Disincentives to work or to search? The effect of increased benefit generosity during the COVID-19 pandemic on unemployment

ERASMUS UNIVERSITY ROTTERDAM

Erasmus School of Economics

Bachelor Thesis [programme Economics & Business Economics]

Name Student: Bob Bremer

Student ID number: 543170

Supervisor: prof.dr. Dinand Webbink

Second Assessor: prof. dr. Anne Gielen

Date final version: 20-07-2022

In this paper I investigate the unemployment effects of the increased benefit generosity of the COVID-19 Unemployment Insurance programs in the United States. After the start of the pandemic, the U.S. government introduced several unemployment insurance programs to help the unemployed. Several states decided to end these programs in June and July 2021, while the other states ended them in September 2021. I exploit this variation in the benefit generosity using a difference-in-difference design. I use the monthly sample of the Current Population Survey from January 2021 until August 2021. I find that the unemployment rate decreased more in the states that terminated the programs early than those that did not. However, there is no significant difference in the employment-to-population ratio, suggesting that the unemployed moved out of the labor force after the programs ended. Using various tests, I show that the parallel trend assumption is likely to hold.

I. Introduction

Unemployment benefits aim to protect workers against unemployment. In times of recession, risk of unemployment is higher and unemployment durations are longer. This is why many governments increase unemployment benefit generosity during recessions. However, unemployment benefits represent a trade-off. While workers are more protected, moral hazard is increasing. Workers can have less incentives to search for a job or become more selective in choosing a job. Investigating this trade-off is difficult as it is generally not possible to observe whether individuals really need a benefit or try to take advantage of benefit schemes. This paper exploits a recent change in benefit schemes to investigate the effect of increased benefit generosity on unemployment.

In March 2020 benefit generosity was increased because of the COVID-19 pandemic. Several Unemployment Insurance (UI) programs were introduced to protect workers against unemployment. As the pandemic continued, the programs were extended. This sparked discussion on whether the programs were too generous and extended too long. As the labor market recovered in the summer of 2021, more states were arguing that the increased benefit generosity led to work disincentives.

I will investigate the effects on unemployment by using a difference-in-difference design exploiting the fact that the federal expiration of the UI programs was in September 2021, but several states decided to end these programs in June or July 2021. The outcome variable in the main analysis is the unemployment rate. I use data from January 2021 to August 2021 from the monthly Current Population Survey (CPS). The identifying assumption of the regression is that the trends would have developed the same in the treatment and control states if the COVID-19 UI programs were not terminated early.

The main finding is that the unemployment rate decreased more in the treatment states than in the control states. The magnitude of this effect is 0.7 percentage point. This result is robust to the inclusion of individual controls, state fixed effects and linear time trends. It is also present when using different groups of control states for which the parallel trend assumption is more likely to hold. Another important finding is that the employment-to-population ratio did not increase. This implies that even though less people were considered unemployed, the number of people that work did not actually increase. This means labor force participation decreased.

Limitations of this study are the design of the data, which is a monthly sample where individuals stay in the sample for four consecutive months. I use the monthly sample to estimate the unemployment rate, which leads to uncertainty regarding the use of standard errors.

Furthermore, the data does not allow to look at individuals over a longer period of time to investigate the total unemployment duration, unemployment flows and the reemployment wages and to investigate which people leave the labor force after treatment. Another limitation is the limited number of observations per state, which leads to less precise results in the additional analyses to provide support for the parallel trend assumption. Finally, the results obtained in this paper on the unemployment effects of benefit generosity are very specific to the pandemic and might not be extrapolated to a different timing and setting.

Much research has been dedicated to the extended benefits in the Great Recession and their macro effect on unemployment (Chodorow-Reich, Coglianese & Karabarbounis, 2019; Boone, Dube, Goodman & Kaplan, 2021). They find small or even insignificant effects. Most of the previous research done in the United States examines the Great Recession. However, several papers look at the introduction of the FPUC program, one of the COVID-19 UI programs. All these papers did not find a significant impact of the FPUC program on employment (Marinescu, Skandalis & Zhao, 2021; Finamor & Scott, 2021; Dube, 2021). The paper closest to mine is from Holzer, Hubbard and Strain (2021). They also look at the early termination of the COVID-19 UI programs and find that the unemployment rate decreased more in the states which terminated the programs early. I add on the literature by providing more insights into the unemployment effect of the COVID-19 programs by using a slightly different approach. Furthermore, I provide more support to show that the parallel trend assumption is likely to hold.

The article is structured as follows. Related literature is discussed in Section II. Section III provides background information on the COVID-19 UI programs. In Section IV the data and methodology are discussed. The results are discussed in Section V. Section VI provides robustness and validity tests. Section VII provides additional analyses. Section VIII discusses the results, and section IX concludes.

II. Related Literature

The effect of unemployment insurance has been studied extensively, especially after the Great Recession, in which unemployment benefits were extended from 26 weeks to up to 99 weeks. Most of these studies look at administrative micro data to investigate the effect of unemployment benefit extensions on the unemployment spell. Johnston and Mas (2018) examine a 16-week cut in potential unemployment duration in Missouri. They compare individuals that enter the benefit program just before and after the cut, using a regression discontinuity design. They find that a one-month reduction in the potential benefit duration leads to a 0.45-month reduction in unemployment spells. They do not find a difference in reemployment earnings conditional on employment. The research implies that policymakers must tradeoff between moral hazard and insurance when determining the duration of the unemployment insurance. It highlights the job search effect of unemployment benefits. Because the benefit increases the outside option of workers, workers reduce their search effort and become more selective. Farber and Valetta (2015) also find negative, but smaller effects of benefit duration on the unemployment spell. They use the increased duration of unemployment insurance during the Great Recession in the United States and exploit variation in timing and size of the UI extension across states to find out what happens to unemployment exits. They find that a one-month increase in benefit duration increases the unemployment duration by 0.06 months. However, the effect of extended benefits on exit from unemployment primarily occurred through a reduction in labor force exits instead of lower job findings. This is the entitlement effect of unemployment benefits: more people participate in the labor force because the value of participation is increased compared to inactivity. Card, Johnston, Leung, Mas and Pei (2015) also look at the effect of unemployment benefits on the duration of unemployment insurance. They use data from Missouri from 2003 to 2013 where they exploit variation around the kink in the UI benefit schedule using a regression kink design. They find that the elasticity of UI duration with respect to the benefit amount is 0.35 in pre-recession periods and between 0.65 and 0.9 in the recession and the years after. Katz and Meyer (1990) examine differences in unemployment spell distributions of UI recipients and non-recipients during 1978 – 1983 in 12 states in the U.S. They find a large increase in the escape rate from unemployment around the time of the benefit exhaustion (end-of-benefit spikes). An increase in the potential benefit duration increases the average unemployment duration by 0.16 – 0.2 weeks. Farber, Rothstein and Valetta (2015) look at the phase-out of benefit extensions after the Great Recession in the United States. They look at the likelihood that an individual exits unemployment and find that the extended benefits decreased the monthly exit rate from unemployment by 15 percent. This effect is mostly driven by exit from labor force rather than exit to employment. They conclude that the phase-out of the extended benefits is not important in explaining the low labor force participation in the recovery of the Great Recession. Another important conclusion is that the extended benefits did not have large moral hazard effects on job-finding rates.

The studies above use micro data on the individual level to investigate the effect of UI. Chodorow-Reich et al. (2019) investigate how an extension of UI affects macroeconomic outcomes like state-level unemployment. They look at the benefit extensions during the Great Recession. An issue is that extensions are usually triggered once unemployment is high. This makes it hard to identify the effect of the program, because of endogeneity issues. They solve this by using measurement error at the timing of implementation to identify the effect of the benefit extension. Using this variation, they find that benefit extensions during the Great Recession had limited influence on state-level outcomes. The unemployment rate increased by at most 0.3 percentage points as a result of the extension from 26 weeks to up to 99 weeks. Boone et al. (2021) also examine the impact of UI on aggregate employment. They use cross-state variation in the benefit duration at the time of the Great Recession using bordering counties in different states. In their first strategy they compare employment outcomes within county pairs using county level employment data from 2007 – 2014 from the Quarterly Census of Employment and Wages. In a second strategy, they use an event study design exploiting variation from national-level policy changes to instrument for the changes in state-level UI duration. They argue that this variation is more likely to be exogenous for the bordering counties than state-level policies. They do not find a significant effect of increasing UI generosity on aggregate employment. This result is not consistent with the negative effect that is often found in studies using micro-data. The results are more consistent with job rationing and aggregate demand channels.

Differences between the micro and macro effects are addressed by the theory of Landais, Michaillat and Saez (2018), who look at the optimal UI in matching models. In earlier models of optimal UI such as Baily (1978) and Chetty (2006), the optimal UI is a trade-off between insurance and incentive. On the one hand it helps workers smooth consumption, but it also discourages job search. However, this only looks at the labor supply side. It may be that UI puts upward pressure on wages because of the higher outside option of workers, which leads to less job creation by firms. There is also the possibility that there is a fixed number of jobs and therefore reduced search efforts by workers improves the opportunities for other workers, which they call the rat race effect. In the first case, the microeconomic effect is smaller than the total effect of UI on unemployment. In the second case, the microeconomic effect is larger than the total effect of UI on unemployment. The total effect of UI on unemployment consists of both the micro-effect and the externalities. This means that increasing UI may decrease or increase labor market tightness depending on which channel dominates. Several papers look at both the micro effect and the market externalities. For example, Lalive, Landais and Zweimüller (2015) investigate equilibrium effects using the Regional Extensions Benefit Program in Austria. They use a difference-in-difference design and exploit the fact that the program was only implemented in certain regions and only a subset of workers was eligible. They look at workers that live in regions where the program is in place but who are not eligible for the program. They find sizable market externalities of UI. For non-eligible workers aged 50 to 54, who are similar to the treated workers, unemployment duration decreases by 6 to 8 weeks compared to similar workers in regions without the program. These results imply that the macro effect of UI on unemployment is smaller than the micro effect and that an increases in the generosity of UI increases labor market tightness. Marinescu (2017) also finds that the macro effect is smaller than the micro effect. She looks at a job platform to investigate the job applications and vacancies around the Great Recession, exploiting state-level variation in potential benefit duration as a result of benefit extension programs. She finds that the increase in the potential benefit duration reduces the number of applications at the state-level but is does not affect the number of vacancies.

Several studies have looked at the effect of the pandemic UI programs by investigating the implementation. One of the programs is the FPUC, which is an additional $600 weekly benefit with the goal of replacing 100% of the mean wage when combined with other UI benefits. Ganong, Noel and Vavra (2020) find that between April and July 2020, 76% of eligible workers had replacement rates above 100%. As a result, they find that the FPUC reverses group-level income patterns. This research shows that the program boosted the income of the unemployed, but it does not focus on the consequences of the programs. Marinescu et al. (2021) investigate the effect of the FPUC on job applications and vacancies using data from an online job platform. They use variation in the proportional increase in benefits and look at the number of applications and vacancy postings each week in each local labor market. They find that a benefit level increase of 10% caused a 3.6% reduction in applications. Vacancy creation was not affected. This implies that the labor market tightness was increasing, consistent with the findings by Marinescu (2017) and Lalive et al. (2015). Finamor and Scott (2020) examine whether changes in UI generosity due to the COVID programs are associated with differential employment outcomes. They focus on the $600 FPUC and use data from Homebase, a firm that provides scheduling and time clock software to firms. They look at changes from week to week. They find that the negative association between replacement rates and employment starts before the program and that workers with more generous benefits did not have different declines in employment during the program. A great limitation of this study is that causal effects cannot be estimated. Dube (2021) instead, focusses on the expiration of the $600 FPUC in July 2020. He exploits the state-level variation in median earnings replacement rates with a difference-in-difference event study design, which estimates macro effects. He uses data from the Census Household Pulse Survey and finds that there is no indication of a substantial impact on employment. This finding is consistent with market-externalities such as job rationing or aggregate demand channels. The paper closest to mine is from Holzer et al. (2021), who use a difference-in-difference and event study design to look at the flow of unemployed workers by exploiting the early termination of the FPUC and PUA in some states. Their outcome of interest is the probability of being unemployed in the previous month and being employed in this month. They use the longitudinal part of the Current Population Survey to look at individuals in consecutive months. They find that early termination is associated with a 14-percentage point increase in the unemployment to employment flow. They use their estimates of the unemployment flows to obtain a counterfactual for the unemployment rate. If the control states would have chosen to end the programs in June, the unemployment rate would have been 0.8 percentage points lower in July and 0.7 percentage points lower in August. The employment rate would have been 0.7 and 0.6 percentage points higher. They also estimate this counterfactual unemployment rate for whole United States. The national unemployment rate would have been 0.3 percentage points lower if all states opted to quit the programs in June. They state that a potential threat to their results is the possibility that treatment and control states were on different paths already before the treatment. Even though they use some tests to support the parallel trend assumption, they cannot be certain the assumption holds.

I add on to the existing literature by looking at the COVID-19 UI programs and their early termination and thus having a credible identification strategy to estimate the effect of UI on the aggregate state-level employment. This gives more insight in the macro effect of an increase in benefit generosity and sheds more light on the effects during the pandemic rather than the Great Recession. Furthermore, I provide more insight into the parallel trend assumption during the time of termination of the programs.

III. Institutional Background

III.A. COVID-19 unemployment insurance programs

When the pandemic hit, the U.S. government implemented several temporary unemployment insurance programs. After their initial implementation, these programs were extended a number of times. The programs include:

-The Federal Pandemic Unemployment Compensation (FPUC). This is $300 a supplement to all regular UI benefits. Initially, the supplement was $600, but this was changed to $300 after the extension on December 26, 2020. During the period of the FPUC, states were not allowed to lower the regular UI benefit amount, to make sure the $300 was a supplement to the regular benefit.

-The Pandemic Emergency Unemployment Compensation (PEUC). This program provides up to 49 extra weeks of UI benefits for those who have exhausted their regular state benefits. Since most states have a normal UI duration of 26 weeks, this leads to a potential benefit duration of up to 75 weeks. This extension is only available to workers who are actively seeking work. There was no phase-out period for this program, meaning that no PEUC benefits were payable after weeks of unemployment that began after the termination date.

-The Pandemic Unemployment Assistance (PUA). For workers who are not eligible for regular UI, this program provides UI up to 75 weeks. These workers have to be unemployed for a COVID-19-related reason and not be able to work from home. This program is targeted at the self-employed, gig workers, and independent contractors. The benefit amount for the PUA was roughly the same as the state's weekly benefit amount. However, because of the lack of UI covered wages, it was based on recent earned income. The PUA did not have a phase-out period either, meaning that no benefits were available in weeks of unemployment starting immediately after the termination date.

-Mixed Earner Unemployment Compensation (MEUC). This provides a $100 supplement for unemployed workers who were both self-employed and employee and are receiving one of the UI programs (but not PUA). This program was added when the programs were extended on March 3, 2021. It was implemented to address a potential difference between the benefits of regular state UI benefits and benefits of the PUA. To qualify for this program, workers needed to have received at least $5,000 in self-employment in the most recent tax year and received a UI benefit other than PUA.

III.B. Termination of the programs

All of the federal COVID-19 UI programs expired September 4, 2021. However, several states decided to terminate the programs early. The red states in figure 1 denote the 20 states that ended all the programs early. The blue states are the states that ended only the FPUC and MEUC, while the other programs remained in place until September 4, 2021. The termination dates range from June 12, 2021, to July 31, 2021. States were required to put in a 30-day notice to the Department of Labor to terminate the programs early. States gave several reasons for the early exit. These include work disincentive effects, decreased state unemployment rates, previous barriers to employment such as industry shutdowns coming to an end, and an increased number of job openings. In two states, the state court issued orders that did not allow the state to terminate the programs early. These states are shown in figure 1 by a light red color. In Indiana, the termination was planned for June 19, 2021, but because of the court ruling, PUA, PEUC, and FPUC remained in place and only the MEUC was terminated July 19, 2021.

According to the court, Indiana Law required the state to accept the benefits. In Maryland all benefit programs continued until September 4, 2021, because of the court prohibiting the planned early termination. There were several other states where the early termination was challenged in court, but those were not successful (Congressional Research Service, 2021).

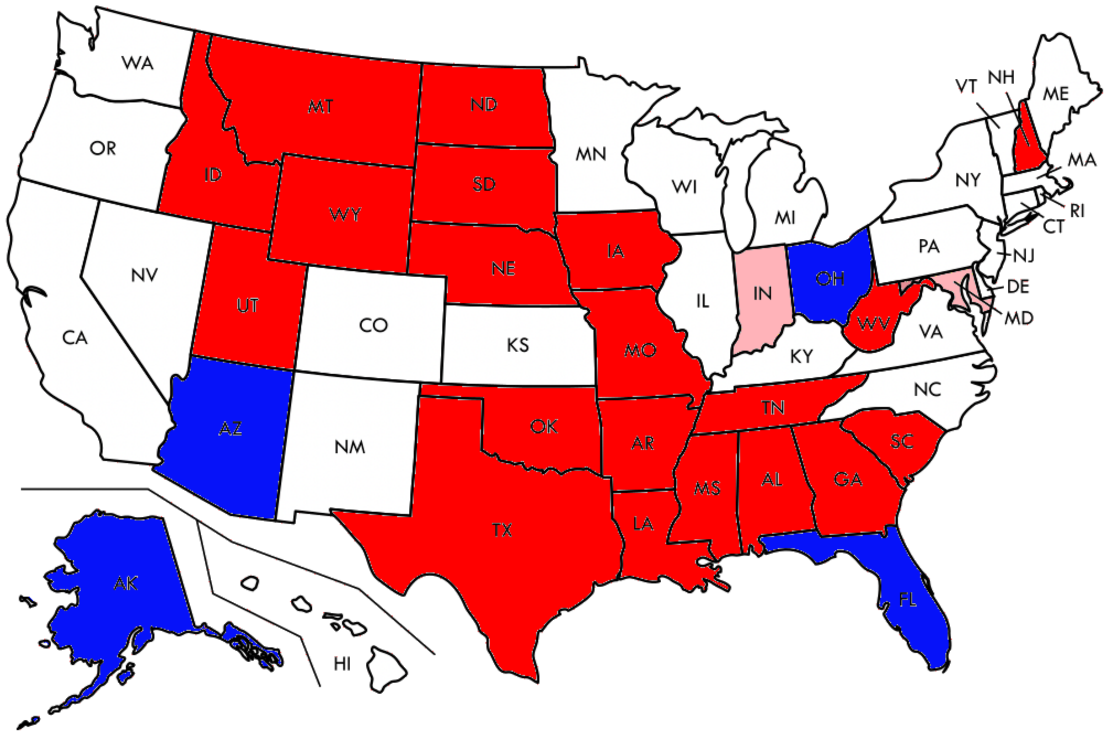

Figure 1: States that terminated the COVID-19 UI programs early

Note: States in red ended all the programs early. States in blue only ended the FPUC/MEUC program early. States in light red wanted to end the programs early but the state court did not allow it.

IV. Data & Methodology

IV.A. Data

I use the Current Population Survey (CPS). This is a monthly survey in the United States done by the U.S. Census Bureau and the U.S. Bureau of Labor Statistics. It is the primary source for labor force statistics for the population of the United States. It has information on the labor force status and other employment outcomes for individuals. About 59,000 households are selected every month for the CPS with an average respond rate of 75%. Households are in the survey for 4 months, then they are not contacted for 8 months, and then they are in the survey again for 4 months after which they leave the sample. This means that for any monthly survey, one-eight is in the sample for the first time. The sample consists of independent samples in each state. I use data from the CPS from January 2021 until August 2021. The monthly survey takes place in the calendar week that includes the 19th of the month, and the questions will be about the calendar week including the 12th of the month. Given that in June 2021, the 12th of the month was a Saturday, the survey questions refer to the week starting on Sunday June 6, 2021, to Saturday June 12, 2021. Given that the earliest termination date was on June 12, 2021, no states were treated yet in the June survey. This means that all of the early exit states terminated the program between the survey of June and the survey of July except Louisiana, which ended the program on July 31, 2021. The sample consists of 624,002 observations for 240,538 individuals. Individuals are observed for an average of 2.6 months in the sample period. This excludes individuals aged 16 or younger and adults that are part of the armed forces because these individuals are not asked about their labor force status in the survey.

The main variable of interest is the unemployment rate in the state. Given that the sample of the CPS consists of independent samples in the different states, the individuals in the sample can be used to investigate the state unemployment rate. The unemployment rate is defined as the number of unemployment people divided by the number of people in the labor force. For the individual, I use a dummy variable that indicates whether an individual is unemployed or not. Being unemployed means not having worked in the reference week but being available to work and having made effort to find a job in the last four weeks. It also includes workers on layoff and waiting to be recalled to the job. These people do not need to look for a job to be considered unemployed.

The treatment group consists of individuals living in a state that ended all the COVID-19 UI programs early. Individuals living in states that only terminated FPUC/MEUC early, are excluded from the sample. These states will be considered in a separate analysis. Individuals living in Louisiana will also be excluded from the sample given that this is the only state that did not terminate the program before the July survey took place. The treatment group is therefore treated in the months July and August. Descriptive statistics of the sample can be found in Table 1. Column 3 in Table 1 shows that the states that ended the programs early and those that did not, are very different. This could cause problems for the identification of causal effects. Fortunately, these are all individual characteristics that can be controlled for. However, it is likely that the states will also differ in unobservable characteristics. This issue will be further discussed in the next section.

Table 1: Summary Statistics for Treatment and Control States

	(1) Early Exit States	(2) Control States	(3) Difference between (1) and (2)
Age	47.94	48.31	0.367*** (0.050)
Male	0.482	0.480	-0.002 (0.001)
Education			
No High School	0.136	0.125	-0.011*** (0.001)
High School Grade or GED	0.470	0.417	-0.053*** (0.001)
Associate Degree – Occupational	0.049	0.039	-0.010*** (0.001)
Associate Degree – Academic	0.056	0.054	-0.002*** (0.001)
Bachelor’s Degree	0.190	0.221	0.031*** (0.001)
Master’s Degree	0.074	0.105	0.032*** (0.001)
Professional School Degree	0.010	0.017	0.007*** (0.000)
Doctorate Degree	0.015	0.022	0.006*** (0.000)
Race			
White	0.839	0.785	-0.054*** (0.001)
Black	0.107	0.095	-0.012*** (0.001)
Native American	0.013	0.011	-0.002*** (0.000)
Asian	0.025	0.083	0.057*** (0.001)
Other	0.016	0.026	0.010*** (0.000)

Column (1) and (2) show the mean of all variables for the early exit states and the control states respectively. Column (3) denotes the t-test of the difference between column (1) and column (2). Standard errors are denoted in parentheses. There are 624,002 observations for 240,538 individuals. Observations are unbalanced panel data from January 2021 to August 2021. The variable unemployment duration is only available for the unemployed observations. Hours worked and hours worked varies are conditional on having a job.

Table 1: Summary Statistics for Treatment and Control States (continued)

	(1) Early Exit States	(2) Control States	(3) Difference between (1) and (2)
Marital Status			
Married	0.524	0.506	-0.019*** (0.001)
Widowed	0.067	0.060	-0.007*** (0.001)
Divorced	0.112	0.100	-0.013*** (0.001)
Separated	0.017	0.015	-0.001*** (0.000)
Never married	0.280	0.320	0.040*** (0.001)
Family Income			
Less than $10,000	0.041	0.034	-0.007*** (0.000)
Between $10,000 and $30,000	0.171	0.140	-0.031*** (0.001)
Between $30,000 and $50,000	0.194	0.163	-0.030*** (0.001)
Between $50,000 and $100,000	0.332	0.313	-0.019*** (0.001)
Above $100,000	0.262	0.349	0.087*** (0.001)
Labor Force Status			
Employed	0.568	0.561	-0.006*** (0.001)
Unemployed	0.027	0.037	0.010*** (0.000)
Not in Labor Force	0.406	0.401	-0.004*** (0.001)
Unemployment duration	23.28	28.84	5.55*** (0.394)
Hours Worked	39.69	39.27	-0.424*** (0.040)
Hours Worked Varies	0.043	0.039	-0.004*** (0.001)
Number of observations	251,573	372,429	

Column (1) and (2) show the mean of all variables for the early exit states and the control states, respectively. Column (3) denotes the t-test of the difference between column (1) and column (2). Standard errors are denoted in parentheses. There are 624,002 observations for 240,538 individuals. Observations are unbalanced panel data from January 2021 to August 2021. The variable unemployment duration is only available for the unemployed observations and hours worked and hours worked varies are conditional on having a job.

IV.B. Methodology

To exploit the variation in benefit generosity across states due to the different termination dates, I use a difference-in-difference design. For the baseline specification I estimate the following Ordinary Least Squares (OLS) regression:

$$Y_{ism} = \alpha_1 + \alpha_2 * EarlyExit_s + \alpha_3 * Post_m + \alpha_4 * EarlyExit_s * Post_m + \gamma_m + \alpha_5 * X_{ism} + \varepsilon_{ism} \quad (1)$$

Where Y_{ist} is the outcome variable of individual i, living in state s, in month m, which will be unemployment. $EarlyExit_s$ is a dummy that equals one if the individual lives in one of the 19 states that ended all the programs between June 12, 2021, and July 12, 2021. α_2 measures the baseline average difference between the treatment and control states. $Post_m$ is a dummy variable that will be equal to 1 in the months July and August. α_3 measures the difference in unemployment rates of the groups compared to the pre-treatment periods. The coefficient α_4 is the coefficient of interest and measures the difference between the control and treatment states compared to their respective pre-treatment values. γ_m are time dummies and X_{ist} is a vector controlling for individual characteristics, which includes gender, age, education, race, and marital status. When using state fixed effects, the regression equation becomes:

$$Y_{ism} = \alpha_1 + \eta_s + \alpha_3 * Post_m + \alpha_4 * EarlyExit_s * Post_m + \gamma_m + \alpha_5 * X_{ism} + \varepsilon_{ism} \quad (2)$$

Where the dummy for early exit states is replaced by η_s, which captures the state fixed effects to control for baseline differences between the states. The coefficient of interest is still α_4.

All individuals will be weighted using the composited final weight that is available in the CPS data. Standard errors are clustered at the state level to account for the fact that individuals reoccur in the monthly sample. Generally, this will result in larger standard errors compared to clustering at the individual or household level or using non-clustered robust standard errors. Other outcome variables that will be used include employment-to-population ratio, unemployment duration, and hours worked.

The regression estimates can be interpreted as causal effects if the parallel trend assumption holds. The states in the treatment and control group are very different in terms of individual characteristics as shown by Table 1, and geographically as shown by Figure 1. Although the difference-in-difference design allows the two groups to differ initially, it could be that the states would have had different developments in the months of treatment even if the programs were still present in both states. Because of the baseline differences between the states, it is important to investigate whether the parallel trend assumption is likely to hold. To test the parallel trend assumption, I shift the intervention one period back and see if the trends were already differing at that moment. I also control for the macroeconomic environment by using linear trends of the unemployment rate. The equation used in the linear time trend regression is shown below. In this specification, the linear time trends are allowed to be different for the treatment and the control states.

$$Y_{ism} = \alpha_1 + \eta_s + \alpha_3 * Post_m + \alpha_4 * EarlyExit_s * Post_m + (\alpha_6 + \alpha_7 * EarlyExit_s) * Month_m + \alpha_5 * X_{ism} + \varepsilon_{ism} \quad (3)$$

Where $Month_m$ is a time variable that measures the linear trend. Coefficient α_6 captures the linear time trend for the control states while α_7 captures the linear time trend for the treatment states. Another specification of the model includes state-specific time trends. The regression equation is shown below.

$$Y_{ism} = \alpha_1 + \eta_s + \alpha_3 * Post_m + \alpha_4 * EarlyExit_s * Post_m + \alpha_8 * \eta_s * Month_m + \alpha_5 * X_{ism} + \varepsilon_{ism} \quad (4)$$

Where α_8 is the coefficient for state-specific time trends. This allows every state to have a different linear trend. The coefficient α_4 will measure the difference in the deviation from the state-specific trend for the treatment and control states.

An issue of the baseline difference-in-difference design might be the endogeneity of early termination. The states provided several reasons as to why they wanted to terminate the programs early. Therefore, it is possible that the early termination could be correlated with the outcome variable. For example, when COVID-restrictions are lifted, and job openings are increasing in states with early termination and not in the other states, it would not be the termination of the program that causes the differences between the treatment and control states, but the lifted COVID-restriction. To limit the endogeneity issue, I also run the baseline regression using only Indiana and Maryland as control states. These states planned on terminating the programs early. However, the court in these states decided that these programs should remain in place. Because the court ruling was based on the state's law, it seems unlikely this ruling in Indiana and Maryland is correlated with unemployment outcomes in these states. This makes the difference in the presence of the UI programs between these two states and the other states, where the program ended early, more likely to be an exogeneous variation. Besides using these two states as the specific control group, I also use a control group constructed in such a way that the treatment and control group have a very similar pre-treatment trend, which makes it more likely that the parallel trend assumption holds.

V. Results

V.A. Baseline specification

Figure 2 shows the fraction of unemployed people in the labor force over time for the states that terminated the COVID-19 UI programs early and those that did not. The trend seems to be quite similar for the two groups and after June, the unemployment rate seems to decline a little more in the states that terminated the COVID-19 UI programs early than in the control states.

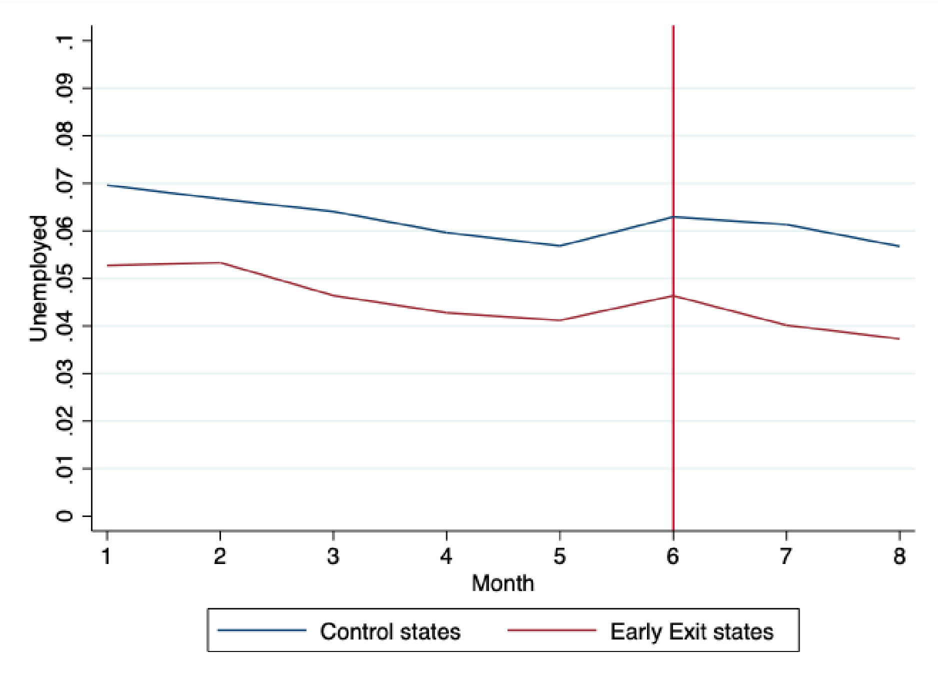

Figure 2: Average unemployment rates in the treatment and control states

Note: The unemployment variable is a dummy indicating whether an individual is unemployed or not. The time period is January 2021 to August 2021, where treatment takes place at the vertical red line. The sample only includes adults that are aged 16 or higher and part the labor force. The early exit states consist of those states that ended the COVID-19 UI programs early. The control states consist of states where all the programs remained in place. The number of observations is 371,525.

However, in the figure it is not possible to control for variables and the figure does not provide any information on the statistical significance. Table 2 shows the estimation of regression equation 1 and 2 given in the Section IV.B. Columns 2 and 4 include state fixed effects and columns 3 and 4 include the individual controls. It can be seen that the coefficient for treatment is negative and statistically significant at the 5%-level. The coefficient is robust to the inclusion of state fixed effects and individual controls. The coefficient of -0.007 indicates that the unemployment rate decreased by 0.7 percentage points more in the Early Exit states than in the control states. My estimate is very similar to the results found by Holzer et al. (2021), who find that the unemployment rate in the control states would have been 0.7 percentage points lower in August if they terminated the programs in June. This result is larger than the result found by other papers investigating the macro-effect of increased benefit generosity, for example, Chodorow-Reich et al. (2019) found that the increase in benefits during the Great Recession increased the unemployment rate by at most 0.3 percentage points.

Table 2: Difference-in-difference estimates of the effect of early termination of COVID-19 UI programs on unemployment

Unemployed	(1)	(2)	(3)	(4)
Early Exit * Post	-0.007**	-0.007**	-0.007**	-0.007**
	(0.003)	(0.003)	(0.003)	(0.003)
State fixed effects	No	Yes	No	Yes
Individual Controls	No	No	Yes	Yes
Observations	371,525	371,525	371,525	371,525
Mean unemployment rate	0.055	0.055	0.055	0.055

OLS estimates for the coefficients and standard errors, which are denoted in parentheses. Standard errors are clustered at the state level. The outcome variable is unemployed, which is a dummy indicating whether a person is unemployed or not. The sample consists of those adults aged 16 or higher that are part of the labor force. Column (1) – (4) show different specifications regarding the state fixed effects and individual controls. Treatment is defined by a state terminating the COVID-19 UI programs early. The mean unemployment rate gives the weighted mean of the outcome variable of all observations in the sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

V.B. Linear Time Trends

To better control for the possibility that the parallel trend assumption might not hold, I extend the difference-in-difference regression with a linear time trend that differs for the treatment and control states. The results are given in Table 3. The first column includes only a general linear time trend, where the time dummies in equation 2 are replaced with a linear time trend. The result is similar to the main analysis. In column 2, the estimate of regression equation 3 is shown, which includes a time trend that differs for the treatment and the control states. The coefficient decreases by 0.001 and the standard error is somewhat larger. Finally, a state-specific linear time trend is added as described in regression equation 4. This also results in a decrease in the coefficient of 0.001 and a slightly larger standard error compared to the main specification. Although adding these time trends makes the coefficients statistically insignificant, the inclusion of the time trend variables might have caused collinearity, which results in less precision and larger confidence intervals. The results shows that the treatment effect is not significantly changed when including the time trends, which provides support for the parallel trend assumption.

Table 3: Difference-in-difference estimates of the effect of early termination of COVID-19 UI programs on unemployment with linear time trends

Unemployed	(1)	(2)	(3)
	Linear Time Trend	Differing Trends for Treatment & Control	State-specific Time Trends
Early Exit * Post	-0.007**	-0.006	-0.006
	(0.003)	(0.004)	(0.004)
State fixed effects	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes
Observations	371,525	371,525	371,525
Mean unemployment rate	0.055	0.055	0.055

OLS estimates for the coefficients and standard errors, which are denoted in parentheses. Standard errors are clustered at the state level. The outcome variable is unemployed, which is a dummy indicating whether a person is unemployed or not. The sample consists of those adults that are aged 16 or higher and part of the labor force. Column (1) – (3) show different specifications with time trends. Treatment is defined by a state terminating the COVID-19 UI programs early. The mean unemployment rate gives the weighted mean of the outcome variable of all observations in the sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

V.C. Indiana and Maryland

Because of the endogeneity issue that might be present in the previous design, I run the same regression as in the baseline specification but now the control group only consists of Indiana and Maryland, because in these states the programs were planned to be terminated early, but the court did not allow it. This variation in the presence of the program is likely to be exogeneous. This could give more insight into the identifying assumption of the baseline regression. If there is no significant difference between these two states and the treatment states, the significant difference found in the baseline case is likely to be caused by endogeneity. A difference between the two control states and the treatment states would provide support for the parallel trend assumption. Panel A in Table 4 shows the results of the regression. The treatment coefficients are very similar to the ones in the baseline regression. Only in column 3 and 4 the coefficient differs 0.001. The coefficients are statistically significant at the 5%-level for column 1 and 2 and at the 10%-level when individual controls are added. It is important to denote the much smaller sample size compared to the baseline case. In this regression, only two states, of which there are 20,789 observations, make up the control states, which means there are only around 2,500 observations per month for the control group. This leads to more imprecise coefficients. However, the coefficients are very similar to the one in the baseline regression and show a difference in the unemployment rate between the treatment and control states. Exploiting this exogeneous variation of the presence of the UI programs which does not change the result, provides support for the parallel trend assumption.

Table 4: Difference-in-difference estimates of the effect of early termination of COVID-19 UI programs on unemployment using different groups of control states

Unemployed	(1)	(2)	(3)	(4)
Panel A: Indiana & Maryland				
Early Exit * Post	-0.007**	-0.007**	-0.006*	-0.006*
	(0.003)	(0.003)	(0.003)	(0.003)
State fixed effects	No	Yes	No	Yes
Individual Controls	No	No	Yes	Yes
Observations	161,909	161,909	161,909	161,909
Mean unemployment rate	0.050	0.050	0.050	0.050
Panel B: Similar pre-trend states				
Early Exit * Post	-0.009**	-0.009**	-0.009**	-0.009**
	(0.003)	(0.003)	(0.003)	(0.004)
State fixed effects	No	Yes	No	Yes
Individual Controls	No	No	Yes	Yes
Mean unemployment rate	0.049	0.049	0.049	0.049
Observations	203,901	203,901	203,901	203,901

OLS estimates for the coefficients, standard errors, which are denoted in parentheses. Standard errors are clustered at the state level. The outcome variable is unemployed, which is a dummy indicating whether a person is unemployed or not. The sample consists of those adults that are aged 16 or higher and part of the labor force. In Panel A, the control states are Indiana and Maryland. In Panel B, the control states are Colorado, Delaware, Indiana, Kansas, Maine, Minnesota, North Carolina, Vermont, and Wisconsin. The treatment states are the states that ended the COVID-19 UI programs early. Column (1) – (4) show different specifications regarding the state fixed effects and individual controls. The mean unemployment rate gives the weighted mean of the outcome variable of all observations in the sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

V.D. Designing the Group of Control States

In the main analysis the control states consisted of all the states where all the COVID-19 UI programs remained in place. However, as figure 2 shows, the level of the unemployment rate before treatment period differs between the treatment states and the control states, although the trend seems to be similar. To provide more support for the parallel trend assumption, I try to construct the control states in such a way that the pre-treatment unemployment rate for the control states is very similar to the treatment states. When only using Colorado, Delaware, Indiana, Kansas, Maine, Minnesota, North Carolina, Vermont, and Wisconsin as control states, the pre-treatment unemployment rates are almost identical, as shown in figure 3. After treatment takes place, the lines start to diverge with the unemployment rate in the treatment states decreasing more. Panel B of Table 4 shows the results of the regression run with these groups. The coefficients are negative and statistically significant at the 5%-level. The coefficients are slightly larger than in the baseline regression, but the standard errors are also slightly larger. The fact that the effect is very similar to the main specification when using the control group with similar pre-treatment unemployment levels, provides more support for the parallel trend assumption.

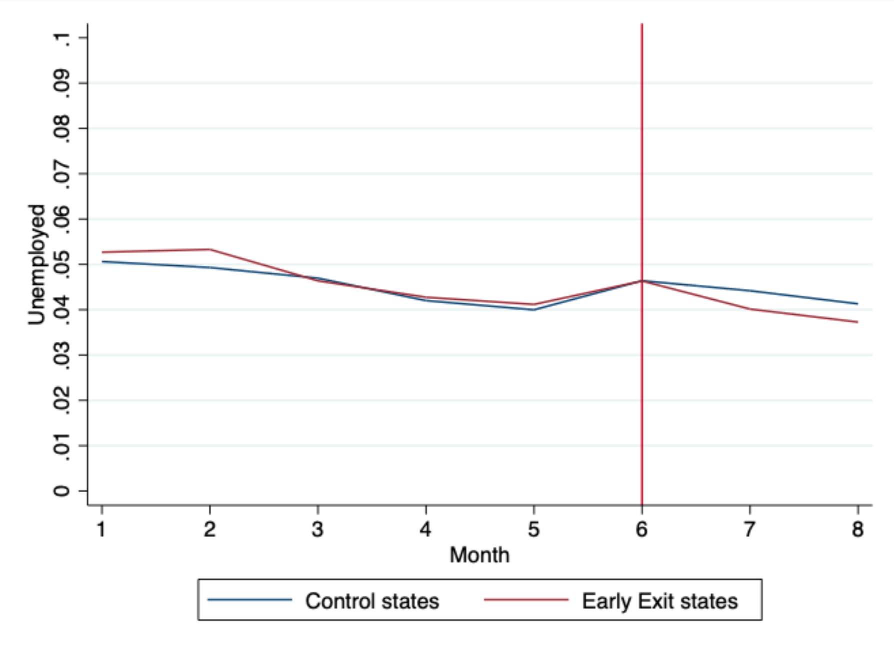

Figure 3: Average unemployment rates in the treatment states and a selection of control states
Note: The unemployment variable is a dummy indicating whether an individual is unemployed or not. The time period is January 2021 to August 2021, where treatment takes place at the vertical red line. The sample only includes adults aged 16 or higher and part of the labor force. The early exit states consist of those states that ended the COVID-19 UI programs early. The control states consist of Colorado, Delaware, Indiana, Kansas, Maine, Minnesota, North Carolina, Vermont, and Wisconsin, which are the states that were best suited to approximate the pre-treatment unemployment rate of the early exit states. The number of observations is 203,901.

VI. Robustness and validity checks

VI.A. Leads of treatment

One way of testing the parallel trend assumption is by including leads of the treatment in the regression. This tests if the trends were already differing before the treatment actually took place. Table 5 shows the baseline regression but replaces the treatment with different leads of the treatment. The results indicate that before treatment took place, the parallel trend assumption seems to hold. The coefficients of leads of treatment are small and not statistically significant. Column 4 of Table 5 shows the result of a regression where the interaction between early exit and September is added. This regression includes data from the September 2021 monthly sample. Because the control states ended the COVID-19 UI programs in September, both groups did not have the UI programs in September. If the parallel trends assumption were to hold, this would mean that there is no longer a significant difference between the two groups in September. The results show an estimation close to zero. This is strong evidence that the parallel trend assumption holds. It is unlikely that the trends start differing after June and stop differing in September, because of something else than the early termination of the programs.

Table 5: Difference-in-difference estimates of the lead of terminating the COVID-19 UI programs early on unemployment

Unemployed	(1)	(2)	(3)	(4)
1^{st} Lead of treatment	-0.000 (0.002)			
2^{nd} lead of treatment		0.003 (0.004)		
3^{rd} lead of treatment			-0.001 (0.003)	
Early Exit * September				0.000 (0.002)
State fixed effects	Yes	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes	Yes
Observations	371,525	371,525	371,525	416,015

OLS estimates for the coefficients and standard errors, which are denoted in parentheses. Standard errors are clustered at the state level. The outcome variable is unemployed, which is a dummy indicating whether a person is unemployed or not. The sample consists of those adults that are aged 16 and higher and part of the labor force. Column (1) – (3) show different inclusions of the lead of treatment. Column (4) shows the coefficient for the interaction between Early Exit and the month September, when all states have ended the programs. Treatment is defined by a state terminating the COVID-19 UI programs early. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

VI.B. Effect on COVID-vulnerable industries

One concern with the baseline difference-in-difference design is the reason why the states decided to terminate the UI programs early. One of the reasons given by the states was that previous barriers to employment such as industry shutdowns were coming to an end. This could mean that people are going back to their jobs in the treatment states once industries open up and as a result, the unemployment rate goes down. If in the control states these businesses were not fully opened up yet and people are therefore not returning to work, the difference in unemployment rates is not caused by the UI programs. One way to provide some more insight into this problem is by splitting the sample in jobs in industries that are likely to be suffering from shutdowns, COVID-restrictions, and other COVID-related issues, and jobs that are not.

To split jobs on their likelihood of being restricted by COVID-19, I use the Remote Labor Index created by Del Rio-Chanona, Mealy, Pichler, Lafond and Farmer (2020). This index measures how much work activities can be done from home. From the twenty industry classifications, I define the ten industries with the highest Remote Labor Index as jobs that are not vulnerable to COVID-19 and the ten industries with the lowest Remote Labor Index as the jobs that are vulnerable to COVID-19. Table A.1 in the Appendix provides a list of vulnerable and non-vulnerable industries. Table 6 shows the result of the baseline difference-in-difference for the two subsamples. It should be noted that the industry variable for the unemployed individuals is only available when the unemployed individual has done paid work before. If the individual is unemployed and has never done paid work before, this variable is not available. However, this is only the case for a small part of the sample (around 6% of unemployed observations).

For the non-vulnerable subsample, the coefficient is 0.001 larger than the baseline specification and is statistically significant at the 10%-level. The coefficient for the vulnerable subsample is smaller and not statistically significant. If it was the case that the lifted COVID-restrictions in the treatment states were causing the unemployment rate to drop, I would have expected the unemployment rate to mostly drop for the subsample that is working in COVID-19-vulnerable jobs. Instead, the lower unemployment rate in the treatment states is more concentrated in the industries that are not vulnerable to COVID-related issues. This result is robust to using different cut-off points in the Remote Labor Index for determining whether an industry is vulnerable or not. This takes away some of the concern that the lifting of COVID-restrictions in the treatment states is what caused the lower unemployment rate.

Table 6: Difference-in-difference estimates of the effect of early termination of COVID-19 UI programs on unemployment for COVID-vulnerable and non-vulnerable industries

Unemployed	(1)	(2)
	Non-vulnerable jobs	COVID-vulnerable jobs
Early Exit * Post	-0.008*	-0.005
	(0.004)	(0.003)
State fixed effects	Yes	Yes
Individual Controls	Yes	Yes
Observations	178,523	170,878
Mean unemployment rate	0.054	0.064

OLS estimates for the coefficients and standard errors, which are denoted in parentheses. Standard errors are clustered at the state level. The outcome variable is unemployed, which is a dummy indicating whether a person is unemployed or not. The sample consists of those adults that are aged 16 or higher and are either employed or unemployed but have worked in the past. Column (1) only includes individuals that are working in industries that are unlikely to be restricted by COVID-19 as measured by the Remote Labor Index. Column (2) only includes individuals that are working in industries with a higher probability of being restricted by COVID-19 as measured by the Remote Labor Index. Treatment is defined by a state terminating the COVID-19 UI programs early. The mean unemployment rate gives the weighted mean of the outcome variable of all observations in the sample. * p < 0.10, ** p < 0.05, *** p < 0.01.

VI.C. Effect on Low-income Families

Another way to divide the sample is by income. This can be done by using the family income variable, which is the combined income of all family members during the last 12 months. This also includes money from social security payments. Exact family income is not given, but brackets are used. I use this variable to see how the unemployment effect is different for low-income households compared to high-income households. The sample is split by using different cut-off points that determine whether a family is low-income or high-income. The first two columns in Table 7 use $74,999 as a cut-off point since this is just below the median family income in the US in 2021. The analysis shows the unemployment effect is larger for individuals living in low-income households compared to the full sample. For the above-median households, the effect is much smaller than for the full sample and not statistically significant. In column 3 and 4, $24,999 is used as the cut-off point because this is close to the poverty threshold for an average family. The unemployment effect is now even larger in the low-income sample. The high-income sample is still small and statistically insignificant. Thus, the difference in unemployment rates in the treatment and in the control states is mostly concentrated in low-income households. Given that low-income households are more vulnerable and dependent on unemployment benefits, this provides more evidence for the termination of the UI programs being the causal effect of the different unemployment rates in the control and treatment states.

For the low-income households, the UI programs are an important part of income when an individual become unemployed while high-income households can probably use the high earnings of the other household members or their built-up savings to pay their expenses. This means that for the low-income households the termination of UI benefits leads to a larger change in the reservation wage, which could result in a higher search effort and a larger acceptance for jobs. This would lead to a decrease in the unemployment rate. Although this provides some more evidence for the termination of the programs to be the cause, there are two points to be noted here. Firstly, it is specified that this family income variable should be used with caution because it is not as accurate as the other variables in the survey. However, I only use this variable to distinguish between low-income and high-income households and therefore a small measurement error is unlikely to change the results. Secondly, there still could be other reasons that the effect is larger for individuals in low-income households, which cannot be attributed to the UI programs. It is possible that macroeconomic conditions were already improving more in the treatment states than in the control states outside of the UI programs. An increase in economic activity could lead to more job creation and therefore more people becoming employed. It is very likely that individuals in low-income families have jobs that are more dependent on economic activity. This additional analysis cannot exclude that possibility.

Table 7: Difference-in-difference estimates of the effect of early termination of COVID-19 UI programs on unemployment for high-income and low-income families

Unemployed	(1)	(2)	(3)	(4)
	Cut-off at $74,999	Cut-off at $74,999	Cut-off at $24,999	Cut-off at $24,999
	Low income	High income	Low income	High income
Early Exit * Post	-0.010**	-0.002	-0.030**	-0.003
	(0.005)	(0.002)	(0.011)	(0.003)
State fixed effects	Yes	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes	Yes
Observations	169,937	201,588	32,335	339,190
Mean unemployment rate	0.092	0.034	0.178	0.049

OLS estimates for the coefficients and standard errors which are denoted in parentheses. Standard errors are clustered at the state level. The outcome variable is unemployed, which is a dummy indicating whether a person is unemployed or not. The sample consists of those adults that are aged 16 or higher and part of the labor force. Column (1) & (2) split the sample in low-income and high-income households using $74,999 as the cut-off. Column (3) & (4) use $24,999 as the cut-off. Column (1) & (3) include the low-income families and column (2) & (4) include the high-income families. Treatment is defined by a state terminating the COVID-19 UI programs early. The mean unemployment rate gives the weighted mean of the outcome variable of all observations in the sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

VI.D Different Types of Unemployment

The labor force status distinguishes between two types of being unemployed. The first one is being unemployed and looking for a job and the second type is waiting to be recalled to a job from which a worker has been laid-off. Instead of taking them together as being unemployed, both can be examined separately. This is shown in Table 8. The coefficient for the first type of unemployment is larger than the baseline specification and statistically significant, while the coefficient for the second type is positive and statistically significant. This means that the number of people that were unemployed on layoff decreased less in the treatment states than in the control states. This could provide more support for the parallel trend assumption. In case the macroeconomic conditions were starting to improve in treatment states relative to control states during the treatment period and that caused a decrease in the unemployment rate for the treatment states, these macroeconomic improvements would influence the demand for labor rather than the supply. If the demand for labor would increase, I would expect the unemployment rate for individuals on layoff to decrease more in the treatment states. The concentration of the decrease of the unemployment rate seems to be caused by the supply side, which provides support for the effect being caused by the change in unemployment benefit generosity.

Table 8: Difference-in-difference estimates of the effect of early termination of COVID-19 UI programs on different types of unemployment

Unemployed	(1) Looking for job	(2) Layoff
Early Exit * Post	-0.012***	0.005***
	(0.003)	(0.001)
State fixed effects	Yes	Yes
Individual Controls	Yes	Yes
Observations	371,525	371,525
Mean unemployment rate	0.047	0.013

OLS estimates for the coefficients and standard errors which are denoted in parentheses. Standard errors are clustered at the state level. The outcome variable in column (1) is unemployed and looking for a job, which is a dummy indicating whether a person is unemployed and searching for a job. In column (2) the outcome variable is unemployed on layoff, which is a dummy indicating whether a person is unemployed while waiting to be recalled to a job from which they had been laid off. The sample consists of those adults that are aged 16 or higher and part of the labor force. Treatment is defined by a state terminating the COVID-19 UI programs early. The mean unemployment rate gives the weighted mean of the outcome variable of all observations in the sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

VI.E. Standard errors

The CPS data that I use in this research has a longitudinal character. The CPS is not panel data, but many individuals appear in the sample for multiple months. This is part of the random sampling and should not be an issue for the calculation of the unemployment rate of a month, which is confirmed by the fact that this sample is used as the primary source in the labor statistics of the U.S. Bureau of Labor Statistics. However, the fact that individuals reoccur in the monthly sample can cause problems for the calculation of standard errors. Throughout the paper, standard errors are clustered at the state level. This generally gives larger standard errors than clustering at the household or individual level or not clustering at all. Table A.2 shows the main specification with standard errors clustered at the individual level and standard errors that are robust but not clustered. The coefficients are similar and the standard errors are slightly smaller than the standard errors clustered at the state level.

VII. Additional analyses

VII.A. The effect of the FPUC/MEUC on unemployment

In the previous section, I looked at the effect of early termination of all the COVID-19 UI programs. It is not possible to disentangle the unemployment effects of the separate programs in that design. However, in four states only the FPUC and MEUC were terminated early. This means I can estimate the effect of the FPUC/MEUC program by using the same difference-in-difference design but only use Alabama, Arizona, Florida, and Ohio as treatment states, while using the same control states as in the baseline specification. This should measure the effect of a $300 increase in weekly UI benefits and an additional $100 for mixed earners. The results of Table 9 show that the coefficients are much smaller than the coefficients in the main analysis, 0.2 percentage points instead of 0.7 percentage points. Also, these coefficients are not statistically significant. This suggests that the additional UI benefit did not increase the unemployment rate. This is conflicting with evidence from Card et al. (2015) who found that an increase in the weekly benefit amount should increase UI duration and therefore increase the unemployment rate. It is more in line with the evidence from Marinescu et al. (2021) who found that FPUC did not increase unemployment because it increased labor market tightness. In other words, companies did not decrease vacancies but because less people search (intensively) for jobs, there are positive externalities for those who are (intensively) searching, which decreases the total unemployment effect. It is also consistent with the finding by Dube (2021), who did not find a significant impact on employment when the $600 FPUC program expired in the summer of 2020.

Table 9: Difference-in-difference estimates of the effect of early termination of the FPUC/MEUC programs on unemployment

Unemployed	(1)	(2)	(3)	(4)
Early Exit * Post	-0.005	-0.005	-0.002	-0.002
	(0.003)	(0.003)	(0.003)	(0.003)
State fixed effects	No	Yes	No	Yes
Individual Controls	No	No	Yes	Yes
Observations	262,722	262,722	262,722	262,722
Mean unemployment rate	0.063	0.063	0.063	0.063

OLS estimates for the coefficients and standard errors which are denoted in parentheses. Standard errors are clustered at the state level. The outcome variable is unemployed, which is a dummy indicating whether a person is unemployed or not. The sample consists of those adults that are aged 16 or higher and part of the labor force. Column (1) – (4) show different specifications regarding the state fixed effects and individual controls. Treatment is defined by a state terminating the FPUC/MEUC program early. The mean unemployment rate gives the weighted mean of the outcome variable of all observations in the sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

VII.B. Changes in the labor force

The main analysis shows a larger decrease in the unemployment rate for the states that terminated the programs early. The unemployment rate could have decreased because of people finding employment, but it could also be because labor participation decreased as a result of the lower benefit generosity (Farber & Valetta, 2015). To see whether more people actually found a job or whether they stopped being unemployed because they left the labor force, I run the same regression as in the baseline specification, but the outcome variables is a dummy for being employed or not and the sample consists of all adults in the survey that are not in the armed forces. This estimates the employment-to-population rate. The results of this regression can be found in Table 10. The treatment coefficient is small and not statistically significant. When the state fixed effects and individual controls are added, the coefficient is slightly higher but still small and insignificant. There seems to be no significant difference between the treatment and control states in the employment-to-population rate. This analysis shows that the lower unemployment rate in the treatment states as found in Table 2 cannot be explained by more people finding a job assuming the population remained constant. This implies that more people moved out of the labor force in the treatment states.

Although from the analysis in Table 10 it is unclear where the decrease in the labor force is coming from, it could be from both more unemployed people moving out of the labor force of because of more employed workers moving out of the labor force, it seems likely that the unemployed moved out of the labor force. This is evidence for the presence of the entitlement effect. When generosity is high, it rewards to be unemployed and search for a job. When generosity decreases, it is no longer rewarding for these people to search for a job, and then they stop searching for a job and are no longer considered unemployed, which leads them to leave the labor force. This is inconsistent with the findings of Holzer et al. (2021), who calculate that the employment-to-population rate increased in the treatment states compared to the control states. They estimated this using the employment flows while assuming the level of the labor force and population stays constant.

Table 10: Difference-in-difference estimates of the effect of early termination of COVID-19 UI programs on the employment-to-population ratio

Employed	(1)	(2)	(3)	(4)
Early Exit * Post	-0.001	-0.001	-0.002	-0.002
	(0.004)	(0.004)	(0.003)	(0.003)
State fixed effects	No	Yes	No	Yes
Individual Controls	No	No	Yes	Yes
Observations	614,527	614,527	614,527	614,527
Mean employment-to-population ratio	0.582	0.582	0.582	0.582

OLS estimates for the coefficients and standard errors which are denoted in parentheses. Standard errors are clustered at the state level. The outcome variable is employed, which is a dummy indicating whether a person is employed or not. The sample consists of all individuals aged at least 16 and not part of the armed forces. Column (1) – (4) show different specifications regarding the state fixed effects and individual controls. Treatment is defined by a state terminating the COVID-19 UI programs early. The mean employment-to-population ratio gives the weighted mean of the outcome variable of all observations in the sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

VII.C Effect on Unemployment Duration

Another aspect of unemployment to look at, is the unemployment duration. Theoretically, an increase in the UI generosity should increase the unemployment duration because of the job search effect (Bailey, 1978; Chetty, 2006; Landais et al., 2015). Much research has been done on the UI effect on the unemployment duration (Johnston & Mas, 2018; Card et al., 2015; Farber & Valetta, 2015). In most of these studies, individuals are followed for their whole unemployment spell. The data I use does not allow for this and as a result, I can only look at the average unemployment duration at a certain time in the sample. Panel A of Table 11 compares the average unemployment duration between the treatment and control states using the difference-in-difference design. The sample only includes individuals that are unemployed. The results show that the average unemployment duration is about two weeks shorter in the states where the programs were ended early. However, the coefficients are not statistically significant. This effect is much smaller than that found by Johnston and Mas (2018), who find that a one-month-reduction of UI leads to an unemployment spell decrease of 0.45 months. However, due to the data design, it is very hard to compare these results to other literature and to actually assess the effect on the total length of the unemployment duration.

Table 11: Difference-in-difference estimates of the effect of early termination of COVID-19 UI programs on unemployment duration and hours worked

	(1)	(2)	(3)	(4)
Panel A: Unemployment duration				
Early Exit * Post	-2.133	-2.107	-2.058	-2.054
	(1.504)	(1.494)	(1.448)	(1.439)
State fixed effects	No	Yes	No	Yes
Individual Controls	No	No	Yes	Yes
Observations	20,452	20,452	20,452	20,452
Mean of unemployment duration	27.580	27.580	27.580	27.580
Panel B: Hours worked				
Early Exit * Post	0.168	0.163	0.091	0.084
	(0.131)	(0.130)	(0.120)	(0.120)
State fixed effects	No	Yes	No	Yes
Individual Controls	No	No	Yes	Yes
Observations	325,410	325,410	325,410	325,410
Mean of hours worked	39.465	39.465	39.465	39.465

OLS estimates for the coefficients and standard errors which are denoted in parentheses. Standard errors are clustered at the state level. The outcome variable in Panel A is unemployment duration, which measures how many weeks an unemployment person has been unemployed. The sample consists of those adults that are aged 16 or higher and unemployed. The outcome variable in Panel B is number of hours worked. The sample consists of employed individuals that are aged 16 or higher. Column (1) – (4) show different specifications regarding the state fixed effects and individual controls. Treatment is defined by a state terminating the COVID-19 UI programs early. The mean unemployment duration and the mean of hours worked give the weighted mean of the outcome variable of all observations in the sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

VII.D Effect on Number of Hours Worked

In Section V, I investigated unemployment by looking at whether an individual was unemployed or not and whether this unemployment rate changed differently for the treatment states. It is also possible that there were changes along the intensive margin, that is, that the people that had to find a job due to the decreasing UI benefit generosity, found jobs with a low number of hours. To investigate this, I run the same regression but change the outcome variable to number of hours worked. This variable is conditional on being employed. Panel B in Table 11 shows the result of this regression. It shows that the number of hours worked is not significantly different between the treatment and control states. Respondents could also answer that their working hours vary from week to week. To check whether varying hours may play a role, I ran the same regression but used a dummy for whether the number of hours vary per week as the outcome variable. In this regression there were no significant differences either.

VIII. Discussion

VIII.A. Identifying Assumption

The results have shown that the unemployment rate decreased more in the states that decided to terminate the COVID-19 UI programs early. I estimate causal effects under the parallel trend assumption. This means that in the absence of the early termination of the programs, the two groups of states would have developed in the same way. It is not possible to know with certainty whether this assumption holds or not. However, I provided several tests that support the parallel trend assumption. Firstly, adding linear time trends that allow the treatment states and control states to have different linear time trends still results in similar coefficients with slightly less statistical significance. Secondly, when using Indiana and Maryland, where early termination was intended but prohibited by the state's court, the coefficients are almost unchanged. Thirdly, using a group of control states with an almost identical pre-treatment unemployment rate, the estimated effect is very similar to the baseline case. Fourthly, leads of the treatment did not have significant coefficients, which means that the trend between the treatment and control states did not differ in the months before July. Furthermore, the trend between the two groups did not differ in September, when all states had eliminated the programs. Only in the two months when there was variation in the presence of the programs, was there a difference in the unemployment rate. It seems unlikely that something other than the UI programs would cause this difference only during these two months. Fifthly, the effect is concentrated at the lower income families, who are more likely to be dependent on the benefit generosity and will therefore search (more intensively) for a job once these benefits disappear. However, if macroeconomic conditions were improving more in the treatment states than in the control states, it is also possible that low-income families benefit more from that, since low-income jobs are more vulnerable to economic activity. However, the number of unemployed workers on layoff did not decrease, while those workers are expected to be recalled to work when economic activity increases, and labor demand rises. Finally, the effect is unlikely to be driven by the lifting of COVID-restrictions in the treatment states because the unemployment effect is concentrated in jobs in industries that are not vulnerable to COVID-19 according to the Remote Labor Index. All of this makes it likely that the parallel trend assumption holds, but it is not possible to state this with certainty.

VIII.B. Limitations

There are a few limitations to this study. First of all, the survey is designed in such a way that individuals are in the sample for multiple months but only for a maximum of four months. Consequently, the sample can be considered panel data but unbalanced. I do not use this longitudinal aspect of the survey since I only use the monthly sample to estimate the unemployment rate. It is therefore unclear what this means for the standard errors, even though the results are robust to different types of clustered standard errors. The design of the data also makes it difficult to look at employment flows, duration of unemployment, and reemployment earnings, because individuals cannot be followed for their entire unemployment spell. This information could be useful for assessing the micro effects of the COVID-19 programs and to assess which people leave the labor force after treatment takes place and this should be assessed by future research. Another limitation is the limited number of observations per state, which resulted in large standard errors in some of the additional analyses to provide support for the identifying assumption. It would have been useful to have more observations and therefore more precisely estimated coefficients in these analyses. Finally, the effects of the increase in benefit generosity as investigated in this paper, are very specific to the pandemic. COVID-19 adds complications to the labor market such as lockdowns, restrictions, illness, or concerns of infection. All of this could have played a role in the decision to search for a job. For this reason, it is difficult to extrapolate the results of this paper to another timing and setting.

VIII.C. Comparison with Previous Literature

The results found here are inconsistent with most literature on the effect of UI on unemployment spells. Many papers found a significant reduction in the unemployment spell when benefit generosity is decreased (Johnston & Mas, 2018; Card et al., 2015; Farber & Valetta, 2015; Farber et al, 2015; Katz & Meyer, 1990). Although I find a lower unemployment rate, this is not caused by more people finding a job. Johnston and Mas (2018) find that a one-month reduction in benefit duration leads to a 0.25-month reduction in nonemployment. This means that reducing the benefit generosity would lead to more employment. The results of Farber and Valetta (2015) and Farber et al. (2015) are more in line with my results. While they find that an increase in benefit duration leads to longer unemployment durations, they also find that there is an increase in labor force attachment but no impact on job finding. This suggests that increased benefit generosity leads to more people in the labor force but no decrease in employment. My results are also consistent with the findings of Boone et al. (2021), who do not find a statistical effect of UI on the aggregate employment. Johnston and Mas (2018) also use their micro findings to calculate the decrease in the unemployment rate. They find that the

16-week cut in potential benefit duration decreased the unemployment rate by 0.8 percentage points, assuming no market-level externalities. This result is quite similar to the result found in this paper. This is larger than other macro-findings such as Chodorow et al. (2019) who found that the benefit extension in the Great Recession increased the unemployment rate by at most 0.3 percentage points. My results do therefore not necessarily point to market externalities as documented by Lalive et al. (2015) and Marinescu (2017).

Most of this literature is focused on the Great Recession. Compared to other papers examining the COVID-19 UI programs, the results are quite similar. I estimate almost the same decrease in unemployment rate as Holzer et al. (2021) although they find a significant increase in the employment-to-population rate. Similar to Finamor and Scott (2021), Marinescu et al. (2020), and Dube (2021), I do not find an effect of the PFUC program on employment. Due to the specific circumstances of the pandemic, it might be difficult to compare the results of research focused on the COVID-19 pandemic to research on benefit generosity during the Great Recession.

IX. Conclusion

The main result of this paper is that the unemployment rate decreased more in states where the COVID-19 UI programs were terminated compared to states where the programs were still in place. The unemployment rate decreased by 0.7 percentage point more in the treatment states than in the control states. This effect holds up when including linear time trends into the regression or using control states that are more likely to have similar trends to the treatment states. The early termination of the programs had no significant impact on the employment-to-population rate. Assuming the population remained constant, this must mean that the labor force decreased. It seems that the people that were considered unemployed and searching for a job, stop searching for a job when the unemployment benefits disappear. Consequently, they are no longer considered unemployed and leave the labor force. Even though the early termination of the program led to a lower unemployment rate, it could be argued that the termination did not reach its desired goal of increasing employment. The effect is a causal effect under the identifying assumption that the trends of the two groups would have developed in the same way if the UI programs were not terminated early. I provide evidence that this assumption is likely to hold. The findings of this paper are important for policymakers. One of the goals of terminating the programs early was to reduce the work disincentive effects. The analysis shows that the increased benefit generosity due to the COVID-19 UI programs did not disincentivize employment since terminating them did not lead to more employment. This can help policymakers in designing the optimal unemployment benefit generosity, keeping the special circumstances of the pandemic in mind.

Literature

Baily, M. N. (1978). Some aspects of optimal unemployment insurance. *Journal of public Economics*, *10*(3), 379-402.

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

Card, D., Johnston, A., Leung, P., Mas, A., & Pei, Z. (2015). The effect of unemployment benefits on the duration of unemployment insurance receipt: New evidence from a regression kink design in Missouri, 2003-2013. *American Economic Review*, *105*(5), 126-30.

Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics*, *90*(10-11), 1879-1901.

Chodorow-Reich, G., Coglianese, J., & Karabarbounis, L. (2019). The macro effects of unemployment benefit extensions: a measurement error approach. *The Quarterly Journal of Economics*, *134*(1), 227-279.

Congressional Research Service (2021) Unemployment Insurance (UI) benefits: Permanent-Law Programs and the COVID-19 Pandemic Response R46687. Last updated January 31, 2022.

Del Rio-Chanona, R. M., Mealy, P., Pichler, A., Lafond, F., & Farmer, J. D. (2020). Supply and demand shocks in the COVID-19 pandemic: An industry and occupation perspective. *Oxford Review of Economic Policy*, *36*(Supplement_1), S94-S137.

Dube, A. (2021). *Aggregate employment effects of unemployment benefits during deep downturns: Evidence from the expiration of the Federal Pandemic Unemployment Compensation* (No. w28470). National Bureau of Economic Research.

Farber, H. S., Rothstein, J., & Valletta, R. G. (2015). The effect of extended unemployment insurance benefits: Evidence from the 2012-2013 phase-out. *American Economic Review*, *105*(5), 171-76.

Farber, H. S., & Valletta, R. G. (2015). Do extended unemployment benefits lengthen unemployment spells? Evidence from recent cycles in the US labor market. *Journal of Human Resources*, *50*(4), 873-909.

Finamor, L., & Scott, D. (2021). Labor market trends and unemployment insurance generosity during the pandemic. *Economics Letters*, *199*, 109722.

Ganong, P., Noel, P., & Vavra, J. (2020). US unemployment insurance replacement rates during the pandemic. *Journal of Public Economics*, *191*, 104273.

Holzer, H. J., Hubbard, R. G., & Strain, M. R. (2021). *Did Pandemic Unemployment Benefits Reduce Employment? Evidence from Early State-Level Expirations in June 2021* (No. w29575). National Bureau of Economic Research.

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

Katz, L. F., & Meyer, B. D. (1990). The impact of the potential duration of unemployment benefits on the duration of unemployment. *Journal of public economics*, *41*(1), 45-72.

Lalive, R., Landais, C., & Zweimüller, J. (2015). Market externalities of large unemployment insurance extension programs. *The American Economic Review*, 3564-3596.

Landais, C., Michaillat, P., & Saez, E. (2018). A macroeconomic approach to optimal unemployment insurance: Theory. *American Economic Journal: Economic Policy*, *10*(2), 152-81.

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

Marinescu, I., Skandalis, D., & Zhao, D. (2021). The impact of the federal pandemic unemployment compensation on job search and vacancy creation. *Journal of Public Economics*, *200*, 104471.

Appendix

Table A.1: List of industries that are vulnerable and not vulnerable to COVID-19

Not vulnerable industries	Vulnerable industries
Finance and Insurance	Admin Support, Waste Management and Remediation
Information	Health Care and Social Assistance
Professional, Scientific, and Technical Services	Wholesale Trade
Utilities	Manufacturing
Educational Services	Mining, Quarrying, and Oil and Gas Extraction
Retail Trade	Management of Companies and Enterprises
Transportation and Warehousing	Construction
Real Estate and Rental and Leasing	Accommodation and Food Services
Other Services (except public administration)	Unclassified
Arts, Entertainment, and Recreation	Agriculture, Forestry, Fishing and Hunting

The industries are classified as vulnerable or not vulnerable to COVID-restrictions using the Remote Labor Index constructed by Del Rio-Chanona, Mealy, Pichler, Lafond and Farmer (2020). From the twenty industry classifications, the ten with the highest Remote Labor Index are classified as not vulnerable. The ten with the lowest Remote Labor Index are classified as vulnerable.

Table A.2: OLS regression with clustered and robust standard errors

Unemployed	(1)	(2)	(3)	(4)
Panel A: clustered at the individual level				
Early Exit * Post	-0.007***	-0.007***	-0.007***	-0.007***
	(0.002)	(0.002)	(0.002)	(0.002)
State fixed effects	No	Yes	No	Yes
Individual Controls	No	No	Yes	Yes
Observations	371,525	371,525	371,525	371,525
Panel B: robust standard errors				
Early Exit * Post	-0.007***	-0.007***	-0.007***	-0.007***
	(0.002)	(0.002)	(0.002)	(0.002)
State fixed effects	No	Yes	No	Yes
Individual Controls	No	No	Yes	Yes
Observations	371,525	371,525	371,525	371,525

OLS estimates for the coefficients and standard errors are denoted in parentheses. Standard errors are clustered at the individual level in Panel A. In Panel B robust standard errors are used. The outcome variable is unemployed, which is a dummy indicating whether a person is unemployed or not. The sample consists of those adults that are aged 16 or higher and part of the labor force. Column (1) – (4) show different specifications regarding the state fixed effects and individual controls. Treatment is defined by a state terminating the COVID-19 UI programs early. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.