

Do Unemployment Insurance Extensions Reduce Employment?

John Coglianesi*

October 31, 2015

Abstract

Unemployment insurance (UI) extensions can have broad effects on labor markets by changing search effort, creating or destroying jobs, and boosting aggregate demand. I use a Federal UI extension in the United States during the Great Recession as a natural experiment to measure the effect on state-level employment. I exploit a feature of this program whereby random sampling error in a national survey altered the duration of unemployment insurance in several states to isolate random variation in the number of weeks of unemployment insurance available. I find evidence that UI extensions may raise employment growth substantially, and although I cannot conclusively rule out an elasticity of zero I can rule out substantial negative effects. I also show several problems with previous attempts to measure the total effect of UI extensions on employment.

1 Introduction

Unemployment insurance (UI) extensions are a common countercyclical tool in the United States. Since the late 1950s, the Federal government has typically responded to recessions by increasing the duration for which

*Harvard University: coglianesi@fas.harvard.edu. I would like to thank Pascal Noel for his work on an early iteration of this project and for his innumerable comments, suggestions, and insights over the course of this project. I thank Evan Storms as well for his work on an early iteration of this project. I am grateful to Mitra Akhtari, Raj Chetty, Gabriel Chodorow-Reich, Peter Ganong, Simon Jager, Larry Katz, Rohan Kekre, Jessica Laird, Greg Mankiw, Jesse Rothstein, Heather Sarsons, and participants at the Harvard macro lunch and Harvard labor lunch for their feedback.

laid off workers can collect unemployment insurance as part of efforts to stimulate the economy. During the Great Recession, UI duration was extended up to 99 weeks in some states, substantially longer than any prior extension. The slow recovery since the recession ended in 2009 raises the question of whether such a large increase in the duration of unemployment insurance held back job growth during this period. Prior research has shown that unemployment insurance can have disincentive effects on job search behavior, so a large extension of unemployment insurance could exacerbate this problem.

At the same time, extending unemployment insurance could have many effects beyond lowering search effort. During a recession, extending unemployment insurance is a particularly straightforward way for the government to increase spending and boost aggregate demand. UI extensions likely make it easier for non-UI-eligible individuals to find jobs, since they face less competition when search effort among individuals collecting unemployment insurance decreases. Additionally firms may face different incentives to either create new vacancies or possibly cut back on hiring.

The broader effects of unemployment insurance cannot be estimated using individual-level variation in UI duration. Research by Rothstein (2011) and Schmieder et al. (2012) has measured the disincentive effects of unemployment insurance by comparing two individuals within the same labor market who differ in the duration of UI they are eligible for. By comparing these two individuals within the same labor market, such analyses net out any broader effects that unemployment insurance has on the labor market. Both individuals face the same amount of aggregate demand, the same number of non-UI-eligible workers, and the same number of available vacancies. These studies are informative about how much an individual's probability of employment changes when UI duration increases, which I call the *micro* effect of unemployment insurance, but they do not capture the market-wide change in employment, which I refer to as the *macro* effect.

In this paper, I estimate the macro effect of extending unemployment insurance at the state-level in the United States using a natural experiment during the Great Recession. An emergency UI extension enacted by the Federal government in 2008 raised the potential duration for which states could offer unemployment insurance up to 53 additional weeks, with different amounts of weeks available based on the unemployment

rate in the state. The measure of the unemployment rate that was used to set duration was based largely on data from a national household survey (the Current Population Survey). I compare this measure to an estimate of the unemployment rate constructed from administrative UI records. In doing so, I isolate the variation in the survey-based unemployment rate that is due to random sampling error and show that in many cases this sampling error led to states being assigned to different UI durations even though they had the same underlying unemployment rate. I compare subsequent employment growth in states randomly assigned to different durations of unemployment insurance and find suggestive evidence that UI extensions raise employment growth. While I cannot conclusively rule out an elasticity of zero, I am able to reject substantially negative elasticities. I also find evidence consistent with unemployment insurance boosting aggregate demand, as employment in non-tradeable industries grows more in response to a UI extension than employment in tradeable industries.

The prior literature on the macro effect of unemployment insurance is largely divided. A series of papers separately estimate the effects of UI extensions on each possible channel and combining these estimates suggests that the macro effect is either less negative than the micro effect or possibly positive (Lalive, Landais, Zweimüller, 2015; Marinescu, 2014; DiMaggio & Kermani, 2015). In contrast, recent work that examines the total effect on employment finds substantial negative elasticities (Hagedorn et al., 2013, 2015). My results are consistent with the former literature, but sharply contradict the latter papers. I examine these studies in more detail to uncover the reason for the discrepancy between our results. I am able to replicate their estimates and find that much of the discrepancy is due to a confluence of three factors: non-parallel trends, measurement error in data based on surveys, and non-smoothness of economic geography.

The remainder of the paper proceeds as follows. Section 2 gives an overview of unemployment insurance extensions in the United States during the Great Recession and describes prior research on the effects of unemployment insurance extensions. Section 3 describes the data and Section 4 gives an overview of the research design I employ, explaining how I identify random sampling error in the CPS. Section 5 discusses the results of this natural experiment. Section 6 compares these results to previous estimates from the literature

and attempts to explain the discrepancies between my results and prior estimates. Section 7 concludes with a discussion of these findings.

2 Background

2.1 Literature

Analysis of unemployment insurance benefits comprises a long literature in labor economics and public finance. Work by Solon (1985), Moffitt (1985), and Katz and Meyer (1990) examined the effects of increased benefit generosity on the incentives that unemployed individuals face in searching for a new job. These studies used cross-state differences in unemployment insurance generosity and generally found that a 10 percent increase in UI benefits raises unemployment duration by between 4 and 8 percent. One potential drawback of these early estimates is that cross-state variation in unemployment insurance is likely endogenous. Several subsequent studies, such as Card and Levine (2000), Card et al. (2007), Rothstein (2011), Schmieder et al. (2012), and Farber and Valletta (2013) have used more plausibly exogenous sources of variation in unemployment insurance benefits at the individual level and find more modest impacts of UI generosity on the probability of re-employment for job seekers.

By estimating the effect of unemployment insurance generosity using only variation in UI eligibility within labor market areas, Rothstein (2011) and others are estimating what I will refer to as the *micro* effect of unemployment insurance. This micro effect is equal to the effect of a benefit change for one worker on that worker's own probability of employment, holding the behavior of all other economic agents fixed. In contrast, this paper will focus on estimating a parameter that I will call the *macro* effect of unemployment insurance, which measures the broad effect of changing unemployment insurance for the whole workforce on the level of employment. This macro effect includes not only the micro effect from increased moral hazard but also includes the effects of unemployment insurance on non-UI-eligible workers, firms' choices to create or destroy jobs, and changes in aggregate demand associated with increased government outflows. Since the macro

effect captures both the micro effect and these additional channels, the macro effect cannot be measured by simply scaling up well-identified estimates of the micro effect.

One approach to estimating the macro effect of unemployment insurance would be to separately estimate the effect of unemployment insurance on each factor included in the macro effect and simply sum these estimates. As mentioned before, the micro effect measured by Rothstein (2011) and Farber and Valletta (2013) during the Great Recession was found to be negative but very small, with an implied elasticity for employment with respect to UI duration on the order of -0.001. Research by Lalive, Landais, Zweimüller (2015) using a natural experiment in Austria where unemployment insurance was extended for certain workers finds that UI extensions have large positive effects on the employment of workers ineligible for the extensions. For vacancy creation, Marinescu (2014) measures the response of job postings on CareerBuilder.com to unemployment insurance extensions in the United States during the Great Recession, finding slightly negative but insignificant responses to the change in benefits. DiMaggio & Kermani (2015) use variation in the generosity of UI benefits across counties in the United States to measure the automatic stabilizer effects of unemployment insurance and find evidence that more generous UI benefits raise aggregate demand. They estimate that higher UI benefits reduce the sensitivity of car sales and non-tradeable employment to labor demand shocks, but find no effect for tradeable employment. Putting these together, the micro effect and the vacancy creation effect both appear to be small but slightly negative, while estimates of the spillover effects on non-UI-eligible workers and the increase in aggregate demand from unemployment insurance are both positive and substantial. Based on the combination of these estimates, one might expect that the macro effect would either reduce employment by less than the micro effect, or possibly lead to greater employment.

A different approach to estimating the macro effect is to identify a natural experiment where unemployment insurance generosity was determined exogenously and examine the effect on employment. In doing so, this estimates the total effect of all channels simultaneously. Research by Hagedorn et al. (2013) and Hagedorn et al. (2015) follow this approach using variation in the maximum potential duration of unemployment insurance induced by Federal policies in the US during the Great Recession and both studies conclude

that more generous unemployment insurance reduces employment. Both studies rely on a county-border-pair research design, where UI duration and employment are compared between two counties on opposite sides of a state border. Based on the assumption that economic geography is smooth, both counties would be subject to the same set of economic shocks in the absence of any changes in unemployment insurance. If unemployment insurance goes up in one state due to a shock to a different part of the state, the potential duration of UI in one county would go up relative to the other county and any difference in subsequent employment growth could be attributed to the change in UI duration. Hagedorn et al. (2013) uses variation in the rollout of UI extensions during the Great Recession and employs a regression specification based on the predictions from a standard search and matching model, while Hagedorn et al. (2015) uses the surprise expiration of benefit extensions at the end of 2013 in a differences-in-differences specification. Both studies find similarly negative effects of UI extensions, with a 1 percent increase in UI duration reducing employment growth by 0.016 percent to 0.032 percent.¹

Estimates of the macro effect in the literature are therefore conflicted, with the approach estimating channels separately suggesting that the total macro effect is less negative than the micro effect and quite possibly positive while the approach estimating all channels together estimates the macro effect to be much more negative than the micro effect. In theory these different approaches should return similar estimates, so this conflict represents somewhat of a puzzle.

2.2 UI Extensions during the Great Recession

Unemployment insurance in the US is operated jointly by State and Federal governments working together to provide benefits to workers who become unemployed through no fault of their own. The benefits funded by State governments are typically provided for up to 26 weeks, with some states offering slightly more or fewer weeks of benefits. The Federal government provides additional benefits for the average worker through two

¹The 0.032 percent reduction refers to the effect of a permanent 1 percent increase in UI duration as reported in Hagedorn et al. (2013). Their results also imply smaller reductions for temporary changes, with a 1 percent increase lasting a year reducing employment by 0.012 percent.

sources: the Extended Benefits (EB) program and emergency measures enacted during recessions. The EB program was enacted in 1970 to allow for 13 additional weeks of benefits in states experiencing temporarily high local unemployment rates with funding split half and half between the Federal and State governments. In addition, the Federal government has adopted a practice of extending unemployment insurance using emergency measures during national recessions as a part of its broader countercyclical policy dating back to the 1958 recession. In 1974, a fully-Federal-funded extension of about 13 weeks was put in place in almost all states. Similar emergency measures were enacted in 1982, 1991, 2002, varying somewhat in the structure but generally providing about 13 extra weeks of fully-Federal-funded benefits for most states with additional weeks available for the hardest hit states. The combination of the EB program and these emergency measures allowed states severely impacted by these recessions to extend benefit duration up to 40 weeks in some cases beyond the regular duration.

During the Great Recession, the Federal government expanded upon previous extensions in two ways. First, the American Reinvestment & Recovery Act of 2009 provided 100 percent Federal funding for the EB program in addition to the emergency measures. Second, while the initial form of the Emergency Unemployment Compensation of 2008 (EUC08) program was similar to past recessions, it was subsequently expanded far beyond prior UI extensions. The initial structure provided an extension of 13 weeks for all states, but was later revised to include four tiers of benefits with unemployment thresholds tied to the higher tiers. The thresholds and number of weeks available for each tier were revised several times over the course of the program, but a state with an unemployment rate exceeding the thresholds for all tiers could provide up to 53 additional weeks of benefits funded by the Federal government at the height of the program. In combination with regular unemployment insurance and the EB program, the maximum duration in several states reached 99 weeks during the deepest part of the Great Recession.

In the EUC08 program, the definition of “high unemployment” hinged on upon two different measures of unemployment, with each tier having a separate threshold for each measure. One measure, which I’ll refer to as the administrative measure of unemployment, used administrative data from the UI system to measure

unemployment by computing the ratio of individuals collecting regular unemployment insurance within a state divided by the total number of jobs paying into the UI system within that state. The second measure, which I'll call the survey-based measure of unemployment, comes from the Bureau of Labor Statistics' Local Area Unemployment Statistics (LAUS) program. This measure uses a time series model to combine a count of unemployed persons from the Current Population Survey with the administrative measure of unemployment from the UI system. In practice, this measure follows the CPS data fairly closely, which is why I refer to it as the "survey-based" measure. The time series model smooths out some of the month-to-month variation in the raw survey data, but otherwise more or less tracks the CPS.

While the EUC08 program had thresholds for both the administrative and survey-based measures of unemployment, in reality only the thresholds on the survey-based measure were binding. Almost all of the cases where a state triggered onto a higher level of benefits involved the survey-based measure crossing a threshold. Only about 4 percent of the cases met the administrative unemployment rate threshold for the second tier without meeting the survey-based threshold, and similar rates for the higher tiers were even lower. This phenomena arose because the thresholds for the administrative measure were set relatively high compared to the survey-based thresholds. When the survey-based unemployment rate was equal to 6 percent, which was the initial threshold for tier two, the administrative measure was on average 2.4 percent, significantly lower than its threshold of 4 percent. Since these two measures are strongly correlated, it was rare for a state to have an administrative measure above a particular tier's threshold without the survey-based measure also meeting its threshold.

These two different measures of unemployment tend to move closely together, as evidenced by a correlation of 0.63. However, there are several factors that can lead to differences between the two, including random sampling error, changes in multiple job holding, cross-state commuting, and changes in non-UI-claiming unemployment. Since the survey-based measure incorporates data from the CPS, random sampling error will affect its estimates of the unemployment rate since each state has a small sample of unemployed individuals. In addition, changes in multiple job holding could drive a wedge between the two estimates

since the administrative measure divides the number of UI claimants by the total number of UI-covered jobs, while the denominator of the survey-based measure counts the number of people in the labor force. Similarly, cross-state commuting would prevent the estimates from being identical since an individual would be counted in the administrative measure in the state where he or she worked, but in the survey-based measure would be counted in the state where he or she lived. Finally, changes in non-UI-claiming unemployment, which includes changes in the number of unemployed people who have exhausted regular UI benefits, would affect the survey-based measure but not the administrative unemployment rate. In the discussion of the research design below, I argue that random sampling error is the most substantial factor leading to a discrepancy between these two measures.

3 Data

I collect data on both the administrative and survey-based unemployment rates from weekly trigger notices released by the Department of Labor. These data represent the real-time unrevised estimates used by the Department of Labor to determine which tiers and programs each state was eligible for. Using the unrevised data is important for my research design as it allows me to reconstruct the unemployment rate tests each state was subject to and determine which eligibility criteria were met and which were not. The real-time administrative measure is equal to the 13-week moving average of the Insured Unemployment Rate computed by dividing the total number of individuals claiming regular UI benefits by the total number of UI-covered jobs. The real-time survey based measure is the three-month moving average of the seasonally-adjusted LAUS unemployment rate. In practice, the data were not quite “real-time” in that eligibility in a particular week was based on data from prior weeks. The administrative-based measure was almost always the value from the prior week, but the survey-based measure sometimes reflected data collected many weeks or even a few months prior. I aggregate both of these measures to the quarterly level and seasonally adjust them on a state-by-state basis using an X-11 procedure applied to the log unemployment rate.

I construct data on the maximum duration of unemployment insurance between 2005 and 2014 from

two sources. First, I tabulated the UI extension eligibility for each state from the Department of Labor’s weekly trigger notices along with the duration of UI associated with each EUC tier. Using this I compute the amount of extra weeks of UI available in each state from both EUC and EB at a weekly frequency and then aggregate to a quarterly frequency. I smooth through several periods where benefits lapsed temporarily but were subsequently re-authorized and benefits were paid for the lapsed period retroactively. I add this to the number of weeks available under the regular unemployment insurance system taken from Farber & Valletta (2013). The quarterly average of UI duration I use is very close to the point-in-time durations listed in Table 1 of Hagedorn et al. (2015), Table 1 of Bradbury (2014), and available on the Department of Labor’s website.²

To measure the effects on employment, I use data on state-by-state counts of UI-covered jobs at a quarterly frequency from the Quarterly Census of Employment and Wages (QCEW). The QCEW data is built from administrative data on quarterly unemployment insurance filings by employers. The use of administrative data means the QCEW is a less noisy measure of state-by-state employment than sources based on survey data, such as the LAUS data. The LAUS measure of employment combines data from a nationwide survey of employers (the Current Employment Survey, or CES) and the CPS household survey. While the LAUS data has the advantage of being more timely, with data released sometimes a year in advance of the QCEW, its reliance on survey data poses potential problems for state-level analysis. I discuss this point further in Section 6 when comparing my results from the QCEW data with other analyses relying on LAUS data. I seasonally adjust the QCEW measure of employment on a state-by-state basis using an X-11 procedure in logs.

²The most recent data released by the Department of Labor can be found at the following link: http://www.ows.doleta.gov/unemploy/docs/potential_weeks_map.pdf

4 Research Design

I divide the research design into two phases. First, I demonstrate how to identify random sampling error in the survey-based measure of unemployment. Second, I use this random sampling error as part of a natural experiment, comparing employment growth among states with the same economic conditions but assigned to different extensions of unemployment insurance.

4.1 Random Sampling Error

The survey-based unemployment rate and administrative unemployment rate frequently differ in both absolute and relative terms. As an example, Figure 1 shows both unemployment rates for the state of New York over the 2005-2014 period. The survey-based unemployment rate has a higher level than the administrative measure and also seems to exhibit more volatility. During the Great Recession, the survey-based unemployment rate rose from about 5 percent at the end of 2007 to nearly 9 percent by 2010 and hovered near 8 percent for several years until falling rapidly in 2013 and 2014. Over the same time period the administrative measure rose from 2 percent to 4.5 percent at the onset of the recession, but steadily fell after peaking at the beginning of 2009. Figure 1 also shows the raw unemployment rate as computed in the CPS, which has been seasonally adjusted to be comparable to the other two measures. The Bureau of Labor Statistics uses a time series model to estimate the survey-based unemployment rate, based on both the administrative measure and the raw survey data. Although it uses both data sources, in practice the survey-based measure appears to predominantly follow the CPS data, not only sharing a similar level but also exhibiting similar dynamics. In New York, the survey-based measure fell in 2011 before rising in 2012 along with the CPS series, while the administrative unemployment rate fell steadily over this time period and did not increase in 2012. While these two series could potentially differ for several reasons, I argue that the main reason for discrepancies like this case is random sampling error. Since the survey-based measure is predominantly based on data from the CPS it will inherit variation in measured unemployment due to sampling error.

The Bureau of Labor Statistics releases standard errors for the survey-based measure of unemployment.

While these standard errors vary across states and over time, they typically produce 95% confidence intervals for the estimates that are between 1 and 2 percentage points (p.p.) wide. In the case of New York state, the confidence intervals around the unemployment rate estimates are not able to reject a consistent decline between 2010 and 2013, instead of the decline and rise and decline suggested by the point estimates. Even in as populous a state as New York, the sample size of the CPS is simply too small at a state level to obtain precise estimates of the unemployment rate. The uncertainty around the survey-based estimates is too large to distinguish between New York suffering somewhat of a double-dip recession in 2012 or experiencing a consistent decline in unemployment over the same period. However, the administrative measure of unemployment is not subject to random sampling error, and indicates precisely how unemployment changed in New York state in 2012. The administrative measure shows no signs of a double-dip recession, exhibiting a consistent decline over this entire period.

More direct evidence of random sampling error is available by looking at the autocorrelation of the discrepancy between these two series. Since the structure of the CPS is that households are interviewed for four months and then rotate out for eight months before rotating in for another four months, I examine the autocorrelation to see whether it picks up at a horizon of about a year. This prediction is driven by the fact that two samples of the CPS taken a year apart should have about half of the households in common (without accounting for attrition), while there will be no overlap between samples six months apart. I estimate the following regression

$$\log(\text{Survey UR}_{s,t}) = \theta_s + \sum_{k=0}^8 \log(\text{Admin. UR})_{s,t-k} + \epsilon_{s,t}$$

where I regress the survey-based measure on the administrative unemployment rate and eight lags of the administrative measure with state fixed effects. I then compute the residuals from this regression and calculate the autocorrelation of these residuals at a quarterly frequency. The autocorrelations are shown in Figure 2, along with autocorrelations from repeating the same exercise with the raw CPS data. Both data sources show autocorrelations that decline with the horizon used, which likely reflects both changes in unemployment among households remaining in the sample over time as well as attrition of households from month-to-month

in the CPS. The unusually high autocorrelation at a yearly horizon is especially clear in the CPS data, and is somewhat more muted in the survey-based estimates, owing to the smoothing effects of the time series model used to produce these estimates.

Another way to test the random sampling error explanation for the discrepancy between the two unemployment rates is to measure whether the discrepancy is informative about labor market conditions. If the discrepancy is due to random sampling error, it will not be predictive of future labor market conditions, while if the discrepancy were due to changes in non-UI-claiming unemployment it could predict future growth in employment. The survey-based unemployment rate on its own conveys substantial information about labor market conditions. Figure 3 shows a binned scatterplot of the log survey-based unemployment rate against the employment growth in the subsequent quarter with a downward-sloping line of best fit, indicating that higher unemployment rates are associated with slower subsequent employment growth. This relationship is statistically significant at the 1 percent level. Figure 4 shows a similar plot, with the only difference being that both the log survey-based unemployment rate and the subsequent employment growth rate were residualized on the log of the administrative measure and eight lags prior to plotting. This graph shows no systematic relationship between the survey-based measure and employment growth and although the line of best fit is slightly downward sloping it is not close to being statistically significant. These plots taken together suggest that conditioning on the administrative measure removes any informative content of the survey-based measure, leaving only variation due to random sampling error.

4.2 Measuring the Effect on Employment Growth

The variation in the survey-based measure of unemployment due to random sampling error was substantial enough to drive variation in unemployment insurance policy during the Great Recession. In New York state, the survey-based measure shot up to 8.8 percent in 2009, triggering on all four tiers of EUC extensions. When the survey-based measure fell below the 8.5 percent threshold in 2010, it switched off of Tier 4 and the potential duration of unemployment benefits fell by 6 weeks. When the survey-based measure rose again in

2012 and crossed the 8.5 percent threshold for a second time, New Yorkers became eligible for Tier 4 benefits again until the survey-based measure fell back below the threshold at the end of 2012. The double-dip pattern seen in the survey-based measure was likely due to random sampling error, but led directly to changes in the duration of unemployment insurance available in the state at different times, creating a natural experiment.

I measure the effect of UI duration on employment by comparing employment growth in states with similar underlying economic conditions but randomly assigned to different levels of EUC benefits. This experiment can be thought of in two different ways. The first is as a random experiment where I compare states with the same administrative unemployment rate but assigned to different UI durations through random sampling error. The second is as a regression discontinuity approach where I exploit the sharp thresholds of the EUC program and compare states with survey-based unemployment rates slightly above and slightly below the threshold. These two approaches are conceptually related, but rely on distinct sources of variation and so I present them separately.

The experimental approach involves comparing states with the same measured administrative unemployment rate, but assigned to different UI durations by random sampling error. Since the administrative unemployment rate is measured without any sampling error, it should be a sufficient statistic for a state's underlying economic conditions. Two states with the same administrative unemployment rate may have very different survey-based unemployment rates because of random sampling error, so they could end up on different sides of the threshold for a particular EUC tier and offer different durations of UI as a result. The confidence intervals reported by the Bureau of Labor Statistics for the survey-based unemployment rates are typically between 1 and 2 p.p. wide, so states frequently ended up on different sides of the EUC thresholds when the survey-based measure was near the threshold even though they had the same underlying economic conditions. To estimate the elasticity of employment with respect to UI duration from this experiment, I estimate the following regression

$$\Delta \log(\text{Employment})_{s,t+1,t} = \theta_s + \delta_t + \beta \text{UI Duration}_{s,t} + \sum_{k=0}^8 \gamma_k \text{Administrative Unemp. Rate}_{s,t-k} + \epsilon_{s,t} \quad (4.1)$$

where I instrument for UI duration with dummy variables for whether a state was triggered onto a each tier of EUC. I include state and time fixed effects in the regression and cluster at the state level. Conditional on the administrative measure, whether a state was triggered onto a particular tier of EUC is dependent only on random sampling error and so this instrument is valid when conditioning on the administrative measure.

The regression discontinuity design uses the sharp cutoff at the thresholds for each EUC tier as a form of random assignment. States just above the cutoff have unemployment rates that are nearly the same as states just below the cutoff, but different durations of UI since the state above the threshold is eligible for the higher EUC tier. Comparing their subsequent employment growth estimates the effect of the additional weeks of UI on employment. This takes the form of the following regression

$$\Delta \log(\text{Employment})_{s,t+1,t} = \theta_s + \delta_t + \beta \text{UI Duration}_{s,t} + \sum_{k \in \{2,3,4\}} \gamma_k \mathbf{1}(|\text{Survey Unemp.} - \text{Threshold}| \leq 1) + \epsilon_{s,t} \quad (4.2)$$

where I instrument for UI duration with dummy variables for whether a state is above or below the threshold for each tier k , conditional on its unemployment rate being close to the threshold (I use 1 p.p. as the bandwidth). This gives three instruments for UI duration, since tiers two, three, and four each had a threshold for at least one quarter, while states were always eligible for tier one benefits at any unemployment rate. I control for quarter and state fixed effects as well. Conditional on being close to the threshold, states have similar economic conditions so whether a state ends up above or below the threshold could be as-if randomly assigned.

Both the natural experiment and regression discontinuity designs compare employment growth between state with similar conditions but with different durations of unemployment insurance. The natural experiment compares states with the same administrative unemployment rate assigned to different levels of UI through random sampling error. In contrast, the regression discontinuity compares states with similar survey-based unemployment rates assigned to different levels of UI based on the sharp cutoff. Both designs identify the same parameter and should return similar estimates, the difference between them lies in which set of states

is being compared. I also present results where I combine both approaches by estimating the regression discontinuity design conditional on the administrative unemployment rate, which compares states slightly above and slightly below the threshold with the same underlying economic conditions according to the administrative measure.

5 Results

I start by outlining the results from the natural experiment and regression discontinuity designs before discussing several robustness checks.

The effect on employment growth can be examined non-parametrically by looking at Figures 5 and 6. Each graph plots annualized employment growth one period into the future against the present unemployment rate and shows separately states on a higher tier of EUC (orange dots) and states not eligible for the higher tier (blue dots). Figure 6 shows this plot for the regression discontinuity research design, with the difference between survey-based unemployment rate and the nearest EUC threshold shown on the x-axis. As the survey-based unemployment rate crosses the threshold, there's a clear change in the tier of EUC as evidenced by the triggered-off states falling entirely to the left of 0 and the triggered-on states falling almost entirely on the right.³ However, at the threshold, there does not appear to be a clear break in employment growth. The orange and blue lines show a local kernel-weighted mean for states triggered on and off the higher EUC tier, respectively. The estimates suggest that if anything, there could be a slight increase in employment growth associated with the unemployment rate crossing the threshold and enabling higher benefits, although the 95% confidence intervals of the two lines overlap.

In contrast, Figure 5 shows the natural experiment design with the predicted survey-based unemployment rate on the x-axis. Here there is substantial overlap of the orange and blue dots since random sampling error can substantially alter the estimated survey-based unemployment rate and assign states with the same

³A few states are shown to be triggered onto EUC despite having unemployment rates below the threshold. This is due to a provision in the program that a state triggering onto a new tier of EUC would remain on the new tier for at least 13 weeks before it could trigger off.

administrative measure to different levels of the EUC program. There is some limit to the extent of random sampling error, however, as the density of the opposite color dots dissipates as the predicted survey-based unemployment rate gets farther away from the threshold, with few orange dots at the far left side and few blue dots on the far right side. In the middle, though, there are substantial numbers of states whose survey-based unemployment rate was predicted to fall below the threshold but the actual estimate was above the threshold, or vice versa. As before, the orange and blue lines show a local kernel-weighted mean for states triggered on and off the higher EUC tier, respectively. Comparing the lines vertically shows the difference in expected employment growth for two states with the same administrative unemployment rate but assigned to different levels of UI benefits through random sampling error. As in the regression discontinuity design, this plot indicates that higher UI durations correspond to higher rates of employment growth, although the lines intersect at some points and the 95% confidence intervals overlap for much of the data.

As described before, each of these research designs can be estimated in a regression framework. Table 1 shows the results from estimating Equations 4.1 and 4.2 in the first two columns. As expected from the non-parametric plots, both regressions suggest a positive elasticity of employment with respect to unemployment insurance duration of around 0.01 to 0.015. This is a sizeable elasticity, implying that a doubling of unemployment insurance duration would raise the employment growth rate between 1 and 1.5 percentage points annually relative to the counterfactual. However, the standard errors are quite wide and the 95 percent confidence interval for both regressions does include an elasticity of zero. Nevertheless, these estimates do conclusively rule out substantial negative elasticities. The third column of Table 1 shows the results from a regression combining both sources of variation in a regression discontinuity design conditional on the administrative unemployment rate, which produces a similar estimate of the elasticity. This estimate is statistically significant at the 5 percent level, although the confidence interval only barely excludes an elasticity of zero. Columns four and five add polynomials of the survey-based measure as additional controls to the regression discontinuity specifications from columns two and three respectively. These produce similar point estimates, although the results are quite noisy and the standard errors are substantially wider.

In Table 2 I show the results of applying the combined regression discontinuity and natural experiment designs to employment in different industries. The first column uses only employment in manufacturing as a measure of employment in tradeable industries, while the second column uses employment in trade (including retail), transportation, utilities, leisure, and hospitality as a measure of non-tradeable employment. In theory, since demand for tradeable goods and services is nationwide, effects on aggregate demand at the state-level should not substantially affect tradeable employment. Therefore, the difference between the response of non-tradeable employment and tradeable employment should give an estimate of the aggregate demand effects of UI extensions. The estimates in column one show that the elasticity of manufacturing employment is only 0.001, which is an order of magnitude smaller than the estimates of the effect on overall employment found in Table 1. However, the elasticity of non-tradeable employment is 0.009 as shown in column two, which is consistent with the effects on overall employment. I present separately the elasticity for construction employment in column three and for all other industries in column four. Construction employment appears to respond an order of magnitude more to UI extensions than overall employment, while the effect on all other industries is largely in line with the effects measured in Table 1.

These estimates offer evidence about the effectiveness of the Federal UI extensions put in place during the Great Recession, but the effects of other programs in other time periods could differ. The Great Recession was a period in which the fiscal multiplier of government spending has been documented to be high (Chodorow-Reich et al., 2011; Leduc and Wilson, 2012). Since interest rates in the US were up against the zero lower bound for much of the period used in this analysis and conventional monetary policy was constrained, unemployment insurance extensions may have been unusually effective at stimulating employment growth (Kekre, 2015). The evidence comparing the effects on manufacturing employment versus food, leisure, and trade is consistent with UI extensions raising aggregate demand during this period, but a UI extension enacted during a different point in the business cycle could have less of an effect on aggregate demand. Furthermore, these extensions were meant to be temporary measures targeted to particularly distressed areas, so the effect could be different for permanent changes if forward-looking firms and workers respond more strongly to a

permanent policy change.

6 Comparison with Prior Literature

A positive effect of UI extensions on employment is consistent with the literature that estimates separately the effects on UI-covered workers, non-UI-covered workers, vacancy creation, and aggregate demand. Taken together, the estimates from the literature of each of these channels are consistent with a positive effect of UI on employment, mainly coming from increased aggregate demand and spillovers on non-UI-covered workers. The results on differential effects across industries provide support for the substantial aggregate demand effects uncovered in DiMaggio & Kermani (2015). However, the results from my analysis are not consistent with prior estimates of the macro elasticity from Hagedorn et al. (2013) and Hagedorn et al. (2015), which find substantial negative elasticities. Importantly, these papers analyze the same setting and time period I examine, ruling out external validity as an explanation for this discrepancy.

Hagedorn et al. (2013, 2015) represent a substantial advance in the literature on unemployment insurance for highlighting the difference between the macro and micro elasticities and for using state-of-the-art empirical techniques to address potential threats to identification. While much of the literature has focused on individual-level variation in UI duration, Hagedorn et al. (2013, 2015) correctly argue that this misses many effects of unemployment insurance and an estimate of the full macro effect of UI extensions is necessary for determining UI policy. The authors take many steps to ensure that their estimates are well identified, including robustness checks, placebo tests, and estimates narrowing in on the channels through which the effect occurs. Both studies have brought attention to the macro effects of unemployment insurance and raised awareness about the importance of using smaller-scale natural experiments to study policy changes at an aggregate level in macroeconomics. I carefully reviewed both studies and present below an attempt to reconcile the differences in estimates with my results from the natural experiment above.

I am able to replicate the estimates produced in both papers. Additionally, their estimates are not substantially altered by using fixed effects and clustered standard errors instead of the interactive effects

structure of Bai (2009) and block bootstrapping which the authors used in both studies. I find that the results using the quasi-difference estimator employed in Hagedorn et al. (2013), which controlled for the impact of future benefit changes, are similar to estimates using the dependent variable in levels but instead directly controlling for leads of the difference in UI benefits across counties in a border pair. I find that the discrepancy between the positive elasticity estimated from the natural experiment above and the negative elasticity estimated by Hagedorn et al. (2013, 2015) is best explained by a combination of three factors: violation of parallel trends assumptions, statistical noise in data sources based on surveys, and non-smoothness of economic geography.

In the first paper, Hagedorn et al. (2013) use the rollout of UI extensions during the onset of the Great Recession through the beginning of 2012 to examine the effect on many labor market measures, including employment. In particular, they use changes in the difference of UI duration between counties on opposite sides of a state border so that any nationwide increase in UI benefits is netted out and thus they use only increases in UI duration that happen in one state but not another in a given quarter. For the UI extensions active during the Great Recession, these changes in UI duration only occur when a state crosses an unemployment rate threshold. In this case, a potential worry is that a negative shock to labor markets could trigger an increase in UI duration and lead to lower employment, which the research design could pick up as an effect of unemployment insurance even if the UI extension had no causal effect. To test the extent of this concern, I estimate a differences-in-differences regression at the state level using the specification:

$$\log(\text{Employment}_{s,t}) = \theta_s + \delta_t + \sum_{k=-8}^3 \beta_k \log(\Delta \text{UI Duration}_{s,t-k}) + \nu_{s,t}$$

This regression shows the measured effect of an increase in UI duration flexibly over time by estimating the effect separately for each quarter before and after the change in benefits for up to eight quarters before the change and four quarters afterwards (the change in UI occurs between period -1 and period 0). I use employment as measured in the LAUS data to be consistent with the original paper. I plot the coefficients β_k in Figure 7 to examine how employment changes around the change in UI benefits. The graph shows

that employment begins to decrease on average about a year before the change in UI duration occurs and continues to decrease for several quarters after the change. However, employment appears to decline at about the same rate before and after the change in UI benefits. A differences-in-differences regression that only examined the change in employment after the change in UI duration might conclude that this is evidence that UI extensions reduce employment, while in fact employment tended to begin declining well before a state triggered onto a higher UI extension. This is evidence that the parallel trends assumption inherent in the differences-in-differences approach would not hold for this source of variation and motivates the use of a more exogenous source of variation.

The specification used by Hagedorn et al. (2013), where the dependent variable is expressed as a quasi-forward difference, could potentially correct for differing parallel trends if county-level labor markets behave as predicted by a standard search and matching model. The specification controls for the equilibrium level of unemployment one period into the future, discounted by the discount factor and the probability of separation, which accounts for the effect that future benefit changes have on the current equilibrium. To evaluate the plausibility of this technique, I conduct a placebo test during a period in which UI policy was constant. Between the beginning of 1996 and the end of 2001, no changes to UI duration occurred in the continental United States. I create a placebo UI extension by pretending that UI duration was increased 13 weeks whenever the survey-based unemployment rate rose above 5 percent. When I regress the quasi-forward difference of unemployment on the placebo UI duration differencing between counties on either side of state borders, I estimate a positive and statistically significant coefficient on the same order of magnitude as the coefficient reported in the original paper. Based on this finding, it seems that the quasi-forward difference specification is not correctly accounting for labor market dynamics and is picking up movements in unemployment that are not related to UI duration. The failure of the quasi-forward difference specification in this placebo test also implies that it is not an appropriate method for correcting the non-parallel trends problem.

The paper by Hagedorn et al. (2015) focuses on a surprise shock to UI benefits created by the lapse in funding for the EUC program at the end of 2013. The benefits paid out by the EUC program needed to

be intermittently re-authorized by Congress and were successfully re-authorized more than a dozen times before the end of 2013. On January 1st, 2014, the EUC program abruptly ended and Congress did not act to re-authorize the program, causing UI duration to fall nationwide back to 26 weeks in most states. Although the change occurred all at once, there was substantial variation in the size of the change across states due to different states being eligible for different tiers of the EUC program. Before the lapse in funding, Illinois, Michigan, Nevada, and Rhode Island had still been eligible for all four tiers of EUC benefits since their survey-based unemployment rates still exceeded 9%, while the unemployment rates in some other states had fallen so much that they were only eligible for the first tier of benefits (which all states received automatically). Using the variation in the change in UI duration caused by the lapse in funding, I conduct the same differences-in-differences design as above and plot the β_k coefficients in Figure 8. As before, I use the LAUS measure of employment to be consistent with the original paper. The plot shows the elasticity of employment with respect to an increase in UI duration, so in the discussion that follows I flip the sign of each coefficient to examine the effect of a decrease like the lapse in EUC funding. Consistent with the results in the original paper, the differences-in-differences regression estimates that employment increased over the four quarters following the benefit cut. However, examining the coefficients before the change in UI shows that employment was also higher about a year before the benefit cut. This plot is consistent with a mean-reverting shock lowering employment over the course of 2013 relative to states with lower UI duration before employment returns to its higher value over the course of 2014. Since UI duration was set according to the survey-based measure of unemployment, a mean-reverting shock to a certain set of states in the form of random sampling error could have lowered measured employment and raised measured unemployment in 2013, triggering a higher tier of benefits in those states (or preventing higher tiers from triggering off), which subsequently reverted in 2014 as new participants rotated into the survey and the survey-based measure of employment returned to its prior level relative to other states.

Hagedorn et al. (2013, 2015) use the measure of employment from the LAUS, which combines counts of employment from a nationwide household survey (CPS) and a survey of employers (CES). The CPS is also

used to construct the survey-based measure of unemployment, so a potential concern is that this measure of employment could be correlated with UI duration mechanically due to random sampling error affecting both. As shown in the differences-in-differences plot in Figure 8, this seems like a plausible concern with estimating a differences-in-differences design around the lapse in benefits at the end of 2013. In contrast, the measure of employment from the QCEW should not be subject to such a concern, since it is built from administrative data and would not be correlated with random sampling error in the CPS. Figures 9 and 10 show the same differences-in-differences plots as Figures 7 and 8 but using QCEW employment as the dependent variable instead of LAUS employment. The QCEW data does little to solve the problematic pre-trends in the first analysis which uses the variation in UI duration caused by states triggering onto different tiers, but it does show much better pre-trends in the analysis using the EUC expiration. In this case, using the QCEW data and the EUC expiration as the source of variation produces an elasticity estimate of -0.009, although this is not statistically different from zero. However, even here there is a concern about the parallel trends assumption being violated, since a national increase in demand in 2014 could have led to higher employment growth in states with more cyclical slack, which would correspond to the states with higher UI duration since UI duration was higher in states with higher unemployment. Although the states that experienced a large drop in UI in the EUC expiration had mostly similar employment growth in 2012 and 2013 as states that experienced smaller drops as a result of the expiration, it is reasonable to expect that counterfactual growth could have been different in 2014 even absent a change in unemployment insurance policy.

These concerns that benefit changes were possibly endogenous to pre-existing trends in economic conditions or measurement error could be alleviated by using a county border pair research design if economic conditions are smooth. Hagedorn et al. (2013, 2015) use a county border pair research design in both analyses. Figure 11 shows the 1,143 counties in the US that share a land border with another state. Assuming for the moment that two neighboring counties on opposite sides of a state border are subject to the same set of economic shocks in the absence of any change in UI policy, using either of the differences-in-differences designs outlined above could be made valid by running the regression on county-level employment and conditioning

on employment in the county across the border. Even if UI duration changed in a state due to a declining trend in economic conditions, if the county across the border were subject to all of the same economic shocks as the county in the state with changing UI policy then the differences-in-differences design would estimate the causal effect of UI extensions. The key assumption that allows this design to be valid is that economic geography is smooth enough that conditioning on the county across the border fully removes the influence of state-level shocks. This results in UI being the only factor driving a wedge between employment in counties on the opposite sides of state borders. If instead there are non-UI state-level economic shocks that stop at the border, the county border pair design is not going to be able to distinguish between an effect of UI and state-level shocks correlated with UI duration.

One way to test this assumption that counties are not affected by non-UI state-level shocks after conditioning on the county across the border is to run a regression of county employment on state-level employment controlling for employment in the county across the border. I do this in the QCEW where the data for counties are constructed directly from the administrative data without using state-level aggregates to impute county data. To further prevent any mechanical correlation between states and counties, I use a leave-one-out measure of employment in the state where I compute for each county the total employment in all other counties in the same state. Both the state and county data are seasonally adjusted using the same procedure. Formally, I estimate the regression

$$\log(Y_{i,t}) = \theta_{ij} + \delta_t + \beta_1 \log(Y_{s(i),t}) + \beta_2 \log(Y_{j,t}) + \epsilon_{ijt}$$

where counties i and j share a border across state lines and $s(i)$ is the state that county i is located in. I include border-pair and time fixed effects. Figure 12 shows a binned scatterplot of this regression using yearly growth rates of employment with state-level data on the x-axis and the county measure on the y-axis. Both variables were residualized on the yearly growth rate of employment in the county across the border as well as border-pair and time fixed effects prior to plotting. Even so, there is a strong relationship between state employment growth and county employment growth. The dotted line in the figure shows the

45 degree line, which would correspond to perfect pass-through of state level shocks to counties. Ignoring the outliers, the relationship is only slightly weaker than perfect pass-through, and even including the outliers the relationship is highly statistically significant. Table 3 shows the coefficients corresponding to this regression and also shows the results of using log employment in levels or one quarter differences instead. Across all specifications, the coefficient on the state-level measure is highly statistically significant, and greater than the coefficient on employment in the county across the border. It is important to note that for this analysis I exclude any years in which a UI extension program was active and thus this relationship must have been driven by non-UI state-level shocks.⁴

These results imply that the county border pair research design does not adequately solve the non-parallel trends problem in this case. Conditioning on the county across the border is not sufficient to control for economic shocks, since there are non-UI state-level shocks that significantly affect county-level employment. This evidence of non-smooth economic geography is potentially consequential, since county border pair research designs have been used in many contexts, including the examination by Holmes (1998) of the effects that right-to-work laws have on manufacturing and the study by Dube et al. (2010) measuring how minimum wage laws affect labor markets. However, the reverse causality problem in this setting is much more mechanical than either of those analyses. In this case, state-level shocks directly lead to a change in UI duration as soon as the Bureau of Labor Statistics determines that the unemployment rate exceeded the threshold. With right-to-work laws and minimum wage laws, the response to a shock is less immediate. Legislatures may respond to an increase in unemployment when deciding on what laws to pass, but they are unlikely to strictly follow a threshold rule as the Bureau of Labor Statistics does. In the case of UI extensions, the mechanical nature of the policy makes any connection between state- and county-level employment particularly problematic, but the evidence presented here is less informative about the validity of other studies relying on county border pair research designs.

⁴Specifically, I drop years 1994-1995, 2002-2004, and 2008-2014.

7 Conclusion

In this paper, I have used a natural experiment created by a Federal UI extension during the Great Recession to estimate the state-level impact on employment. A particular feature of this UI extension was that the program set UI duration based on the unemployment rate in each state, which was measured with a national household survey. By using administrative data I was able to isolate random sampling error in this survey, which allowed me to compare employment growth in states with the same underlying conditions that were randomly assigned to different durations of unemployment insurance. I estimate that higher UI duration raised the growth rate of employment with an elasticity of about 0.01. Although the estimates have large standard errors and I'm unable to rule out an elasticity of zero, I am able to reject substantial negative elasticities.

The prior literature on the macro effect of UI extensions is split on this topic. Papers that measure the different channels of the macro effects separately imply that the elasticity of employment to UI duration is likely positive due to substantial aggregate demand effects of UI extensions. In contrast, two studies that attempt to capture all channels of the macro effect at once estimate substantial negative elasticities. My results support the former set of estimates, both the finding that the total effect is positive and that the aggregate demand effects are substantial. After examining the latter set of papers, I find that the negative effects they estimate are possibly confounded by a confluence of three factors: violations of the parallel trends assumption, measurement error in data based on surveys, and non-smoothness of economic geography.

While this analysis sheds light on the macro effect of the EUC extension, the results may differ in other contexts. The EUC program existed only during a deep recession when conventional monetary policy was constrained and the multiplier on increased government spending was possibly quite high. Furthermore, the effects measured in this paper concern temporary increases in UI duration and may or may not be informative about permanent changes to the unemployment insurance system. Nevertheless, these results show that UI extensions are a positive force for stimulating the economy during recessions and motivate their continued use as a standard element of countercyclical fiscal policy in the United States.

References

- Bai, Jushan. “Panel Data Models with Interactive Fixed Effects.” *Econometrica* 77.4 (2009): 1229–1279.
- Chodorow-Reich, Gabriel et al. “Does State Fiscal Relief during Recessions Increase Employment? Evidence from the American Recovery and Reinvestment Act.” *American Economic Journal: Economic Policy* (2012): 118–145.
- Di Maggio, Marco, and Amir Kermani. “The Importance of Unemployment Insurance as an Automatic Stabilizer.” (2015).
- Dube, Arindrajit, T. William Lester, and Michael Reich. “Minimum Wage Effects across State Borders: Estimates Using Contiguous Counties.” *The review of economics and statistics* 92.4 (2010): 945–964.
- Farber, Henry S., and Robert G. Valletta. *Do Extended Unemployment Benefits Lengthen Unemployment Spells? Evidence from Recent Cycles in the US Labor Market*. National Bureau of Economic Research, 2013.
- Hagedorn, Marcus et al. *Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects*. National Bureau of Economic Research, 2013.
- Hagedorn, Marcus, Iourii Manovskii, and Kurt Mitman. *The Impact of Unemployment Benefit Extensions on Employment: The 2014 Employment Miracle?* National Bureau of Economic Research, 2015.
- Holmes, Thomas J. “The Effect of State Policies on the Location of Manufacturing: Evidence from State Borders.” *Journal of Political Economy* 106.4 (1998): 667–705.
- Katz, Lawrence F., and Bruce D. Meyer. “The Impact of the Potential Duration of Unemployment Benefits on the Duration of Unemployment.” *Journal of public economics* 41.1 (1990): 45–72.
- Kekre, Rohan. “Unemployment Insurance in Macroeconomic Stabilization.” (2015).

- Lalive, Rafael, Camille Landais, and Josef Zweimüller. “Market Externalities of Large Unemployment Insurance Extension Programs.” (2013).
- Leduc, Sylvain, and Daniel Wilson. “Roads to Prosperity or Bridges to Nowhere? Theory and Evidence on the Impact of Public Infrastructure Investment.” NBER Macroeconomics Annual 2012, Volume 27. University of Chicago Press, 2012. 89–142.
- Marinescu, Ioana. The General Equilibrium Impacts of Unemployment Insurance: Evidence from a Large Online Job Board. mimeo, University of Chicago, 2014.
- Moffitt, Robert. “Unemployment Insurance and the Distribution of Unemployment Spells.” *Journal of Econometrics* 28.1 (1985): 85–101.
- Rothstein, Jesse. “Unemployment Insurance and Job Search in the Great Recession.” *Brookings Papers on Economic Activity* 2 (2011).
- Schmieder, Johannes F. et al. “The Effects of Extended Unemployment Insurance Over the Business Cycle: Evidence from Regression Discontinuity Estimates Over 20 Years.” *The Quarterly Journal of Economics* 127.2 (2012): 701–752.
- Solon, Gary. “Work Incentive Effects of Taxing Unemployment Benefits.” *Econometrica: Journal of the Econometric Society* (1985): 295–306.

Tables

Table 1: Natural Experiment and Regression Discontinuity Results

	(1)	(2)	(3)	(4)	(5)
Dependent Variable	Annualized Employment Growth in the Next Quarter (QCEW)				
Log UI Duration (Weeks)	0.00805* (0.00484)	0.0166 (0.0102)	0.0183** (0.00871)	0.0146 (0.0255)	0.0446 (0.0274)
Instruments	EUC Status	Above/Below Thresholds			
Controlling for Administrative Measure and Eight Lags	YES	NO	YES	NO	YES
Controlling For Polynomials	NO	NO	NO	YES	YES
First Stage F Statistic	1364.85 ⁺	76.13 ⁺	77.71 ⁺	13.53 ⁺	13.06 ⁺
N	1,989	1,989	1,989	1,989	1,989
Number of Clusters	51	51	51	51	51
R^2	0.738	0.692	0.736	0.702	0.733

This table reports coefficient estimates and standard errors from five instrumental variables (IV) regressions. The dependent variable in all specifications is the growth rate of QCEW employment at an annual rate over the subsequent quarter and the coefficient of interest corresponds to the log of UI duration. Instruments vary across specifications, as indicated under the instruments section. Each specification includes state and time fixed effects. The statistic testing for weak identification in the IV regressions is the first stage F statistic and a “+” denotes whether the test statistic exceeds the 5% Stock-Yogo critical value assuming 10% bias. The weak identification statistics are computed under the assumption of homoskedasticity. Standard errors on regression coefficients are clustered by state. Triple, double, and single asterisks denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 2: Heterogeneity by Industry

	(1)	(2)	(3)	(4)
Dependent Variable	Annualized Employment Growth in the Next Quarter (QCEW)			
Industry	Tradeables	Non-Tradeables	Construction	All Others
Log UI Duration (Weeks)	0.000965 (0.0199)	0.00932 (0.00727)	0.117** (0.0500)	0.0127 (0.00958)
Instruments	Above/Below Thresholds			
Controlling for Administrative Measure and Eight Lags	YES	YES	YES	YES
First Stage F Statistic	77.48 ⁺	77.71 ⁺	77.59 ⁺	77.71 ⁺
N	1,981	1,989	1,985	1,989
Number of Clusters	51	51	51	51
R^2	0.624	0.668	0.624	0.498

This table reports coefficient estimates and standard errors from four instrumental variables (IV) regressions. The dependent variable in all specifications is the growth rate of QCEW employment in the specified industry at an annual rate over the subsequent quarter and the coefficient of interest corresponds to the log of UI duration. Instruments vary across specifications, as indicated under the instruments section. Each specification includes state and time fixed effects. The statistic testing for weak identification in the IV regressions is the first stage F statistic and a “+” denotes whether the test statistic exceeds the 5% Stock-Yogo critical value assuming 10% bias. The weak identification statistics are computed under the assumption of homoskedasticity. Standard errors on regression coefficients are clustered by state. Triple, double, and single asterisks denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table 3: Correlation Between State and County Employment

	(1)	(2)	(3)
Units	Logs	Quarterly Growth Rate	Annual Growth Rate
Dependent Variable	County-Level Employment (QCEW)		
State Leave-One-Out Employment (QCEW)	0.489*** (0.0407)	0.0538** (0.0213)	0.120*** (0.0439)
Opposite County Employment (QCEW)	0.200*** (0.0243)	0.0248*** (0.00319)	0.0623*** (0.00585)
N	130,360	127,764	120,259
R^2	0.314	0.002	0.011

This table reports coefficient estimates and standard errors from three least squares regressions on county border pairs. The dependent variable in each specification is specified unit of QCEW county-level employment and the coefficients of interest correspond to state-level employment and employment in the county across the state border. Each specification includes county border pair fixed effects. Standard errors on regression coefficients are clustered by county border pair. Triple, double, and single asterisks denote statistical significance at the 1%, 5%, and 10% level, respectively.

Figures

Figure 1: Unemployment Rates for New York State




Figure 2: Residual Autocorrelation, CPS and Survey-Based Measure

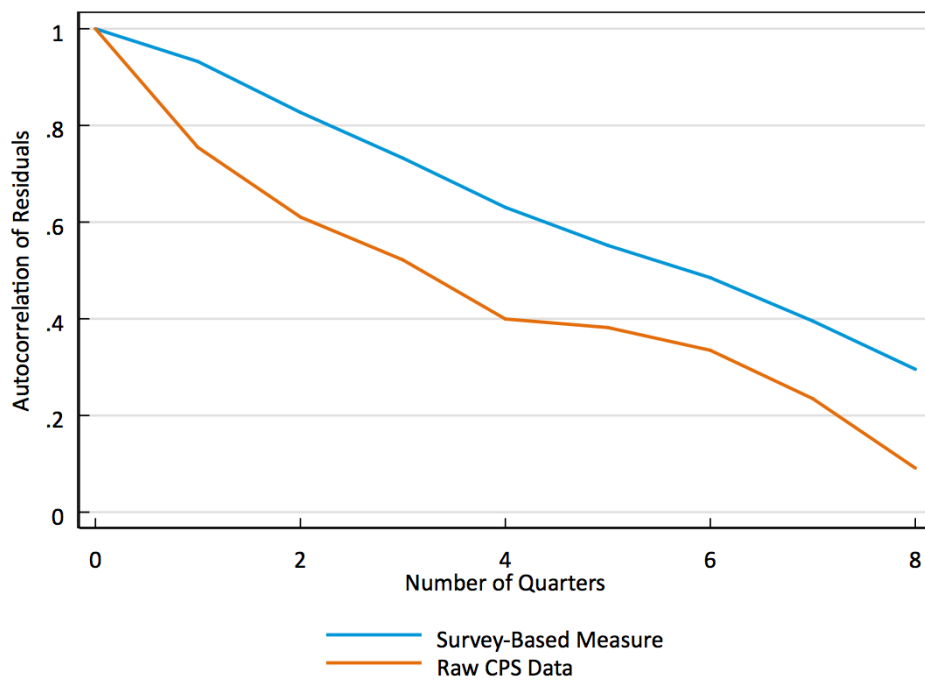


Figure 3: Employment Growth vs. Survey-Based Measure




Figure 4: Employment Growth vs. Survey-Based Measure, Controlling for Administrative Measure


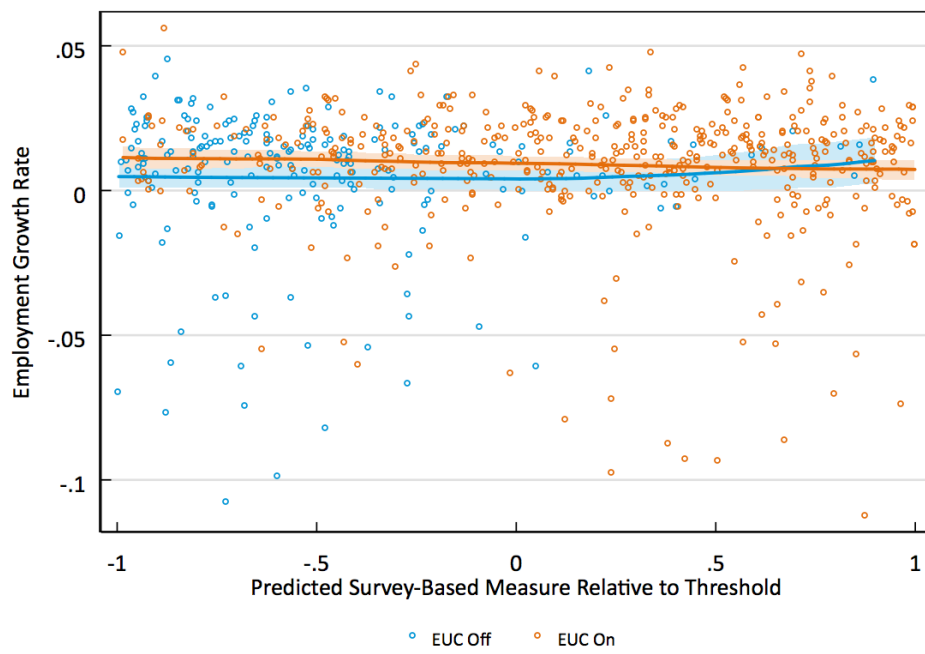


Figure 5: Comparing Employment Growth: Natural Experiment



N = 796.

Figure 6: Comparing Employment Growth: Regression Discontinuity

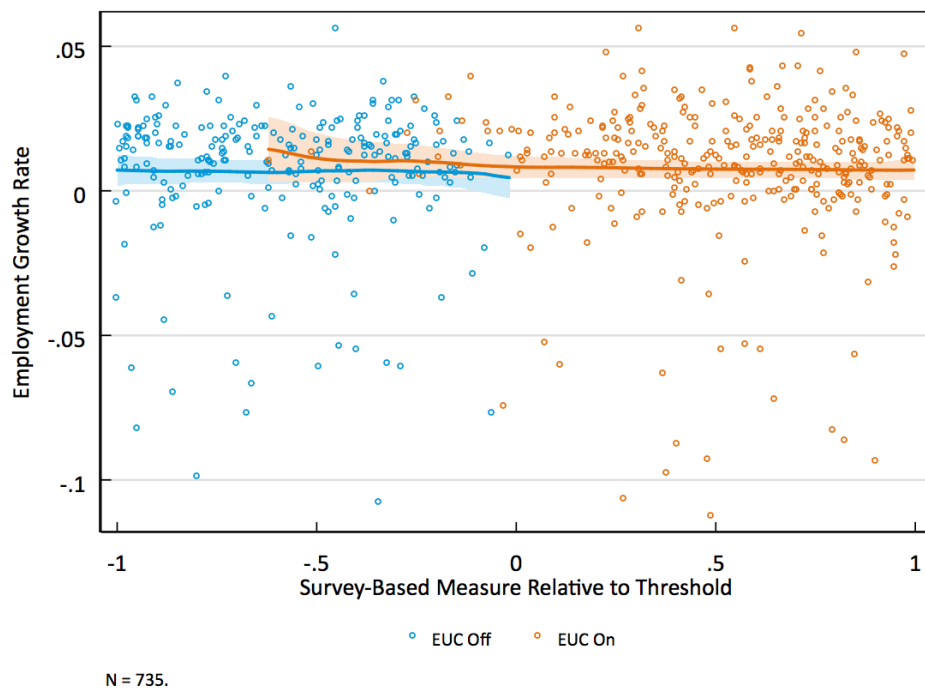


Figure 7: Differences-in-Differences: All Changes in UI Duration




Figure 8: Differences-in-Differences: EUC Expiration




Figure 9: Differences-in-Differences: All Changes in UI Duration (QCEW)




Figure 10: Differences-in-Differences: EUC Expiration (QCEW)




Figure 11: Counties Located on State Borders

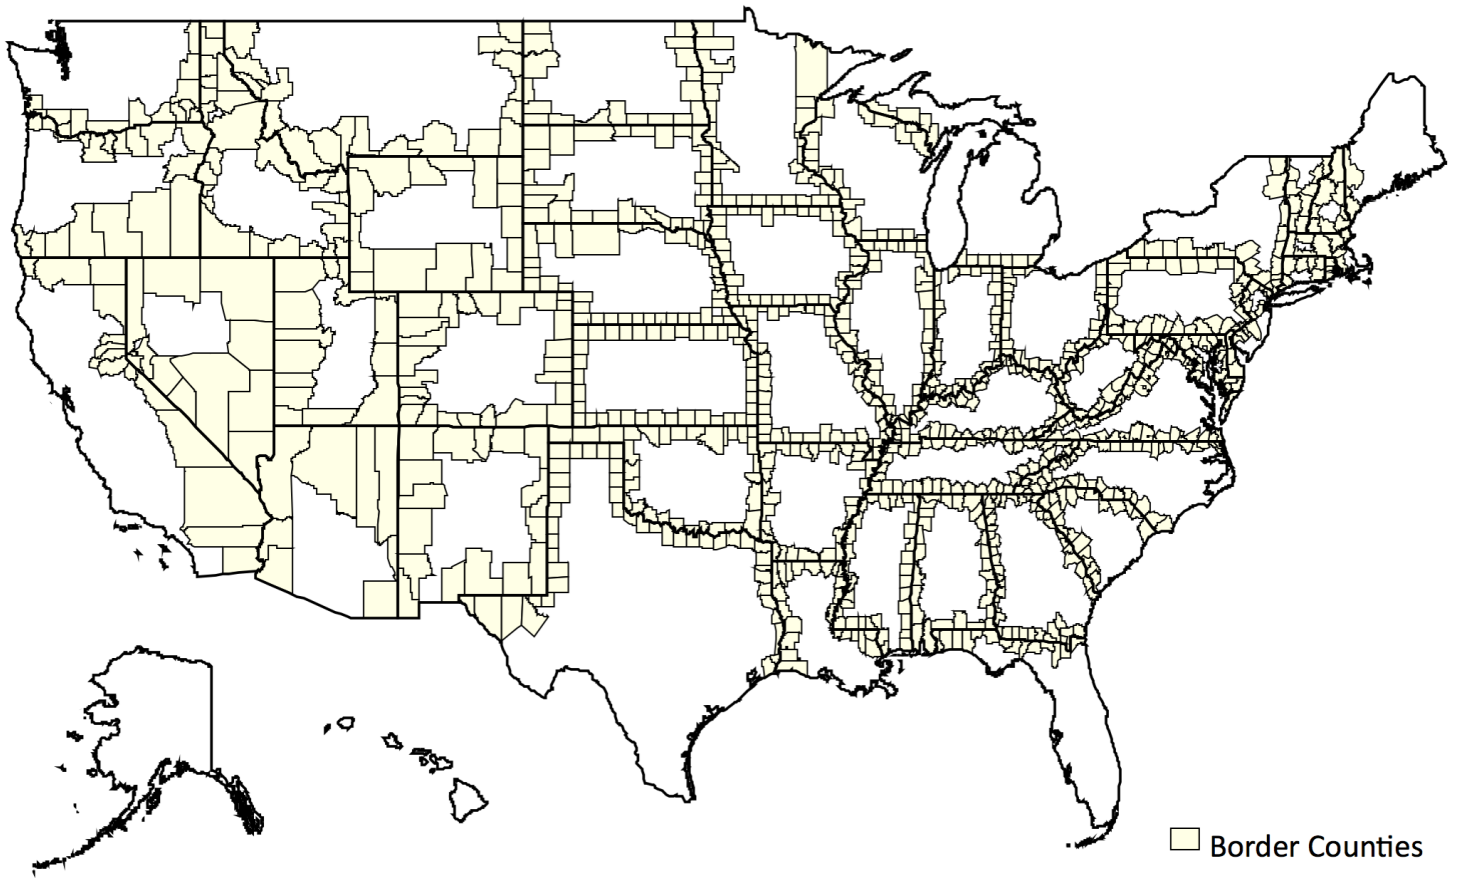


Figure 12: County Employment Growth vs. State Employment Growth

