The Jobs Effect of Ending Pandemic Unemployment Benefits: A State-Level Analysis *

Iris Arbogast and Bill Dupor†

September 6, 2022

Abstract

This paper uses the asynchronous cessation of emergency unemployment benefits (EUB) in 2021 to investigate the jobs impact of ending unemployment benefits. While some states stopped providing EUB in September, others stopped as early as June. Using the cessation month as an instrument, we estimate the causal effect on employment of reducing unemployment rolls. In the first two months following a state's program termination, for every 100 person reduction in beneficiaries, state employment causally increased by about 26 persons. The effect is statistically different from zero and robust to a wide array of alternative specifications.

JEL Code: J65

Keywords: pandemic emergency unemployment benefits

1 Introduction

Between January and April of 2020, U.S. employment declined by 24 million people (about 15 percent) as a result of the COVID pandemic. In response, the federal government introduced several temporary programs to lessen the negative impact of this decline. These emergency unemployment benefits (EUB) included providing program eligibility to many individuals who would not otherwise be covered, such as contract and gig

*The analysis set forth does not reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System. The authors thank Serdar Birinci, Tim Conley, Victoria Gregory and Amanda Michaud for useful comments. Comments are welcome. Authors' emails: william.d.dupor@stls.frb.org and irisarbogast@gmail.com.

† Federal Reserve Bank of St. Louis.

workers, extensions of benefit durations and a $600 weekly add-on for recipients.[1] These initial programs were extended by the federal government through September 2021, with the only major adjustment being a reduction of the add-on to $300 per week.

By spring 2021, following a major but incomplete rebound of employment, job growth began to slow and job openings were steadily rising, reaching 9.6 million in May 2021, up 2.5 million from their pre-pandemic level. At the same time, some state governors voiced concerns that the generosity of EUB was contributing to businesses' difficulty filling job vacancies. They announced plans to halt either some or all of these benefits several months before the programs' completion in September. In total 26 governors ended benefits before September, with 24 halting participation between June 12 and July 3.[2] In this paper, we use this asynchronous benefit termination across states to estimate the causal effect of losing unemployment benefits on employment.[3]

There are several important things to note about difference across states. First, a key indicator for each governor's decision to become an early halter or remain a late halter was the governor's political party. Our identification uses this political party as a source of exogenous variation, aligning with other economics research that uses political party to identify causal effect.[4]

Despite the important role of political party for determining which states were early halters, the two types of states tended to differ on another dimension. On average, early halters had recovered much closer to their pre-pandemic employment levels by the spring of 2020 than late halters. That is, early halters in general had relatively smaller remaining "employment gaps" immediately before treatment. With smaller remaining gaps, employment growth for early halters had already slowed relative to late halters and EUB recipient rolls in early halting were generally lower relative to late halting states. In May 2021, for example, EUB per capita in late halting states was nearly double that of early

[1]Throughout the paper, EUB will abbreviate both "emergency unemployment benefits" as well as "emergency unemployment beneficiaries." The meaning that is appropriate in each instance will be clear. For a description of these programs, enabled by U.S. Congress (2020), see U.S. Department of Labor (2020a) and U.S. Department of Labor (2020b).

[2]To fix language, we shall refer to states that ended EUB in September as late halters and state that ended EUB before September as early halters.

[3]It is worth noting the tremendous amount of national news regarding the quasi-natural experiment afforded by the end of EUB and the differential timing of benefit cessation. See for example Horsley (2021), Smialek (2021) and Mollica and Santilli (2021).

[4]For example, Knight (2002) uses political power of state congressional delegations as a source of exogeneity to identify potential crowd-out of federal highway grants on state government spending. Benhabib and Spiegel (2022) use the comparison of the president's political party to the share of a state's congressional delegation as an instrument to estimate the effect of sentiments on economic activity.

halters. We address this difference by explicitly using the reduction in the number of EUB recipients as our treatment variable, while maintaining exogeneity by instrumenting using the month of program termination.

Finally, we address endogenity concerns by including both early and late halting states in our analysis. Including late halting states strengthens our case for exogeneity. While early halters made the explicit—although politically influenced—decision of when to halt benefits, late halters did not individually choose the timing of benefit termination. The late-halting states' September 2021 termination date was set in December 2020 by federal legislation. Federal lawmakers could not have foreseen states' economic conditions when the termination date was determined.

Our main outcome variable is the state-level three month change in employment. In our benchmark specification, we estimate that for every 100 person reduction in beneficiaries driven by program termination, employment increased by roughly 26 persons comparing the second month following a state's EUB termination to the month before EUB termination. Thus, there was a strong, rapid employment response to the reduction in EUB rolls resulting from the end of emergency benefits.

Next, we show that our results are robust to a battery of controls. In the benchmark specification we control for the mask usage rate and an index for lockdown stringency. We further show that the results are robust to many alternative specifications, such as changing the controls, adjusting the start month of the sample and modifying the sample states. Robustness checks control for pandemic-related health variables such as COVID deaths and COVID cases as well as the leisure and hospitality share of employment. We employ placebo tests to show that our regression results do not reflect spurious findings.

We forgo an extensive literature review due to our intention to present our findings concisely and the number of excellent surveys on the topic of UI and labor markets (e.g., Krueger and Meyer (2002)) and describe only a few most closely related papers.

Holzer, Hubbard and Strain (2021) is perhaps the closest study to ours. They use the June 2021 halt month to conduct an event study using CPS data from roughly the same period as this paper. They find that the national employment-population ratio would have been 0.1-0.2 percentage points higher in July and August had all states ended FPUC and PUA in June. They do not use beneficiaries data in their analysis. Our estimates share the same sign, but are larger in magnitude than theirs.

Marinescu, Skandalis and Zhao (2021) use disaggregated data from an online jobs site to estimate the effect of the additional $600 per week benefit on job applications filed and on vacancy creation. They study the March to July 2020 period. They find that higher benefits led to a small decrease in applications and no change in vacancy creation. Their results are not comparable to ours for three reasons. (i) Over their period, jobs were scarce, whereas over our period jobs were plentiful. (ii) They look at micro-level data. Cross-individual spillovers (either positive or negative) will imply that micro estimates differ from macro estimates.[5] (iii) Whereas Marinescu, Skandalis and Zhao (2021) look at variation in benefits levels that are positive before and after treatment, the benefits amounts change from positive to zero for many beneficiaries during the episode we consider.

Finamour and Scott (2021) study the $600 per week EUB add-on program (Federal Pandemic Unemployment Compensation, FPUC) and its effect on employment and unemployment insurance (UI) replacments rates. They analyze labor market trends for the period before FPUC, during introduction of FPUC and after FPUC ended. In relation to our paper, the most relevant finding is that the relationship between UI generosity and employment does not change upon expiration of FPUC at the end of July 2020.[6] This result is suggestive that removing EUB had little effect on employment.

Bartik, et.al. (2020) study many aspects of the labor market during the early phase of the pandemic. This includes estimating a state-level event study using the distribution of payment initiation for two of the EUB programs. They motivate exogeneity by appealing to heterogeneous delays in program implementation. They find no evidence that the number of hours worked changed following the payment initialisation. However, similarly to Marinescu, Skandalis and Zhao (2021), jobs were scarce during their period of study compared to our period of study.

2 EUB Payouts Before and After Program Cessation

In May 2021, over 12 million non-working individuals were receiving some form of EUB in the U.S. When a state halted EUB, a natural question at the time was: How many recipients losing benefits would take up work within a short time span? Our approach is motivated by that question. In this section, we discuss two issues that complicate this approach. First, as mentioned in the introduction, early halting states had substantially fewer beneficiaries per capita than late halters. This difference alone means that one should interpret benefit cessation across states with caution. A state with zero benefi-

[5] This is due to a violation of the SUTVA assumption.

[6] FPUC was reintroduced in December 2020 at the lower amount of $300 per week.

ciaries would necessarily have no beneficiaries returning to work.

Second, when states ceased participation in EUB, they often continued to pay benefits to some recipients for unemployment spells that occurred before EUB cessation. From the perspective of individuals looking forward after the state's cessation date, benefits would no longer de-incentivize employment. That is, the relative cost of not working would rise. However, if they were still receiving or anticipating receiving benefits for past EUB, then they might have an income stream that could finance consumption and allow them to remain out of the workforce. Therefore, we will take into account how many people actually stopped receiving benefits when a state halted EUB.

Figure 1 shows the path for the number of beneficiaries by month in twelve large states.[7] The left-hand panel of the figure plots the ratio of EUB recipients to pre-pandemic employment by month for six large early halting states. The horizontal axis marks the number of months before and after the respective state halted benefits. For example, during its cessation month, Arizona (red solid line) had a recipients share equal to roughly 2.3 percent of its pre-pandemic employment. The number of benefit recipients in that state was relatively stable before halting; it began to decline following halting. The right panel presents the analogous data for six large late halters.

Many states saw very high EUB take up. For example, the right panel of the figure indicates that California had the number of beneficiaries equal to about 15 percent of pre-pandemic employment in the month before it halted benefits (August 2021). We discuss the reasons for and confirm, using other data, these high beneficiary rates in Appendix A.

Relative to early halters, late halters tended to have higher EUB-employment ratios in the halting month and saw steeper declines in EUB following halting. This implies that the large majority of the variation in EUB reductions comes from late, rather than early halting states. From the month preceding to the second month following the halt month, non-working EUB recipients fell by 2.4 million persons in early halting relative to 7.5 million persons in September halting states.

Importantly, EUB for two of the main programs, PEUC and PUA, did not collapse to zero in the month following cessation.[8] This may indicate that some EUB beneficiaries were still receiving EUB for weeks of unemployment preceding EUB program cessation. Given the slow decline, our benchmark model looks at the two month employment

[7] The federal government reports the number of weeks of benefits paid in a month. We approximate the number of beneficiaries by dividing the weeks reported by four.

[8] PUA and PEUC stand for Pandemic Unemployment Assistance and Pandemic Emergency Unemployment Compensation.

Figure 1: Ratio of number of non-working EUB recipients to pre-pandemic employment level, by state and months since benefit halt date

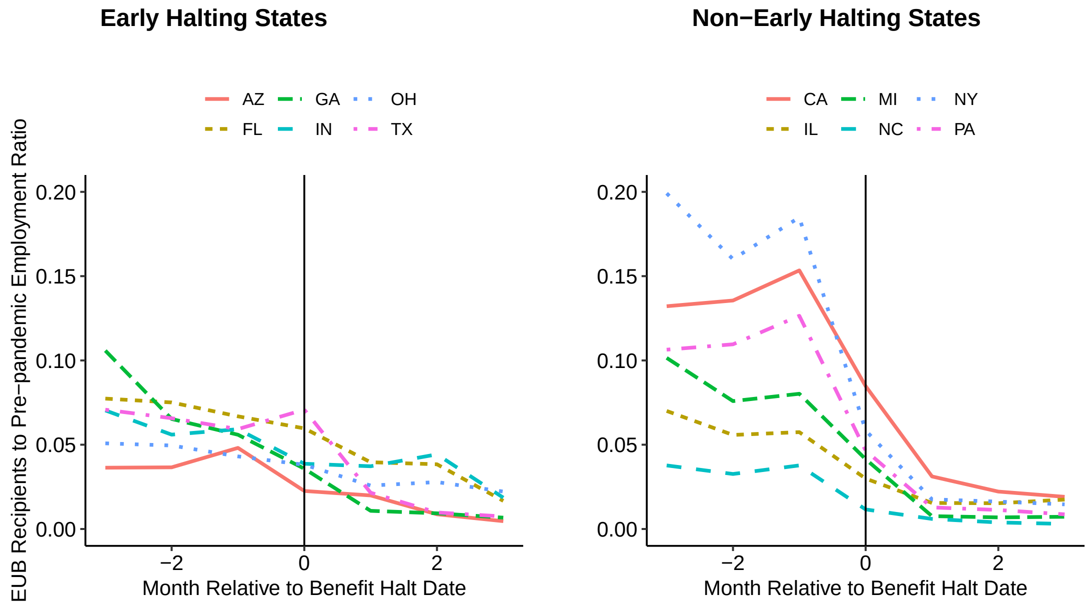

Note: EUB recipient data are from monthly state reports to the Bureau of Labor Statistics and include PUA, PEUC and regular beneficiaries. Pre-pandemic employment levels are from CPS micro data aggregated to the state level for January 2020.

change following the halting of EUB rather than the impact effect in the first month.

3 Data and Econometric Model

We study monthly data from every state and the District of Columbia covering October 2020 through June 2022. Let $Y_{i,t}$ denote the number of employed persons in state i at month t aged 16+. Our employment data are from the Current Population Survey (CPS) individual-level data, which we aggregate to the state level. We consider two other employment data sources, the CES and the LAUS, in the robustness section and in Appendix Sections E and F. We seasonally adjust our data using the Census Bureau's X-13 procedure in order to control for seasonal changes in employment.

Let $H_{i,t} = 1$ if state i halted in month t, and zero otherwise. Using this measure, we define t_i^* to equal the halt month of state i. To assign a specific cessation month to a state, we record the calendar date that EUB halted in that state. Next, we note that the survey occurs during the calendar week (Sunday through Saturday) containing the twelfth day of the respective month. We choose the treatment month as the one where the twelfth day of the month is closest in time (either before or after) to the date at which the state ended benefits. For example, the earliest four halting states ended benefits on June 12. Thus in these states, we choose June as the treatment month. Appendix B lists the halt month for each state.[9]

Next, let $C_{i,t}$ denote the number of non-working emergency unemployment beneficiaries in state i in month t, measured in hundreds of persons. We construct this variable using Bureau of Labor Statistics (BLS) data on the number of weeks of emergency unemployment benefits paid, which we call $B_{i,t}$. First, we assume that a state resident collecting EUB receives four weeks of benefits per month. Second, from Census Pulse surveys taken during the pandemic, we note that nationwide about 20 percent of beneficiaries were working while collecting benefits. In terms of the effect of program cessation on net employment changes, its important to exclude this 20 percent from the treatment.[10] Finally, we measure the final variable in hundreds of persons to ease interpretation of our estimated coefficients. Thus,

$$C_{i,t} = \frac{1}{4}\left(\frac{4}{5}\right)\left(\frac{1}{100}\right)B_{i,t}$$

Before presenting our the full econometric model, we plot the pre-treatment trends in employment for the early halting and late halting groups and estimate a simple event study regression. Figure 2 plots employment as an index (equal to 100 in January 2020) both before and after treatment for June and September halting states. The vertical lines indicate the months immediately preceding the two most common halting months (June and September). Note that in the immediate response to the pandemic (March and April 2020), employment fell more in early halting states relative to late halting states. Employment in both groups increased quickly after April.

For April 2020 through June 2021, the two groups follow roughly parallel trends in which employment growth slows as time progresses. The parallel trends observed in Figure 2 provide additional support for the exogeneity assumption we maintain in order

[9] A few states temporarily delayed execution of their cessation plans. Appendix C discusses these cases. Also, in our robustness analysis we show that changing the halt date for these states to reflect the delays does not materially affect our results.

[10] This adjustment is necessary because if an individual losing benefits had already been working some hours pre-cessation, then their survey response would be "employed" both before and after the state terminated benefits. Throughout the paper, beneficiaries refers to non-working beneficiaries.

Figure 2: Employment index for early and late halting states (index equals 100 in January 2020)

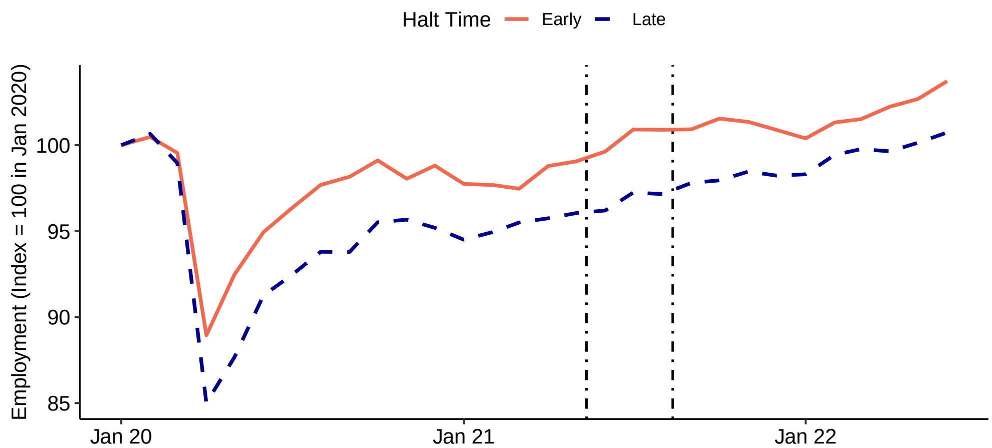

Note: Employment in each group is the sum of seasonally adjusted state level employment calculated from the micro-CPS dataset in each month divided by the value in January 2020.

to achieve identification.

We adopt the local projections method, described in Dube, et.al. (2022), as a straightforward way to estimate dynamic causal effects in the presence of staggered treatment as a way to avoid potential bias. However, as a first step, we demonstrate our basic finding in a more traditional event study analysis. Note that we cannot estimate the causal effect via a standard event study approach using all of our data without potentially introducing bias because of the staggered EUB cessation.[11]

To construct an event study that does not suffer from the above potential bias from staggered treatment, we limit the sample by: (i) including only June and September halting months in the regression; (ii) ending the estimation sample with August 2021. Since benefits had not yet ended for the September halters in August, in absence of anticipation effects, the September states form a valid untreated control group. The June halters make up the treatment group. By excluding states that halt in July and August, we avoid the issue of asynchronous treatments and can apply the basic event study technique.

[11]See for example Goodman-Bacon (2021) on this point. See Appendix D for a discussion of alternative methods for addressing this potential bias.

First, define $\tilde{y}_{i,t} = 100 \times \left(Y_{i,t}/Y_{i,Jan20}\right)$. The estimation equation is:

$$\tilde{y}_{i,t} = \psi_i + \mu_t + \sum_{h=-H_1}^{H_2} \alpha_h 1\left(H_{i,t-h} = 1\right) + \phi' W_{i,t} + \eta_{i,t} \tag{1}$$

We estimate the model over the period June 2020 through August 2021. $W_{i,t}$ contains two variables: log COVID deaths and a government COVID-19 policy stringency index in state i during month t. The deaths variable controls for cross-state differences in the course of the pandemic. The stringency index controls for differential state COVID-19 policy response unrelated to EUB policies. ψ_i and μ_t are state and time fixed effects. We weight regressions by pre-pandemic employment and compute standard errors clustering by census division.

Figure 3 plots the regression coefficient α_h for each time-to-treatment h as dots, enveloped by 90 percent confidence intervals. It indicates a strong employment effect that builds gradually in the months following benefit termination. Two months following cessation, the halting of benefits increases employment by 2.0 percent in the treated relative to the untreated group. To express this effect as a total number of jobs, consider a counterfactual in which all states and DC halted in June relative to another counterfactual in which they halted in September. The point estimate implies that employment would be 3.1 million persons larger in August in the former case relative to the latter.

Next, note that in the twelve months before treatment, the coefficient oscillates around zero. For $h < 0$, the coefficients are statistically indistinguishable from zero. This indicates the same absence of differential pre-trends as suggested by Figure 2 and thus bolsters our case for our exogeneity assumption used to achieve identification.

While the event study regression is instructive, it has several limitations. First, we are not using data after August 2021, nor are we using the July or August halting states. Second, the September halting states have a control phase and a treatment phase, the latter of which we are not exploiting above. Third, although we measure the employment effect of terminating benefits, we are not measuring precisely how many people were taken off of EUB as a result of the halting. To address these three concerns, our main econometric specification is a local projections instrumental variables regression. Our estimation equation is:

$$\frac{Y_{i,t+k} - Y_{i,t-1}}{Y_{i,Jan20}} = \theta_{i,k} - \gamma_k \frac{C_{i,t+k} - C_{i,t-1}}{Y_{i,Jan20}} + \beta_k' X_{i,t-1} + \varepsilon_{i,t+k}^k \tag{2}$$

Figure 3: Effect of halting EUB on employment pre and post treatment, 90 percent pointwise confidence intervals

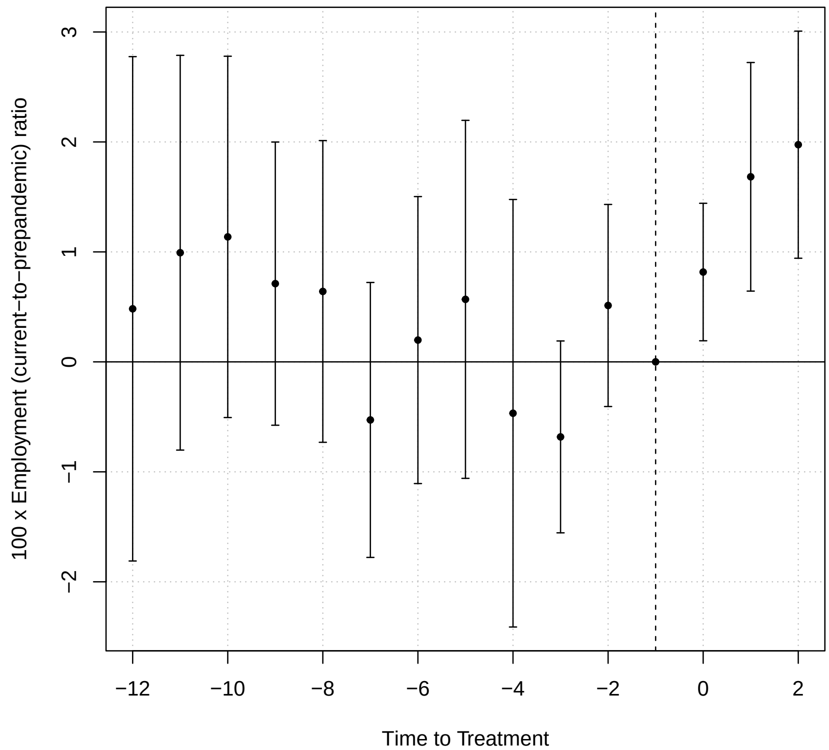

Note: Outcome variable=100 $\times$ Ratio of employment level to pre-pandemic employment level. We weight regressions by pre-pandemic employment and cluster standard errors by census division.

where we estimate the model at various horizons $k = 0, \ldots, M$.

Here, $\theta_{i,k}$ is a state fixed effect, and $X_{i,t-1}$ is a vector of conditioning variables. We estimate (2) using $H_{i,t}$ as an instrumental variable. We assume that beyond period $t + M$ there is no additional effect of EUB cessation on the outcome (where t is the time of benefits halting). This allows us to expand our sample to use some post-treatment months as control months in our estimation. This assumption implies that state i employment following termination adjusts within M month following t_i^*. In the robustness section, we explore how estimates change as we vary M.

We draw our sample using the procedure given in Dube, et.al. (2022) in order to ensure that, for periods close to t_i^*, we do not contaminate our estimates of the horizon k response with observations that are influenced by the treatment. Specifically, for each state i and horizon k, we drop observations

$$t \in \{t_i^* + g\} \text{ for } g \neq 0 \cap g \in \{-k, \ldots, M\} \tag{3}$$

Dropping the observations up-to k months before treatment immunizes against using a contaminated control period resulting from overlap induced by the k-month ahead dependent variable. Dropping the up-to M month post-treatment observations immunizes against using a contaminated control period resulting from the delayed effects of the treatment on the outcome. We weight regressions by state pre-pandemic employment. We compute robust standard errors adjusted with a finite-sample correction and also cluster at the census-division level. Restricting the sample via equation (3), ensures that each observation is in either a treatment period or a "clean control" period.

The main coefficient of interest is γ_k, which gives the causal impact on employment of a 100 person reduction in benefit recipients driven by program termination. In most of our specifications, we report the estimate at horizon $k = 2$ and report the estimate at other horizons in the robustness section. Our benchmark conditioning variable vector $X_{i,t-1}$ contains three lags of the one month change in EUB and three lags of the one month change in employment, where each variable is scaled by pre-pandemic employment. The benchmark also controls for the mask usage rate measured by the percent of the population reporting always wearing a mask when leaving their home reported by the Institute of Health Metrics and Evaluation. It controls for COVID-19 lockdown intensity using an index from the Oxford COVID-19 Government Response Tracker. The index is measured as the average over indicator variables for containment policies, closure policies and public information campaigns. For both controls, we construct a 3-month change ending before $t - 1$ and scale the variables to have mean zero and unit variance.

4 Results

This section reports and interprets estimates of equation (2) and its variants. The results deliver a consistent message: Terminating EUB had a quantitatively and statistically significant positive impact on employment, which we show is robust to many sensible alternative statistical specifications.

First, column 1 of Table 1 contains our benchmark two-stage leasts squares estimate. The coefficient equals 25.64 (SE=10.43), indicating that employment increases by about

Table 1: Jobs effect of benefit termination: employment change (over three months) per 100 person reduction in the number of beneficiaries

	Dep Var: Emp Change			Dep Var: Decl in claims
	(1) 2SLS (Bmark)	(2) Reduced Form	(3) OLS	(4) First Stage
Decl in EUB	25.64**		3.86	
	(10.43)		(4.38)	
Halt month		1.51**		5.88**
		(0.61)		(1.76)
Partial F				11.2**
Num. obs	663	663	663	663

Notes: The dependent variable is the 3 month change in employment per 100 pre-pandemic employees. Instrument is an indicator for halting month of EUB. Standard errors are robust, finite-sample corrected and clustered by census division. Regressions include state fixed effects and are weighted by state pre-pandemic employment level (1/2020). See text for additional details.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

26 persons for each 100 person reduction in EUB after two months following a state's EUB cessation (compared to the month prior to cessation). The estimate is statistically different from zero at a five percent level. Note that using state-level data implies that we cannot follow individuals losing EUB benefits over time to see if they specifically boosted employment following EUB termination. Rather, our outcome variable is the net change in total state employment. Parsing the extent to which the employment increase comes from the group of those losing EUB is not possible using our approach.

Column 2 contains the reduced-form estimate, which directly regresses the outcome on the instrument (as well as the additional control variables).[12] The coefficient equals 1.51 (SE=0.61), indicating that benefit cessation causes employment over a three month horizon (inclusive of the impact month) to increase by 1.51 percent of pre-pandemic employment. Since pre-pandemic employment was roughly 157 million, the estimate implies that a nationwide cessation policy would have increased employment by 2.37 million jobs.

This reduced form estimate should be of particular interest to a reader concerned that

[12] The regression weights and standard error methodology follow those for the two-stage least squares estimates described above.

there may be other channels—–besides labor supply increasing through incentive effects on individuals—operating to influence observed employment growth. For example, an observer could argue that labor demand might increase as firms see their probability of filling vacancies increase as labor supply shifts out. In our view, quantitatively the overwhelming channel for the jobs effect is the labor supply impact of millions losing benefits, although this is—–of course—untestable without further econometric structure. Nonetheless, our reduced form estimate provides the total jobs effect of halting benefits, independent of the relative importance of the labor supply versus demand channels. The large, positive reduced form jobs effect is also confirmed by the two-way fixed effects results presented earlier (see Figure 3).

Column 3 presents the least squares analog of the benchmark two-stage least squares coefficient from column 1. Without instrumenting, there is no significant relationship between changes in the number of EUB and the number of employed.

Column 4 contains the first-stage regression estimate, i.e., regressing the scaled change in beneficiaries on the halt indicator, along with the control variables. The point estimate equals 5.88 (SE=1.76), indicating that over a three-month horizon EUB falls by 5.88 percent of pre-pandemic employment. This implies that, had every state simultaneously halted EUB, the number of beneficiaries would have fallen by 9.23 million persons over this horizon, relative to a no-halt counterfactual.

If one sums the number of non-working EUB recipients across states in each state for the month following each state's respective cessation, this total would equal 3.0 million persons. This indicates that, despite a substantial reduction in beneficiaries, many continued to collect benefits. Two potential reasons account for this effect: (i) As evidenced by Figure 1, some individuals continued to receive benefits for months after program termination perhaps because of payment delays; (ii) Many individuals on regular state benefits remained on the program despite losing the $300 EUB add on. Note that the point estimate is statistically different from zero at a five percent level, indicating that the halt month is a strong instrument.

Figure 4 presents a visualization of the jobs effect we identify. It consists of ten panels, each of which plots four data points. There are panels for the national total, each of the eight largest states and the sum of the remaining states. The vertical axis indicates the three-month change in either employment or decline in non-working beneficiaries per 100 pre-pandemic workers. The horizontal axis for each panel assigns the value one in the month of cessation and zero to the relevant non-halting months.[13]

Consider Texas (the third panel from the left). The dashed blue line connects two points. The left point on this line gives the average three month change in employment, per 100 persons, across all clean control months in the sample in which Texas did not initiate the cessation of benefits. The right point gives the three month change in employment starting from the month before cessation. For each state, the left point is an average and the right point is a single observation. The two values for Texas are 0.62 and 3.35 respectively. Thus Texas saw an increase in its 3-month employment change of 2.73 (=3.35-0.62) between its halting month and its average non-halting month. Texas employment increased more rapidly in the halting month relative to the average non-halting months, which is immediately evident from the line's upward slope.

The yellow solid line is constructed similarly except it represents the reduction in the number of beneficiaries in the state between the halting month and the average non-halting months. The figure indicates that the three month decline in beneficiaries was 4.95 in the halting month and 0.70 in the average non-halting month for Texas. The difference equals 4.25. Taking the ratio of the employment change effect and the (reduction in) beneficiary effect, we have 0.64.

The figure's leftmost panel presents the same information as described above, except it presents the weighted sum of the state specific values across the entire U.S. Both lines are upward-sloping, indicating that halting was associated with both faster employment growth and a faster decline in beneficiaries. The ratio of the two slopes equals 0.30. Thus, according to this tabulation, employment increased by 30 persons for every 100 persons losing benefits as a result of program termination.[14]

Table 2 presents 15 sets of estimates of equation (2). First, we restate the benchmark estimate in entry 1. We drop the state fixed effects in entry 2 and the lag control variables in entry 3. The estimate size decreases slightly to 22.30 (SE = 10.35) when we remove state fixed effects and increases slightly to 30.31 (SE = 9.41) when we drop the lagged controls.

Recall that the benchmark estimate corresponds to the employment change two months following cessation (i.e., $k = 2$). Entries 4 and 5 report the jobs effect at two other horizons: $k = 1$ and $k = 3$. At $k = 1$, the jobs effect is 19.10 (SE=8.15) and at $k = 3$ the jobs effect is 23.62 (SE=25.07). This indicates an increasing employment effect of benefit cessation which begins to dissipate at the last horizon. For the impact month $k = 0$, in

[13] The relevant non-halting months are given by equation (3).

[14] The difference between the jobs effect from this tabulation and our benchmark estimate is due to our conditioning variables. We discuss the role of conditioning variables below.

Figure 4: 3-month change in non-working beneficiaries and employment comparing halting and non-halting months

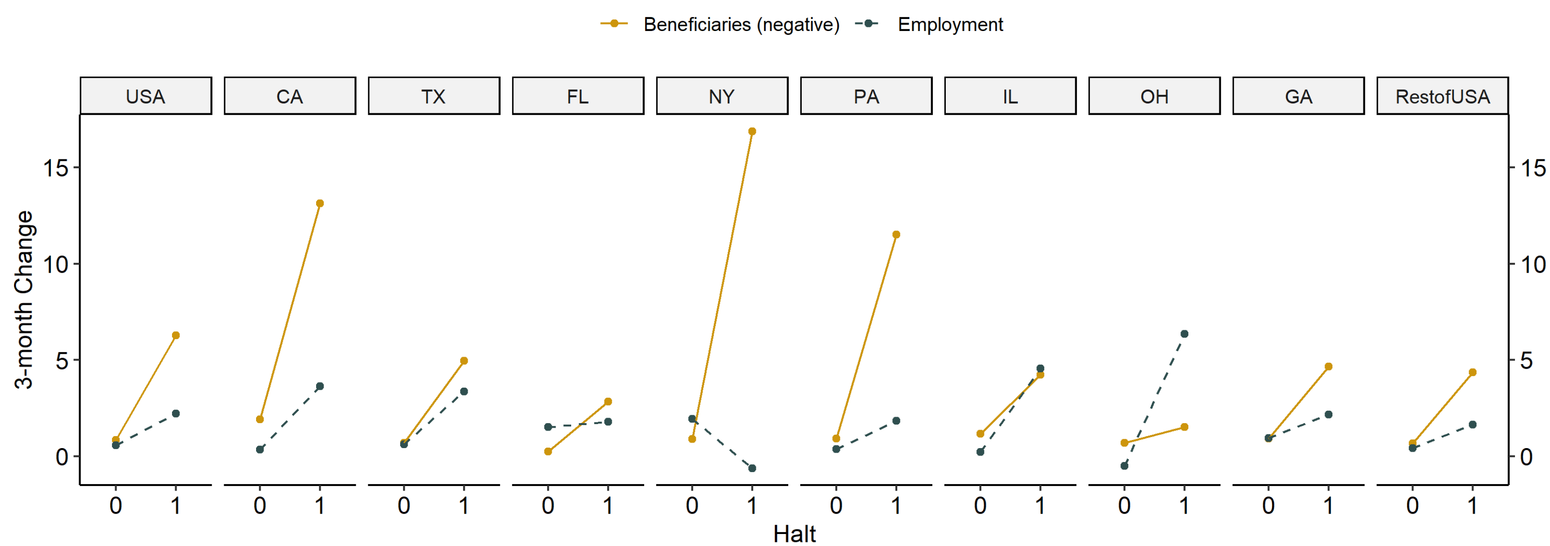

Note: USA and RestofUSA are aggregates of state level data using pre-pandemic employment weights. Halt=1 indicates month of program termination; Halt=0 indicates relevant non-halting months, determined by equation (3). 3-month changes are measured per 100 pre-pandemic workers. See text for additional details.

results not reported here, the coefficient is small and not statistically different from zero, indicating that the jobs effect of halting benefits occurs with delay.

Entry 6 is identical to our benchmark case except that we construct our outcome variable without a seasonal adjustment. This change results in no material difference from our benchmark case. In our benchmark specification we set $M = 4$, the final horizon for which the treatment influences the outcome. In entries 7 and 8 we change this assumption, using $M = 3$ and $M = 5$ instead, and do not find substantive changes in the estimates. Entries 9 and 10 vary the start date of the sample (using either November 2020 or December 2020 instead of our benchmark start date of October 2020). The point estimates change slightly, equaling 25.91 (SE=11.08) and 21.79 (SE=9.18), respectively.

Next, because Indiana and Maryland (announced early halters) each delayed their EUB cessation, entry 11 drops these two states from the sample and entry 12 changes $H_{i,t}$ to assign these two states to be late halters. As seen in the table, neither adjustment has an appreciable effect on the point estimate or standard error relative to our benchmark.

In Figure 4, we observe that New York had a large decrease in beneficiaries and a decrease in employment comparing halting and non-halting 3-month changes. In Ohio, employment increased more than beneficiaries decreased in the same comparison. Because these states are large and potentially influential, entry 13 drops both states from the sample. The estimate increases by about 6 people to 31.44 (SE = 9.84). The final way that we change the states included in our estimation is by running our benchmark specification with only June-halting states and then running the benchmark specification with only September-halting states. The samples are likely too small to find precise estimates, but the estimates are positive and larger for June-halting states.

Table 3 reports estimates from several placebo specifications. Each column uses our reduced form specification which regresses our employment change variable on the instrument (halt month). In the placebo specifications, halt months are artificially assigned to different periods.

First, in column 2, we run our reduced form specification with halt dates pushed backward by three months (e.g., September halting states counterfactually halt in June and June halting states counterfactually halt in March). This method keeps the groupings of halting and non-halting states the same, but assigns the treatment to a time period in which we expect to find no effect. In column 3, we randomly reassign halt months for each state between January and March. For both of these first two placebo regressions we use the same data and set of cleaned controls as in the reduced form benchmark.

Table 2: Jobs effect of benefit termination: employment change (over three months) per 100 person reduction in beneficiaries, various alternative specifications

Vary Controls/horizon	(1) Bmark	(2) Drop St FE	(3) Drop Lag Cntrls	(4) k = 1	(5) k = 3
Decl in EUB	25.64**	22.30*	30.31**	19.10**	23.62
	(10.43)	(10.35)	(9.41)	(8.15)	(25.07)
Partial F	11.2**	11.0**	13.5***	10.9**	18.3***
Num. obs	663	663	663	714	612
NSA/Vary sample months	(6) NSA	(7) M = 3	(8) M = 5	(9) Start Nov	(10) Start Dec
Decl in EUB	27.93*	26.33**	24.56*	25.91**	21.79**
	(12.76)	(10.57)	(11.41)	(11.08)	(9.18)
Partial F	11.2**	11.4***	11.3***	9.4**	11.6***
Num. obs	663	714	612	612	561
Vary sample states	(11) Drop IN, MD	(12) IN, MD in Sept	(13) Drop OH, NY	(14) June Halt	(15) Sept Halt
Decl in EUB	25.73**	26.92**	31.44**	39.99	16.32
	(10.61)	(10.83)	(9.84)	(40.31)	(8.78)
Partial F	12.1***	10.3**	14.2***	21.4***	11.2**
Num. obs	637	663	637	286	325

Notes: The dependent variable is the 3 month change in employment per 100 pre-pandemic employees. Instrument is an indicator for halting month of EUB. Standard errors are robust, finite-sample corrected and clustered by census division. Regressions include state fixed effects unless otherwise described and are weighted by state pre-pandemic employment level (1/2020). See text for additional details.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Jobs effect of benefit termination: employment change (over three months) per 100 person reduction in beneficiaries, placebo tests

	(1) Bmark Reduced Form	(2) Push Back 3 Months	(3) Random Months	(4) Push Back 3 Years
Halt Month	1.51**	0.29	0.37	−0.08
	(0.61)	(0.51)	(0.59)	(0.58)
Num. obs	663	663	663	663

Notes: The dependent variable is the 3 month change in employment per 100 pre-pandemic employees. All columns are variations of our reduced form specification which estimates the impact of our instrument (halt month) on the dependent variable. Standard errors are robust, finite-sample corrected and clustered by census division. Regressions include state fixed effects and are weighted by state pre-pandemic employment level (1/2020). See text for additional details.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

For a third placebo test we push the halt dates backward 36 months to 2018, allowing us to use a set of dates unaffected by COVID-19 to check for effects from seasonality. We construct clean controls using the same method as the benchmark for this new time period. This specification controls for three lags of employment, but not the lagged beneficiary variables or stringency and mask usage indexes. The coefficients for all three placebo specifications are small in size and statistically insignificant.

Table 4 alternatively removes and adds a series of controls that may potentially have impacted employment levels during the COVID-19 pandemic. For each control, we construct a 3-month change ending before $t-1$ and scale the variables to have mean zero and unit variance. In column 2 we present the benchmark specification estimate, except we drop the mask usage variable and stringency index. In columns 3 and 4, we control for log normalized COVID-19 cases and deaths from *The New York Times*, based on reports from state and local health agencies. Our results are robust to controlling for these measures of pandemic intensity. In column 5 we control for the leisure and hospitality share of employment to ensure that our results are not driven by differences in labor market dynamics of these sectors across states during the pandemic. The non-farm employment and leisure and hospitality employment data used to construct this share come from the CES (Current Employment Statistics) dataset.

Next, note that we do not include month fixed effects among our specifications. Month fixed effects are highly collinear with $H_{i,t}$. Since most states halt in either June or September, adding month fixed effects will result in an imprecisely estimated jobs effect. When we add month fixed effects to the benchmark specification, the jobs effect is 31.74 (SE=36.38) with a partial F-statistic equal to 28.4. Thus, while the point estimate changes little with the addition of month fixed effects, the standard error increases dramatically.

We conclude this section by discussing alternative state-level employment measures that one could have used in this study. We use the CPS instead of the CES because the latter excludes gig and contract workers. A large fraction of individuals collecting EUB did so under the Pandemic Unemployment Assistance program — which was established specifically to cover these types of workers. If individuals on this program returned to gig and contract work upon reentry into employment, this would be missed if we had used the employer-based data.

Second, we use the CPS instead of the state-level Local Area Unemployment Statistics (LAUS). Roughly speaking, LAUS employment attempts to control for outliers at the regional level; however, over this two year period, the difference between the LAUS and

Table 4: Jobs effect of benefit termination: employment change (over three months) per 100 person reduction in beneficiaries, varying the conditioning variables

	(1) Bmark	(2) Drop All Controls	(3) Cases	(4) Deaths	(5) Leis Emp
Decl in EUB	25.64^{**}	30.07^{*}	25.78^{**}	25.71^{**}	26.77^{**}
	(10.43)	(16.14)	(10.67)	(10.56)	(10.89)
Cases			-0.11		
			(0.42)		
Deaths				-0.10	
				(0.33)	
Leis Emp Share					0.32
					(0.75)
Partial F	11.2**	13.2***	11.0**	10.4**	12.0***
Num. obs	663	663	663	663	663

Notes: The dependent variable is the 3 month change in employment per 100 pre-pandemic employees. Instrument is an indicator for halting month of EUB. Standard errors are robust, finite-sample corrected and clustered by census division. Regressions include state fixed effects unless otherwise described and are weighted by state pre-pandemic employment level (1/2020). See text for additional details.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

state-aggregated CPS measures are systematically related to whether or not a state has halted. If the LAUS were simply "controlling for outliers," there is no reason to believe that this noise would be correlated with states' halt months.[15]

5 Conclusion

Using the cessation of EUB across states in 2021, we establish a strong positive casual link from a reduction in the number of beneficiaries to state employment growth. We find that for every 100 people in a state that lost unemployment benefits, 26 people became employed two months after halting in that state. The effect is statistically different from zero and robust to a wide array of alternative specifications. Note that for every 100 people in a state that lost unemployment benefits, there was not a 100 person increase in employment - i.e. there was a gap of 74 people between the 100 person reduction in beneficiaries and the increase in employment. Future research might examine the effect of cessation on other labor market outcomes, such as labor force participation, as well as other economic indicators, such as consumption, income and financial distress.

Note that by focusing on a state-level outcome variable, our estimates are closer to "macro responses" than related studies which instead look at an individual-level outcomes. While interesting from a decision theoretic perspective, using individual-level data may miss important cross-individual spillovers that could be either positive or negative. For example, if a person increases consumption upon losing emergency benefits and then takes a job (e.g., from spending on clothes, fuel and car maintenance for travelling to and from work), then this may drive up demand for goods in the rest of the economy. This may in turn stimulate employment in the state indirectly. This indirect positive effect would be missed in individual-level regressions and thus bias downward the macroeconomic jobs effect of halting unemployment benefits. One could envision negative spillovers, on the other hand, that would reverse the direction of the bias. By working with data aggregated to the state-level, we are at least in part immunized against this concern.[16]

[15] Appendix E and F explain these and other reasons why using either the CES or LAUS datasets would be inappropriate for the question our paper answers.

[16] Note that our approach does not take cross-state spillovers into account. If one state halts unemployment benefits there may be positive (or negative) effects on other states' employment, which would downwardly (or upwardly) bias our estimates.

References

Bartik, A., M. Bertrand, F. Lin, J. Rothstein and M. Unrath. 2020. "Measuring the Labor Market at the Onset of the COVID-19 Crisis," *Brookings Papers on Economic Activity*.

Benhabib, Jess and Mark Spiegel. 2022. "Sentiments and Economic Activity: Evidence from U.S. States," *Economic Journal*.

Dube, A., D. Girardi, O. Jorda and A. Taylor. 2022. "A Local Projections Approach to Difference-in-Difference Event Studies," working paper.

Finamour, L. and D. Scott. 2021. "Labor Market Trends and Unemployment Insurance," *Economic Letters* 199.

Goodman-Bacon, A. 2021. "Difference-in-differences with Variation in Treatment Timing," *Journal of Econometrics*, 225 (2): 254-277.

Knight, Brian. 2002. "Endogenous Federal Grants and Crowd-out of State Government Spending: Theory and Evidence from the Federal Highway Aid Program," *American Economic Review* 92 (1): 71-92.

Holzer, H., R. Hubbard and M. Strain. 2021. "Did Pandemic Unemployment Beefits Reduce Employment from Early State-Level Expirations in June 2021," working paper.

Horsley, Scott. 2021. "Millions Lose Jobless Benefits Today. It Doesn't Mean They'll Be Rushing Back to Work," *National Public Radio*, Sep. 6.

Krueger, Alan B., and Bruce D. Meyer. 2002. "Labor Supply Effects of Social Insurance." *In Handbook of Public Economics,* edited by Alan J. Auerbach and Martin Feldstein, 2327–92. Amsterdam: Elsevier

Marinescu, I., D. Skandalis and D. Zhao. 2021. "The Impact of the Federal Pandemic Unemployment Compensation on Job Search and Vacancy Creation," *Journal of Public Economics* (forthcoming).

Mollica, Andrew and Peter Santilli. 2021. "Half of U.S. States Ended Federal Covid-Related Jobless Benefits Early. Here Is How They Compare with the Other Half," *The Wall Street Journal*, July 16.

Smialek, Jeanna. 2021. "Will Cutting Off Federal Unemployment Benefits Shake Up the Job Market?" *New York Times*, Aug. 5.

U.S. Congress. *Coronavirus Aid, Relief and Economic Security Act*. 2020.

U.S. Department of Labor, Employment and Training Administration, *Unemployment Insurance Program Letter No. 16-20.* 2020a.

U.S. Department of Labor, Employment and Training Administration, *Unemployment Insurance Program Letter No. 17-20.* 2020b.

Appendix for *The Jobs Effect of Ending Pandemic Unemployment Benefits: A State-Level Analysis*

A Were as many people really on EUB as suggested by Figure 1?

Figure 1 indicates that in California, EUB was roughly 15 percent of prepandemic employment in August 2021. In August 2021, the seasonally adjusted California unemployment rate was 7 percent with 1.34 million unemployed persons according to the LAUS. Is our 15 percent calculation too high?

We compare this estimate to U.S. Department of Labor data on the number of continuing claims (which is distinct from the weeks paid out data we use to construct Figure 1 in the paper.). For the week ending August 14, 2021, there were 1.74 million PUA continuing claims and 1.07 million PEUC continuing claims. Together, there were 2.81 million PUA and PEUC claims, more than double the number of unemployed persons in California in that month. This does not include the regular program state claims. In short, our 15 percent calculation for California aligns closely with federal claims data.

One reason that the number of unemployed people was less than the number of recipients is that, in August 2021, some emergency benefit recipients had several ways to collect benefits without conducting a job search (which is a flag for being classified as unemployed in the CPS). For example, according to California Economic Development Department (2021), PUA recipients could meet PUA requirements by enrolling "in training or education courses that will help the business and does not interfere with an ability to return to full-time self employment." Also, California Economic Development Department (2021) discusses "pandemic-related exemptions to the work search requirement for PUA recipients such as providing primary care to a person diagnosed with COVID-19 or a child who cannot attend school because of a COVID-19 health emergency."

B For each state and DC, what are the halt months given by your algorithm?

Table 5: Halt dates by state based on our algorithm

Month	States
Jun	AK, AL, AR, FL, GA, IA, ID, IN, MO, MS, MT, ND, NE, NH, OH, OK, SC, SD, TX UT, WV, WY
Jul	AZ, MD, TN
Aug	LA
Sep	CA, CO, CT, DC, DE, HI, IL, KS, KY, MA, ME, MI, MN, NC, NJ, NM, NV, NY, OR PA, RI, VA, VT, WA, WI

Notes: Cutoff dates are the closest CPS survey reference date (the 12th of each month) to the state's official EUB halting dates.

C Didn't a few states delay their implementation of halting EUB?

In our analysis, we set the halt month for Indiana and Maryland as the month each governor's actions initially took place, June 19 and July 3 respectively. In both states, legal challenges and other impediments that followed slowed the actual implementation of the termination of benefits. Both states' actual termination took place in September as the federal dollars ran out.

For robustness, we rerun the regressions using September as the two states' halt month and alternatively drop the two states in the paper's main text. There was no material difference in the estimates relative to the benchmark ones.

D Aren't there other methods (besides local projections) to deal with the staggered treatment issue?

Yes, there are. Unless properly accounted for, using the differences-in-differences method with staggered treatment can induce biased estimates. See Goodman-Bacon (2021) on this point. As such, in the main text of our paper, we adopt the approach of Dube, et.al. (2022), who use local projections. Dube, et.al. (2022) shows that the above bias is not present when local projections are applied correctly.

There are other papers that develop approaches to deal with staggered treatment. These include Sun and Abraham (2021) and Callaway and Sant'Anna (2021) as well as a few others not cited here. In our view, Dube, et.al. (2022) is the most straightforward, simplest to implement and easiest to understand.

E Why not construct the outcome variable using the LAUS instead of the CPS?

To estimate our state-level model, one alternative to using the CPS is using the Local Area Unemployment Statistics (LAUS) panel. While the LAUS panel includes a monthly state-level employment measure, as we explain here, the model-based nature of that measure makes it inapplicable in our context.

We start by plotting the data. Figure 5 has the same general setup as Figure 4 from the paper's main body, except here we show 3-month employment changes only and omit the changes in beneficiaries. For each state (or group of state), we plot the employment changes using the CPS, CES and LAUS measure. While the CPS shows substantial changes, the other two do not. For example, the solid light blue line corresponds to the LAUS values. For the US, RestofUS and the plotted states, the LAUS changes are much flatter than the corresponding CPS values. The LAUS sees a much smaller difference across halt and relevant non-halt months.

Next, we examine what drives the difference between the CPS and LAUS measures. In particular, we explain how the smoothing procedure used for the CPS systematically removes employment fluctuations that are correlated with employment when states terminate EUB. First, BLS (2020) explains how the LAUS measures are imputed: "Estimates for states are derived from signal-plus-noise models that use the monthly employment

Figure 5: 3-month change in employment comparing halting and relevant non-halting months, three employment measures (seasonally adjusted)

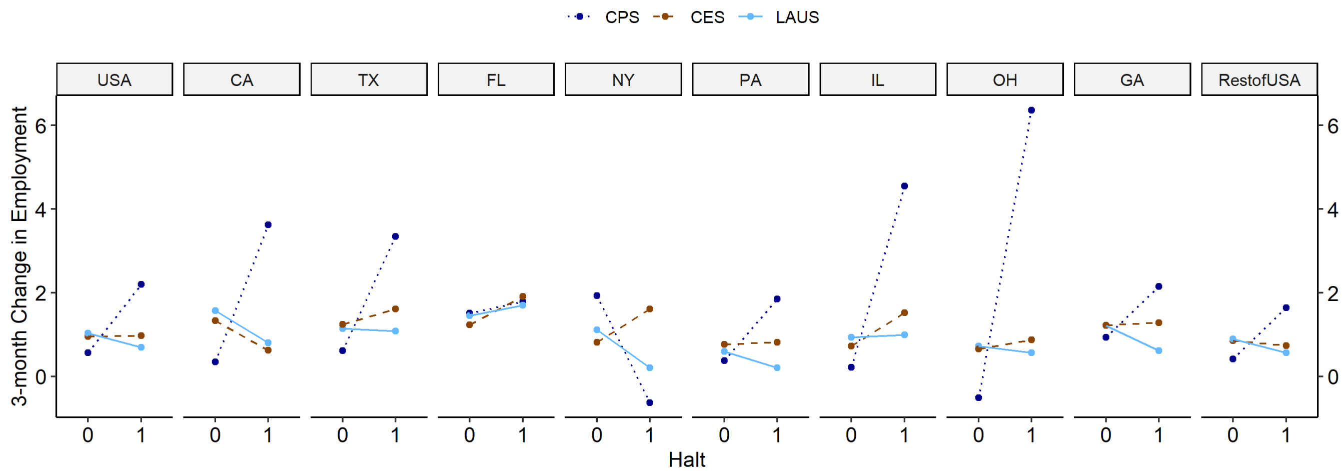

Note: USA and RestofUSA are aggregates of state level data using pre-pandemic employment weights. Halt=1 indicates month of program termination; Halt=0 indicates relevant non-halting months, determined by equation (3). See text for detailed description. 3-month changes are measured per 100 pre-pandemic workers.

and unemployment measures tabulated from the Current Population Survey as the primary inputs."[17]

BLS (2022) explains that the LAUS estimates are "model-based," rather than tabulated from direct sampling. BLS (2022) states:

> The signal-plus-noise model postulates that the observed Current Population Survey estimate consists of a true, but unobserved, labor force value (the signal) plus noise that reflects the error arising from taking a probability sample rather than a complete census of the population.

If the LAUS were simply smoothing out noise that results from taking a probability sample rather than a complete census of working and non-working adults, then we would expect this noise to be uncorrelated with our instrument: each state's halt month.

To assess whether this is the case, we construct the monthly deviation of the LAUS

[17] BLS (2020) also notes that payroll employment estimates from the CES survey are also used as model inputs.

Table 6: Effect of halt month on CPS "noise" (LAUS employment - CPS employment)

	(1)	(2) State FEs	(3) State and month FEs
Halt month	-0.98^{***}	-0.98^{**}	-0.83^{***}
	(0.36)	(0.37)	(0.29)
Num. obs.	816	816	816
N Clusters	51	51	51

$^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$

employment from its comparable state-level-aggregated CPS value. That is, define

$$noise_{i,t} = 100 \times \left(Y_{i,t}^{laus} - Y_{i,t} \right)$$

where we use the seasonally adjusted values of both variables. We run the following regression

$$\frac{noise_{i,t+2}}{Y_{i,Jan20}} = \alpha + \beta H_{i,t} + \epsilon_{i,t+2}$$

We run the regression for the same sample as our benchmark specification. We weight regressions and compute standard errors as in our benchmark specification.

Column 1 of Table 6 reports a coefficient of -0.98, which implies that two months following a state's halt month, the noise term is reduced by 0.98 percent. Rather than eliminating noise due to sampling, the coefficient indicates that the "noise" may be resulting from the large employment effect of the cessation of EUB.

This suggests that the signal-to-noise procedure is masking the causal effect of employment that can be seen directly if one uses the (state-level-aggregated) survey-based CPS employment panel. This effect is statistically significant at a one percent level. Column 2 adds state fixed effects. Column 3 adds both state and month fixed effects. In both cases, the effect of halt month on "noise" is negative and statistically different from zero.

BLS (2021) detailed difficulties with constructing the LAUS series during this episode, noting that COVID 19 presented "an unprecedented challenge due to its magnitude and scope." For these reason, we use the state totals from direct CPS respondents as our employment measure.

F Why not construct the outcome variable using the CES instead of the CPS?

Figure 6: Effect of halting EUB on CES employment pre and post treatment, 90 percent pointwise confidence intervals

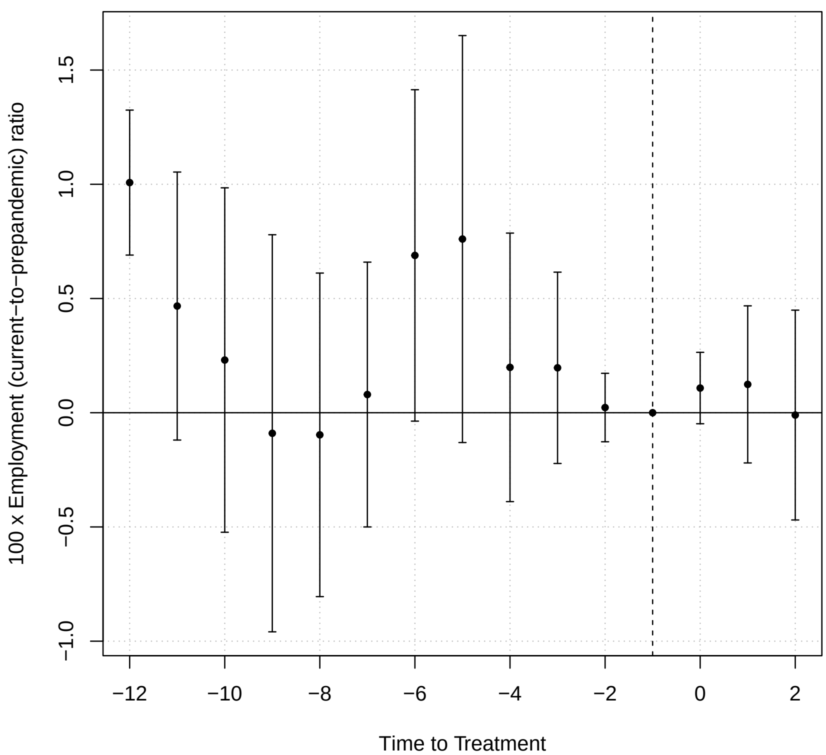

Note: Outcome variable=100 × Ratio of CES employment level to pre-pandemic employment level. We weight regressions by pre-pandemic employment and compute standard errors clustering by census division.

The CES employment data are not suitable for our purposes. First, as explained in the main text, the CES excludes gig and contractor workers. Pandemic Unemployment Assistance benefits were specially established to provide benefits to workers in this category who lost employment. If individuals on this program returned to gig and contract work upon reentry into employment this would be missed if we had used the employer-based CES data.

Second, a previous literature has found that the household employment survey marks changes (turning points) in employment series more rapidly than the CES. For example, Bolwer and Morisi (2006) discusses the historical discrepancies between the household survey and CES found that after the trough of the 2001 recessions: "establishment survey employment continued to decline while household survey employment began to show signs of growth." The authors cite firm births and deaths, "off-the-books" employment and self employment as potential reasons for the discrepancy. Cajner et al. (2022) argue that firm birth and death errors in the CES are typically small but large in recessions using the Great Recession as an example.

Nonetheless, one might expect establishment survey employment to show a muted (relative to CPS employment) but positive response to halting EUB. To investigate this, we present an event study plot similar to Figure 3 except we use state-level CES employment to construct our outcome variable. Figure 6 plots the analogous coefficients and confidence regions. The figure indicates no substantive difference between the treated group (June halters) and untreated group (September halters) in the first few month following treatment. While the above discussion of the CES employment construction likely plays a role in explaining the apparent disconnect between the CES and CPS employment responses, the issue merits further investigation.

References for Appendix

Bowler, M. and T. Morisi. 2006. "Understanding the Employment Measures from the CPS and CES Survey," Monthly Labor Review, February, 23-38.

Bureau of Labor Statistics. 2020. "Local Area Unemployment Statistics: Estimation Methodology," https://www.bls.gov/lau/laumthd.htm, extracted July 25, 2022.

Bureau of Labor Statistics. 2021. "Questions and Answers on Changes to Model-Based Estimation in the Local Area Unemployment Statistics (LAUS) Program"

Bureau of Labor Statistics. 2022. "Local Area Unemployment Statistics: Frequently Asked Questions," https://www.bls.gov/lau/laufaq.htmQ21, extracted July 25, 2022.

California Economic Development Department. 2021. "News Release: EDD to Resume Work Search Requirement to Obtain Unemployment Benefits," June 17.

Cajner, T., L. Crane, R. Decker, A. Hamins-Puertolas, and C. Kurz. 2022. "Improving the Accuracy of Economic Measurement with Multiple Data Sources: The Case of Payroll

Employment Data," NBER Chapters, in: Big Data for Twenty-First-Century Economic Statistics, pages 147-170, *National Bureau of Economic Research, Inc.*

Callaway, B. and P. Sant'Anna. 2021. "Differences-in-Difference with Multiple Time Periods," *Journal of Econometrics*, 225(2), 200-230.

Dube, A., D. Girardi, O. Jorda and A. Taylor. 2022. "A Local Projections Approach to Difference-in-Difference Event Studies," working paper.

Goodman-Bacon, A. 2021. "Difference-in-differences with Variation in Treatment Timing," *Journal of Econometrics*, 225(2), 254-277.

Sun, L. and S. Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogenous Treatment Effects," *Journal of Econometrics*, 225(2), 175-199.