are salary reports: 5,566 treated and 324,252 control. Panel C rows are firm-occupation-county-quarter cells: 608,388 treated and 30,554,839 control.

4. Empirical Strategy

4.1. Two-Way Fixed Effects

The baseline specification for the state-level outcomes is a two-way fixed effects regression:

$$Y_{st} = \alpha_s + \lambda_t + \delta D_{st} + \varepsilon_{st} \tag{1}$$

where Y_{st} is the recipiency rate or seasonally adjusted unemployment rate for state s in quarter t, α_s are state fixed effects, λ_t are quarter fixed effects that absorb aggregate time trends common to all states, and D_{st} is a binary treatment indicator equal to one if state s has implemented a benefit duration reduction by quarter t. Standard errors are clustered at

the state level.

The recipiency rate is mechanically tied to the maximum benefit duration, since weeks compensated is bounded above by the number of weeks a state pays, so its change captures both the mechanical effect of a shorter maximum and behavioral responses such as changes in take-up or exit, which I do not observe separately. I therefore show the effective dose of the policy rather than a behavioral treatment effect. For the unemployment rate, Y_{st} is the seasonally adjusted unemployment rate described in Section 3, collapsed from months to state-quarters so that both outcomes of Equation (1) are estimated on the same quarterly panel. Quarters with fewer than three observed months are dropped, so a partial quarter cannot enter as though it were a full one. The unemployment rate as a measure is not deterministically linked to the maximum duration of benefits. Therefore, one would interpret δ as the effect of that dose on the share of the labor force recorded as unemployed. This would not coincide with an effect on employment as a δ of zero could indicate offsetting movements in employment and participation or no effect at all.

For the wage outcomes, the specification includes firm-level fixed effects:

$$\log(\text{Pay})_{efst} = \beta D_{st} + \phi_e + \alpha_s + \lambda_t + \varepsilon_{efst} \tag{2}$$

where $\log(\text{Pay})_{efst}$ is the log of base pay for employer e (or employer-occupation f) in state s at time t, ϕ_e are firm fixed effects (or firm-by-job fixed effects for the BGI analysis), and the remaining terms are defined as before. The firm fixed effects absorb permanent differences across employers in pay levels, so the identifying variation comes from within-firm changes in pay across treated and control states surrounding the policy change.

Restricting to multi-state firms is only valid if firms do not set wages nationally. If a firm posts one figure for a job in every state in which it hires, no state policy can move the within-firm contrast and the estimated effect will be zero whether or not UI generosity affects wages. I test this for each employer-job by measuring the posting-weighted spread across states and

its deviation from the employer-job by quarter mean, using pre-reform observations only so the classification cannot itself be an outcome of treatment. The median employer-job's posted wage varies by 0.0682 log points across state lines, with an interquartile range of 0.0291 to 0.1153, and only 14.4% of employer-jobs, from 2.9% of ads, post within 0.010 log points of a single national figure. These results suggest firm wages in this sample do vary by state.

4.2. Event Study Design

To trace the dynamic path of treatment effects, the analysis estimates event study specifications:

$$Y_{st} = \alpha_s + \lambda_t + \sum_{\substack{k=-K \\ k\neq -1}}^{K} \mu_k \cdot \mathbf{1}[\text{EventTime}_{st} = k] + \varepsilon_{st} \tag{3}$$

where EventTime_{st} is the number of periods (quarters for the state-level outcomes and the posted-wage panel; years for Glassdoor) between quarter or year t and the quarter or year in which state s first implemented its benefit duration cut. The coefficients of interest $\{\mu_k\}$ trace out the effect of the policy change at each lead and lag relative to the omitted reference period ($k = -1$, the quarter immediately before treatment). The event window is $K = 8$ quarters for quarterly data and $K = 4$ years for annual data. The establishment-level LEHD analysis (Model 3) is monthly and uses $K = 10$ months around treatment and is estimated with the stacked design of Cengiz et al. (2019) that interacts cohort-specific "stack" fixed-effects with every included FE rather than Equation (3). The recipiency, unemployment-rate and posting-volume event studies are estimated with TWFE. The posted-wage and Glassdoor estimates in Section 5.3 use the imputation estimator of Borusyak et al. (2024), which omits no reference period. Section 4.4 sets out the choice. Each figure and table states its estimator.

4.3. Identification Assumptions and Threats

The standard identifying assumption for event study models is parallel trends; absent the benefit duration cuts, treated and control states would have experienced the same trends in labor market outcomes. It is conventionally assessed by testing whether the coefficients prior to treatment are jointly indistinguishable from zero or each other.

A state that cuts duration differs from a comparison state in ways that may also move its labor market, whether the motive is a preference for a less generous safety net or trust fund insolvency. State and employer-job fixed effects absorb these differences where they are permanent. The design therefore requires that adopting states were not on a different trend from non-adopting states. The within-employer-job comparison I make supports this assumption. Holding the employer and the job title fixed removes permanent pay differences across employers and with it the concern that the firms hiring in reforming states are systematically different firms. A multi-state design removes effects correlated with a firm's decision to operate in a treated state compared to a control state. The remaining variation identifies a state-by-quarter shock that moves relative pay across establishments of the same employer posting the same job. I also assume that the *timing* of enactment is idiosyncratic conditional on selection into treatment, and two features of the reforms support this. The first is that duration-cut bills recur across sessions and mostly fail. Missouri, for example, considered and rejected benefit reduction bills in three separate years. For states that did enact, the gap between passage and the date benefits fell is highly variable, running from 15 days in Iowa to 31 months in Tennessee. Which session finally passes a bill therefore appears to depend more on legislative factors than on the state's contemporaneous labor market. Where the motive is fiscal, the argument is the one Dahl and Knepper (2022) make for their own setting, since cuts spurred by UI insolvency reflect the benefits paid in the previous downturn rather than conditions in the quarter the cut takes effect.

I test whether timing responds to the labor market with a discrete-time hazard model among the nine treated states. Each state is observed monthly from January 2014 until its

bill is announced, dated to the first public text of the cut, and leaves the sample afterward. The hazard is regressed on the change in the state unemployment rate over the prior six months, with calendar-month fixed effects and log months at risk. The coefficient is 0.008 (s.e. =0.130), and it remains insignificant when announcement is dated to the introduction of the enacting bill instead. This suggests recent changes in the state unemployment rate do not predict when states announce a cut (Table A9). Dating the event to the governor's signature instead gives a coefficient of -0.277 (s.e. =0.120), which is significant at the 5% level under clustered standard errors but not under a permutation test ($p = 0.138$). With nine events the test has limited power, so the timing evidence is suggestive rather than conclusive.

The cuts also did not follow a deterioration in treated labor markets. Over the 8 quarters before each cut, the unemployment rate fell from 4.01 to 3.37 percent in treated states and from 5.02 to 3.62 percent in control states over the same calendar quarters. Unemployment therefore fell in both groups, by 0.75 percentage points less in treated states, which is the gap closing that the unemployment-rate leads show (Section 5.2).

I also assume that states and their labor markets do not adjust in advance of the cut. Extended legislative deliberation makes some anticipation possible, and Tennessee is an example of this. Its cut was enacted in May 2021 but did not fully come into effect until December 2023, giving 31 months of advance notice. Pre-treatment movement would be one signal of such adjustment of anticipated cuts. I therefore re-run the regressions with treatment dated to enactment rather than the operative date as a check (Section 5.3). Finally I assume no cross-state spillovers. Cross-state labor mobility could violate this, and the county identifiers in the posted-wage extract allow the assumption to be tested rather than asserted. Section 7.4 re-estimates the posted-wage specification separately on counties that sit on a treated/control state line, where cross-border search should bias the estimate toward zero, and on counties that do not. In both cases I find the results insensitive to these checks.

The sample period also overlaps the COVID-19 pandemic and the federal UI expansions

(PUA, PEUC, FPUC) that accompanied it, which temporarily overrode state-level variation in duration by providing additional weeks of benefits and higher weekly benefit amounts. Quarter fixed effects absorb the common component of pandemic-era disruption but not a shock that fell differently on treated and comparison states. I therefore re-estimate the model removing the pandemic quarters (2020Q2–2021Q4), dropping 294 of the 2,058 state-quarter observations. The TWFE recipiency effect is -5.99 percentage points on the remaining 1,764 observations, against -7.26 on the full panel which suggests the estimate is not an artifact of the pandemic quarters.

Finally, states that reduce UI generosity may simultaneously implement other labor market policies that independently affect recipiency or wages. No treated state raised employer taxes to compensate for program insolvency, and state fixed effects absorb permanent differences in program and regulatory infrastructure (the ways benefits are applied for, claimed, program requirements, capacity, etc.). They do not absorb a policy that changes when the duration cut does. One treated state is an example of this: Arkansas bundles its cut with a work-search mandate and an employer tax cut on a single effective date.[2] I therefore re-estimate the model dropping the Arkansas cohort rather than recoding it as a control and the result is insensitive to the change (Sections 5.3 and 7).

4.4. Estimators and Inference

The TWFE estimator in Equation (1) can be biased in staggered-adoption settings when treatment effects are heterogeneous across cohorts or over time since treatment (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Roth et al., 2023). The concern is limited here. The Goodman-Bacon (2021) decomposition places 91.5% of the TWFE weight on comparisons of treated with never-treated states, so already-treated states rarely serve as controls.

[2] Act 196 of 2023 cut the maximum duration for claims filed on or after 1 January 2024. On that same date Act 587 raised the work-search requirement to five employer contacts per week, and employer unemployment taxes fell. A work-search mandate is a labor supply shock pointing the same way as a duration cut, so the Arkansas cohort's contribution to the average effect is not attributable to UI duration alone.

For the recipiency rate, the unemployment rate and posting volume I therefore report TWFE. I report robust estimates that account for staggered adoption using de Chaisemartin and D'Haultfœuille (2020) in Appendix A.2. For posted wages and Glassdoor starting pay, whose models include employer-job or firm fixed effects, I use the imputation estimator of Borusyak et al. (2024). I report its native average, which estimates $e = -1$, and its standardized average. Every standard error is clustered by state.

Four of the five outcomes estimated below have pre-trends that reject: the recipiency rate (Section 5.1), the unemployment rate (Section 5.2), Glassdoor starting pay (Section 5.3), and posting volume (Section 5.4). The posted-wage leads drift smoothly before reform (Section 5.3), and 0 of 8 reject individually. Section 5 reports the conventional estimates, which assume parallel trends. Given these violations, Section 6 reports bounds that relax that assumption.

I also employ a wild cluster bootstrapping procedure to handle the small number of treated states. With 9 treated states, asymptotic standard errors clustered at the level of policy variation are known to substantially over-reject (Cameron et al., 2008; MacKinnon and Webb, 2018). This affects every outcome, since treatment is assigned by state and every estimate clusters by state. I report restricted wild cluster bootstrap p-values for the TWFE estimates of the recipiency rate, the unemployment rate and posting volume in Section 7.1. The imputation estimates for posted wages and Glassdoor starting pay use analytic clustered standard errors. The Glassdoor effect is identified by few treated states (Section 5.3), so its standard error likely overstates precision. Appendix A.3 reports the bootstrap in full.

5. Results

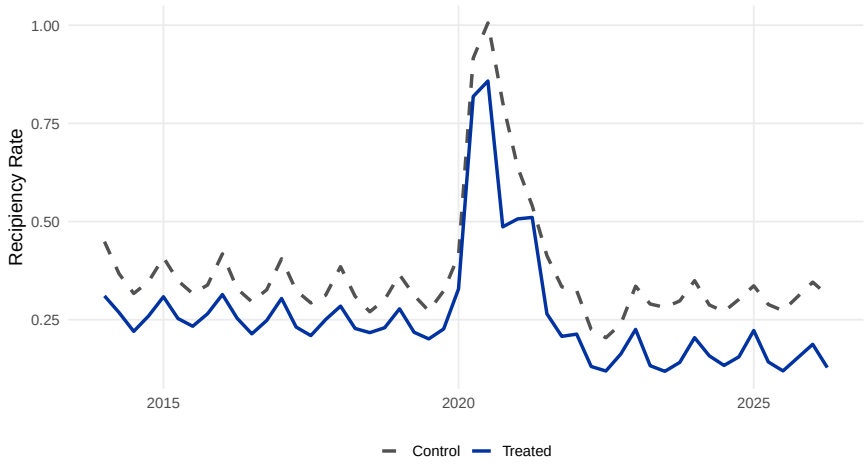
5.1. UI Recipiency

The first question is whether the reductions in maximum benefit duration affected the weeks of UI actually paid. Figure 1 plots the raw series for treated and control states. Both

groups trend similarly over the sample period, with the treated group below the control throughout. These trends both spike during the COVID-19 pandemic before widening from 2021 on. TWFE returns an average effect of -7.26 percentage points, an associated decrease in the short-term recipiency rate (s.e. = 2.31).[3] The decrease is significant at the 1% level under analytic clustered standard errors and at the 5% level under the restricted wild cluster bootstrap (Cameron et al., 2008; MacKinnon and Webb, 2018). Within the ±8-quarter window of the event study, the estimate is -8.27 percentage points (s.e. = 4.48), significant at the 10% level, so the stronger rejection relies on the quarters furthest from the cut. However, recipiency falls in 6 of 7 treatment cohorts, and dropping any one cohort leaves an estimate between -5.80 and -8.87 (Appendix A.4). A cut can only lower the weeks a state pays, so a one-sided test is also appropriate. It rejects a non-negative effect with p <0.001 for the static estimate and $p = 0.046$ for the post-period event-study average. Figure 2 plots the event study. Post-treatment effects turn negative immediately, with the largest decreases three to four quarters after the cut, a path similar to previous literature (Dahl and Knepper, 2022). I report robust alternative estimates in Appendix A.2.

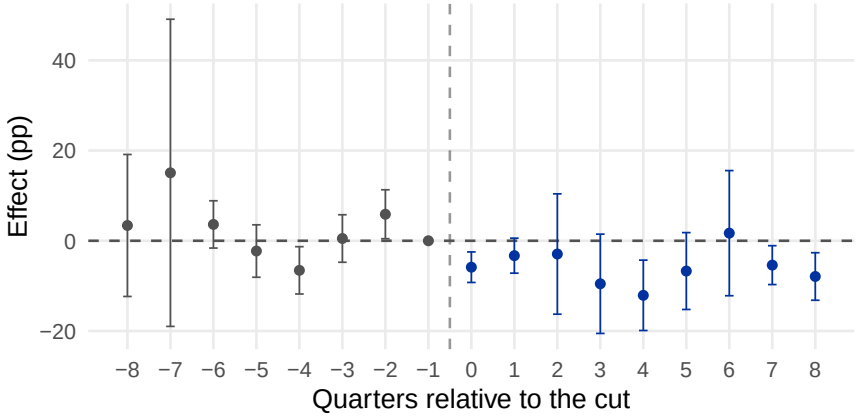
[3]I report recipiency effects in percentage points, the rate multiplied by 100.

Figure 1: UI Recipiency Rate in Treated and Control States



Notes: Quarterly mean recipiency rate, defined as average weekly UI claimants (weeks compensated divided by the weeks in the quarter) per short-term unemployed worker, for treated states (solid) and control states (dashed), in calendar time.

Figure 2: Event Study: UI Recipiency Rate



Notes: TWFE event study (Equation (3)) of the effect of a duration cut on the recipiency rate, in percentage points, by quarter relative to the cut, with state and quarter fixed effects. Coefficients are relative to $k = -1$, which is omitted and set to zero. Bars are 95% confidence intervals; standard errors clustered by state.

The TWFE leads jointly reject (p <0.001), whereas the placebos of de Chaisemartin and D'Haultfœuille (2020) on the same window do not ($p = 0.268$). The far leads are imprecise,

and the CPS denominator is noisy. I relax the parallel trends assumption in Section 6. I treat the recipiency rate here as evidence that the policy binds, since it moves mechanically with the maximum duration a state pays.

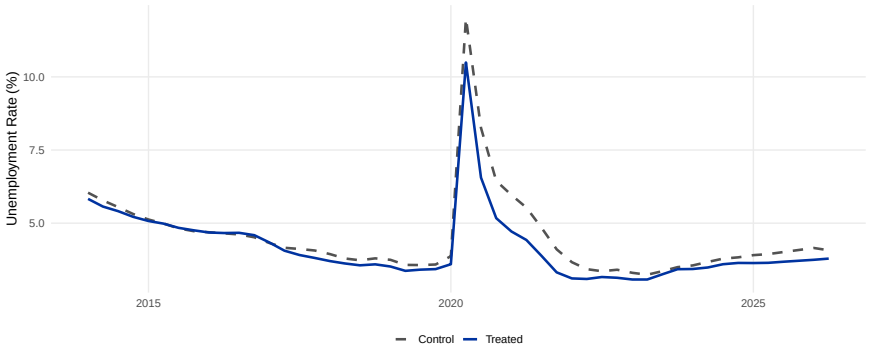
The decrease comes from the numerator. In log points multiplied by 100, the effect on weekly claimants is -43.7 (s.e. =10.4) and on the short-term unemployed -2.7 (s.e. =5.9), an implied effect on the ratio of -41.0 (Appendix Figure A3). Fewer weeks of benefits are therefore paid, while the number of short-term unemployed does not change detectably.

5.2. The State Unemployment Rate

Section 5.1 establishes that the legislated cuts bind. The second question is whether binding produced a response large enough to appear in the state unemployment rate. On the seasonally adjusted LAUS unemployment rate over 2,058 state-quarter observations, the TWFE average effect is -0.05 percentage points (s.e. = 0.31). The estimate is statistically indistinguishable from zero, and its 95% confidence interval of [-0.67, 0.57] excludes the previous estimates from Dahl and Knepper (2022).

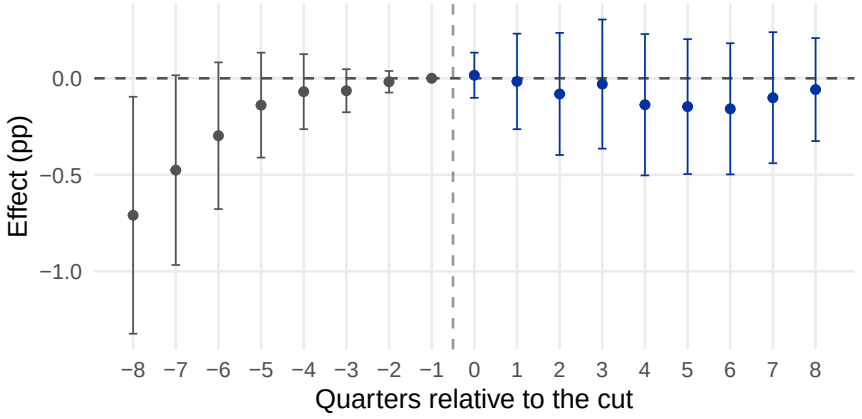
Figure 3 plots the treated and control series in calendar time, and Figure 4 presents the TWFE event study. The leads rise toward zero in the quarters before treatment, and they jointly reject (p <0.001). Treated states sit below control states in levels throughout, so the rise reflects the gap closing as control states recovered from a higher pandemic peak, not rising unemployment in treated states. After treatment no coefficient is individually distinguishable from zero. The sign of the static estimate depends on Alabama, whose cut took effect in the quarter before the pandemic; without it the estimate is 0.26 percentage points (s.e. = 0.16), and no leave-one-cohort-out estimate is distinguishable from zero (Appendix A.4). I still apply the bounding exercise in Section 6 which relaxes the parallel trends assumption and re-estimates intervals for post-treatment estimates.

Figure 3: Unemployment Rate in Treated and Control States



Notes: Quarterly mean of the seasonally adjusted LAUS unemployment rate, in percent, for treated states (solid) and control states (dashed), in calendar time.

Figure 4: Event Study: Unemployment Rate



Notes: TWFE event study (Equation (3)) of the effect of a duration cut on the seasonally adjusted LAUS unemployment rate, in percentage points, by quarter relative to the cut, with state and quarter fixed effects. Coefficients are relative to $k = -1$, which is omitted and set to zero. Bars are 95% confidence intervals; standard errors clustered by state.

These estimates do not rule out an effect on the level of employment since a rise in employment can be offset by a fall in participation. I therefore estimate Equation (1) for the LAUS employment-population ratio and labor force participation rate on the same quarters, through 2026Q2. The employment-population ratio changes by 0.57 percentage points (s.e.

= 0.43) and participation by 0.57 (s.e. = 0.35). Neither is distinguishable from zero, and their leads do not reject ($p = 0.420$ and 0.515, respectively). I therefore find no evidence that a fall in participation offsets a change in employment (Appendix Figure A4).

In the CPS, the long-term unemployment rate, the share of the labor force unemployed for more than 26 weeks, changes by 0.05 percentage points (s.e. = 0.14), and the prime-age (25–54) employment-population ratio by 0.39 (s.e. = 0.37). However, the leads reject for both (p <0.001 and $p = 0.002$), so I report these as descriptive (Appendix Figure A5).

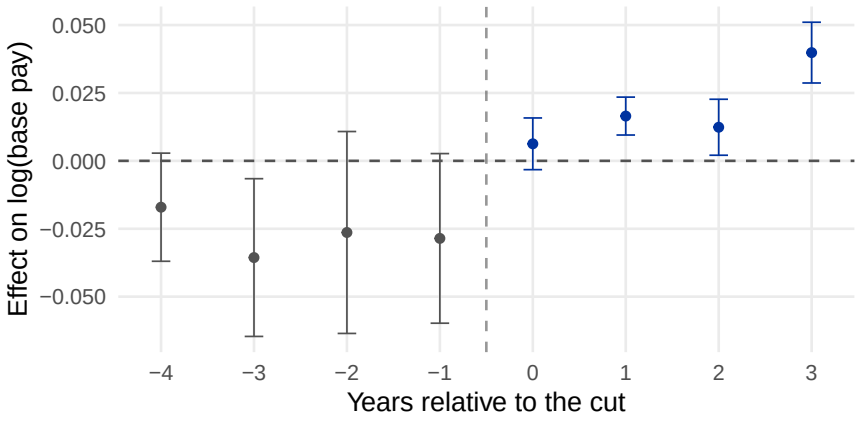
5.3. Wage Results

5.3.1 Reported Starting Salaries

I estimate the effect on log reported starting salaries in Glassdoor with the imputation estimator of Borusyak et al. (2024), with firm, state and year fixed effects and standard errors clustered by state. After the cut the average effect is an increase of 1.51 percent (s.e. = 0.32). Reported starting pay therefore rises after the cuts, which contrasts with previous literature in direction of effect (Dahl and Knepper, 2022). By horizon the effect is 0.0063 log points in the year of the cut and 0.0399 three years after.

I do not read this rise as a wage response to the cuts, for two reasons. First, the leads jointly reject ($p = 0.0015$). Figure 5 shows reported pay in treated states below that of the comparison group in every pre-treatment year, so the post-treatment rise cannot be separated from the closing of that gap. Relative to $e = -1$ the average is larger, 0.0436 log points, because the estimated pre-treatment coefficients are negative. Second, few states identify the effect. In the year of the cut, 5 treated states contribute. After the first year only Alabama and Idaho identify it, and Alabama supplies most of those observations (Appendix Table A5).

Figure 5: Event Study: Starting Salaries (Glassdoor)



Notes: Borusyak et al. (2024) imputation estimates of the effect of a duration cut on log base pay of new hires (0–2 years of experience), by year relative to the cut. No period is omitted; the pre-treatment coefficients, including $k = -1$, are estimated. Firm, state, and year fixed effects. Bars are 95% confidence intervals; standard errors clustered by state.

5.3.2 Posted Wages

Table 3 reports the main posted-wage estimates. I compare postings made by the same employer for the same occupation and pay frequency between reform and non-reform states. Here, I hold employer-job composition fixed which lets me compare posted salaries for the same job within the same company across state lines so that the remaining variation is the change in policy. Estimation uses the Borusyak et al. (2024) imputation estimator on 26.1 million establishment-job × quarter cells, with standard errors clustered by state.

The average treatment effect on the treated suggests an increase of 0.44 percent, which is significant at the 10% level but not at 5%.[4] In the event study specification relative posted wages in treated states sit below the control group two years out and converge smoothly by the quarter before reform. Therefore the post-treatment path has no clear separation from a continuation of that drift. This continues into a slight increase in post reform estimates that are statistically significant at the 5% level for three quarters before flattening out and

[4]This is the native imputation average, which estimates $e = -1$ rather than omitting it.

becoming statistically insignificant after.

The bounds in Section 6 require a confidence set relative to $e = -1$, and therefore a change of basis, since the estimator omits no reference period (Borusyak et al., 2024). Re-expressing the event-study vector relative to $e = -1$ gives an average of 0.0009 (s.e. $=$ 0.0058), which is not statistically different from zero. This suggests there is no discernible effect on posted wages due to a change in the maximum duration of UI benefits.

The estimate is insensitive to the robustness checks in Appendix A.6. Dropping the Arkansas cohort, whose cut shares an effective date with a work-search mandate and an employer tax cut, or dating treatment to enactment does not meaningfully move the ATT estimate. Given these specifications are still subject to pre-trends, I apply the same trend-robust bounding exercise in Section 6, which requires no specification choice. The multiple robustness checks do not meaningfully move any estimates. Further tests and specifications can be found in Appendix A.6.

Table 3: Effect of UI Duration Cuts on Posted Wages

Quarters relative to reform	Coefficient	(SE)
$e = -8$	-0.0107	(0.0104)
$e = -7$	-0.0123	(0.0101)
$e = -6$	-0.0110	(0.0083)
$e = -5$	-0.0025	(0.0069)
$e = -4$	-0.0040	(0.0074)
$e = -3$	-0.0007	(0.0069)
$e = -2$	0.0022	(0.0061)
$e = -1$	0.0035	(0.0069)
$e = 0$	0.0039**	(0.0020)
$e = 1$	0.0059**	(0.0024)
$e = 2$	0.0061**	(0.0026)
$e = 3$	0.0040	(0.0026)
$e = 4$	0.0035	(0.0029)
$e = 5$	0.0032	(0.0032)
ATT	0.0044*	(0.0025)
relative to $e = -1$	0.0009	(0.0058)
robust to a linear trend	[-0.0251, 0.0176]	

Notes: Borusyak et al. (2024) imputation estimator on establishment-job × quarter cells (employer-job by county within state); employer-job, state and quarter fixed effects; SEs clustered by state. No period is omitted, so $e = -1$ is estimated rather than set to zero. Pre-period coefficients are placebos; 0 of 8 reject zero. The ATT row is the estimator's own output. The last two rows re-express it relative to $e = -1$, the basis the Rambachan and Roth (2023) framework requires, and report its bound at $\bar{M} = 0$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Sections 6 and 8 examine the drift from previous literature and what I can definitively say about these new reforms. My estimate differs in sign and is insignificant at the 5% level, and the Dahl and Knepper (2022) estimate falls outside my conventional confidence set of [-0.0106, 0.0124]. I attribute the difference to a different policy environment. Previous literature focused on cuts after the Great Recession, while these cuts look at an entirely new set of states, mainly after the COVID-19 recession. Additionally, the previous reforms take effect at a mean state unemployment rate of 8.4% whereas the second wave takes place at a mean state unemployment rate of 3.4%. The two waves therefore differ by 5.0 percentage points in mean state unemployment rates for treated states.

The reforms also differ in size. Weighting each treated state by its contribution of treated

observations to the post-reform window, the average reform in this sample reduces maximum duration by 42.6%, compared to 23.1% in the six moderate-reform states from which Dahl and Knepper (2022) estimate the corresponding parameter. Both sets of reforms primarily cut duration, and weekly benefit amounts moved little (Section 6), so the comparison is between two duration shocks of different size.

5.4. Job-Posting Volume

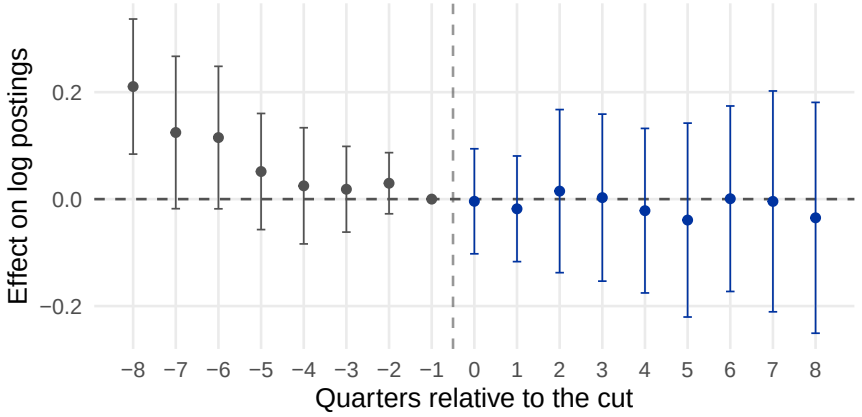
The posted-wage analysis measures the intensive margin of a change in posted wages. The final outcome asks whether the cuts moved the extensive margin: the number of postings. Equilibrium search predicts more postings in response to maximum duration cuts, since a lower bargained wage raises the surplus a filled vacancy returns and induces entry (Hagedorn et al., 2019), although Marinescu (2017) finds that longer benefits reduce applications without significantly changing vacancies. Withdrawing benefits from households with high marginal propensities to consume lowers local spending and the derived labor demand (Ganong and Noel, 2019). Kroft and Notowidigdo (2016) find that the search response to UI generosity is larger when unemployment is low so a worker-side response should be stronger in the tighter labor markets. This conflicts with the demand channel: where UI is a small share of local income, cutting it should do little to local spending (Ganong and Noel, 2019). The effect of a reduction in UI benefits on job posting volume is therefore ambiguous.

I reconstruct the full BGI/Lightcast sample and aggregate postings that carry a usable posted salary to state-by-quarter cells. The outcome is the log count of postings. The quarter effects absorb the growth in Lightcast coverage over the sample, which is equivalent to expressing each cell as a state's share of national postings. Standard errors are clustered at the state level.

The TWFE estimate suggests an associated change of -10.61 percent (s.e. = 7.82). Under parallel trends, its 95% confidence interval of [-24.70, 6.11] rules out increases larger than 6.11 percent. The post-period average of the event-study coefficients is much smaller, -0.0115

log points (Section 6), likely because the static estimate also absorbs the pre-treatment gap. It is not statistically significant, including under the wild cluster bootstrap, and the estimate does not scale with the size of the cut (Table 4). Figure 6 plots the event study. The leads jointly reject ($p = 0.001$), which suggests treated states were already losing relative posting share before their cuts. Section 6 shows that allowing for that pre-trend does not yield an informative bound. I therefore find no evidence that the cuts moved posting volume, although the data cannot rule out moderate effects in either direction.

Figure 6: Event Study: Job-Posting Volume



Notes: TWFE event-study estimates of Equation (3) for log postings in state × quarter cells, by quarter relative to the cut. Coefficients are relative to $k = -1$, which is omitted and set to zero. State and quarter fixed effects. Bars are 95% confidence intervals; standard errors clustered by state.

Table 4: Effect of UI Duration Cuts on Job-Posting Volume

	Log postings
TWFE treated × post	-0.1122
	(0.0875)
Clustered p-value	0.207
Wild cluster bootstrap p-value	0.242
TWFE treated × post × weeks cut	-0.0149
	(0.0179)

Notes: Outcome is log postings in state × quarter cells, counting postings that carry a usable posted salary. TWFE estimates Equation (1) with state and quarter fixed effects. The dose row interacts treatment with weeks cut. Standard errors clustered by state in parentheses. The wild cluster bootstrap p-value imposes the null and resamples Rademacher weights over state clusters (1,999 replications).

6. Sensitivity to Parallel Trends

Every estimate in Section 5 assumes that treated and comparison states would have moved together absent the cuts. However, the pre-period rejects this for four of the five outcomes, and the posted-wage leads drift smoothly before reform. In this section I relax the parallel trends assumption and report bounds in place of the point estimates.

I use the partial-identification framework of Rambachan and Roth (2023). Each event-study coefficient in Equation (3) is the treatment effect plus a violation of parallel trends,

$$\mu_k = \tau_k + \nu_k, \qquad \nu_{-1} = 0, \tag{4}$$

and the pre-treatment coefficients measure that violation before the cuts, since $\tau_k = 0$ for $k < 0$. The framework restricts how the violation may evolve after treatment and returns a confidence set under each restriction. The target in every case is the equal-weighted post-

period average relative to $e = -1$,

$$\theta = \frac{1}{\bar{K}+1} \sum_{k=0}^{\bar{K}} \tau_k, \tag{5}$$

where $\bar{K} = K$ for the state-level outcomes and $\bar{K} = 5$ for posted wages. I apply two restrictions to every outcome. The first, relative magnitudes, bounds each post-treatment change in the violation by $\bar{M}$ times the largest change between consecutive pre-treatment periods:

$$|\nu_{k+1} - \nu_k| \leq \bar{M} \max_{-K \leq j \leq -2} |\nu_{j+1} - \nu_j| \qquad \text{for all } k \geq -1. \tag{6}$$

$\bar{M} = 0$ returns the conventional interval. The second bounds the change in the slope of the violation,

$$|(\nu_{k+1} - \nu_k) - (\nu_k - \nu_{k-1})| \leq M \qquad \text{for all } k, \tag{7}$$

which Rambachan and Roth (2023) denote Δ^{SD}. With $M = 0$ it allows the violation to continue a linear pre-treatment trend of any slope. For this restriction I report the conditional (C-LF) confidence set, which also tests whether the leads fit a line, and Table 6 reports the fixed-length interval (FLCI) that Rambachan and Roth (2023) recommend. The inputs are the TWFE event studies for the recipiency rate, the unemployment rate and posting volume, and the imputation coefficients re-expressed relative to $e = -1$ for posted wages.

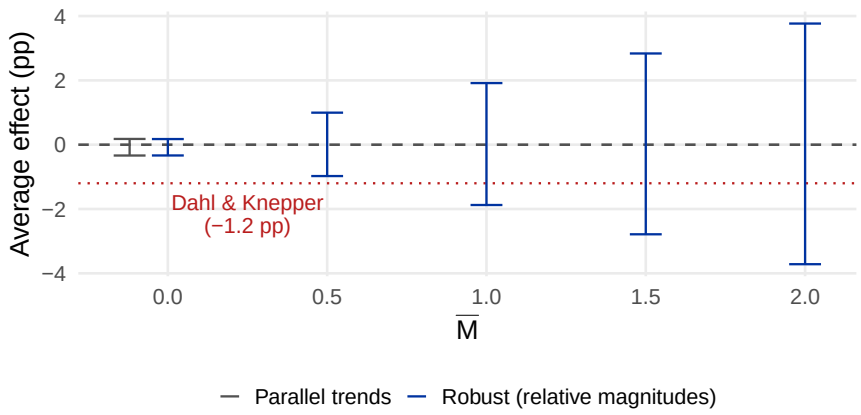
For an estimate that is already indistinguishable from zero, the smallest $\bar{M}$ at which the robust set admits zero is uninformative. I therefore report the smallest $\bar{M}$ at which the relative-magnitudes set admits the estimate of Dahl and Knepper (2022). This value measures how large a trend violation is needed to reconcile these data with their finding. Recipiency has no breakdown value because its set contains zero even at $\bar{M} = 0$, and Dahl and Knepper (2022) report no posting-volume estimate. Their moderate reforms remove 23.1% of potential duration, against 42.6% here. A level comparison therefore holds the effect fixed while my dose is larger, which favors finding a larger effect here. Table 5 collects

the bounds.

Recipiency Under exact parallel trends, the post-period average of the TWFE event-study coefficients is -5.78, with a confidence set of [-12.50, 0.93] that contains zero. The event-study average is therefore smaller and less precise than the static TWFE estimate of Section 5.1. The linear-trend set, [-6.0, 53.2], also contains zero. The relative-magnitudes restriction is uninformative for this outcome. Its benchmark is the largest change between consecutive pre-treatment coefficients (11.67), which the imprecise far leads set. The robust set at $\bar{M} = 1$ is therefore [-198.7, 183.6] which is uninformative given the bounds are set by an outlying lead estimate.

The Unemployment Rate Figure 7 plots the bounds. The confidence set for the post-period average runs from -0.337 to 0.179 percentage points under exact parallel trends. Permitting a differential linear trend widens it to [-0.44, 0.10], which still excludes the -1.2 percentage point effect that Dahl and Knepper (2022) report. The relative-magnitudes set is wider, [-1.88, 1.92] at $\bar{M} = 1$, and first admits -1.2 at $\bar{M} = 0.63$. For the first post-treatment quarter the breakdown value is >3. The null therefore survives a differential linear trend. However, it does not survive post-treatment changes in the violation of 0.63 times the largest pre-treatment change.

Figure 7: Sensitivity to Parallel-Trends Violations: Unemployment Rate



Notes: 95% confidence sets for the average post-period effect on the unemployment rate, in percentage points. Gray: exact parallel trends. Blue: Rambachan and Roth (2023) relative-magnitudes sets by $\bar{M}$. Dotted line: the -1.2 percentage point estimate of Dahl and Knepper (2022).

Posted Wages Under exact parallel trends, the confidence set for the average posted-wage effect relative to $e = -1$ is [-0.0106, 0.0124]. Permitting a differential linear trend of arbitrary slope widens it to [-0.0251, 0.0176], and relative magnitudes at $\bar{M} = 1$ give [-0.0503, 0.0596]. All three sets contain zero. The Dahl and Knepper (2022) estimate lies outside the conventional set but inside both robust ones, and the relative-magnitudes set admits it at $\bar{M} = 0.15$. Once parallel trends is relaxed, therefore, the data cannot reject their estimate in levels.

The confidence interval of the native imputation average, 0.0044, also excludes their estimate. The imputation estimator fits its pre-treatment coefficients on untreated observations only, so they measure the violation less directly than the leads of a conventional event study.

The FLCI for the linear trend is narrower, [-0.0151, -0.0002], and excludes zero. However, it extrapolates the leads as an exact line, and the leads reject a line ($J = 53.1$, $p < 0.001$). The FLCI admits zero once the trend may bend by 0.00066 log points per quarter, a small fraction of the 0.0099 the leads already bend. Posted wages therefore fall only under strict linear extrapolation, and that decline does not survive the curvature already present in the

pre-period.

To compare this bound with the wider literature, I express it as an elasticity of posted wages with respect to UI generosity. Rinz and Wasser (2026) are the only benchmark that estimates this elasticity. I therefore use the definition of Rinz and Wasser (2026): a regression of log wages on log simulated maximum potential benefits, the weekly benefit amount times the maximum number of weeks. The matching denominator is the log change in maximum potential benefits, 0.516 on average across treated observations.[5] On this definition, the Rambachan and Roth (2023) bound is [-0.034, 0.049] under a differential linear trend of arbitrary slope and [-0.116, 0.097] under relative magnitudes at $\bar{M} = 1$. I report the bound rather than the point estimate. These data are difficult to reconcile with elasticities larger than roughly 0.049 in absolute value under a linear trend.

Rinz and Wasser (2026) estimate the elasticity for CPS new hires out of unemployment. Their baseline includes state-year fixed effects, so their variation comes from differences across worker cells within a state and year rather than from reforms. The comparison is therefore between two different sources of variation and two wage measures, realized and posted. They report approximately 0.21 over 2005–2019 and 0.04 over 2020–2024, which covers much of my sample. My linear-trend and relative-magnitudes bounds exclude the 2005–2019 elasticity and contain the 2020–2024 one. These bounds also contain zero, so the 2020–2024 estimate is consistent with these data. The FLCI is narrower and excludes zero and both Rinz and Wasser (2026) estimates. Dahl and Knepper (2022) find that posted and reported salaries move in the same direction. Here reported starting pay rises while posted wages do not move. However, the reported-pay leads reject parallel trends (Section 5.3), so I do not read the two measures as diverging.

[5] five of the nine treated states show a higher weekly maximum in the semester their cut took effect. For four of them the increase is the state's routine indexation step, which comparison states receive as well and which the difference-in-differences removes. Tennessee is the exception: its maximum rose independently, although its maximum potential benefits still fall overall. The denominator attributes every one of those increases to the reform, which is the conservative reading. In weeks alone, the average reform cuts potential duration by 42.6%.

Posting Volume Under exact parallel trends, the post-period average of the TWFE event-study coefficients is only -0.0115, and its confidence set, $[-0.1544, 0.1314]$, contains zero. The relative-magnitudes set at $\bar{M} = 1$ is $[-0.753, 0.750]$, and the linear-trend set is $[-0.150, 0.424]$. All three sets contain zero. The FLCI for the linear trend, [0.065, 0.189], is positive, which is the direction equilibrium search predicts. However, the leads reject a line ($J = 12.7$, $p = 0.048$), and the interval admits zero once the trend may bend by 0.0021, against 0.076 in the leads. I therefore find no robust evidence that the cuts changed the volume of job postings.

Across the four outcomes, the bounds are informative for two. For recipiency the exercise returns nothing, so the evidence that the cuts bind rests on the mechanical link between recipiency and the maximum weeks a state pays. For posting volume the interval is too wide to rule out a moderate job-creation response. For the unemployment rate, a differential linear trend still excludes the effect Dahl and Knepper (2022) report, although relative magnitudes admit it at $\bar{M} = 0.63$. For posted wages, the bound excludes the 2005–2019 wage elasticity of Rinz and Wasser (2026) under both restrictions and admits their 2020–2024 one. Equilibrium search predicts that a benefit cut lowers posted wages and raises postings (Burdett and Mortensen, 1998; Hagedorn et al., 2019). I observe both margins in the same data, and neither moves detectably. However, the relaxed sets admit moderate effects in levels on both margins, and only the wage margin yields a bound tight enough to rule out a previous estimate.

Table 5: Sensitivity to Violations of Parallel Trends

Outcome	Prior estimate	Exact parallel trends	Relative magnitudes	Linear trend	Breakdown $\bar{M}$
Recipiency rate (pp)	—	[-12.50, 0.93]	[-198.7, 183.6]	[-6.0, 53.2]	—
Unemployment rate (pp)	-1.2	[-0.337, 0.179]	[-1.88, 1.92]	[-0.44, 0.10]	0.63
Posted wages (log)	-0.0135	[-0.0106, 0.0124]	[-0.0503, 0.0596]	[-0.0251, 0.0176]	0.15
Posted-wage elasticity	0.21		[-0.116, 0.097]	[-0.034, 0.049]	
Posting volume (log)	—	[-0.1544, 0.1314]	[-0.753, 0.750]	[-0.150, 0.424]	—

Notes: 95% robust confidence sets (Rambachan and Roth, 2023). Relative magnitudes: $\bar{M} = 1$. Linear trend: a differential linear trend of any slope ($M = 0$). Breakdown $\bar{M}$: the smallest $\bar{M}$ at which the set admits the Dahl and Knepper (2022) estimate. Elasticities divide the posted-wage sets by the log change in maximum potential benefits (a decrease of 0.516); the prior elasticity is Rinz and Wasser (2026), 2005–2019.

Table 6: Linear-Trend Bounds: Fixed-Length Confidence Intervals

Outcome	FLCI, linear trend	Linear fit J (p)	M admitting zero	Lead curvature
Recipiency rate (pp)	[-9.2, 4.2]	39.2 (<0.001)	—	23.0
Unemployment rate (pp)	[-0.18, 0.08]	28.1 (<0.001)	—	0.089
Posted wages (log)	[-0.0151, -0.0002]	53.1 (<0.001)	0.00066	0.0099
Posting volume (log)	[0.065, 0.189]	12.7 (0.048)	0.0021	0.076

Notes: FLCI: Rambachan and Roth (2023) fixed-length confidence interval under a differential linear trend ($M = 0$). J tests that the leads lie on a line (6 degrees of freedom). M admitting zero: the smallest smoothness bound at which the FLCI contains zero, shown only where the $M = 0$ interval excludes it. Lead curvature: the largest absolute second difference of the leads. Posted wages: Borusyak et al. (2024), relative to $e = -1$; other outcomes: TWFE.

7. Robustness

This section reports checks that vary the estimator and the outcome margin. Checks that vary the specification or sample within the posted-wage design appear with that result, in Section 5.3 and Appendix A.6.

7.1. Alternative Estimators

The state-level headline estimates are TWFE (Section 4.4). Appendix A.2 reports de Chaisemartin and D'Haultfœuille (2020) as an alternative. For recipiency, it agrees with TWFE on sign but is less precise. For the unemployment rate, it agrees on sign and is more precise (Table A2). The Goodman-Bacon (2021) decomposition places 91.5% of the TWFE weight on comparisons of treated with never-treated states. Neither the sign of the recipiency estimate nor the unemployment-rate null is therefore an artifact of negative weighting. Under TWFE the recipiency decrease is significant. A restricted wild cluster bootstrap (Cameron et al., 2008; MacKinnon and Webb, 2018), which is robust to the small number of treated clusters, raises its p-value from 0.00313 to 0.033 over 1,999 replications. The TWFE estimate also does not rest on Arkansas, the one cohort whose reform bundles three policies (Section 4). Dropping it returns -7.35 against -7.26 on the full panel (Table 7).

Table 7: Effect of UI Duration Cuts on the Recipiency Rate: TWFE by Sample

	(1)	(2)	(3)
Sample	Full	Ex. pandemic	Ex. Arkansas
Treated × post	-7.26	-5.99	-7.35
	(2.31)	(2.08)	(2.54)
Clustered p-value	0.00313	0.00629	0.00616
Wild cluster bootstrap p-value	0.033		
Clean-comparison weight (%)	91.5		92.2
State-quarters	2,058	1,764	2,009

Notes: Effects in percentage points (the recipiency rate multiplied by 100). TWFE estimates Equation (1) with state and quarter fixed effects. Column (2) drops 2020Q2–2021Q4; column (3) drops the Arkansas cohort. The wild cluster bootstrap p-value imposes the null and resamples Rademacher weights over state clusters (1,999 replications). The clean-comparison weight is the Goodman-Bacon (2021) share of the TWFE weight on treated-versus-never-treated comparisons. Standard errors clustered by state in parentheses.

Table 8: Effect of UI Duration Cuts on the Unemployment Rate

	TWFE
Treated × post	-0.0511
	(0.3146)
95% confidence interval	[-0.668, 0.565]
Clustered p-value	0.872
Wild cluster bootstrap p-value	0.920
Clean-comparison weight (%)	91.5
State-quarters	2,058

Notes: Outcome is the seasonally adjusted LAUS unemployment rate in percentage points, state-quarters from 2014-Q1 to 2026-Q2. TWFE estimates Equation (1) with state and quarter fixed effects. The wild cluster bootstrap p-value imposes the null and resamples Rademacher weights over state clusters (1,999 replications). The clean-comparison weight is the Goodman-Bacon (2021) share of the TWFE weight on treated-versus-never-treated comparisons. Standard errors clustered by state in parentheses.

Table 9 reports the three state-panel outcomes under five estimators, with never-treated controls and state clusters. Recipiency falls under every estimator, from -7.26 to -5.70 percentage points. It is significant at the 5% level under TWFE, Callaway and Sant'Anna (2021) and imputation, but not under Sun and Abraham (2021) or de Chaisemartin and D'Haultfœuille (2020). No estimator finds a significant effect on the unemployment rate or on posting volume.

Table 9: State-Panel Outcomes by Estimator

	Recipiency (pp)		Unemp. rate (pp)		Volume (log pts)	
Estimator	ATT	SE	ATT	SE	ATT	SE
TWFE	-7.26	2.31	-0.05	0.31	-11.22	8.75
Callaway and Sant'Anna (2021)	-6.42	3.18	-0.06	0.13	-3.15	6.46
Sun and Abraham (2021)	-5.70	3.51	-0.09	0.13	-1.34	6.89
de Chaisemartin and D'Haultfœuille (2020)	-5.70	3.57	-0.09	0.13	-1.34	6.65
Imputation (Borusyak et al., 2024)	-7.15	1.91	-0.04	0.31	-11.98	8.08

Notes: Never-treated controls; standard errors clustered by state. Sun and Abraham (2021) and de Chaisemartin and D'Haultfœuille (2020): average of the effects 0–8 quarters after the cut, which coincide in point estimates with never-treated controls on the same window. Sun and Abraham (2021) and imputation standard errors: state cluster bootstrap, 499 draws, since each cohort holds one or two states. Volume in log points multiplied by 100.

7.2. Matched Control States

Treated states had lower unemployment and recipiency than the never-treated states before the cut. I therefore re-estimate Equation (1) on control groups matched to the treated states. For each treated state I keep the k never-treated states nearest by Mahalanobis distance over its 8 quarters before the cut, on the mean unemployment rate and its change, the change in the employment-population ratio, mean recipiency and the maximum weekly benefit. With $k = 2$, 14 control states remain, and their pre-cut unemployment rate and recipiency are 3.5 and 26.5 percent, against 3.6 and 26.2 in treated states, respectively. The recipiency effect is -6.91 percentage points (s.e. =2.56) and the unemployment-rate effect -0.16 (s.e. =0.30), whose leads no longer reject at the 5% level ($p = 0.077$). However, the match uses pre-cut values of the outcomes themselves, so the improved balance is partly built in (Table 10).

Table 10: Matched Control States

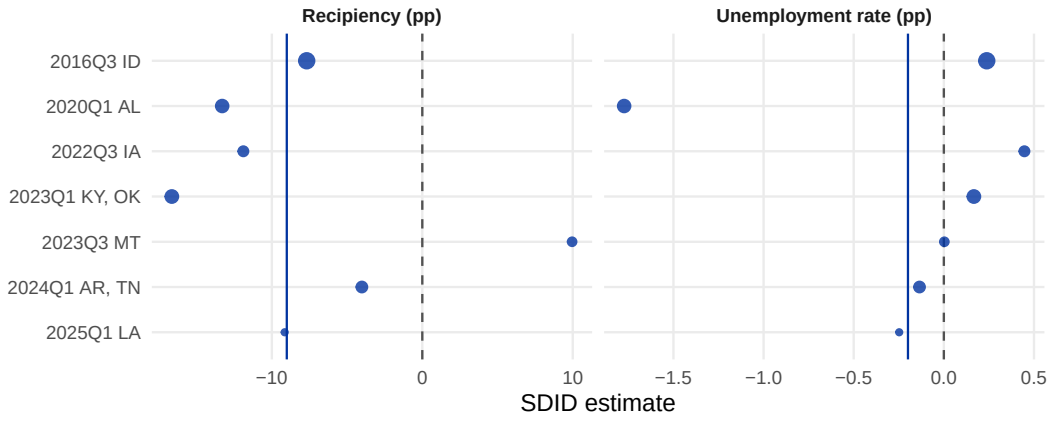
		Recipiency (pp)			Unemp. rate (pp)		
Controls	States	TWFE	SE	Leads p	TWFE	SE	Leads p
All never-treated	33	-7.26	2.31	<0.001	-0.05	0.31	<0.001
Nearest 1	8	-9.56	2.43	0.003	-0.18	0.32	0.128
Nearest 2	14	-6.91	2.56	<0.001	-0.16	0.30	0.077
Nearest 3	19	-6.83	2.46	<0.001	-0.14	0.29	0.073
Before the cut							
	Treated	All controls			Nearest 2		
Unemployment rate (%)	3.6	4.1			3.5		
Change in unemployment rate	-0.65	-1.40			-0.84		
EPOP (%), not matched	58.5	61.2			63.4		
Change in EPOP	0.99	1.22			0.93		
Recipiency (%)	26.2	34.2			26.5		
Maximum weekly benefit ($)	447	536			489		

Notes: Static TWFE (Equation (1)), state and quarter fixed effects, standard errors clustered by state. Leads p: joint test of the event-study leads. Nearest k: the k never-treated states closest to each treated state over its 8 quarters before the cut, pooled with replacement. Balance rows average over treated states, each on its own window.

7.3. Synthetic Difference-in-Differences

I also estimate the state-panel effects with synthetic difference-in-differences (Arkhangelsky et al., 2021), which reweights control states and pre-period quarters to match each treated state's pre-period path. The estimator takes one adoption date, so I estimate each cohort against the 33 never-treated states and average the cohort estimates, weighted by treated state-quarters after the cut. Standard errors come from 200 placebo draws. The recipiency effect is -9.01 percentage points (s.e. = 2.43), against -7.26 under TWFE, and the unemployment-rate effect -0.20 (s.e. = 0.22), against -0.05. The cohort estimates vary (Figure 8).

Figure 8: Synthetic Difference-in-Differences by Adoption Cohort



Notes: Arkhangelsky et al. (2021) estimates for each adoption cohort against the never-treated states, using every quarter of the balanced panel. Point size: treated state-quarters after the cut. Solid line: the average weighted by those state-quarters. Percentage points.

7.4. Cross-State Spillovers: Border-County Design

Section 4 assumes the cuts do not spill across state lines. If workers search across a border, a control county adjacent to a treated state is partially treated. This compresses the treated-control contrast and biases the estimate toward zero. Contamination of this kind should be strongest in counties on a treated/control border and weakest in the interior. I therefore test whether the estimate in the border band is attenuated relative to the interior. Relative to $e = -1$, the band returns 0.0007 against 0.0009 in the interior, with confidence sets of [-0.0073, 0.0087] and [-0.0078, 0.0097], respectively (Table A7). The band holds 22.3% of the employer-jobs, against 83.6% in the interior, yet its confidence set is marginally narrower. The two estimates are approximately the same. I therefore find no evidence that cross-border search compresses the contrast or produces the null of Section 5.3.

Of the 2,616 cross-state county adjacencies, 309 counties fall in the band. I estimate both halves with the Borusyak et al. (2024) imputation estimator on establishment-job × quarter cells, as in Section 5.3, so the split changes only the sample. This is not the contiguous

county-pair design of Dube et al. (2010), which estimates a different parameter. Because one employer-job can post in both a band and an interior county, the comparison is a diagnostic rather than a formal test.

8. Conclusion

Nine states cut their maximum UI benefit duration below the 26-week standard between 2016 and 2025. Recipiency changes by -7.26 percentage points under TWFE. However, the TWFE leads reject and the bounds are uninformative, so the evidence that the cuts bind is mainly mechanical. On wages, I use a multi-state firm design that compares postings and salary reports for the same job within the same employer across state lines. I find no detectable change in posted wages. Reported starting pay rises by 1.51 percent, but its leads reject and after the first year only two states identify it. Bounds that permit a differential linear trend of arbitrary slope rule out elasticities of posted wages larger than 0.049 in absolute value. The state unemployment rate did not move either. A differential linear trend still excludes the effect Dahl and Knepper (2022) report, although relative magnitudes admit it at $\bar{M} = 0.63$.

I argue that the contrast with Dahl and Knepper (2022) may reflect labor market conditions rather than the policy itself. The reforms studied here took effect when the mean state unemployment rate was 3.4%, compared to 8.4% for the reforms Dahl and Knepper (2022) study. The pool of unemployed workers whose search behavior a benefit cut could change was therefore smaller in my sample. Kroft and Notowidigdo (2016) find that the search response to UI is larger when unemployment is low, so the response per worker may have been larger here. However, far fewer workers were exposed, and the aggregate effect depends on both. The motivation for the cuts is similar across the two waves: fiscal pressure on state trust funds. However, given the timing of reforms, the economic environment, and my data, I cannot discern whether I would estimate similar effects as the previous literature.

My results have four limitations. First, the pre-period rejects parallel trends for four of

the five outcomes and drifts for the fifth. For this reason the bounds in Section 6, rather than the conventional confidence sets, carry the substantive claims. Second, the wage data are proprietary and carry no demographic fields, so whether the reforms changed the composition of new hires remains open. Third, the unemployment-rate null concerns a ratio. The employment-population ratio and participation also show no detectable change, but their intervals admit increases of up to 1.41 and 1.25 percentage points, respectively. Fourth, my design speaks to posted and reported wages, not to the welfare consequences of the cuts. Where a cut returns workers to the labor market sooner, I observe nothing about the quality of the matches they form.

Two extensions remain for future work. First, heterogeneity across industry and firm size is available in these data and unexploited here. Second, recent evidence that UI generosity moves the wages of job changers and continuously employed workers, not only new hires from unemployment (Rinz and Wasser, 2026), suggests that wage spillovers to non-displaced workers within treated states are a more promising margin than the one studied here.

References

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118.

Boone, C., Dube, A., Goodman, L., and Kaplan, E. (2021). Unemployment insurance generosity and aggregate employment. *American Economic Journal: Economic Policy*, 13(2):58–99.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6):3253–3285.

Burdett, K. and Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 39(2):257–273.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics*, 134(3):1405–1454.

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy*, 116(2):173–234.

Dahl, G. B. and Knepper, M. M. (2022). Unemployment insurance, starting salaries, and jobs: Evidence from multi-state firms. NBER Working Paper 30152, National Bureau of Economic Research.

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996.

Dube, A., Lester, T. W., and Reich, M. (2010). Minimum wage effects across state borders: Estimates using contiguous counties. *Review of Economics and Statistics*, 92(4):945–964.

Ganong, P. and Noel, P. (2019). Consumer spending during unemployment: Positive and normative implications. *American Economic Review*, 109(7):2383–2424.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Econometrica*, 89(5):2261–2290.

Hagedorn, M., Manovskii, I., and Mitman, K. (2019). The impact of unemployment benefit extensions on employment: The 2014 employment miracle? *American Economic Review*, 109(12):4224–4264.

Hazell, J. and Taska, B. (2025). Downward rigidity in the wage for new hires. *American Economic Review*, 115(12):4183–4217.

Imbens, G. W. and Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction.* Cambridge University Press, Cambridge.

Johnston, A. C. and Mas, A. (2018). Potential unemployment insurance duration and labor supply: The individual and market-level response to a benefit cut. *Journal of Political Economy*, 126(6):2480–2522.

Kroft, K. and Notowidigdo, M. J. (2016). Should unemployment insurance vary with the unemployment rate? theory and evidence. *Review of Economic Studies*, 83(3):1092–1124.

MacKinnon, J. G. and Webb, M. D. (2018). The wild bootstrap for few (treated) clusters. *Econometrics Journal*, 21(2):114–135.

Marinescu, I. (2017). The general equilibrium impacts of unemployment insurance: Evidence from a large online job board. *Journal of Public Economics*, 150:14–29.

McCall, J. J. (1970). Economics of information and job search. *Quarterly Journal of Economics*, 84(1):113–126.

Mortensen, D. T. (1977). Unemployment insurance and job search decisions. *Industrial and Labor Relations Review*, 30(4):505–517.

Rambachan, A. and Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591.

Rinz, K. and Wasser, D. (2026). Unemployment insurance generosity and wage determination. Working Paper 26-11, Federal Reserve Bank of Cleveland.

Roth, J., Sant'Anna, P. H. C., Bilinski, A., and Poe, J. (2023). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics*, 235(2):2218–2244.

Schmieder, J. F., von Wachter, T., and Bender, S. (2012). The effects of extended unemployment insurance over the business cycle: Evidence from regression discontinuity estimates over 20 years. *Quarterly Journal of Economics*, 127(2):701–752.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

U.S. Department of Labor (2024). Unemployment compensation: Federal-state partnership. Technical report, Office of Unemployment Insurance, Division of Legislation.

U.S. Government Accountability Office (2015). Unemployment insurance: States' reductions in maximum benefit duration have implications for federal costs. Report GAO-15-281, GAO.

Whittaker, J. M. (2024). Unemployment insurance: Programs and benefits. Report RL33362, Congressional Research Service.

A. Appendix

A.1. TWFE Event-Study Coefficients

Table A1 reports the TWFE event-study coefficients plotted in Figures 2 and 4.

Table A1: TWFE Event-Study Coefficients: Recipiency and Unemployment Rates

	Recipiency rate	Unemployment rate
$e = -8$	3.39 (7.79)	-0.709 (0.304)
$e = -7$	15.07 (16.86)	-0.476 (0.243)
$e = -6$	3.63 (2.59)	-0.298 (0.188)
$e = -5$	-2.27 (2.88)	-0.139 (0.134)
$e = -4$	-6.55 (2.60)	-0.069 (0.096)
$e = -3$	0.50 (2.61)	-0.064 (0.055)
$e = -2$	5.88 (2.69)	-0.018 (0.028)
$e = 0$	-5.86 (1.68)	0.016 (0.058)
$e = 1$	-3.30 (1.92)	-0.016 (0.123)
$e = 2$	-2.94 (6.61)	-0.081 (0.157)
$e = 3$	-9.54 (5.45)	-0.030 (0.166)
$e = 4$	-12.09 (3.86)	-0.137 (0.181)
$e = 5$	-6.71 (4.22)	-0.147 (0.173)
$e = 6$	1.69 (6.87)	-0.158 (0.168)
$e = 7$	-5.40 (2.13)	-0.101 (0.168)
$e = 8$	-7.92 (2.61)	-0.059 (0.132)
Joint test of leads, p-value	<0.001	<0.001
State-quarters	1,764	1,764

Notes: Coefficients μ_k of Equation (3), by quarter relative to the cut, with state and quarter fixed effects; $e = -1$ is omitted. Recipiency in percentage points (the rate multiplied by 100); unemployment rate in percentage points. Treated observations outside $e = -8, \ldots, 8$ are dropped. The joint test is a Wald test that the seven leads are zero. Each cell gives the estimate and, in parentheses, its standard error clustered by state.

A.2. Alternative Estimator: de Chaisemartin and D'Haultfœuille

As an alternative to TWFE for the state panels, I estimate the effects with the estimator of de Chaisemartin and D'Haultfœuille (2020), using the Stata package `did_multiplegt_dyn` with 9 effects ($e = 0, \ldots, 8$), 8 placebos and only never-treated states as controls. Its average

effect weights each horizon by the number of treated states observed at it. Table A2 reports the estimates.

For recipiency, the average effect is -5.70 percentage points (s.e. = 3.57), which is not statistically significant. Its placebos do not jointly reject ($p = 0.268$), whereas the TWFE leads on the same window do ($p < 0.001$). For the unemployment rate, the average effect is -0.090 percentage points (s.e. = 0.128), and its placebos reject ($p = 0.002$).

Table A2: State Panel: TWFE and de Chaisemartin–D'Haultfœuille

	Recipiency rate	Unemployment rate
TWFE treated × post	-7.26	-0.0511
	(2.31)	(0.3146)
dCdH average effect	-5.70	-0.090
	(3.57)	(0.128)
dCdH 95% confidence interval	[-12.69, 1.30]	[-0.340, 0.161]
dCdH placebo test, p-value	0.268	0.002
TWFE joint test of leads, p-value	<0.001	<0.001

Notes: Recipiency in percentage points (the rate multiplied by 100); unemployment rate in percentage points. TWFE is the static estimate of Equation (1). The de Chaisemartin and D'Haultfœuille (2020) average effect is from `did_multiplegt_dyn` with 9 effects ($e = 0, \ldots, 8$), 8 placebos and never-treated states as controls; it weights each horizon by the number of treated states observed at it. The placebo test is the package's joint test of the placebos. Standard errors clustered by state in parentheses.

A.3. Inference with Few Treated Clusters

With 9 treated states, asymptotic standard errors clustered at the level of policy variation over-reject substantially (Cameron et al., 2008; MacKinnon and Webb, 2018). The de Chaisemartin and D'Haultfœuille (2020) standard errors are analytic and rest on the same few clusters. I therefore apply the restricted wild cluster bootstrap of MacKinnon and Webb (2018) to the TWFE recipiency effect of Section 7.1, imposing the null and resampling Rademacher weights over the 42 state clusters (1,999 replications). The bootstrap p-value is 0.033, against an analytic clustered p-value of 0.00313. The TWFE effect therefore remains significant at the five percent level, but not at the one percent level the analytic p-value implies. I apply the same bootstrap to the unemployment rate and posting volume (Tables 8

and 4). The posted-wage estimates also cluster by state, but the bootstrap is not feasible for the imputation estimator on 26,060,534 observations, so the same caution applies to them.

A.4. Event Window and Treatment Cohorts

The static TWFE estimates use every quarter of the panel, and treated states differ in how long they are observed after the cut. Idaho (2016Q3) and Alabama (2020Q1) contribute most of the post-treatment quarters, and 46.04% of treated post-treatment observations fall more than eight quarters after the cut. I therefore check whether the estimates depend on the far quarters, on any single cohort, or on how cohorts are weighted.

Table A3 restricts the treated states to quarters within $\pm K$ of the cut, while never-treated states keep every quarter. The recipiency estimate is -6.29, -8.27 and -6.50 percentage points for $K = 4$, 8 and 12, respectively, against -7.26 on all quarters. The point estimate is stable, but it is less precise inside the window. The unemployment-rate estimate is not distinguishable from zero in any window.

Table A3: Static TWFE by Event Window

	Recipiency rate (pp)				Unemployment rate (pp)			
Treated quarters	Est.	(s.e.)	p	p_{WCR}	Est.	(s.e.)	p	p_{WCR}
± 4	-6.29	(3.03)	0.044	0.039	-0.02	(0.15)	0.885	0.891
± 8	-8.27	(4.48)	0.072	0.087	0.14	(0.23)	0.547	0.572
± 12	-6.50	(3.70)	0.087	0.140	0.28	(0.28)	0.310	0.344
All quarters	-7.26	(2.31)	0.003	0.028	-0.05	(0.31)	0.872	0.918

Notes: TWFE estimates of Equation (1) with state and quarter fixed effects. Rows restrict treated states to quarters with event time in $[-K, K]$; never-treated states keep every quarter. Standard errors clustered by state. p_{WCR}: restricted wild cluster bootstrap with Webb weights, 9,999 draws.

Table A4 estimates each cohort against the never-treated states alone and drops one cohort at a time from the pooled estimate. Figure A1 plots the first. Recipiency falls in 6 of 7 cohorts; Montana is the exception. The simple average across cohorts is -6.75 percentage points and the average weighted by treated observations is -7.17, against -7.26

pooled. Dropping Alabama gives -6.60 (s.e. = 2.60).

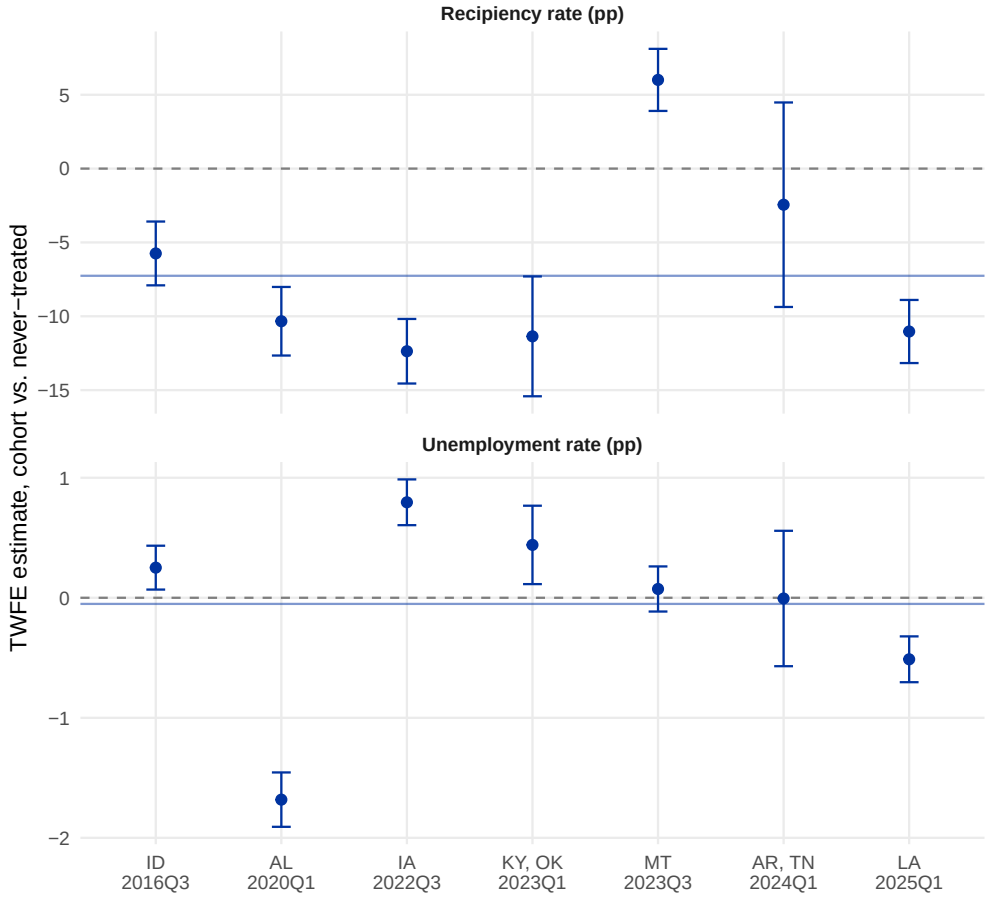
The unemployment-rate cohort estimates differ in sign. Alabama's is -1.68 percentage points, and the average of the other cohorts is 0.17. Alabama's post-period is the pandemic and recovery. Its unemployment rate peaked at 10.93% in 2020Q2, against 12.01% in control states, and stood at 3.97% in 2021Q1, against 5.96%. The Alabama estimate therefore reflects a smaller pandemic shock and a faster recovery, which the cut cannot be separated from. Dropping any one cohort leaves an estimate between -0.20 and 0.26, none of them distinguishable from zero. The cohort-alone standard errors rest on one or two treated clusters and understate the uncertainty, so I read those estimates for their sign and size only.

Table A4: Static TWFE by Treatment Cohort

		Recipiency rate (pp)				Unemployment rate (pp)			
Cohort	Treated obs.	Cohort alone		All but cohort		Cohort alone		All but cohort	
ID (2016Q3)	39	-5.75	(1.10)	-7.36	(2.54)	0.25	(0.09)	-0.08	(0.35)
AL (2020Q1)	25	-10.34	(1.18)	-6.60	(2.60)	-1.68	(0.12)	0.26	(0.16)
IA (2022Q3)	15	-12.37	(1.12)	-6.43	(2.51)	0.80	(0.10)	-0.18	(0.33)
KY, OK (2023Q1)	26	-11.36	(2.07)	-5.80	(2.70)	0.44	(0.17)	-0.20	(0.38)
MT (2023Q3)	11	6.01	(1.07)	-8.87	(1.82)	0.07	(0.10)	-0.07	(0.35)
AR, TN (2024Q1)	18	-2.45	(3.53)	-8.43	(2.54)	-0.01	(0.29)	-0.06	(0.38)
LA (2025Q1)	5	-11.03	(1.09)	-7.02	(2.44)	-0.51	(0.10)	-0.03	(0.33)
All cohorts (pooled)	139			-7.26	(2.31)			-0.05	(0.31)
Simple average of cohorts		-6.75		-0.09					
Weighted by treated obs.		-7.17		-0.08					

Notes: "Cohort alone" estimates each cohort against the never-treated states only. "All but cohort" drops that cohort from the pooled sample. Treated obs. is the cohort's number of treated state-quarters. Standard errors (in parentheses) clustered by state; with one or two treated clusters the cohort-alone standard errors understate the uncertainty.

Figure A1: Static TWFE Estimate for Each Cohort Against Never-Treated States



Notes: Each point is a TWFE estimate from one cohort and the never-treated states. The solid line is the pooled estimate. Bars are 95% confidence intervals from state-clustered standard errors, which understate the uncertainty with one or two treated clusters.

A.5. Glassdoor Regression Sample by Treated State

Table A5 reports the Glassdoor estimation sample by treated state and the event-time horizons at which each state contributes post-reform observations.

Table A5: Glassdoor Regression Sample by Treated State

State	First full treated year	Obs.	Pre	Post	Years post
ID	2017	1,810	2	1,808	7
AL	2020	4,567	446	4,121	4
IA	2023	3,211	2,635	576	1
KY	2023	4,593	3,699	894	1
OK	2023	4,086	3,333	753	1
AR	2024	2,611	2,611	0	0
MT	2024	676	676	0	0
TN	2024	8,941	8,941	0	0
LA	2025	3,561	3,561	0	0

Notes: Worker-level Glassdoor observations in multi-state firms. The extract ends in 2023, so states whose first fully-treated year falls at or after 2024 contribute no post-period observations and do not identify the treatment effect.

A.6. Posted Wages: Alternative Specifications and Pre-Trends

Table A6 re-estimates Table 3 with treatment dated to each state's enactment rather than its operative date, and Table A7 drops the Arkansas cohort and splits the sample into border and interior counties (Section 7.4). Every confidence set contains zero in both the native and the normalized convention.

Table A6: Posted Wages: Alternative Specifications

Specification	ATT (native)	ATT (norm.)	Robust to linear trend
Baseline, firm-clustered	0.0044	0.0009	[-0.0155, 0.0080]
Excluding Montana	0.0017	0.0008	[-0.0148, 0.0109]
+ Federal withdrawal × quarter	0.0085	-0.0013	[-0.0140, 0.0085]
Treatment dated to enactment	0.0026	0.0083	[-0.0085, 0.0229]
Enactment dating and withdrawal	0.0087	0.0055	[-0.0151, 0.0127]

Notes: each row re-estimates Table 3 under one change. "Native" is the Borusyak et al. (2024) output with no omitted period; "norm." subtracts the $e = -1$ coefficient. The final column is the Rambachan and Roth (2023) smoothness bound at $M = 0$. Every row clusters by firm, including the baseline; the state-clustered headline is in Table 3.

Table A7: Posted Wages: Sample Restrictions and Alternative Design

Specification	ATT, native	ATT, rel. $e = -1$	95% confidence set	Leads rejecting
Baseline, firm-clustered	0.0044 (0.0044)	0.0009 (0.0045)	[-0.0079, 0.0097]	3/8
Employers that price by state	0.0057 (0.0046)	0.0006 (0.0057)	[-0.0106, 0.0118]	3/8
Excluding Idaho and Alabama	0.0083 (0.0049)	0.0004 (0.0049)	[-0.0091, 0.0100]	5/8
Excluding Arkansas	0.0054 (0.0046)	0.0008 (0.0041)	[-0.0073, 0.0089]	4/8
Cells with at least two ads	0.0071 (0.0052)	0.0001 (0.0060)	[-0.0116, 0.0118]	1/8
Border counties	0.0107 (0.0058)	0.0007 (0.0041)	[-0.0073, 0.0087]	6/8
Interior counties	0.0043 (0.0044)	0.0009 (0.0045)	[-0.0078, 0.0097]	3/8
Stacked DiD	0.0094 (0.0074)			

Notes: Each row re-estimates Table 3 with the Borusyak et al. (2024) imputation estimator on the stated sample. "Native" is the estimator's output with no omitted period; "rel. $e = -1$" subtracts the $e = -1$ coefficient, and the 95% confidence set is for that average. Leads rejecting counts pre-period coefficients that reject zero at 5%. Price by state keeps employer-jobs whose pre-reform cross-state spread is at least 0.010 log points. The stacked row follows Cengiz et al. (2019) on state × quarter cells (Appendix A.9), with standard errors clustered by state. All other rows, including the baseline, cluster by firm; the state-clustered headline is in Table 3.

Pre-treatment coefficients. The leads in Table 3 are not zero, although none rejects under state clustering. The pattern is a smooth convergence rather than a discrete jump. Relative posted wages in treated states sit below the comparison group two years before reform, with a far-lead average of -0.0113 log points, and reach parity in the quarter before reform, with no break at $e = 0$. I tested five explanations.

First, treatment coding. I verified every treatment date and dose against statute amendment histories and enrolled bill text. Three states were mis-coded, one because the recorded reform was a bill that died in committee. Correcting them did not change the pattern. Second, the sample window. Moving the start of the panel back from 2018 to 2014 did not change it. Third, contamination of the comparison group. Removing the eight states whose duration changed within the window did not change it. Fourth, the unit of observation. Rebuilding the panel from state-pooled cells to establishment-jobs removed the displacement in the five quarters immediately before reform. That portion was therefore an artifact of averaging across establishments within a state. Fifth, anticipation. Several reforms were enacted well before they bound, Tennessee's by 31 months. Re-dating treatment to each

state's enactment quarter pushes the near leads negative but leaves the far leads in place.

A differential trend of approximately 0.24% per quarter therefore remains, concentrated in the six to eight quarters before reform and only partly explained. For the estimates in Table 3, the linear-trend C-LF set contains zero, although the FLCI under the same restriction excludes it (Section 6).

A.7. Dose Response in the Recipiency Rate

The treatment indicator in Equation (1) is binary, but the cuts range from 4 to 14 weeks. If the recipiency decline is a mechanical consequence of a shorter maximum duration, it should scale with the size of the cut. This tests the mechanical reading directly, at the level of the identifying variation.

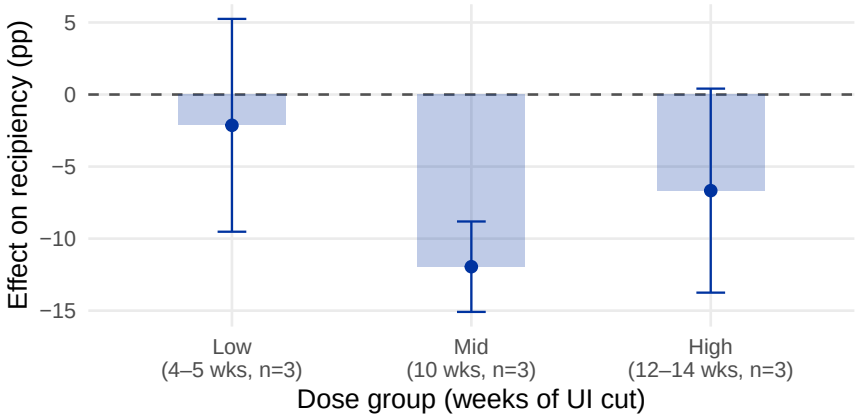
I split treated states into terciles of weeks cut and estimate a separate post-treatment coefficient for each. The effect is -2.14 percentage points (s.e. =3.66) in the smallest-cut tercile, -11.95 (s.e. =1.55) in the middle, and -6.67 (s.e. =3.51) in the largest. Table A8 reports the estimates and Figure A2 plots them.

The ordering only partly matches the mechanical account. The smallest cuts produce the smallest and least precise effect, as expected. However, the largest-cut tercile returns a smaller effect than the middle one, so the relationship is not monotone. I do not read this as evidence against the mechanism. With 9 treated states, each tercile holds a third of them, so each coefficient rests on very few clusters, and the middle-tercile standard error is implausibly small for the reason given in Section 4.4. The sign is the same in all three terciles, and the ranking among them is not identified at this sample size.

I also estimate a continuous version, interacting the treatment indicator with weeks cut centered at the mean cut of 9.2 weeks (range 4–14). The interaction is the change in the effect per additional week cut. It is -0.62 percentage points (s.e. =0.68) for recipiency, -0.10 (s.e. =0.05) for the unemployment rate and -1.49 (s.e. =1.79) log points multiplied by 100 for posting volume. No slope differs from zero at the 5% level ($p = 0.368$, 0.073 and 0.409,

respectively).

Figure A2: Recipiency Effect by Size of Duration Cut



Notes: TWFE post-treatment coefficients on the recipiency rate, in percentage points, for treated states grouped into terciles by weeks cut, relative to never-treated states. Axis labels give each group's range of weeks cut and number of states. Bars are 95% confidence intervals. State and quarter fixed effects; standard errors clustered by state.

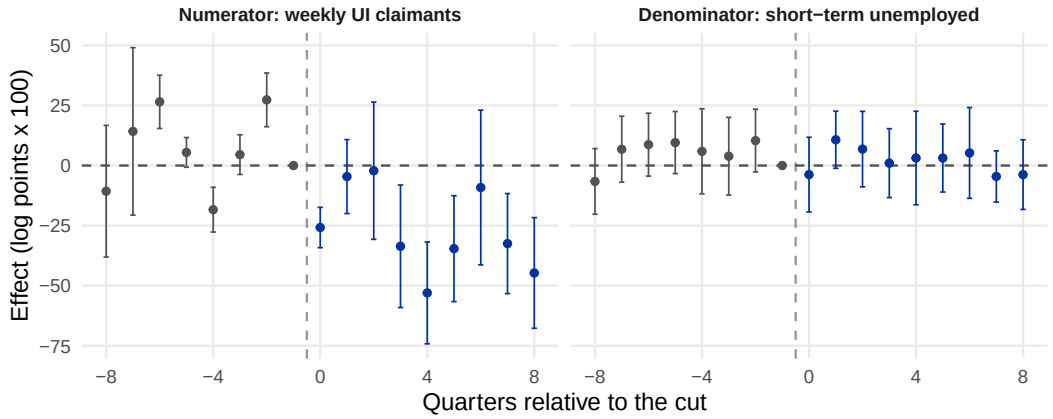
Table A8: Recipiency Effect by Tercile of Weeks Cut

	Smallest cuts	Middle	Largest cuts
Treated × post	-2.14	-11.95	-6.67
	(3.66)	(1.55)	(3.51)

Notes: Separate post-treatment coefficients for each tercile of the size of the duration cut, in weeks (4 to 14), in percentage points. Never-treated states are the omitted category. State and quarter fixed effects; standard errors clustered by state in parentheses.

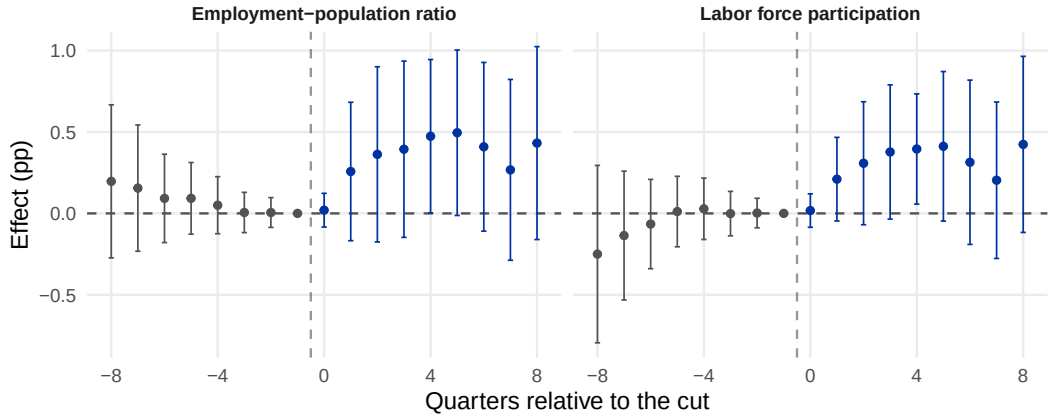
A.8. Recipiency Components, Employment and Long-Term Unemployment

Figure A3: Event Studies: Recipiency Numerator and Denominator



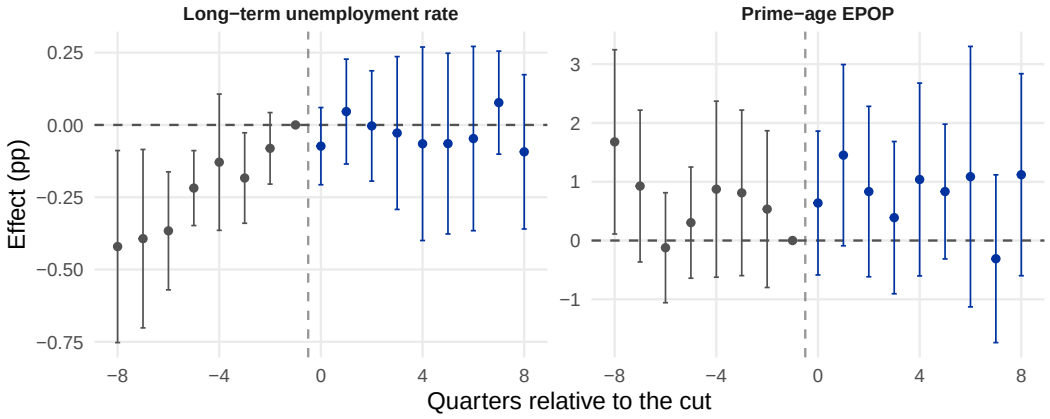
Notes: TWFE event studies (Equation (3)) of log average weekly UI claimants (left) and log short-term unemployed (right), in log points multiplied by 100, by quarter relative to the cut. Coefficients are relative to $k = -1$. State and quarter fixed effects. Bars are 95% confidence intervals; standard errors clustered by state.

Figure A4: Event Studies: Employment-Population Ratio and Participation



Notes: TWFE event studies (Equation (3)) of the seasonally adjusted LAUS employment-population ratio and labor force participation rate, in percentage points, by quarter relative to the cut. Coefficients are relative to $k = -1$. State and quarter fixed effects. Bars are 95% confidence intervals; standard errors clustered by state.

Figure A5: Event Studies: Long-Term Unemployment and Prime-Age Employment (CPS)



Notes: TWFE event studies (Equation (3)) of the long-term unemployment rate (unemployed more than 26 weeks, share of the labor force) and the employment-population ratio for ages 25–54, from CPS basic monthly files weighted by the final person weight, in percentage points. Coefficients are relative to $k = -1$. State and quarter fixed effects. Bars are 95% confidence intervals; standard errors clustered by state.

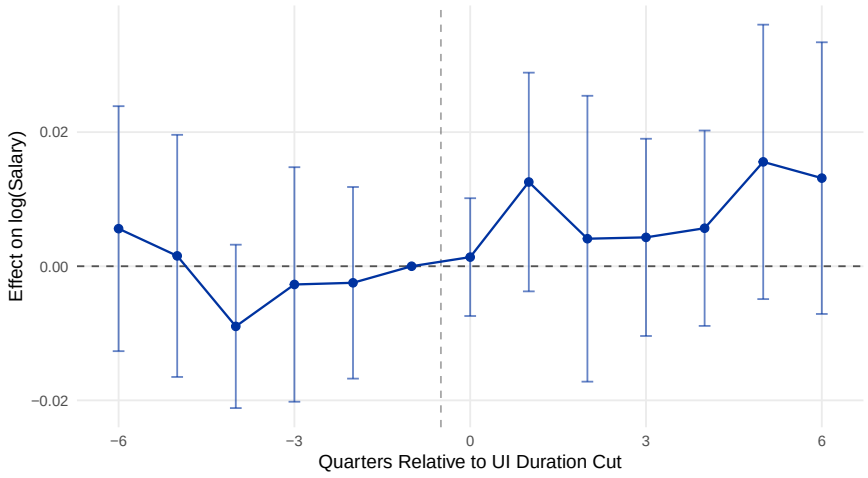
A.9. Stacked Difference-in-Differences (BGI/Lightcast)

The estimates in Section 5.3 come from the Borusyak et al. (2024) imputation estimator on establishment-job × quarter cells. To check that the near-zero result is not an artifact of that estimator, I estimate a stacked difference-in-differences following Cengiz et al. (2019).

This is a different design, not a re-estimation. For each treatment cohort, I construct a comparison window of ±6 quarters containing the treated cohort and the never-treated and not-yet-treated states. I *aggregate postings to the state × quarter level*, weighted by observation count, and stack the cohort panels with cohort-interacted state and quarter fixed effects. Aggregating to the state discards the within-employer-job contrast that identifies the main specification. This design therefore asks whether average posted wages in treated states diverge from those in comparison states, not whether the same employer posts different wages for the same job across state lines.

The stacked estimate is 0.0094 (s.e. =0.0074). Figure A6 presents the event study, which is also flat, with slightly positive estimates several quarters after treatment.

Figure A6: Event Study: Posted Wages, Stacked Difference-in-Differences



Notes: Stacked difference-in-differences estimates, following Cengiz et al. (2019), of the effect of a duration cut on log posted salary (BGI/Lightcast), by quarter relative to the cut. Each cohort is compared with never-treated and not-yet-treated states over ±6 quarters. Postings are aggregated to the state × quarter level, weighted by observation count. Cohort-specific state and quarter fixed effects. Bars are 95% confidence intervals; standard errors clustered by state.

A.10. Timing of Announcement

Table A9: Timing of Announcement and Prior Labor Market Conditions

	First public text				Bill introduced			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Cloglog	Cloglog	Logit	LPM	Cloglog	Cloglog	Logit	LPM
ΔUR, 6 months	0.008		-0.011	-0.0013	-0.187		-0.247	-0.0027
	(0.130)		(0.133)	(0.0012)	(0.316)		(0.395)	(0.0019)
Mean UR, 6 months		-0.419				0.021		
		(0.264)				(0.310)		
log(months at risk)	3.308	3.091	3.231	0.0120***	2.978	3.108	2.919	0.0123***
	(2.186)	(2.136)	(2.226)	(0.0011)	(1.922)	(1.902)	(2.056)	(0.0014)
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State-months	353	353	353	815	140	140	140	800
Events	9	9	9	9	9	9	9	9

Notes: Discrete-time hazard of announcement among the nine adopting states. Each state is observed monthly from January 2014 through its announcement month and leaves the sample afterward. Columns (1)–(4) date announcement by the first public text of the duration cut: bill introduction, or for Tennessee (May 2021) and Oklahoma (March 2022) the amendment that inserted the cut into a shell bill, and for Iowa (January 2022) the study bill. Columns (5)–(8) date it by the introduction of the enacting bill number. ΔUR is the change in the seasonally adjusted LAUS unemployment rate from $t-7$ to $t-1$; mean UR is its average over $t-6$ to $t-1$. All models include log months at risk. Month fixed effects drop calendar months without an event from the nonlinear models. Standard errors clustered by state (nine clusters) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.