Increasing Employment by Halting Pandemic Unemployment Benefits

Iris Arbogast and Bill Dupor

In mid-2021, 26 states halted participation in all or some federal emergency unemployment benefits (EUB) programs before those programs' federal funding lapsed. This article uses this asynchronous EUB cessation between early- and late-halting states to estimate the causal impact of benefit cessation on employment. We find that cessation increased employment by 29 persons for every 100 (pre-halt) EUB recipients. Expressed as a number of jobs, if all states had halted EUB in June, September employment would have been 3.4 million persons higher relative to a no-halt counterfactual. Late-halting states could have significantly accelerated their states' jobs recoveries in the second half of 2021 through early program cessation. (JEL J65, E24)

Federal Reserve Bank of St. Louis *Review*, forthcoming 2022.

As part of the 2020 CARES Act, the federal government augmented regular state unemployment insurance (UI) programs with several temporary measures: a $600 weekly add-on for UI recipients, extended eligibility to persons who otherwise would not have been covered by their state programs (e.g., gig and contractor workers), and the extension of benefits beyond the duration of those provided by regular state programs. These emergency unemployment benefits (EUB) were renewed in later legislation, with the only major change being a reduction in the add-on from $600 to $300 per week.

In late winter and spring 2021, job vacancies in the United States soared to near historic highs while employment growth slowed. This pattern for vacancies and employment was observed across regions and sectors. The Federal Reserve Beige Books from those months provide accounts of business owners lamenting difficulty in filling job vacancies. Moreover, business owners nationwide linked this difficulty with the historic generosity of unemployment benefits.

Over a period of several weeks, 26 governors announced that their states would halt participation in these EUB programs either partially or fully.[1] Twenty states halted program participation between June 19 and July 3. Four states did so on June 12 as did two states later in July. After September 9, the remaining states ended participation because the programs' federal funding lapsed.

Iris Arbogast is a research associate and Bill Dupor is vice president and economist at the Federal Reserve Bank of St. Louis.

Figure 1

Aggregate Employment Gap Path, January 2020-December 2021, and State-Level May 2021 Employment Gaps Placed Along the Aggregate Gap Path

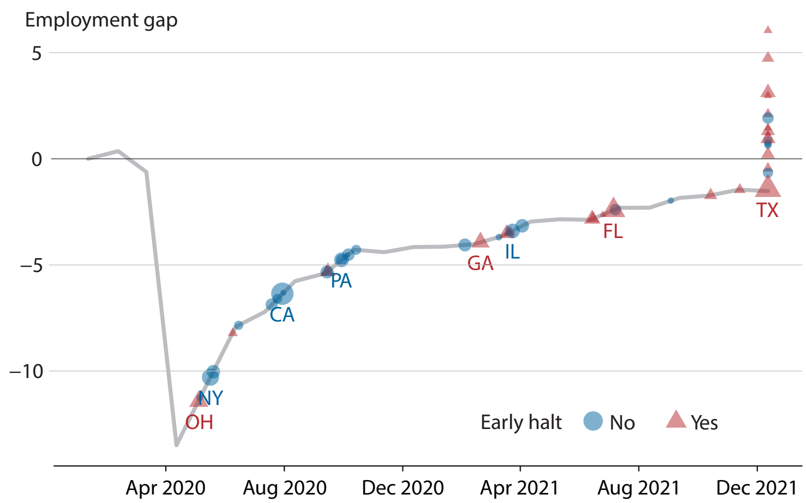

NOTE: The employment gap equals the deviation of employment at a point in time from employment in January 2020. The four largest early- halting and late-halting states are labeled with postal abbreviations. The sizes of the circles and triangles represent the relative pre-pandemic employment levels.

While both Republican- and Democrat-led states saw high job openings and slowed employment growth,[2] with one exception, only Republican governors implemented the policy change before funding lapsed. This politically driven policy variation provides a source of identification to assess the jobs effect of terminating EUB. Our outcome variable is the four-month change in employment, scaled by the lagged number of EUB recipients. We run a panel least-squares regression of the outcome on a dummy indicator for the halt month and include several alternative sets of control variables. The model is identified by our assumption that, conditional on our control variables, the regression error term is orthogonal to the halt-month indicator.

We find that, as a result of cessation, employment increased by 29 persons for each 100 individuals receiving these benefits pre-cessation. We show this effect is statistically significant and robust to controlling for a battery of additional covariates. Expressed as an aggregate jobs effect, our estimate implies that had all states terminated benefits in June 2021, employment would have been 3.4 million persons higher in September 2021 relative to a no-halt counterfactual.

1 STATE-LEVEL EMPLOYMENT DYNAMICS BEFORE HALTING BEGINS

In this section, we document an important difference between early- and late-halting states: pre-cessation employment dynamics. Figure 1 demonstrates that early halters were further along in their employment recoveries than late halters in the period immediately preceding the first set of states halting benefits. First, the solid gray path plots the log of the ratio of national employment in each month to national employment in January 2020 multiplied by 100. Thus, a value of –10 implies that, in the respective month, national employment was 10 percent below its pre-pandemic level. The line gives a sense of the "average employment trajectory" along the recovery path. The figure indicates that in the first few months following the employment trough, national employment recovered very quickly, after which the pace of the recovery slowed as the nation moved closer and closer to closing its employment gap.

On the same figure, we also plot the corresponding variable for each state in May 2021. For each state, we locate its point on the national path to give a sense of how far the state had progressed in its recovery up to that date. States that are off the national path had recovered further by May 2021 than the corresponding national value for December 2021. The sizes of the blue circles (late halters) and the red triangles (early halters) indicate the relative pre-pandemic employment levels across states. We label the four largest states of each type with their postal codes.

The figure starkly shows that, immediately prior to the first wave of benefit terminations, early halters were much closer to their pre-pandemic employment levels on average than late halters. The early halters had much smaller employment gaps even before the policy interventions began than the late halters did.

The immediate takeaway from the figure is that one should not interpret states' asynchronous cessation timing as a pure natural experiment. In this article, we take the view that some conditioning is required to infer a causal treatment effect from a cross-state comparison of employment for this episode. Our analysis will address systematic differences across states in two ways we describe below.

Our outcome variable will be the scaled four-month change in employment. We use the four-month change because many continued to receive benefits several months past cessation (due to delayed filing and administrative delays) and also because a share of recipients may have built up a few months of savings to finance not working following their final benefit payment. Also, we will include several one-month employment changes to control for the pattern observed in Figure 1: Along the transition path, employment growth slows as states close their remaining employment gaps.

Next, another implication of Figure 1 is that—since early halters had smaller employment gaps—the early halters also had smaller per capita shares of the population collecting benefits. In fact, EUB recipients per capita in May 2021 in late-halting states were nearly double those in early-halting states. To adjust for this difference across states, we use a novel form of scaling of the outcome variable. We scale the change in employment in a state by the lagged number of EUB recipients in that state, as described in the next section.

2 DATA, MODEL, AND RESULTS

The sample consists of 46 U.S. states and Washington, D.C. The underlying data are monthly and cover December 2020 through December 2021. First, let $Y_{i,t}$ equal the number of employed persons 16 years of age and older in state *i* at month *t*. These non-seasonally adjusted data are from Current Population Survey (CPS) microdata, which we sum to the corresponding state levels. We discuss alternative employment measures after presenting our benchmark findings. We exclude Alabama, Georgia, South Carolina, and Vermont because the U.S. Bureau of Labor Statistics (BLS) does not provide complete EUB data for these states.

To construct our outcome variable, we divide the change in employment by the number of EUB recipients. To construct our recipient measure, we first calculate for each state the sum of weeks paid on regular state programs, Pandemic Unemployment Assistance (PUA), and Pandemic Emergency Unemployment Compensation (PEUC), which are available from the Employment and Training Administration (ETA) 5159 reports.[3] To map the number of weeks paid to the number of recipients, we divide the number of weeks by four. Next, we multiply this raw recipient number by 0.8 to construct our recipient number. This adjustment is necessary because, according to Census Pulse Surveys from the time, about 20 percent of surveyed UI recipients reported engaging in some work while collecting UI. 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 article, "recipients" refers to non-working recipients.

Our outcome variable is then

$$y_{i,t+\delta} = 100 \times \left(\frac{Y_{i,t+\delta} - Y_{i,t-1}}{Recipients_{i,t-5}} \right). \tag{1}$$

Importantly, we use the fifth lag of the number of beneficiaries in the denominator of the dependent variable. For the response horizon of interest, the denominator is (in every period) predetermined with respect to the halt month. Thus, the response to the halt shock arises from the dependent variable's numerator, that is, the change in employment.

The independent variable of interest, $h_{i,t}$, equals 1 if state i ended participation in all or some of the EUB programs in month t and zero otherwise. The CPS survey occurs during the calendar week containing the 12th day of the respective month. With this in mind, for each state we choose the halt month as the month in which the 12th day is closest to the date the state ended benefits. For example, 24 states and Washington, D.C., terminated benefits on September 9; therefore, we set September as the halt month in these cases.

Our use of lagged recipients in the denominator of the dependent variable is motivated by the following overarching question in public discourse during the episode: When a state halts benefits, how many recipients losing benefits will become employed? The outcome variable we choose is an "employment yield" for those losing benefits, which squarely addresses this question. If instead of the change in employment per beneficiary we were to use employment growth as the outcome variable, we would be answering a *different* question than the one posed above.

To further see the usefulness of our construct, consider the following extreme example. If a state had zero beneficiaries when it announced a halt to benefits, then that state's "return to work" channel would be absent.

Our regression equation is

$$y_{i,t+3} = \gamma_i + \psi_t + \phi h_{i,t} + \beta' X_{i,t-1} + \eta_{i,t+3}. \tag{2}$$

As explained in the previous section, following the employment trough in spring 2020, states generally followed a similar dynamic path for employment, although their initial gaps and transition rates differed substantially; therefore, we condition on these dynamics in our estimation.

The coefficient ϕ is the causal impact of halting benefits on the employment change per 100 EUB recipients between $t-1$ and $t+3$. γ_i and ψ_t are state and time fixed effects, respectively. In our

Table 1

Jobs Effect of Benefit Termination: Four-Month Employment Change Per 100 Emergency Benefit Recipients

	(1) Benchmark	(2) Drop lags	(3) Drop FE	(4) Placebo
Termination	28.82** (11.19)	22.03* (11.30)	29.02*** (9.37)	5.13 (7.00)
Emp change, 1-month lag	–0.69*** (0.08)		–0.54*** (0.08)	–0.68*** (0.08)
Emp change, 2-month lag	–0.55*** (0.11)		–0.44*** (0.12)	–0.54*** (0.11)
Emp change, 3-month lag	–0.52*** (0.15)		–0.38*** (0.14)	–0.52*** (0.15)
Number of observations	470	470	470	470

NOTE: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. FE, fixed effects; emp, employment. Regressions include state and time fixed effects (not reported) and are weighted by the pre-pandemic (1/2020) employment level. Standard errors are computed using state-level clustering. Sample covers 12/2020 to 12/2021. Sample excludes Alabama, Georgia, South Carolina, and Vermont because of unavailability of EUB recipient data.

benchmark specification, $X_{i,t-1}$ contains three lags of the one-month change in employment scaled by the lagged number of recipients. We estimate equation (2) using least squares in which we weight each observation by its pre-pandemic (January 2020) employment level. We compute standard errors using state-level clustering.[4]

Table 1 reports estimates of equation (2). Column 1, our benchmark specification, indicates that halting EUB increased employment by 28.8 persons for each 100 (pre-halt) recipients. The estimate is statistically different from zero at the 5 percent level. Thus, halting benefits provided a substantial boost to state employment. By delaying the end of EUB, late-halting states delayed part of their states' employment recoveries dramatically.

To express this effect as the number of jobs created, note that in January 2021, five months prior to the treatment, the total number of EUB recipients was about 11.8 million. Thus, had all 47 states in our sample halted EUB in June, national employment in September would have been 3.4 million persons higher relative to a "no early halt" counterfactual.[5] In actuality, the 23 early-halting states for which we have data had about 3.3 million EUB recipients five months prior to their halt dates. Thus, the jobs effect of the 23 early-halter initiatives was to increase those states' combined employment by 950,000 persons.

Column 2 of Table 1 modifies the benchmark by dropping the lagged employment changes as controls. This modification reduces the employment effect from 28.8 jobs to 22.0 jobs. This occurs because being an early-halting state is correlated with having a smaller remaining employment gap, and having a smaller remaining employment gap implies that future employment growth tends to be low. Thus, failing to condition on pre-cessation state employment dynamics downwardly biases the jobs effect estimate.

In our benchmark specification, we follow the common practice of adding fixed effects to the panel regression. For example, state fixed effects would control for different trend behavior in

Table 2

Jobs Effect of Benefit Termination: Four-Month Employment Change Per 100 Emergency Benefit Recipients, Impact of Changing Beginning Month of Sample

	(1) Start 12/2020 (benchmark)	(2) 9/2020	(3) 10/2020	(4) 11/2020	(5) 1/2021
Termination	28.82** (11.19)	28.20** (10.88)	28.30** (10.97)	28.60** (11.04)	29.00** (11.31)
Number of observations	470	610	564	517	423

NOTE: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions include state and time fixed effects as well as three lags of the one-month change in employment (not reported) and are weighted by the pre-pandemic (1/2020) employment level. Standard errors (in parentheses) are computed using state-level clustering. Each sample ends in 12/2021. Sample excludes Alabama, Georgia, South Carolina, and Vermont because of unavailability of EUB recipient data.

employment growth. Strict exogeneity of the halt-month variable would mean that removing state fixed effects should have little effect on our results. In Column 3 we report results without month and state fixed effects and find little change in the employment effect.

Next, as a placebo exercise, we randomly reassign halt months across states, with a uniform distribution from June to September, to create a new termination month variable. Column 4 reports the benchmark results except we replace the actual halt variable with this placebo variable, suggesting that our instrument is not estimating a spurious jobs effect.

The starting month of the sample is dictated by our intent to capture employment dynamics using a sufficiently long sample as well as to build a model for the pandemic period. The latter concern limits how far back our period of estimation can reasonably extend. Recall that we include as controls three lags of the one-month change in employment scaled by the number of recipients. We lag the recipients variable in the denominator such that, for every state in every month, the variable is measured in a pre-treatment month. To accomplish this, the variable is lagged by five periods. As such, we cannot start a sample before September 2020, as data for EUB start in May 2020 in many states.

We start our sample in December for the benchmark results in order to avoid using recipient data from the early months of the pandemic, before recipient numbers had stabilized. For example, California's number of recipients more than doubled between May and August 2020. Table 2 shows that our results are robust to changes in the start dates of the sample used. Estimates from samples starting in September, October, and November 2020 and January 2021 compared with the December 2020 results vary by less than one job per 100 benefit recipients.

Next, Table 3 shows that recognizing cross-state differences in the number of EUB recipients is important in understanding our results; ignoring them biases the jobs effect downward. First, Column 1 restates the benchmark estimates. In Column 2, our dependent variable's definition deviates from the benchmark model. Instead of scaling employment changes by the lagged number of recipients, Column 2 scales by lagged state employment multiplied by the national recipients-to-population ratio.[6] As such, we treat each state as if it had the same number of recipients (after controlling for state population). In this case, the jobs coefficient is 14.43, falling roughly one-half from the benchmark specification.

Table 3

Jobs Effect of Benefit Termination: Four-Month Employment Change Per 100 Emergency Benefit Recipients, Effect of Changing Variable Used to Scale Employment Changes

	(1) Scale by claims (benchmark)	(2) Scale by average claims	(3) Scale by PP emp
Termination	28.82** (11.19)	14.43** (6.41)	1.17** (0.50)
Number of observations	470	470	470

NOTE: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. PP emp, pre-pandemic employment. Regressions include state and time fixed effects as well as three lags of the one-month change in employment (not reported) and are weighted by pre-pandemic (1/2020) employment level. Standard errors (in parentheses) are computed using state-level clustering. Sample covers 12/2020 to 12/2021. Sample excludes Alabama, Georgia, South Carolina, and Vermont because of unavailability of EUB recipient data. See text for dependent variable construction in columns 2 and 3.

Table 4

Jobs Effect of Benefit Termination: Four-Month Employment Change Per 100 Emergency Benefit Recipients, Effect of Using Alternative Employment Measures

	(1) CPS non-SA (benchmark)	(2) CPS SA	(3) CES non-SA	(4) CES SA
Termination	28.82** (11.19)	25.73* (10.78)	3.17 (3.90)	0.98 (1.35)
Number of observations	470	470	470	470

NOTE: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. CPS, Current Population Survey; CES, Current Employment Statistics; SA, seasonally adjusted. Regressions include state and time fixed effects as well as three lags of the one-month change in employment (not reported) and are weighted by the pre-pandemic (1/2020) employment level. Standard errors (in parentheses) are computed using state-level clustering. Sample covers 12/2020 to 12/2021. Sample excludes Alabama, Georgia, South Carolina, and Vermont because of unavailability of EUB recipient data.

In Column 3, we scale employment changes by pre-pandemic employment, instead of the number of recipients, to construct the dependent variable. This is a second way, beyond the Column 2 specification, to ignore cross-state differences in the number of recipients. The jobs coefficient falls to 1.17. In this case, the coefficient is interpreted, roughly, as a growth rate response rather than a per 100 persons response. The total jobs effect implied by Column 3 is 1.73 million ($= 0.0117 \times 148$ million: January 2020 employment for the 47 states in our sample). This is roughly one-half the corresponding impact from our benchmark specification.

Table 4 reports the effect of changing the employment measure used. The benchmark specification uses non-seasonally adjusted data because we do not want to exclude seasonal hires from the jobs effect of terminating benefits. Column 2 of Table 3 uses the benchmark specification's employment series, except we seasonally adjust using the U.S. Census Bureau's X-13 procedure; the seasonal adjustment has minimal impact on the coefficient of interest. Two alternatives to the state-aggregated CPS microdata are employment from the Current Establishment Survey (CES) and the Local Area Unemployment Statistics (LAUS) datasets.

Table 5

Jobs Effect of Benefit Termination: Four-Month Employment Change Per 100 Emergency Benefit Recipients, Effect of Adding Additional Controls

	(1) (Benchmark)	(2)	(3)	(4)	(5)	(6)	(7)
Termination	28.82** (11.19)	28.95** (11.09)	27.39** (10.63)	28.29** (11.36)	28.82** (11.17)	27.75** (10.56)	27.67** (10.45)
Mask use		1.13 (16.88)				3.75 (16.93)	11.42 (16.55)
Stringency index			–6.10 (5.94)			–6.36 (5.91)	–6.94 (5.88)
COVID-19 cases				–9.80 (17.49)			–14.80 (18.23)
COVID-19 deaths					–0.16 (22.08)		
Number of observations	470	470	470	470	470	470	470

NOTE: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Regressions include state and time fixed effects as well as three lags of the one-month change in employment (not reported) and are weighted by the pre-pandemic employment level (in January 2020). Standard errors (in parentheses) are computed using state-level clustering. Sample covers December 2020 to December 2021. Sample excludes Alabama, Georgia, South Carolina, and Vermont because of unavailability of EUB recipient data.

We use the household-based CPS data rather than the firm-based CES data because the latter excludes contractors and gig workers. More than one-half of EUB recipients in May 2021 were collecting PUA benefits, a program for contractors and gig workers. Using the establishment-based data would likely bias our results downward substantially if pre-recession unemployed contractors and gig workers tended to return to the same work upon taking jobs. Results using non-seasonally adjusted CES data are reported in Column 3 of Table 4, and results for seasonally adjusted data are reported in Column 4.

We use aggregated micro-level data from the CPS instead of the model-based LAUS data from the BLS. The LAUS uses data from the CPS as well as from the CES and state UI systems as inputs to time-series models. The BLS noted that the COVID-19 pandemic created an unprecedented challenge to the BLS's LAUS model estimation due to the magnitude and scope of outliers. According to the BLS (2021), in a BLS FAQ section for the LAUS, the model estimates were biased by the influence of pre-pandemic data during the pandemic.[7]

Table 5 adds a series of alternative control variables that might potentially explain the different outcomes between early and late halters: the mask-usage rate, a lockdown intensity index, COVID-19 cases, and COVID-19 deaths. We introduce scaled versions of each variable expressed as four-month changes between months $t - 4$ and t. We measure lockdown intensity using an index from the Oxford COVID-19 Government Response Tracker that averages over indicator variables for containment policies, closure policies, and public information campaigns. Mask-usage data are from the Institute of Health Metrics and Evaluation and indicate the percent of the population reporting always wearing a mask when leaving home. Case and death data are from the *New York Times*, based on reports from state and local health agencies.

Figure 2

Benefits-to-Pre-pandemic Employment Ratio for 12 States in the Month of and Months Surrounding Benefit Cessation

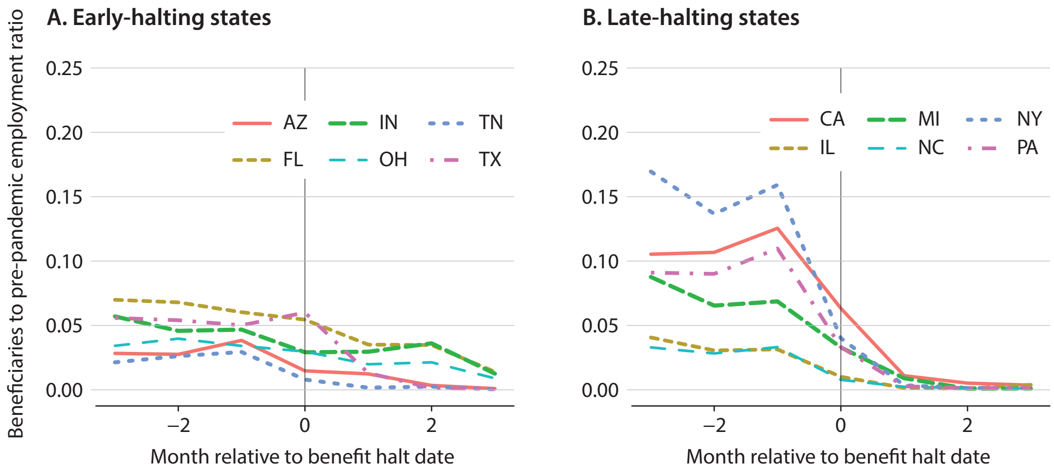

NOTE: Panel A reports the recipient ratio for the six early-halting states with the largest populations, and Panel B reports the recipient ratio for the six late-halting states with the largest populations. States are labeled with postal abbreviations.

We standardize each regressor to have mean zero and unit variance. Thus, the reported coefficient on a given control variable should be interpreted as the four-month change in employment per 100 EUB pre-halt recipients in response to a one-standard-deviation increase in that control variable. For both levels (results available on request) and changes, the benchmark results do not change by more than two persons per 100 emergency benefit recipients. This likely occurs because the covariates at the state level vary only slightly over time and state-level differences are controlled for with fixed effects.

3 BENEFITS PAID AFTER PROGRAM TERMINATION

In this section, we document that many states, even after halting EUB, continued to pay out benefits for several months. For example, Alaska paid about 21,400 weeks of PUA benefits to recipients in May (the month before halting) and roughly 12,200 weeks in August. This persistence was possible because federal law permitted individuals to submit claims if a past unemployment spell occurred when the PUA program was in effect. Also, backlogs in some states may have played a role in payment delays.

Figure 2 plots the ratio of the sum of PUA and PEUC recipients relative to pre-pandemic employment for six early- and six late-halting states. The figure indicates that many states experienced only a gradual decline in EUB payouts following the termination of their program participation.

Even if an individual received benefits from a state for a past unemployment spell, program termination meant that—going forward—a potential disincentive for supplying labor had been lifted. The termination channel should lead to increased employment in the months after the termination date. However, some individuals continued to receive benefits. Liquidity constrained consumers may have been able to finance consumption and use those benefits to shift their time away from work and toward non-market activities.

In another article (Arbogast and Dupor, 2022), we take a different approach. Using the same underlying data as in this article, we regress changes in employment on the reduction in the number of beneficiaries, using the timing of benefit termination as an instrument. In that article, we find that employment increases by 37 persons for every 100 persons who, on net, stop receiving EUB. As such, the approach taken in the current article may understate the stimulative jobs effect of EUB cessation.

4 RELATED RESEARCH

There is a large body of empirical research examining the incentive effects of unemployment on the labor supply going back at least 50 years. A long literature review is beyond the scope of a short article such as this one. Instead, we describe a few particularly relevant articles on UI benefit termination and extension during the two most recent recessions.

During the 2007-09 recession, unemployment benefits were extended through two temporary federal programs. One extended benefits up to 53 weeks depending on the state and expired in 2013. The other extended benefits for between 13 and 20 weeks in states with high unemployment rates. Farber, Rothstein, and Valletta (2015) study changes in benefit lengths resulting from each program's end, using individual-level CPS data. They estimate that availability of UI to an unemployed worker reduced the probability of exiting unemployment by 3.5 percentage points from 2008 to 2011 and 2.7 percentage points from 2012 to 2014 but do not find significant impacts on the likelihood of transitioning from unemployment to employment. They conclude that the primary effect of phasing out the benefits programs was on labor force attachment.

Coombs et al. (2021) study the impact of EUB early termination during the same time period as this article. They use individual-level data on bank transactions from a sample of low-income and credit-constrained workers to examine the impact of the policy change. They use a difference-in-difference approach comparing UI recipients in early-halting and non-early-halting states and reweight this sample for people in non-early-halting states to match the unemployment duration in the early-halting sample. They find a 4.4-percentage-point relative increase in the probability of job finding through the first week of August 2020 in states where EUB were halted early.

Several authors look at how CPS unemployment-to-employment transition rates differ between halting and non-halting states and use these rates as a measure of the treatment effect of early benefit termination (e.g., Holzer, Hubbard, and Strain, 2021). However, there were nearly twice as many EUB recipients measured by the 5159 data as there were unemployed measured by the CPS. This difference is due, at least in part, to the fact that many EUB recipients had no work-search requirements during much of the early pandemic period. Our observation motivates us to use the number of EUB recipients as the denominator in our treatment variable.

In contrast to the above articles that use individual-level data, we use data aggregated to the state level. Using state-level data brings us closer to identifying a macroeconomic employment effect, which is our primary interest, than one would find using individual-level data in the presence of cross-individual spillovers. Using individual-level data may miss important spillovers that could be either positive or negative. For example, if a person increases consumption upon losing emergency benefits and taking a job (e.g., from spending on clothes and fuel and car maintenance for traveling to and from work), then this increase in consumption may drive up demand for goods in the local economy. This increase may in turn stimulate employment in the state indirectly. This indirect positive effect would be missed in an individual-level regression. Thus, the estimate from that regression would provide a downwardly biased estimate of the macro jobs effect of halting UI benefits.

Indeed, the results that we find are larger than those in the individual-level articles we discuss above. Farber, Rothstein, and Valletta (2015) find no increase in the probability of job finding in the 2007-09 recession. Coombs et al. (2021) find only a small effect; however, it is not directly comparable to our results, because they use job-finding probabilities as an outcome.

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 partially immunized against this concern. This is because within-state spillovers are subsumed by aggregating the outcome variable to the state level.[8]

This article estimates that employment increased by 29 persons for each 100 pre-halt recipients in a state. We note, however, that this does not necessarily mean that 71 persons stopped receiving benefits and/or were no longer making claims for past unemployment spells. Our use of state-level employment data implies that we cannot track individuals receiving EUB before cessation to see whether these persons were the ones that boosted employment or alternatively whether employment increased through a more indirect channel. Rather, our outcome variable is the net change in total state employment. Parsing the extent to which the employment increase comes from the pre-treatment number of EUB recipients, the unemployed not collecting EUB, or those out of the labor force is not possible using our methodology.

5 CONCLUSION

The exogenous variation resulting from the decisions of about one-half of state governors to cease providing EUB in June and July 2021 provides a valuable opportunity to identify EUB programs' short-term causal effect on state-level employment. Our estimate indicates a 29-person increase in state employment—within a few months—for every 100 (pre-halt) EUB recipients. The jobs effect is statistically significant and robust to allowing for a number of controls. There remains a great deal of work to be done on this episode. Useful research might examine the effect of halting benefits on, for example, the employment response of older individuals, labor force participation, job openings, quits, and hires. ■

NOTES

1 Montana Governor Greg Gianforte announced his state's plan on May 4, 2021. On May 6, South Carolina Governor Henry McMaster set forth a similar plan. By May 17, 21 total states had laid out plans to halt benefits.

2 Beige Book reports discussing worker shortages and linking them to unemployment benefits span both Federal Reserve Bank regions with primarily non-halting states (e.g., the New York Fed) and those with primarily halting states (e.g., the Atlanta Fed).

3 We include regular state programs in this measure because a regular program recipient also received the $300 add-on from the EUB program in halting states before those states ended participation.

4 Given our panel's short time dimension, we do not attempt to adjust standard errors for potential serial correlation in our reported results. In results available upon request, we compute standard errors—applying a Newey-West correction—and find the adjustment has little effect.

5 Recall from equation (1) that recipients are lagged by five months in the outcome variable; thus, we use January recipient data for our June halting counterfactual.

6 Specifically, the denominator becomes $P_{i,t-5} \Big/ \left(\sum_i Recipients_{i,t-5} \Big/ \sum_i P_{i,t-5} \right)$. Here, P is state population.

7 The FAQ section states that "the regression coefficients showed some flexibility in real time between the CPS and covariate inputs, but not enough to prevent bias due to the influence of past data."

8 Note that our approach does not account for potential cross-state spillovers. These spillovers imply our state-level estimates would provide a biased estimate of the national effects of halting EUB. Thus, using state-level data is not sufficient to completely overcome the spillover issue. For a methodology to conduct causal inference regarding aggregate, spillover, and local effects in a unified framework, see Conley et al. (2021).

REFERENCES

Arbogast, I. and Dupor, B. "The Jobs Effect of Ending Pandemic Unemployment Benefits: A State-Level Analysis." Working Paper 2022-10A, Federal Reserve Bank of St. Louis, 2022; https://doi.org/10.20955/wp.2022.010.

Conley, T.; Dupor, B.; Ebsim, M.; Li, J.C. and McCrory, P. "The Local-Spillover Decomposition of an Aggregate Causal Effect." Working Paper 2021-006A, Federal Reserve Bank of St. Louis, 2021; https://doi.org/10.20955/wp.2021.006.

Coombs, K.; Dube, A.; Jahnke, C.; Kluender, R.; Naidu, S. and Stepner, M. "Early Withdrawal of Pandemic Unemployment Insurance: Effects on Earnings, Employment and Consumption." Working paper, University of Toronto, August 2021.

Farber, H.; Rothstein, J. and Valletta, R. "The Effect of Extended Unemployment Insurance Benefits: Evidence from the 2012-2013 Phase-Out." *American Economic Review*, 2015, *105*(5), pp. 171-76; https://doi.org/10.1257/aer.p20151088.

Holzer, H.; Hubbard, R. and Strain, M. "Did Pandemic Unemployment Benefits Reduce Employment? Evidence from Early State-Level Expirations in June 2021." NBER Working Paper 29575, National Bureau of Economic Research, 2021; https://doi.org/10.3386/w29575.

U.S. Bureau of Labor Statistics. "Questions and Answers on Changes to Model-Based Estimation in the Local Area Unemployment Statistics (LAUS) Program." 2021; https://www.bls.gov/lau/gen-5-changes-in-2021.htm.