

NBER WORKING PAPER SERIES

AGGREGATE EMPLOYMENT EFFECTS OF UNEMPLOYMENT BENEFITS
DURING DEEP DOWNTURNS:
EVIDENCE FROM THE EXPIRATION OF THE
FEDERAL PANDEMIC UNEMPLOYMENT COMPENSATION

Arindrajit Dube

Working Paper 28470
<http://www.nber.org/papers/w28470>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2021

I thank Chris Boone, Gabriel Chodorow-Reich, Peter Ganong, Pascal Michaillat, Jesse Rothstein for thoughtful comments. Ihsaan Bassier provided excellent research assistance. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Arindrajit Dube. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Aggregate Employment Effects of Unemployment Benefits During Deep Downturns: Evidence from the Expiration of the Federal Pandemic Unemployment Compensation

Arindrajit Dube

NBER Working Paper No. 28470

February 2021

JEL No. E24,E62,E65

ABSTRACT

The expiration of the temporary \$600 boost to weekly UI benefits under the Federal Pandemic Unemployment Compensation (FPUC) led to a sharp, unprecedented, 98 percentage point reduction (on average) in the replacement rate during a time when employment was recovering during the Covid recession. Leveraging the considerable variation in this drop across states, I use a difference-in-differences event study design to estimate the macro employment effects. I find little impact of job gains from the benefit reduction, especially when I focus on groups (non-college graduates, and those from non-high-income households) that comprise of most UI recipients. The estimates rule out job gains implied by much of the micro UI duration elasticities from the existing literature.

Arindrajit Dube

Department of Economics

University of Massachusetts

Crotty Hall

412 N. Pleasant Street

Amherst, MA 01002

and NBER

adube@econs.umass.edu

1 Introduction

The Federal Pandemic Unemployment Compensation (FPUC) passed under the CARES Act added a \$600 boost to weekly UI benefits in March 2020. This led to an unprecedented increase in the benefit replacement rate, which (on average) rose from roughly 48 to 145 percent (Ganong, Noel and Vavra 2020). Does generous earnings replacement of this type during a deep downturn impede employment growth, either via a labor supply or via a wage push channel? To answer this, I take advantage of the sharp curtailment in benefits induced by the July 2020 expiration of FPUC to assess whether the generous benefit replacement rate had a negative impact on employment by disincentivizing labor supply (or by reducing job creation). I use an event-study design using the new, high-frequency, Census Household Pulse Survey (HPS), and consider how non-college, adult worker employment evolved differentially by the state level variation in the median earnings replacement rates following the expiration of the FPUC.

While the replacement rates varied markedly across states (because of very different earnings levels of workers), I find that there is no indication that this had a substantial impact on the employment following the expiration of the FPUC, which reduced the replacement rate roughly from 145 to 85 percent. The point estimate is negative (opposite of what is expected from labor supply considerations), and not statistically distinguishable from zero. These conclusions hold after controlling for demographic differences and state-level job loss measures during the pandemic. An even stronger conclusion obtains when I focus only on those two groups that comprise of the vast majority of UI recipients in this period: those without a college degree, and those in households that did not have incomes exceeding \$150,000 in 2019. The preferred estimates rule out an increase in overall employment more than around 1.3%. I compare my macro-based estimates to employment gains implied based on alternative UI duration elasticities from the micro literature. If the macro labor response measured here were driven primarily by micro labor supply considerations, the findings here rule out large and many mid-sized UI duration elasticities found in the literature, especially when I focus on the non-college and non-high-income samples. The findings here are also consistent with the presence of rationing or aggregate demand externalities, which can rationalize more sizably positive micro estimates (from a benefit cut) with a less positive (or

even negative) macro estimate.

My findings of limited sensitivity of employment to UI generosity are similar to what was found in Altonji et al. (2020) and Bartik et al. (2020) using different data and empirical design, focusing on the introduction of the FPUC and rehiring at the outset of the crisis. A challenge for these analyses is that they are analyzing employment changes by different replacement rate regimes, but the concurrent onset of the crisis and the passage of the CARES Act makes this challenging; moreover, states with high earnings like in the Northeast with lower replacement rates under FPUC were hit much harder by the initial wave of the pandemic.¹ In contrast, in this paper I leverage relatively high frequency variation around the time of a sharp, large drop in benefit levels. The policy change was at the national level, so the changes in replacement rates at state levels were not driven by state-level labor market conditions. Moreover, the change is occurring during a time with an overall increase in employment, which arguably is a much more informative about labor supply issues than when many states were under lockdown as in spring. Overall, then, the expiration event provides a much cleaner source of identifying variation, which is based primarily on how the common benefit reduction translated very differently in terms of the replacement rates due to cross-sectional differences in earnings. Moreover, during this period, I show that states with different sized drops in UI replacement rates were following parallel employment trends prior to expiration. The closest to this paper is the study conducted by Finamor and Scott (2021), who also study how individual level employment outcomes changed for workers in Homebase using variation in replacement rates over 2020 within state-industry-week cells. Their identifying assumption is, therefore, quite different than here as they rely on workers' wage differences within local labor market while I am using variation across labor markets. Moreover, this (along with the fact that they only observe employment within Homebase data which is heavily weighted towards small businesses, especially restaurants) means their employment elasticity might be thought of as a "mixed" elasticity between micro and macro effects (in the terminology of Landais, Michaillat, and Saez 2018). In contrast, by using cross-state variation and all employment, I am estimate

¹Another related paper is by Marinescu, Skandalis and Zhao (2020), who find that applications-per-vacancy were somewhat lower in occupation-state groups with a larger increase in the replacement rate; however, they are not able to provide a causal interpretation of this fact given the crisis affected these groups differentially.

the macro effect of the policy, which I can clearly relate to the macro and micro estimates in the existing literature.

While there is a large literature on the effect of UI benefits on unemployment duration, there is little evidence on the macro effects of large changes in the replacement rate during downturns. Most of the evidence on benefit amounts come from micro-level studies (e.g., Card et al. 2015). At the same time, the literature on macro level effects of UI policy (including during downturns) have focused on extension of maximum benefit duration. In contrast, this paper adds evidence on the macro effects of a change in the benefit levels (i.e., replacement rates), which affect a much larger set of UI recipients.

Overall, the weight of evidence seems to suggest that the unprecedentedly generous UI benefit levels during the Covid crisis did not have any substantially negative effect on jobs. These findings are relevant for both policy design in the current environment as well as thinking about designing optimal UI policies during downturns.

2 Data and Research Design

I use the weekly Census Household Pulse Survey, which collects data from an average of around 97,000 respondents per week using email and texts. The details for the dataset can be found here: https://www2.census.gov/programs-surveys/demo/technical-documentation/hhp/2020_HPS_Background.pdf. The survey collection has varied somewhat across the various waves of the survey. It was collected every week between May and July; there was a subsequent gap, followed by collection every two weeks starting late August. I use all 18 survey rounds between June and November in this analysis. I focus on individuals between 18 and 65 years of age, producing a total of around 1 million observations in my sample. There are a number of reasons I use the HPS in this paper. First, and most importantly, the HPS provides high frequency measures of employment rates. This is especially true in the pre-treatment period when it was collected weekly, while it was biweekly in the post-treatment period. This higher frequency measurement with fairly large sample size allows for a more compelling event study analysis, where we can assess both preexisting trends and dynamic effects. Second, the HPS also asks several questions (including whether

anyone in the household has experienced job loss since March 2020, and 2019 income) which are useful for constructing controls and defining likely treated subgroups.

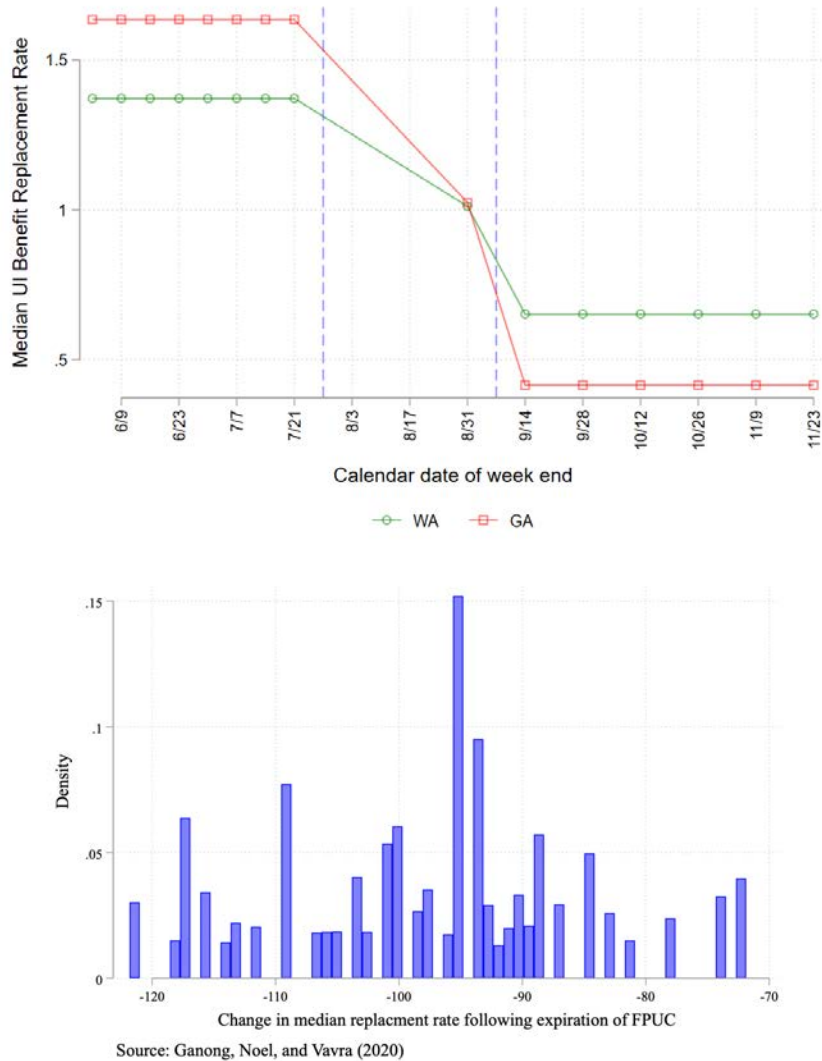
I focus on the following variables: whether a person was working in the previous week (ANY-WORK), their educational credentials (EEDUC), and what share of their income came from unemployment benefits (SPNDSRC5), whether someone in their household experienced job loss since March 2020 (WRKLOSS), their race (RRACE), hispanic origin (RHISPANIC), year of birth (TBIRTH_YEAR), gender (EGENDER), home ownership and rental information (TENURE), and previous year’s household income category (INCOME). Starting in round 13 (week ending August 31), the HPS also collects information about UI receipt (UI_RECV); I use this information estimate likely UI receipt across various subgroups (and hence the exposure to the policy change).

I use the information on the median replacement rate under CARES Act (i.e., what share of one’s pre-pandemic earnings are being replaced using FPUC and regular benefits) as well as prior to the CARES Act from Ganong, Noel and Vavra (2020).

The replacement rate prior to FPUC was $R_{NoFPUC} = \frac{\min(bW, \bar{B})}{W} = \frac{B}{W}$, where B is the benefit level, \bar{B} is the maximum benefit amount, and b is the conditional replacement rate (typically 50 percent), and W is the earnings in the reference period. Under FPUC, there was an additional \$600 in benefits, leading to a replacement rate of $R_{FPUC} = \frac{B}{W} + \frac{600}{W}$. Averaged across our sample, the median replacement rate was around 146 in July, 2020 under the FPUC. Following the expiration of the FPUC, almost all states provided a temporary boost of \$300 under the Lost Wages Assistance (LWA), which lasted for around six weeks in most cases. While the exact timing of the LWA *payments* varied slightly across states (due to administrative delays), the payments *covered* unemployment typically for five to six weeks immediately following the FPUC expiration (i.e., weeks between July 26 and September 5). Therefore, the replacement rate for someone who was unemployed in mid September onward was at the pre-pandemic levels in all states, at $\frac{B}{W}$. Given the timing of the various rounds of the HPS (where there is a 5 week gap in surveying during the month of August), only the replacement rate for the survey ending on the week of August 31 is likely affected by the LWA. For this reason, I will define the post-FPUC period excluding this

transition week in the analysis below, Moreover, this means the post-FPUC replacement rate of R_{POST} . The change in the replacement rate, then, was $\Delta R = R_{POST} - R_{FPUC} = \frac{-600}{W}$. Averaged across all states, the median replacement rate fell to from 146 to 48 percent upon FPUC expiration, a 98 percent point reduction. However, there was substantial variation in the reduction in the replacement rate. As shown in Figure 2, the fall in the expiration ranged between 71 percent point in Washington State and 122 percent point in Georgia.

Figure 1: Distribution of Change in Median UI Replacement Rates Across States Following Expiration of FPUC



Notes. The top panel shows the evolution of the median UI replacement rates over time for the two states with the smallest (WA) and largest (GA) changes from the July expiration of FPUC. The bottom panel shows a histogram of the the change in the median replacement rate following the expiration.

I use the change in the median replacement rate ΔR induced by the expiration of the FPUC

to assess labor supply impacts. Identification requires that the variation induced by different earnings, W , satisfies the parallel trends assumption: that high and low earnings states were not experiencing differential employment growth (e.g., due to different pace of the recovery, etc.), at least conditional on pre-expiration characteristics (including size of job losses experienced in each state).

To test this, I use an event-study approach that plots the evolution of employment differentially in states with high versus low median replacement rate, R_s . I estimate the following regression:

$$Emp_{it} = \alpha + \sum_{\tau=1}^{15} \beta_{\tau} \left(I_t^{\tau} \times \Delta R_{s(i)} \right) + \Gamma X_{it} + \rho_{s(i)} + \phi_t + e_{it} \quad (1)$$

Here Emp_{it} is an indicator for if the person i worked in the previous week at all, I_t^{τ} is an indicator for τ^{th} survey week, $s(i)$ is the state of individual i , X_{it} is a vector controls (worker’s race, hispanic origin, quartic in age, gender, 2019 income, state-level share of households experiencing job loss since March 2020). Since states wages may be correlated with how hard the state was hit by the pandemic-induced job losses in March and April, inclusion of the state-level measure of pandemic job loss may be important.² I estimate the regression both with and without these additional controls, X_{ist} . I plot the coefficients β^{τ} , normalizing β^1 as zero, which trace out changes in employment since the first sample week (ending in May 5, 2020). Regressions use the sampling weight provided in the dataset, and standard errors are clustered by state. Finally, I multiply all coefficients, β_{τ} , by -0.98, which means the coefficients can be interpreted as the impact of a 98 percent point reduction in the replacement rate, which was the average reduction from the expiration of FPUC.

Additionally, I also calculate difference-in-differences estimates as follows:

$$Emp_{it} = \alpha + \beta_{DID} \left(Post_t \times \Delta R_{s(i)} \right) + \Gamma X_{it} + \rho_{s(i)} + \phi_t + e_{it} \quad (2)$$

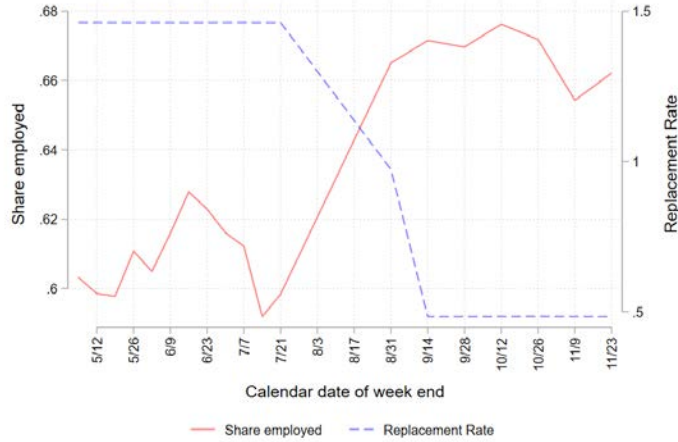
Here $Post_t$ is dummy for the period following the July 31 expiration of the FPUC. Again, I multiply the coefficient (and standard errors) by -0.98 for interpretational ease.

²One may worry that job losses could be affected by the expiration of the FPUC. However, in practice, inclusion of this makes little difference. The DID estimate for all workers excluding (including) this covariate while including other covariates is -0.011, s.e. = 0.015, (-0.017 s.e. = 0.015).

To start with, I focus on are all individuals of age 18-65. This provides the clearest estimate of the aggregate employment impact which is the key goal of this paper. However, I also provide estimates for groups who are more likely impacted by the policy. As one restriction, I focus on non-college grads, who compose of around 79% of UI recipients over the sample period when that data is available (starting in round 13 ending on August 31, 2020); and around 76% of those stating that they have relied on UI for their expenditures (collected over all rounds of the HPS). As a second restriction, I exclude households who report earning high incomes in 2019 (\$150,000 or more), or around 12% of the sample. (To stress, this is based on historical income information and therefore unaffected by events in 2020, including UI policy changes.) The non-high-income group includes around 94% of UI recipients and 94% of those stating they have relied on UI for their expenditures over the same sample periods. By excluding higher income, or college educated, individuals, we are able to more sharply focus on the aggregate employment effects for the population that contain almost all likely UI recipients who were affected by the policy change.

The variation in the replacement rate here is coming from the variation in benefit wage rates for those who are unemployed across states. One concern is that high and low-wage states may have experienced different rates of employment change during the downturn, and hence different (latent) rates of employment recovery during the second and third quarters of 2020. I address this concern in several ways. First, the event study design will allow us to assess whether employment trends were diverging across these states prior to the expiration. Second, my preferred specification controls for state-level measure of pandemic job loss which should mitigate bias. Third, we can also directly assess whether this measure of cumulative pandemic related job loss at the beginning of our sample (week ending May 5, 2020) was correlated with the future change in replacement rate ΔR . As shown in Appendix Figure A1, there is little indication that the severe job losses during March and April were systematically related to the future change in replacement rate in August, ΔR . This provides additional validation for the research design.

Figure 2: Employment Rate and Median Replacement Rate Over Time: U.S. during 2020



Notes. The figure plots the median replacement rate averaged across states over time (left y-axis), and the share employed among 18-65 year olds (right y-axis) from the Census Household Pulse Survey.

3 Findings

First, as shown in Figure 3, employment for non-college educated workers started to recover after mid-May, followed by an employment decline starting in late June—as a second peak in Southern and Western states led to greater restrictions and consumer reactions. Subsequently, we see a sharp rebound in August, which was around the time of the expiration of the FPUC. This naturally raises the question: did the expiration of the FPUC help spur the employment recovery through incentivizing labor supply?


To assess this, in Figure 4, I show the key results from the estimates from equation (1) that use cross-sectional variation in the earnings of UI recipients. The plotted coefficients represent the impact of a 98 point reduction in the replacement rate, R_s , on employment probabilities of individuals, ages 18-65. In the top panel A, I show the results with only the time and state fixed effects as controls, while in the bottom panel B, I include the full set of individual and state-level controls (preferred estimate). I estimate the same regression (with controls) but now for all individuals (ages18-65).

First, in both panels, employment changes in June and July were not systematically different by future replacement rate change, ΔR_s ; states with high and low treatment intensities appear to have had similar employment trajectories prior to the policy change at the end of July. This lack


of systematic pre-existing trends provides an important validation of the research design.

Second there is little impact on the employment. In both panels, the DID estimates are not statistically distinguishable from zero. In my preferred specification with all controls, the magnitude of the DID estimate on employment probability (-0.011 , $s.e. = 0.015$), along with the July sample employment rate of 0.60 , implies a 1.8% *reduction* in jobs from the expiration of FPUC. Recall that labor supply consideration suggests employment should rise from the benefit cut, so the point estimate are of the “wrong” sign. Moreover, the confidence interval rules out more than a 3.1% increase in overall employment in response to the at the 95% confidence level. These estimates are useful for producing an overall macro employment estimates and confidence bounds for the US labor market. Given our in-sample employment level for this group of around 120 million in end of July, 2020, the point estimate suggests a loss of 2.2 million jobs, while the confidence interval rules out employment gains exceeding 3.7 million at the 95% confidence level.

Figure 3: Impact of FPUC Expiration on Employment Probability: All Individuals 18+
 Panel A: Only state and time fixed effects without additional controls



Panel B: Full set of controls



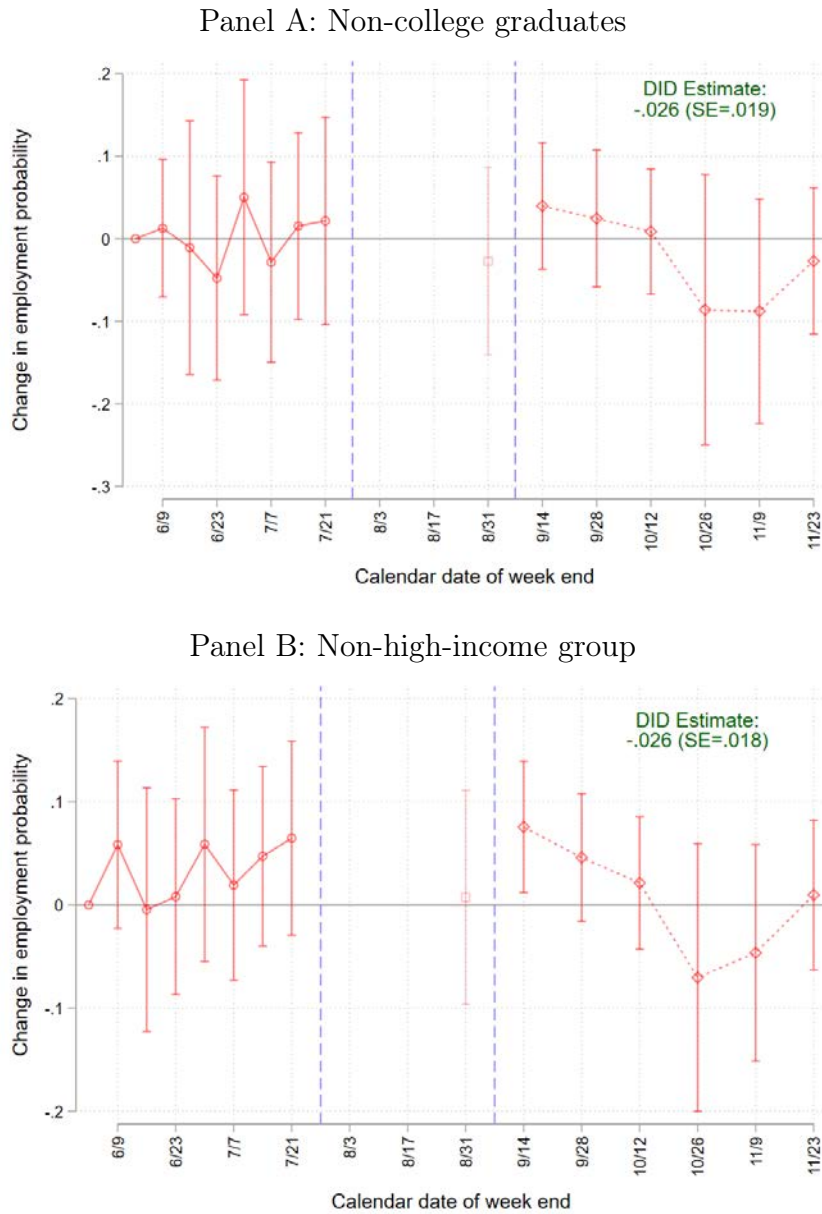
Notes. Coefficient represents the impact of a 89 point reduction in the median replacement rate from the July expiration of FPUC for all individuals ages 18-65. All employment changes normalized to week ending in 6/2. The dashed blue line indicates the expiration of the FPUC. The estimates from the week ending 8/31 (unconnected and in lighter shade) are in the “transition” period, while the subsequent estimates are in the post-treatment period. Regression controls for week and state fixed effects, quartic in age, gender, race, hispanic origin, 2019 income categories, and share of households in state experienced job loss since March 2020. The 95% confidence intervals are based on standard errors clustered at the state level.

We can also specifically focus on higher impact groups. Panel A of Figure 5 reports the impact on non-college graduates. Similar to the overall workforce sample, I find little pre-existing trends, along with a negatively signed but not statistically significant DID point estimate of -0.026 (s.e. 0.019). As before, using a July employment rate of around 0.52 and employment level of 75.3M

suggests a loss of 3.8M non-college graduate employment; and the 95% confidence interval rules out more than 1.6M increase in jobs. (If we wanted to factor in that non-college graduates were 79% of likely UI recipients and scaled up the estimate by $1/0.79$, it would rule out more than 2.1M overall job gains.)

Panel B provides the estimated effects for the non-high-income group. Here, too, employment seems to have been following parallel trends prior to expiration; and there is a negative signed but not statistically significant change on employment. The DID point estimate of -0.025 (s.e. 0.017), coupled with a July employment rate of 0.57 and employment level of 102M suggests a 4.5M loss of jobs, and the confidence intervals rule out job gains among this group exceeding 1.5M . Again, if we wanted to scale this up given that the non-high-income group contain around 94% of likely UI recipients, the confidence bounds would rule out more than 1.6M overall job gains. I take this to be the most informative bound on likely aggregate job gains, since this group excludes only a small group of individuals who are almost entirely unaffected by the UI policy changes.

Figure 4: Impact of FPUC Expiration on Employment Rate: Excluding College Grads and High Income Households



Notes. Coefficient represents the impact of a 98 point reduction in the median replacement rate from the July expiration of FPUC for all individuals ages 16-65 without a college degree (Panel A); individuals ages 16-65 in the households with 2019 income less than \$150,000. All employment changes normalized to week ending in 6/2. The dashed blue line indicates the expiration of the FPUC. The estimates from the week ending 8/31 (unconnected and in lighter shade) are in the “transition” period, while the subsequent estimates are in the post-treatment period. Regression controls for week and state fixed effects as well as controls for: quartic in age, gender, race, hispanic origin, 2019 income categories, and share of households in state experienced job loss since March 2020. The 95% confidence intervals are based on standard errors clustered at the state level.

The estimates here consider the linear relationship between the drop in replacement rate and employment outcomes. To assess possible non-linear relationships, Appendix Figure A2 shows

the evolution of employment in middle and top terciles of the replacement rate change (relative to the bottom tercile). There is no indication of bigger employment gains for larger drops in the replacement rate here.

4 Connecting Macro and Micro Effects of UI

Overall, the findings in this paper provide evidence that the sharp, and unprecedentedly large, reduction in the UI benefit replacement rates due to the expiration of the FPUC did not lead to any sizable increase in employment. As noted above, we can rule total job gains exceeding 1.6 million at the 95% confidence level when we consider our non-high-income sample which contains virtually all likely UI recipients. While these estimates are inconsistent with some claims made in the popular press,³ how do these bounds compare to what we may expect from labor supply effects from previous research?

First, this paper provides macro estimates of employment changes from UI policy, similar to Boone et al. (forthcoming), Chodorow-Reich et al. (2018), and Hagedorn et al. (2016) that considered changes in maximum benefit duration during the Great Recession. A separate, micro, approach comes from considering how individual worker UI duration responds to the benefit changes. There is a substantial range in the estimates for the duration elasticity. Chetty (2008) finds an overall elasticity of around 0.5 using data observational data, though it varies substantially between those who are liquidity constrained (0.7) and those who are not (0.2). At the low-end, Meyer and Mok (2014) find estimates in the 0.1-0.2 range using a change in maximum benefit levels in New York for a subset of recipients using a difference-in-differences design. In contrast, Card et al. (2015) use data from Missouri and a regression kink design (RKD) and find UI duration elasticities of around 0.35 during non-recessionary times, and between 0.65-0.9 during recessions. Applying the RK design to data from 5 states between 1970-1984, Landais (2015) find on overall elasticity around 0.4.⁴ Overall, I take the 0.1 to 0.9 as the likely range for the micro

³For example, in a Wall Street Journal op-ed from December 3, 2020, Casey Mulligan and Stephen Moore argue that the expiration of FPUC helped create between 3 and 5 million jobs. <https://www.wsj.com/articles/unemployment-bonus-proves-its-harm-11607037230>

⁴A related, but distinct, set of studies estimate the elasticity of UI or unemployment duration with respect to the maximum benefit duration. In general, we expect the elasticity of UI duration with respect to benefit level (the

duration elasticity, ϵ^m , with a central estimate of 0.5.

We can use an unemployment duration elasticity (with respect to benefit levels) to calculate the implied change in the number of jobs as follows. (For a similar exercise evaluating the macro and micro estimates of maximum benefit durations during the Great Recession, see Boone et al. (forthcoming)). First, note that for a given elasticity, ϵ^m , the aggregate employment change implied by a particular benefit change $\% \Delta B$ can be approximated as⁵:

$$\Delta E \approx -N_{UI} \times \% \Delta B \times \epsilon^m \quad (3)$$

Here N_{UI} is the number of UI recipients: this was around 17 million at the time of FPUC benefit expiration for those receiving regular benefits, extended benefits, or Pandemic emergency unemployment compensation.⁶ $\% \Delta B = -\frac{0.98}{1.46} = -0.67$ refers to the percentage reduction in UI benefit levels from the policy change, and ϵ_D^m is the (micro) elasticity of UI benefit duration with respect to the benefit level.

If we take the 0.1 to 0.9 as the likely range, and 0.5 as a central estimate of ϵ^m , these imply a range of ΔE between 0.9M and 8.2M, with a central estimate of around 4.6M. These contrast with the negative signed ΔE found in this paper, and the confidence intervals for employment

question of interest in this paper) to be smaller than the elasticity with respect to maximum benefit duration; this is because the former boosts benefits all beneficiaries starting at the beginning of the spell, while the latter only boosts benefits for those whose spell reach the end of regular benefits. Katz and Meyer (1990) finds a range of 0.36-0.43. For a discussion of estimates of how the extension of maximum benefit duration in the Great Recession affected unemployment duration, see Rothstein (2011), whose findings suggest a very small duration elasticity of 0.06. (Note, however, that Landais et al. (2018) argue the estimates in Rothstein likely reflect a mix and micro and macro factors.)

⁵The elasticity of unemployment duration is nearly identical to the elasticity of the unemployment rate with respect to UI. If the job finding rate from unemployment is f_{UE} while the separation rate to unemployment is s_{EU} , the average duration of unemployment is $d_U = 1/f_{UE}$, while the unemployment rate is $u = \frac{f_{EU}}{f_{UE} + s_{EU}}$; with $s_{EU} \ll f_{UE}$, $u \approx \frac{s_{EU}}{f_{UE}} = s_{EU}d_U$, and hence $\frac{\partial \ln u}{\partial x} \approx \frac{\partial \ln d_U}{\partial x}$. (Also see Landais, Michailat, and Saez 2018, who use the same approximation). This, along with the fact that $\Delta E = -\Delta U = -\Delta N_{UI}$ leads to the approximation for ΔE .

⁶Total continuing claims from these 3 sources for the week of July 25 from these sources was 17.1M (see here: <https://oui.doleta.gov/press/2020/081320.pdf>). For my main analysis, I do not use the additional 10.7M continuing claims under the Pandemic Unemployment Assistance (PUA). The PUA was a new program which was designed to help those with too little earnings to qualify for regular UI, and those who were self-employed. The verification system for PUA was substantially different than for regular UI. It is possible that some individuals who were unlikely to be in the labor force even absent the pandemic claimed PUA. Moreover, it is not clear if some individuals continued to claim PUA while working informally (say as babysitters). For these reasons, I take a conservative approach by focusing only on those who were on regular UI. If we were to use the 27.8M UI recipient estimate: as expected, those estimates are more concentrated around zero (and would rule out an even larger part of the range of ϵ^m).

effects rule out a substantial part of this range. To relate the micro-based predictions and macro estimates here more clearly, we can invert equation (3) to obtain the implied macro elasticity (ϵ^M) implied by the confidence bounds of our various macro employment estimates, $\Delta\hat{E}^M$, under the assumption that the micro labor supply considerations are the only force at work:

$$\epsilon^M \approx \left(\frac{\Delta\hat{E}^M}{N_{UI}} \right) \left(\frac{1}{\% \Delta B} \right) = \% \Delta \hat{E}^M \times \left(\frac{E}{N_{UI}} \right) \times \frac{1}{\% \Delta B}$$

Here $\frac{E}{N_{UI}} = 15M/120M = 0.09$. When we focus on specific subgroups (e.g., non-college graduate), we have to modify the formula to account for the shares of that group among those unemployed and among those in the labor force. Namely, based on the macro estimates for a group g , the implied duration (or unemployment) elasticity is:

$$\epsilon_g^M \approx \% \Delta \hat{E}_g^M \times \left(\frac{E_g}{E} \cdot \frac{E}{N_{UI}} \cdot \frac{N_{UI}}{N_{UI,g}} \right) \times \frac{1}{\% \Delta B}$$

$$\epsilon_g^M \approx \% \Delta \hat{E}_g^M \times \left(\frac{E}{N_{UI}} \right) \times \frac{1}{\% \Delta B} \times \frac{s_g^E}{s_g^{UI}}$$

When we focus on a specific group, g , the conversion of the employment change, $\% \Delta \hat{E}_g^M$, to the implied duration elasticity requires adjustment by the relative shares of that group among those who are (1) employed (s_g^E), versus (2) those receiving UI (s_g^{UI}). If, for example, the group composes a greater share of those on UI than those working, the duration elasticity implied by a specific value of the employment change is smaller; and vice versa.

Table 1 reports the employment effects and the implied duration elasticities for our three key groups: all individuals 18-65; those without a college degree; and those in the non-high-income households. The column 4 implied point estimates of ϵ_D^M are of the “wrong sign” given the negative employment estimates found here (recall the micro duration elasticities, ϵ_D^m , are positive in sign). For the overall workforce, the 95% confidence intervals rule out ϵ^M larger than 0.31, which includes large and medium-sized micro elasticities, estimates for which largely range

Table 1: Employment Effects and Implied Duration Elasticities Across Subgroups

	% Δ Employment	Share among UI Recipients	Share among Employed	Implied Duration Elasticity, ϵ_D^M
<i>Overall Workforce</i> (N=963,744)	-0.020 (.025) CI: [-0.069 , 0.029]	100%	100%	-0.208 (0.264) CI: [-0.726, 0.311]
<i>Non-college graduates</i> (N=449,700)	-0.047 (0.036) CI: [-0.118 , 0.024]	79%	67%	-0.399 (0.301) CI: [-0.988, 0.190]
<i>Non-high-income householders</i> (N=802,431)	-0.045 (0.030) CI: [-0.104 , 0.138]	94%	88%	-0.431 (0.294) CI: [-1.001, 0.146]

Notes. Column 1 reports the estimated percentage change in employment from the 98 point reduction in the replacement rate due to the expiration of FPUC. These are based on estimates from regression equation (2). Columns 2 and 3 report the shares of the group among UI recipients (s_g^{UI}), and among employed (s_g^E), respectively. Column 4 reports the implied duration elasticity, ϵ_D , needed to rationalize the employment effect found in column 1 for group g . The estimates are provided for 1) the overall workforce (of ages 18-65), 2) those without a college degree, and 3) those in households that earned less than \$150,000 in 2019. The standard errors are in parentheses; the 95% confidence intervals are shown in brackets.

is roughly between 0.1 and 0.9. When we focus on those without a college degree—who compose of around 79% of the UI recipients, and 67% of labor force participants—the 95% confidence intervals rule out out elasticities larger than 0.19. Finally, when we consider individuals from non-high-income households—who compose of around 94% of the UI recipients, and 88% of labor force participants—the 95% confidence intervals rule out out elasticities larger than 0.15. In other words, once we trim the college educated (or high income) parts of the sample, the implied confidence bounds rule out all but the smallest set of micro elasticities in the literature. It is worth noting that when we focus on the groups that are more likely to include UI recipients, there is no indication of a more positive employment effect from the expiration. These estimates also demonstrate the value of being able to focus on high impact subgroups, where individuals compose of a greater share of UI recipients than the labor force at large. Given the number of likely UI recipients among the non-college-graduate and the non-high-income samples, the overall employment effects in those groups are not consistent with employment effects implied by anything but the smallest set of micro elasticities.

It is possible that the micro labor supply elasticity was depressed during the Covid crisis due to health restrictions; note, however, that the August to November post-treatment period is one where unemployment was steadily falling and more and more people were returning to work. Arguably more importantly, labor supply is not the only (or most relevant) constraint on employment during downturns. One possibility is that the change in UI led to large fall in wages and spurred vacancies, creating a micro-macro wedge (as in Hagedorn et al. 2013); however, this would only worsen the problem of explaining the lack of job growth following the large reduction in UI benefits. More relevant to our case, as Landais, Michaillat, and Saez (2017) and Kekre (2019) point out, macro employment effects of a benefit cut may be more negative than the micro estimates due to search and aggregate demand externalities. In a slack labor market where jobs are rationed, one person finding a job may lead to a reduced chance of finding a job for others, creating a wedge between the micro and macro effects. (For empirical evidence on displacement externality, see Lalive, Landais, Zweimueller 2015). Such rationing can explain why any job gains at the macro level may be more muted than implied by the micro elasticities, though it would not be able to explain a negative signed effect. In addition, it is possible that there are positive local aggregate demand effects from the sharp reduction in unemployment benefits led to an actual decrease in jobs. Kekre (2019) shows this in the context of expanded UI benefits during the Great Recession within a New Keynesian model. Consistent with that channel, as shown by Farrell et al. (2020), consumption rose among UI recipients following the receipt of the benefits, with a marginal propensity to consume of around 0.7, which was not very different from pre-Covid evidence. For both of these reasons, it is possible that there were negative macro-level offsets to positive micro-level labor supply effects from the benefit cuts. Going forward, future research would benefit from better understanding the source of this very modest employment response to UI benefit generosity, and whether it was driven by unusually low micro-level response of unemployment duration, or macro-level wedges affecting aggregate employment.⁷

⁷It is worth noting that if we look at relatively narrow subgroups, the employment effects may capture less of the macro effects than when we consider as a whole. For example, if some of the displacement externality occurs outside the group, looking more narrowly at some groups will not capture those. However, in our case, most of the likely displacements are probably within groups (i.e., within non-college grads, and within non-high-income households—making the employment estimates fairly close to the pure macro elasticities.

To summarize, this paper uses a historically unprecedented and sharp reduction in the UI benefit levels which cut the replacement rate by around 98 points, but with much geographic variation. I find surprisingly little effect on overall employment, and the estimates rule out large and medium-sized employment responses implied by estimates from previous micro-level studies. Importantly, taking well-identified micro elasticities from the literature and using those to predict aggregate job changes would have vastly overstated any likely increase in jobs from the dramatic cutback in the replacement rate during the downturn. These findings suggest the costs of providing very high benefit levels during deep downturns may be fairly small. The findings from this unprecedented change in benefit levels are useful to consider as we contemplate reforming the current UI system in the U.S., including the possibility of raising the benefit replacement rate, and tie it to the state of the business cycle through more rules than discretion.

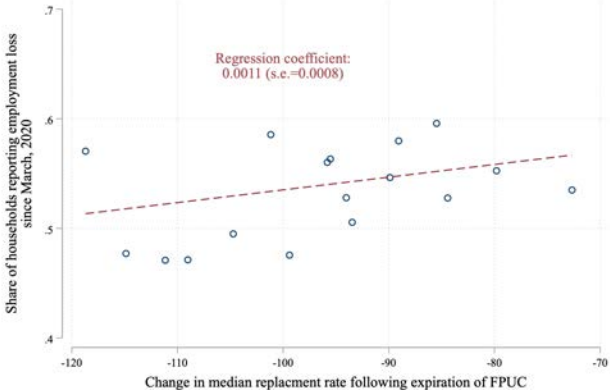
References

1. Altonji, J., Contractor, Z., Finamor, L., Haygood, R., Lindenlaub, I., Meghir, C., O’Dea, C., Scott, D., Wang, L. and Washington, E., 2020. “Employment Effects of Unemployment Insurance Generosity During the Pandemic.” *Mimeo*. Available at: [http://tobin.yale.edu/sites/default/files/files/C-19%20Articles/CARES-UI_identification_vF\(1\).pdf](http://tobin.yale.edu/sites/default/files/files/C-19%20Articles/CARES-UI_identification_vF(1).pdf)
2. Bartik, Alexander W., Marianne Bertrand, Feng Ling, Jesse Rothstein, and Matthew Unrath. 2020. “Measuring the Labor Market at the Onset of the COVID-19 Crisis.” *BPEA Conference Draft*.
3. Boone, Christopher, Arindrajit Dube, Lucas Goodman, and Ethan Kaplan. 2020. “Unemployment Insurance Generosity and Aggregate Employment.” *American Economic Journal - Economic Policy*.
4. Card, David, Andrew Johnston, Pauline Leung, Alexandre Mas, and Zhuan Pei. 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.
5. Chetty, Raj. 2008. “Moral hazard versus liquidity and optimal unemployment insurance.” *Journal of political Economy* 116(2): pp.173-234.
6. Chodorow-Reich G, Coglianesi J, Karabarbounis L. 2019. “The macro effects of unemployment benefit extensions: a measurement error approach.” *The Quarterly Journal of Economics* 134(1): pp.227-79.
7. Farrell, D., Ganong, P., Greig, F., Liebeskind, M., Noel, P. and Vavra, J., 2020. “Consumption Effects of Unemployment Insurance during the Covid-19 Pandemic.” *Mimeo*. Available at SSRN 3654274. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3654274
8. Finamor, Lucas, and Dana Scott. 2021. “Labor market trends and unemployment insurance generosity during the pandemic.” *Economic Letters* 199(2021): 109722.

9. Ganong, P., Noel, P.J. and Vavra, J.S. 2020. "US Unemployment Insurance Replacement Rates During the Pandemic." National Bureau of Economic Research Working Paper No. w27216.
10. Farrell, Diana, Peter Ganong, Fiona Greig, Max Liebeskind, Pascal Noel, and Joseph Vavra. 2020. "Consumption effects of unemployment insurance during the covid-19 pandemic." Available at SSRN 3654274.
11. Hagedorn, Marcus, Fatih Karahan, Iourii Manovskii, and Kurt Mitman. 2013. "Unemployment benefits and unemployment in the great recession: the role of macro effects.." *National Bureau of Economic Research Working Paper* No. w19499.
12. Kekre, Rohan. 2019. "Unemployment Insurance in Macroeconomic Stabilization." *Mimeo*. Available at: https://sites.google.com/site/rohankekre/UI_v0619.pdf?attredirects=0&d=1
13. Lalive R, Landais C, Zweimüller J. 2015. "Market externalities of large unemployment insurance extension programs." *American Economic Review*. 105(12): pp. 3564-96.
14. Landais Camille. 2015. "Assessing the welfare effects of unemployment benefits using the regression kink design." *American Economic Journal: Economic Policy*. 2015 7(4): pp. 243-78.
15. Landais, C., Michaillat, P. and Saez, E. 2018. "A macroeconomic approach to optimal unemployment insurance: Applications." *American Economic Journal: Economic Policy* 10(2), pp.182-216.
16. Marinescu, I.E., Skandalis, D. and Zhao, D., 2020. "Job Search, Job Posting and Unemployment Insurance During the COVID-19 Crisis." *Mimeo*. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3664265
17. Meyer BD, Mok WK. 2014. "A short review of recent evidence on the disincentive effects of unemployment insurance and new evidence from New York State." *National Tax Journal* 67(1): p. 219.

Online Appendix

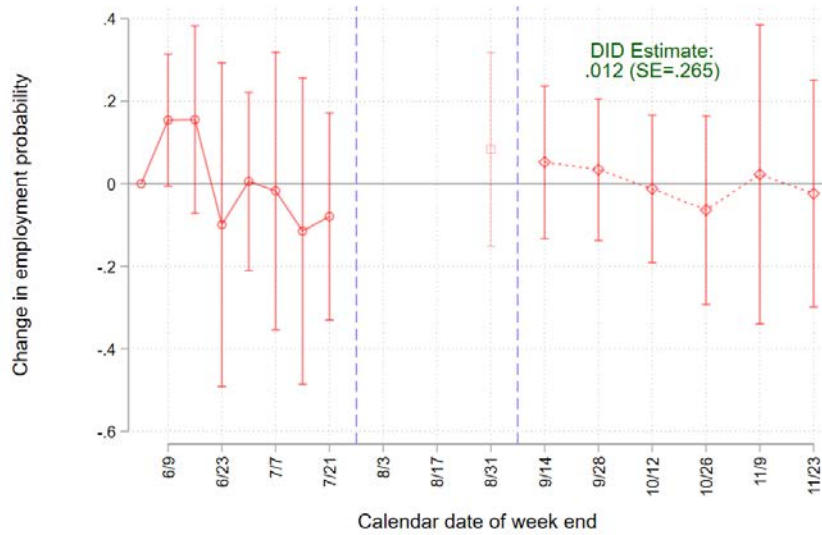
Figure A1: Relationship between Pandemic Job Loss since March, and the Change in Future Median Replacement Rate in July of 2020



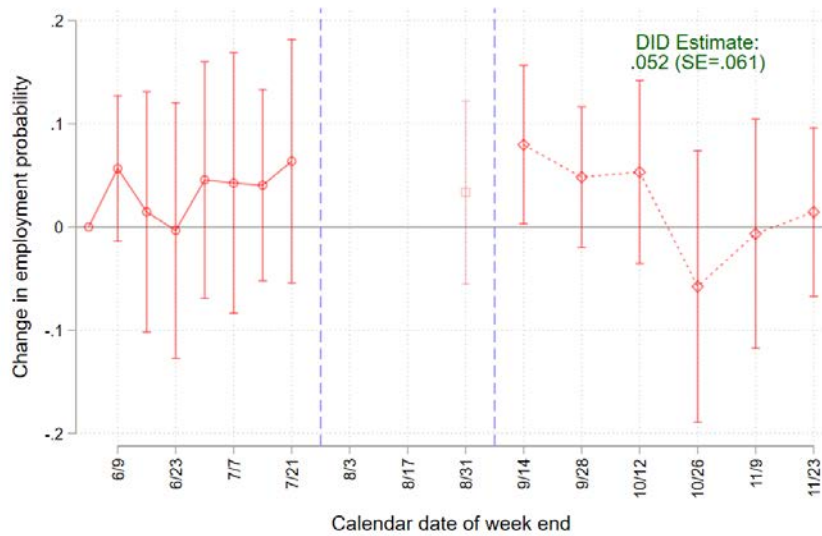
Notes. Binned scatter plot of the change in median replacement rate from the July expiration of the FPUC, and the share of households reporting employment losses since March, 2020 during the week ending May 5, 2020.

Figure A2: Impact of FPUC Expiration on Employment Probability: By Tercile of Change in Replacement Rate

Panel A: Mid-size decrease compared to smallest tercile decrease



Panel B: Large decrease compared to smallest tercile decrease



Notes. Coefficients are scaled such that they represents a 98 point reduction in the median replacement rate from the July expiration of FPUC for all individuals ages 18-65. All employment changes normalized to week ending in 6/2. The dashed blue line indicates the expiration of the FPUC. The estimates from the week ending 8/31 (unconnected and in lighter shade) are in the “transition” period, while the subsequent estimates are in the post-treatment period. Regression controls for week and state fixed effects, quartic in age, gender, race, hispanic origin, 2019 income categories, and share of households in state experienced job loss since March 2020. The 95% confidence intervals are based on standard errors clustered at the state level.