Do Unemployment Insurance Extensions Reduce Employment?

John Coglianese*

November 30, 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 analyze a natural experiment created by a federal UI extension enacted in the United States during the Great Recession and measure the effect on state-level employment. I exploit a feature of this UI extension whereby random sampling error in a national survey altered the duration of unemployment insurance in several states, resulting in random variation in the number of weeks of unemployment insurance available at the state level. Point estimates of the impact of this UI extension imply that unemployment insurance raises employment growth. Although I cannot conclusively rule out an elasticity of zero, I can rule out substantial negative effects. I also document several issues 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 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 in 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 can 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

*Harvard University: coglianese@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 Jaeger, Larry Katz, Rohan Kekre, Jessica Laird, Greg Mankiw, Jesse Rothstein, Heather Sarsons, James Stock, and participants at the Harvard macro lunch and Harvard labor lunch for their feedback.

likely make it easier for non-UI-eligible individuals to find jobs, since they face less competition as the search effort of individuals collecting unemployment insurance decreases. Additionally, firms may face different incentives to either create new vacancies or cut back on hiring.

The broader effects of unemployment insurance cannot be estimated using individual-level variation in UI duration. Rothstein (2011) and Schmieder et al. (2012) have measured the disincentive effects of unemployment insurance by comparing two individuals within the same labor market who differ in the duration for which they are eligible for unemployment insurance. 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 by up to 53 additional weeks beyond the usual duration, with different amounts of weeks available based on the unemployment rate in the state. The measure of the unemployment rate that the federal government 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 different 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 estimate that UI extensions raise employment growth on average. While I cannot conclusively rule out that the elasticity of employment with respect to UI duration is equal to zero, I am able to reject substantially negative elasticities. I also find evidence consistent with unemployment insurance boosting aggregate demand, as employment in non-tradable industries grew faster in response to UI extensions than did employment in tradable industries.

The prior literature on the macro effect of unemployment insurance is largely divided. A series of papers estimate the effects of UI extensions on each possible channel separately, and combining these estimates suggests that the overall macro effect is either less negative than the micro effect or possibly positive (Lalive, Landais, Zweimüller, 2015; Marinescu, 2014; Di Maggio & Kermani, 2015). In contrast, other recent work that examines the total effect on employment estimates substantial negative elasticities (Hagedorn et al., 2013, 2015). My results are consistent with the former literature, but sharply contrast with 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 provides 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 I employ in this analysis. Section 4 details the research design and results, explaining how I identify random sampling error in the CPS and outlining the main empirical findings. Section 5 compares these results to previous estimates in the literature and explains the discrepancies between my results and prior estimates. Section 6 concludes with a discussion of these findings.

2 Background

2.1 Literature

Analysis of UI benefits comprises a long literature in labor economics and public finance. Solon (1985), Moffitt (1985), and Katz and Meyer (1990) examine the effects of increased benefit generosity on the incentives that unemployed individuals face in searching for a new job. These studies use cross-state differences in UI generosity and generally find 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), use more plausibly exogenous sources of variation in UI benefits at the individual level and find more modest impacts of unemployment insurance on the probability of re-employment for job seekers.

By estimating the effect of unemployment insurance using only variation in UI eligibility within labor market areas, Rothstein (2011) and others estimate 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 on the level of employment of changing unemployment insurance for the whole workforce. This overall 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 each component of 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. Looking at spillovers, Lalive, Landais, Zweimüller

(2015) analyze a natural experiment in Austria where unemployment insurance was extended for only certain workers and find that UI extensions have large positive spillover effects on the employment of workers ineligible for the extensions. For the vacancy creation channel, 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 changes in benefits. Di Maggio & 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-tradable employment to labor demand shocks, but find no effect for tradable 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 have been estimated to be 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 approach 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. Both studies conclude that more generous unemployment insurance reduces employment, based on a county-border-pair research design where UI duration and employment are compared between two counties on opposite sides of a state border. Given 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) use variation in the rollout of UI extensions during the Great Recession and employ a regression specification based on the predictions from a standard search and matching model, while Hagedorn et al. (2015) use the surprise expiration of benefit extensions at the end of 2013 with 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 that estimates channels separately suggesting that the total macro effect is at least less negative than the micro effect and 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

[1]The 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.

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 most workers through two 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 that are experiencing temporarily high local unemployment rates, with funding split half and half between the federal and State governments. In addition, the federal government has since 1958 adopted a practice of extending unemployment insurance using emergency measures during national recessions as a part of its broader countercyclical policy. In 1974, a fully-federally-funded extension of about 13 weeks was put in place in almost all states, lasting until 1978. Similar emergency measures were enacted in 1982, 1991, 2002, varying somewhat in the structure but generally providing about 13 extra weeks of fully-federally-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 in some cases up to 40 weeks 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. Second, the federal government enacted an expansive emergency extension of unemployment insurance called the Emergency Unemployment Compensation 2008 (EUC08) program. While the initial form of the 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 in November 2009 to include four tiers of benefits. Each state was eligible for the first tier automatically, but eligibility for each higher tier depended on a state's unemployment rate exceeding a particular threshold. 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 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 upon two different measures of unemployment, with each tier having two separate thresholds, one for each measure. One measure used administrative data from the UI system to measure unemployment. The administrative unemployment rate is equal to 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 comes from the Bureau of Labor Statistics' Local Area Unemployment

Statistics (LAUS) program. This survey-based 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 raw CPS data fairly closely. 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 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 lead to differences between the two estimates since an individual would be counted in the administrative measure in the state where he or she worked, but would be counted in the survey-based measure 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 collected 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 as it allows me to reconstruct the unemployment rate tests each state was subjected 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 during 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. To reduce the influence of differences in timing, I aggregate both of these measures to the quarterly level. I also seasonally adjust each series on a state-by-state basis using an X-11 procedure applied to the log unemployment rate. I also construct a 13-week moving average of the revised administrative unemployment rate for comparison purposes and similarly aggregate this series to a quarterly frequency and then seasonally adjust.

Drawing on two sources, I constructed data on the maximum duration of unemployment insurance between 2005 and 2014. 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 EUC08 tier. Using this, I computed the amount of extra weeks of UI available in each state from both EUC08 and EB at a weekly frequency and then aggregated to a quarterly frequency. I smoothed through several periods where benefits lapsed temporarily but were subsequently reauthorized and benefits from the lapsed period were paid retroactively. I added this to the number of weeks available under the regular unemployment insurance system taken from Farber & Valletta (2013). Figure 1 shows the median UI duration in the United States between 2005 and 2014 along with the 10th and 90th percentiles, weighting by population. Even at the 10th percentile, UI duration more than tripled between 2008 and 2010. 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 online at 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 are 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 estimates 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 have 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

[2]The 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

adjust the QCEW measure of employment on a state-by-state basis using an X-11 procedure in logs.

4 Research Design and Results

This section is divided into three parts. First, I show that random sampling error is responsible for sizeable differences between the administrative unemployment rate and the survey-based unemployment rate. Second, I outline how I use this random sampling error as part of a natural experiment, comparing employment growth among states with the same economic conditions that were randomly assigned to different durations of unemployment insurance. Finally, I detail the results of this natural experiment and examine how much of the effect of unemployment insurance can be explained by aggregate demand effects.

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 2 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 in New York 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 in New York 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 2 also shows the raw unemployment rate for New York 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 CPS 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 with the CPS data, 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, the main reason for discrepancies like these in New York and other states 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 estimated unemployment rate 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 at a quarterly frequency. 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; it 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 consists of households that 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. Using data on all states, 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 3, 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 whether random sampling error is responsible 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 4 shows a binned scatterplot of the log survey-based unemployment rate against 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 highly statistically significant, even controlling for state and time fixed effects. Figure 5 shows a similar plot, with the only difference being that, prior to plotting, 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. Figure 5 shows no systematic relationship between the survey-based measure and employment growth and the line of best fit is in fact slightly upward sloping, although it is not close to being statistically significant. These plots taken together indicate 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, making New York eligible for all four tiers of EUC08 extensions. When the survey-based measure fell below the 8.5 percent threshold in 2010, New York switched off of Tier 4 and the potential maximum 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 EUC08 benefits. This experiment can be thought of in two different ways. The first is as a random experiment in which 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 EUC08 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 in this context. 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 EUC08 tier and offer different UI durations 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 EUC08 thresholds when the survey-based measure lay near the threshold, even though they had the same administrative unemployment rate.

To estimate the elasticity of employment with respect to UI duration from this experiment, I estimate the following regression on employment growth:

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

where I instrument for UI duration with dummy variables for whether a state was triggered onto each tier of EUC08. The dependent variable here is the annualized growth rate of employment in quarter $t+1$, which I use to avoid issues of simultaneous causality. 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 EUC08 is dependent only on random sampling error and so these instruments are valid when conditioning on the administrative measure. I use the revised administrative unemployment rate to improve precision, but the results are virtually unchanged when using the unrevised, "real-time" administrative unemployment rate.

The regression discontinuity design uses the sharp cutoff at the thresholds for each EUC08 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 EUC08 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

$$\begin{aligned} \Delta log(\text{Employment})_{s,t+1,t} = \theta_s + \delta_t &+ \beta \text{UI Duration}_{s,t} \\ &+ \sum_{k \in \{2,3,4\}} \gamma_k \mathbf{1}\left(|\text{Survey Unemp. Rate} - \text{Threshold}_k| \leq 1\right) + \epsilon_{s,t} \end{aligned} \tag{4.2}$$

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.[3] 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. As before, I use subsequent employment growth at an annual rate as the dependent variable. I control for state and time 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 is as-if randomly assigned.

Both the natural experiment and regression discontinuity designs compare employment growth between states 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

[3] I use 1 p.p. as the bandwidth, but the results are similar with smaller and larger bandwidths.

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, as measured by the administrative measure. This involves estimating the following regression

$$\begin{aligned}\Delta log(\text{Employment})_{s,t+1,t} = \theta_s + \delta_t + \beta \text{UI Duration}_{s,t} \\ + \sum_{k=0}^{8} \gamma_k \text{Administrative Unemp. Rate}_{s,t-k} \\ + \sum_{q \in \{2,3,4\}} \psi_q \mathbf{1}\left(|\text{Survey Unemp. Rate} - \text{Threshold}_q| \leq 1\right) + \epsilon_{s,t}\end{aligned} \tag{4.3}$$

where I instrument for UI duration with dummy variables for whether a state is above or below the threshold for each tier q, conditional on the survey unemployment rate being within 1 p.p. of the threshold. As before, I control for state and time fixed effects and cluster at the state level. This specification exploits variation in UI duration between states with both the same administrative unemployment rate and with similar survey-based unemployment rates just above and below an EUC08 threshold. By leveraging variation from both the natural experiment and regression discontinuity designs, this approach should provide more precise estimates of the elasticity.

4.3 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 6 and 7. Each graph plots annualized employment growth one period into the future against the present unemployment rate and separately shows states on a higher tier of EUC08 (orange dots) and states not eligible for the higher tier (blue triangles). Figure 6 shows the natural experiment design with the predicted survey-based unemployment rate relative to the nearest EUC08 threshold on the x-axis. There is substantial horizontal overlap of the two types of points since random sampling error can substantially alter the estimated survey-based unemployment rate and assign states with the same administrative measure to different levels of the EUC08 program. There is some limit to the extent of random sampling error, however, as the density of the opposite type of points 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 triangles on the far left 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. The orange and blue lines show a local kernel-weighted mean for states triggered on and off the higher EUC08 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. These estimates indicate 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.

In contrast, Figure 7 shows this plot for the regression discontinuity research design, with the difference between survey-based unemployment rate and the nearest EUC08 threshold shown on the x-axis. Here there is very little overlap between the orange dots and blue triangles. As the survey-based unemployment rate crosses the threshold, there is a clear change in UI duration 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.[4] However, at the threshold, there does not appear to be much of a break in employment growth; if anything there appears to be a slightly positive break at 0. The orange and blue lines show a local kernel-weighted mean for states triggered on and off the higher EUC08 tier, respectively. These estimates suggest that 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.

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 0.02. This is a sizeable elasticity, implying that a doubling of UI duration would raise the employment growth rate between 1 and 2 percentage points annually relative to the counterfactual. However, the standard errors are quite wide and the 99 percent confidence interval for each regression includes an elasticity of zero. Nevertheless, the confidence intervals for these estimates do 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, as laid out in Equation 4.3, 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 linear and quadratic terms of the difference between the survey-based measure and the threshold, as well as interactions of these polynomial terms with the dummy for the unemployment rate being above the threshold. These polynomial terms act as additional controls to the regression discontinuity specifications from columns two and three respectively. These specifications produce slightly larger 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 approach from Equation 4.3 to employment in different industries. I divide industies into four mutually exclusive categories by NAICS supersector: 1) manufacturing; 2) trade, transportation, utilities, leisure, and hospitality; 3) construction; and 4) all other supersectors. Manufacturing is a clear case of a tradable industry, while the second and third are clear examples of non-tradable industries. In theory, since demand for tradable goods and services is nationwide, effects on aggre-

[4] A few states are shown to be triggered onto EUC08 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 EUC08 would remain on the new tier for at least 13 weeks before it could trigger off.

gate demand at the state-level should not substantially affect tradable employment. Therefore, the difference between the response of non-tradable employment and tradable employment should give an estimate of the change in aggregate demand due to the UI extension. The estimates in column one of Table 2 show that the elasticity of manufacturing employment is -0.002, which is substantially smaller than the estimates of the effect on overall employment found in Table 1. However, the elasticity of non-tradable employment in column two is 0.013, which is consistent with the effects on overall employment. I show construction separately from the other non-tradable category since the dynamics of this industry were very different during the period examined, which falls in the aftermath of the national housing crash. I present the elasticity for construction employment in column three and for employment in 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 other industries is largely in line with the effects measured in Table 1.

An additional dimension of heterogeneity that I consider is how the effect of UI extensions changed throughout the recession. The stimulative effect of UI extensions should be greater during periods in which the zero lower bound was binding and the labor market had a greater amount of slack (Kekre, 2015). I examine the effect of increased UI duration separately in three different intervals: 2006-2008, 2009-2011, and 2012-2014. Table 3 shows the estimates of applying Equation 4.3 to each period separately. Across all periods, estimates of the elasticity are large and positive. While it may be surprising that the elasticity was positive in all periods, monetary policy was constrained by the zero lower bound for the entire period in which the UI extensions were in effect. The estimated elasticity is greatest for the 2009-2011 period, which coincided with the peak in the unemployment rate. However, although the point estimates for 2006-2008 and 2012-2014 are smaller, I cannot reject that the elasticity is the same in all three time periods. Nevertheless, I can reject substantial negative elasticities in latter two time periods.

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). The evidence comparing the effects on employment in different industries 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. Since interest rates in the US were up against the zero lower bound during the entire duration of the EUC08 program and conventional monetary policy was constrained, UI extensions may have been more effective at stimulating employment growth during this period than if conventional monetary policy had not been constrained. 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.

5 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, estimates from the literature on each of these channels suggest 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 Di Maggio & 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), who find substantial negative elasticities. Importantly, Hagedorn et al. (2013) and Hagedorn et al. (2015) analyze the same setting and time period I examine, ruling out external validity as an explanation for this discrepancy. In this section, I consider this discrepancy more carefully by reviewing both of the Hagedorn et al. (2013, 2015) studies and presenting analysis seeking to reconcile the differences in their estimates and the results from the natural experiment presented in the preceding section of this paper.

Hagedorn et al. (2013, 2015) clearly make a substantial contribution to the literature on unemployment insurance by highlighting the difference between the macro and micro elasticities and by 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 that an estimate of the macro effect of UI extensions is necessary for determining UI policy. In both of their studies, Hagedorn et al. (2013, 2015) also take many steps to verify 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 how to use smaller-scale natural experiments to study policy changes at an aggregate level in macroeconomics.

I am able to replicate the estimates produced in both studies. 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 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 while directly controlling for future changes in UI benefits across counties in a border pair. I find that the discrepancy between the positive elasticity estimated from the natural experiment in the present paper and the negative elasticity estimated by Hagedorn et al. (2013, 2015) is best explained by a combination of three factors: non-parallel trends, statistical noise in data sources based on surveys, and non-smoothness of economic geography.

In their 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. Thus, the variation they are left with consists of 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 the state crosses an unemployment rate threshold. In this case, a potential worry is that a persistent negative shock to labor markets could trigger an increase in UI duration and simultaneously lead to lower employment, which the research design would 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:

$$log(\text{Employment}_{s,t}) = \theta_s + \delta_t + \sum_{k=-8}^{3} \beta_k log(\Delta\text{UI Duration}_{s,t-k}) + \nu_{s,t} \tag{5.1}$$

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 in order to be consistent with the original paper. I plot the coefficients β_k in Figure 8 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), in which 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. This 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 conducted 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 occured in the continental United States. I created a placebo UI extension by assuming that UI duration was increased 13 weeks whenever the survey-based unemployment rate rose above 5 percent. When I regressed the quasi-forward difference of unemployment on the placebo UI duration, differencing between counties on either side of state borders, I estimated a positive and statistically significant coefficient on the same order of magnitude as the coefficient reported in Hagedorn et al. (2013). 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 implies that it is not an adequate method for correcting the non-parallel trends problem.[5]

The paper by Hagedorn et al. (2015) focuses on a surprise shock to UI benefits created by the lapse in funding for the EUC08 program at the end of 2013. The benefits paid out by the EUC08 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 1, 2014, the EUC08 program abruptly ended and Congress did not act to re-authorize the program, causing UI duration to fall nationwide back to the regular UI duration. Although the change occured 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 EUC08 program. Before the lapse in funding, Illinois, Michigan, Nevada, and Rhode Island had still been eligible for all four tiers of EUC08 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 variation in the change in UI duration caused by the lapse in funding, I conduct the same differences-in-differences design as outlined above in Equation 5.1 and plot the β_k coefficients in Figure 9. As before, I use the LAUS measure of employment to be consistent with the original paper. Each coefficient in 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, such as the lapse in EUC08 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. This would trigger a higher tier of benefits in those states (or prevent higher tiers from triggering off), but then the measurement error shock 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

[5] When I apply the same placebo test to my research design, I estimate an elasticity close to zero that is not statistically significant.

9, this seems like a plausible concern around the time of 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 10 and 11 show the same differences-in-differences plots as Figures 8 and 9 but using QCEW employment as the dependent variable instead of LAUS employment. The QCEW data do 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 EUC08 expiration. In this case, using the QCEW data and the EUC08 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 because EUC08 raised UI duration in states with higher unemployment. Although the states that experienced a large drop in UI duration in the EUC08 expiration had mostly similar employment growth in 2013 compared to states that experienced smaller drops when EUC08 expired, counterfactual growth could have been different between these states 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 12 shows the 1,143 counties in the US that share a land border with another state. The research design used by Hagedorn et al. (2013, 2015) involves comparing employment in counties on opposite sides of a state border. 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 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 an effect of non-UI 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 with QCEW data, where counts of employment for counties are constructed directly from the administrative data without using state-level aggregates to impute county-level 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 13 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 the relationship is highly statistically significant even when including the outliers. Table 4 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.[6]

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 the context on unemployment insurance extensions is much more mechanical than in the context of right-to-work or minimum wage laws. With UI extensions, 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.

[6] Specifically, I drop years 1994-1995, 2002-2004, and 2008-2014.

6 Conclusion

In this paper, I have used a natural experiment created by a federal UI extension during the Great Recession to estimate the impact on state-level employment. A key feature of this UI extension was that the program set UI duration based on the unemployment rate in each state, which was measured in a national household survey. By using administrative data I was able to isolate random sampling error in the household 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 estimated that higher UI duration raised the growth rate of employment with an elasticity of between 0.01 and 0.02. Although the estimates have large standard errors and I cannot 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 has been split. Papers that measure the different channels of the macro effect 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 in the finding that the total effect is positive and in that the aggregate demand effects are substantial. After examining two studies that report negative elasticities, I identify three factors as being responsible for the discrepancy between our respective results: non-parallel trends; measurement error in data based on surveys; and non-smoothness of economic geography.

While the analysis I present sheds light on the macro effect of the EUC08 extension, and perhaps by implication the effect of unemployment insurance on employment more generally, the results may differ in other contexts. The EUC08 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, which would 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, Laura Feiveson, Zachary Liscow, and William Gui Woolston. "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." mimeo, University of California, Berkeley, 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, Fatih Karahan, Iourii Manovskii, and Kurt Mitman. "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." mimeo, Harvard University, 2015.

Lalive, Rafael, Camille Landais, and Josef Zweimüller. "Market Externalities of Large Unemployment Insurance Extension Programs." mimeo, London School of Economics, 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., Till M. von Wachter, and Stefan Bender. "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*	0.0229**	0.0232**	0.0382	0.0588**
	(0.00484)	(0.0111)	(0.00952)	(0.0272)	(0.0267)
Instruments	EUC08 Status	Above/Below Thresholds			
Controlling for Administrative Measure and Eight Lags	YES	NO	YES	NO	YES
Controlling For Linear and Quadratics of Running Variable	NO	NO	NO	YES	YES
First Stage F Statistic	1364.85^{+}	71.57^{+}	73.51^{+}	21.84^{+}	21.01^{+}
Hansen J-Statistic	1.179	1.187	0.0204	0.524	0.624
(p-value)	(0.555)	(0.552)	(0.990)	(0.770)	(0.732)
N	1,989	1,989	1,989	1,989	1,989
Number of Clusters	51	51	51	51	51
R^2	0.738	0.689	0.733	0.689	0.722

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. Regressions 4 and 5 include linear and quadratic terms of the difference between the survey-based measure of unemployment and the nearest EUC08 threshold, as well as interactions between these terms and a dummy variable for being above the threshold. 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. The J-statistic of overidentification is computed for the GMM estimator using clustered long-run variance weight matrix. 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	Manufacturing	Trade/Transport./ Utilities and Leisure/Hospitality	Construction	All Others
Log UI Duration (Weeks)	-0.00223 (0.0228)	0.0133 (0.00860)	0.137** (0.0546)	0.0191** (0.00932)
Instruments	Above/Below Thresholds			
Controlling for Administrative Measure and Eight Lags	YES	YES	YES	YES
First Stage F Statistic	73.22^{+}	73.51^{+}	73.36^{+}	73.51^{+}
Hansen J-Statistic (p-value)	5.313 (0.0702)	0.514 (0.773)	0.283 (0.868)	0.555 (0.758)
N	1,981	1,989	1,985	1,989
Number of Clusters	51	51	51	51
R^2	0.624	0.667	0.619	0.494

This table reports coefficient estimates and standard errors from four instrumental variables (IV) regressions covering mutually exclusive industry categories. 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. The J-statistic of overidentification is computed for the GMM estimator using clustered long-run variance weight matrix. 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: Heterogeneity by Time Period

	(1)	(2)	(3)
Dependent Variable	Annualized Employment Growth in the Next Quarter (QCEW)		
Time Period	2006-2008	2009-2011	2012-2014
Log UI Duration (Weeks)	0.0225 (0.0463)	0.0299*** (0.00993)	0.0207* (0.0122)
Instruments	Above/Below Thresholds		
Controlling for Administrative Measure and Eight Lags	YES	YES	YES
First Stage F Statistic	137.6^{+}	46.10^{+}	15.30^{+}
Hansen J-Statistic (p-value)	–	1.568 (0.456)	0.801 (0.670)
N	612	612	561
Number of Clusters	51	51	51
R^2	0.845	0.858	0.074

This table reports coefficient estimates and standard errors from three instrumental variables (IV) regressions covering mutually exclusive time periods. 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. The J-statistic of overidentification is computed for the GMM estimator using clustered long-run variance weight matrix. 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 4: 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: Distribution of UI Duration, 2005-2014

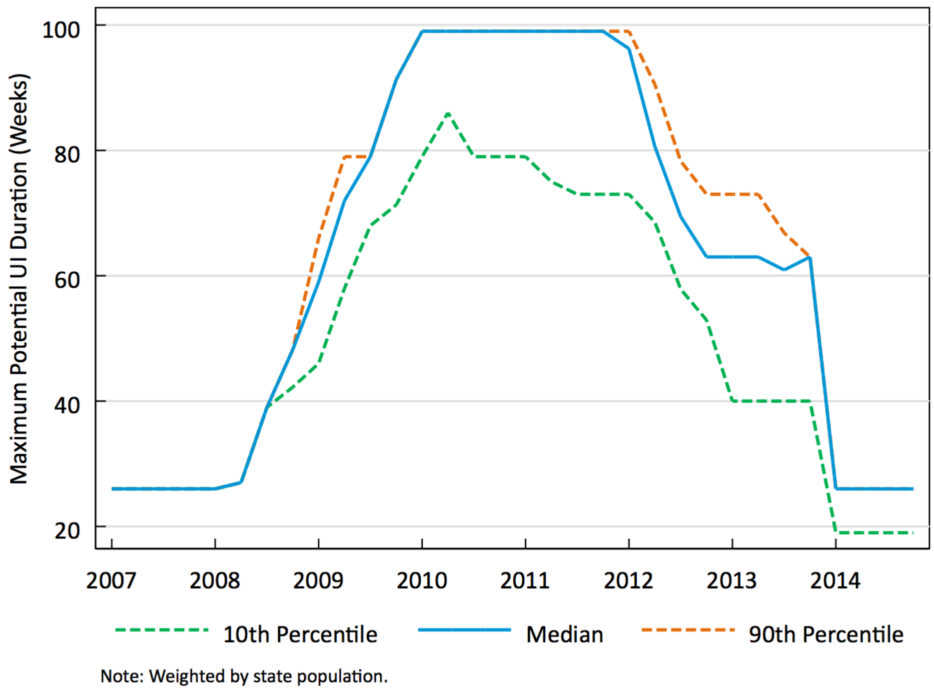

This figure shows median UI duration in weeks across US states between 2007 and 2014, along with the 10th and 90th percentiles. In each period, the median and percentiles are computed weighting by state population. I construct a weekly measure of the extra benefits available in each state from UI extensions, and add this to data on baseline UI duration in each state from Farber & Valletta (2013).

Figure 2: Unemployment Rates for New York State

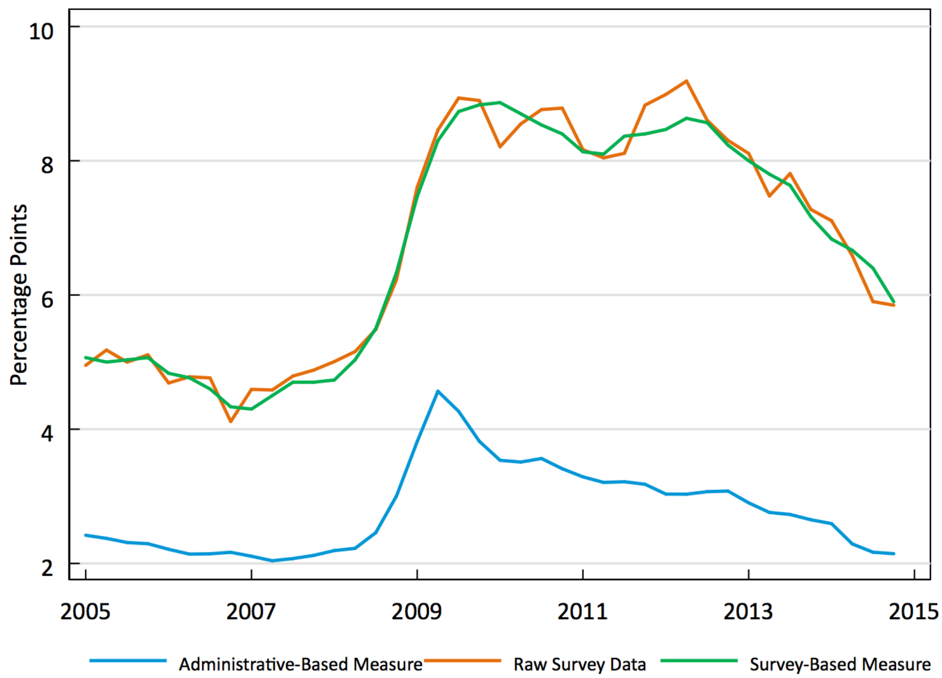

This figure shows several different measures of unemployment for the state of New York between 2005 and 2015. The administrative-based measure corresponds to the Insured Unemployment Rate, which is computed as the ratio of regular UI claimants to total UI-covered employment. The raw survey data is the unemployment rate computed from the CPS basic monthly public use files. The survey-based measure is the unemployment rate estimate produced by the Local Area Unemployment Statistics program. All three series have been seasonally adjusted using an X-11 procedure in logs.

Figure 3: Residual Autocorrelation, CPS and Survey-Based Measure

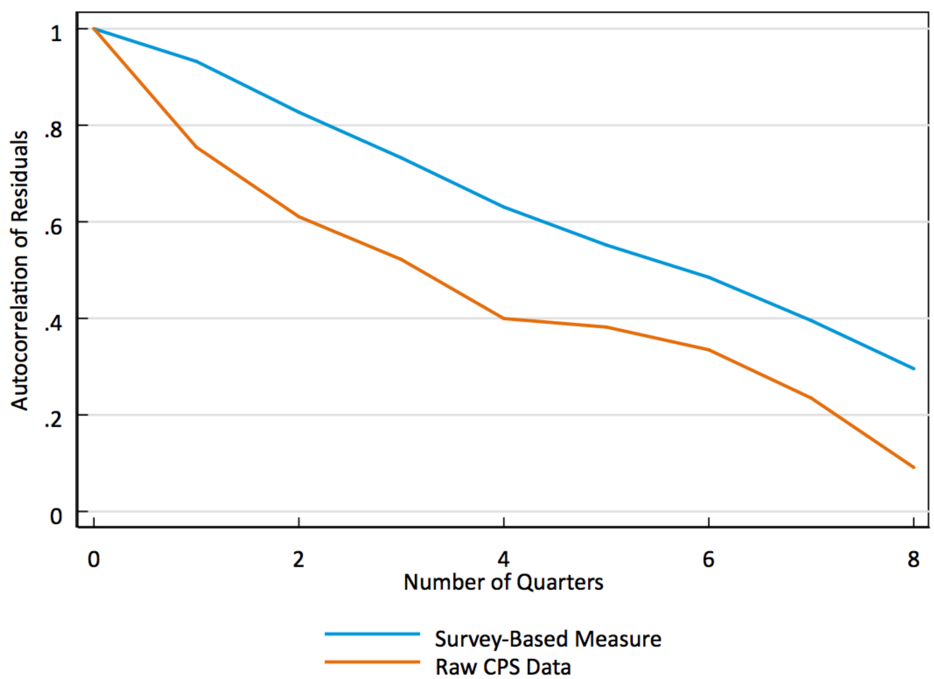

This figure shows the autocorrelation coefficients of residuals over horizons ranging from one quarter to eight quarters. I computed the residuals for two different series, the survey-based unemployment rate from LAUS and the unemployment rate computed from the CPS public use files, by regressing each on the administrative unemployment rate and eight lags of the administrative unemployment rate along with state fixed effects. All series were seasonally adjusted using an X11 procedure in logs prior to regressing. While the autocorrelation declines monotonically with horizon for both series, the autocorrelations decline less steeply between four and five quarters out, consistent with the structure of the CPS.

Figure 4: Employment Growth vs. Survey-Based Measure

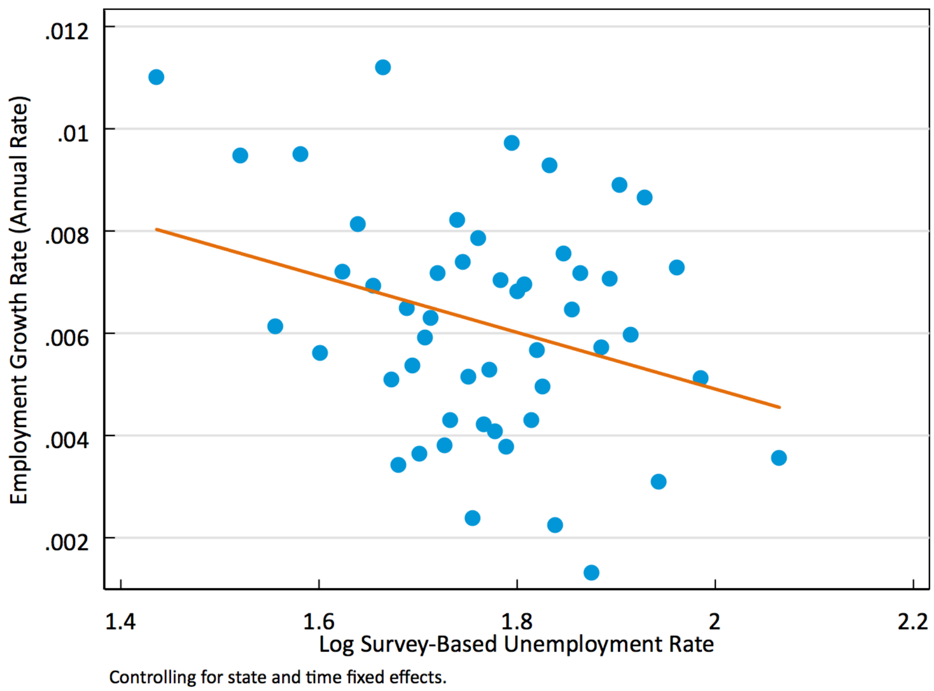

This figure shows a binned scatterplot of the log survey-based unemployment rate against employment growth over the subsequent quarter. Each dot represents the mean of about 50 observations along both dimensions. The line of best fit is shown, with a slope of -0.006, which is statistically significant at the 5% level. The observations were residualized on state and time fixed effects before plotting.

Figure 5: Employment Growth vs. Survey-Based Measure, Controlling for Administrative Measure

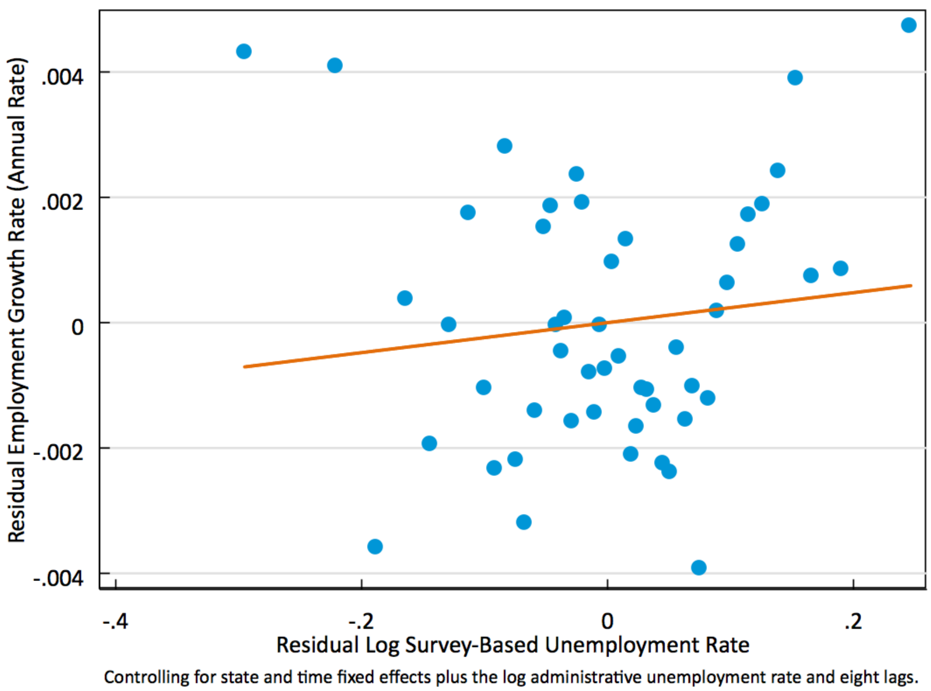

This figure shows a binned scatterplot of the log survey-based unemployment rate against employment growth over the subsequent quarter. Prior to plotting, both variables were residualized on the log administrative unemployment rate and eight lags, along with state and time fixed effects. Each dot represents the mean of about 50 observations along both dimensions. The line of best fit is shown, with a slope of 0.002, which is not statistically significant.

Figure 6: Comparing Employment Growth: Natural Experiment

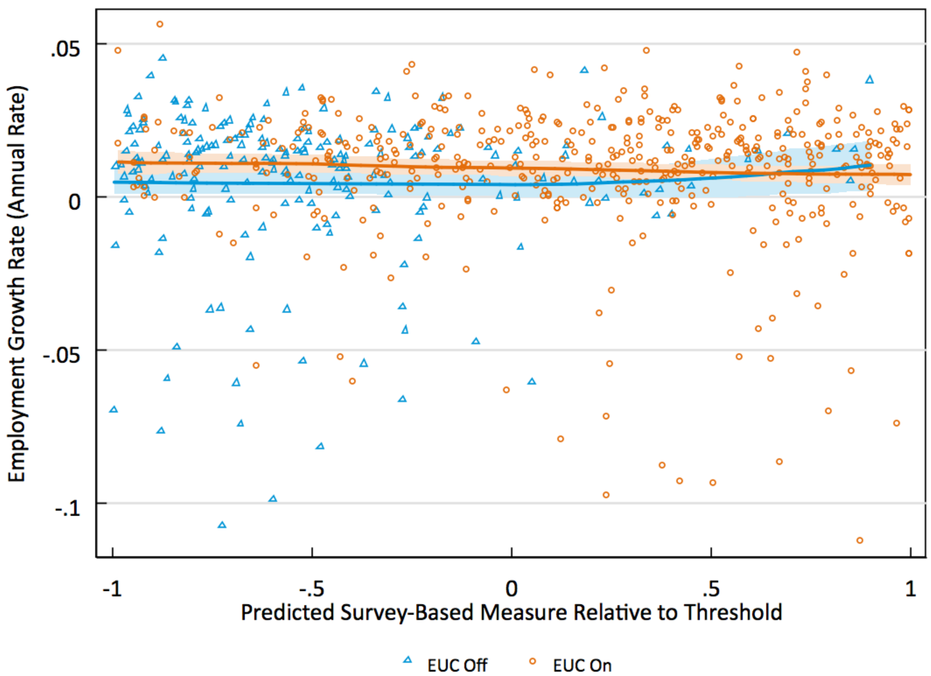

This figure shows employment growth over the subsequent quarter against the difference between the predicted survey-based measure and the nearest EUC threshold. The predicted survey-based measure was computed from a regression of the survey-based unemployment rate on the administrative measure and eight lags, along with state and time fixed effects. The two lines show local kernal-weighted means for the two types of observations, using an Epanechnikov kernel and a rule-of-thumb bandwidth.

Figure 7: Comparing Employment Growth: Regression Discontinuity

This figure shows employment growth over the subsequent quarter against the difference between the actual survey-based measure and the nearest EUC threshold. The two lines show local kernal-weighted means for the two types of observations, using an Epanechnikov kernel and a rule-of-thumb bandwidth. There are a few observations where the survey-based unemployment rate was below the threshold but EUC was triggered on, due to a rule that once a state triggered onto a new tier of EUC, they would remain eligible for that tier of benefits for at least 13 weeks, even if the unemployment rate fell below the threshold during those 13 weeks.

Figure 8: Differences-in-Differences: All Changes in UI Duration

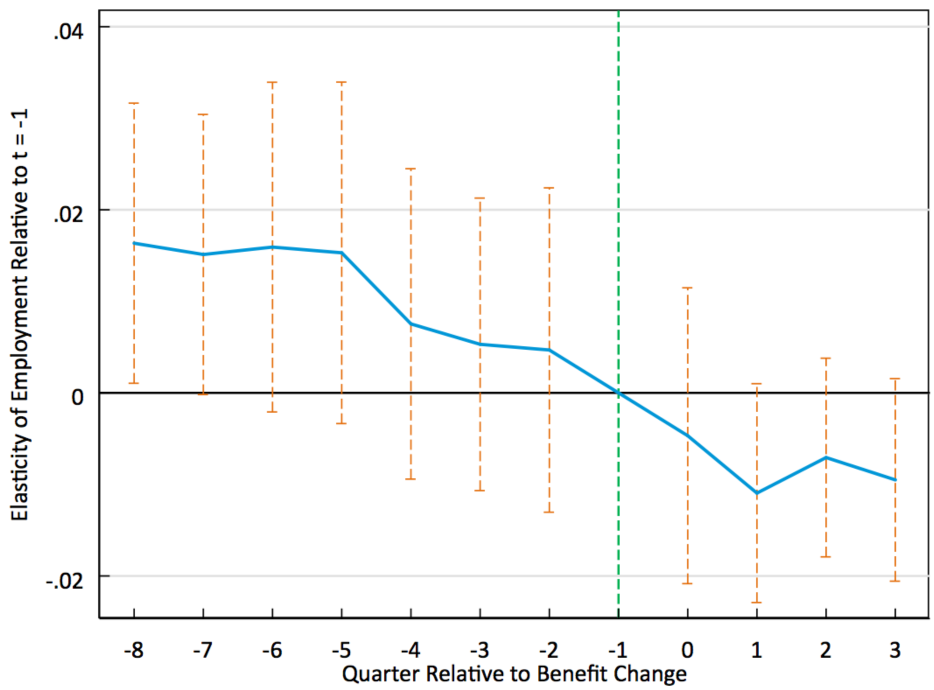

This figure shows coefficients from an event study regression. I regress log LAUS employment on leads and lags of the change in log UI duration, along with state and time fixed effects. I use all changes in UI benefits between 2006 and 2014. Each coefficient represents the elasticity of employment in that quarter relative to employment in the quarter immediately before the change in UI duration, for a 1% increase in UI duration.

Figure 9: Differences-in-Differences: EUC08 Expiration

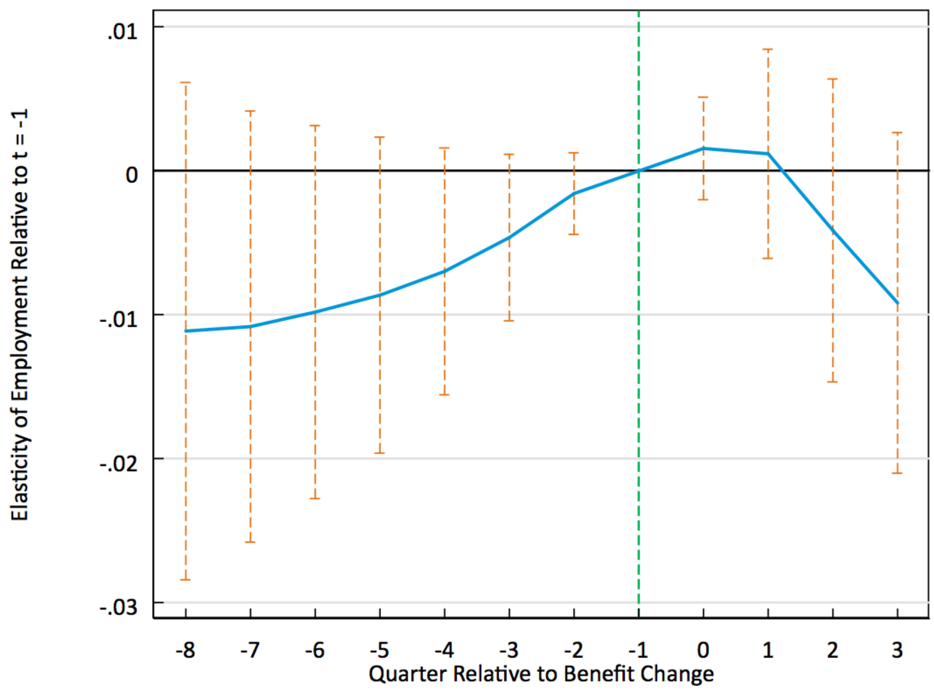

This figure shows coefficients from an event study regression. I regress log LAUS employment on leads and lags of the change in log UI duration, along with state and time fixed effects. I use only the change in UI duration caused by the expiration of the EUC program. Each coefficient represents the elasticity of employment in that quarter relative to employment in the quarter immediately before the change in UI duration, for a 1% increase in UI duration.

Figure 10: Differences-in-Differences: All Changes in UI Duration (QCEW)

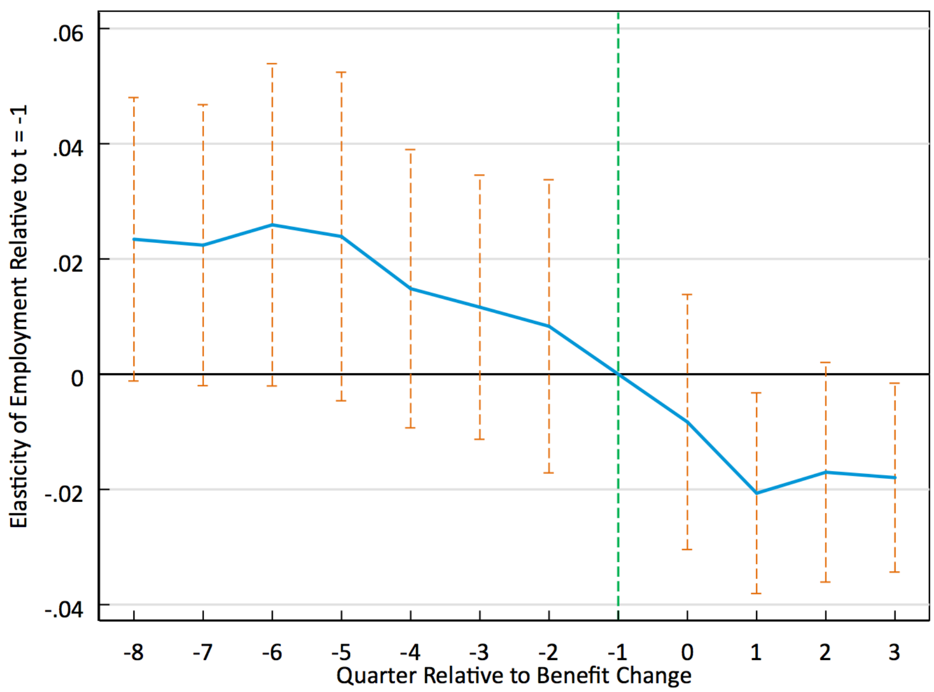

This figure shows coefficients from an event study regression. I regress log QCEW employment on leads and lags of the change in log UI duration, along with state and time fixed effects. I use all changes in UI benefits between 2006 and 2014. Each coefficient represents the elasticity of employment in that quarter relative to employment in the quarter immediately before the change in UI duration, for a 1% increase in UI duration.

Figure 11: Differences-in-Differences: EUC08 Expiration (QCEW)

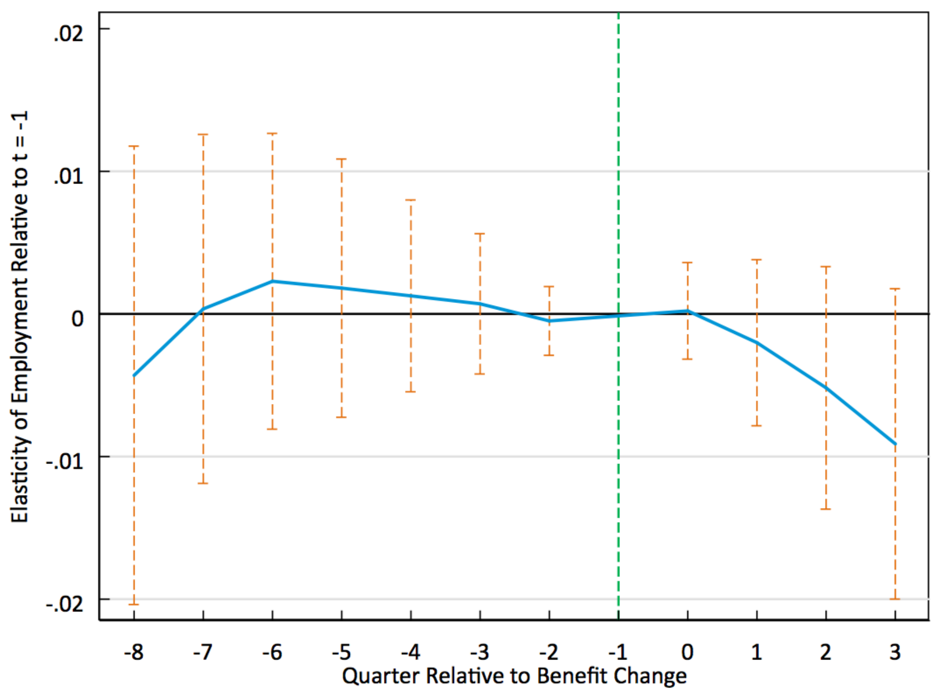

This figure shows coefficients from an event study regression. I regress log QCEW employment on leads and lags of the change in log UI duration, along with state and time fixed effects. I use only the change in UI duration caused by the expiration of the EUC program. Each coefficient represents the elasticity of employment in that quarter relative to employment in the quarter immediately before the change in UI duration, for a 1% increase in UI duration.

Figure 12: Counties Located on State Borders

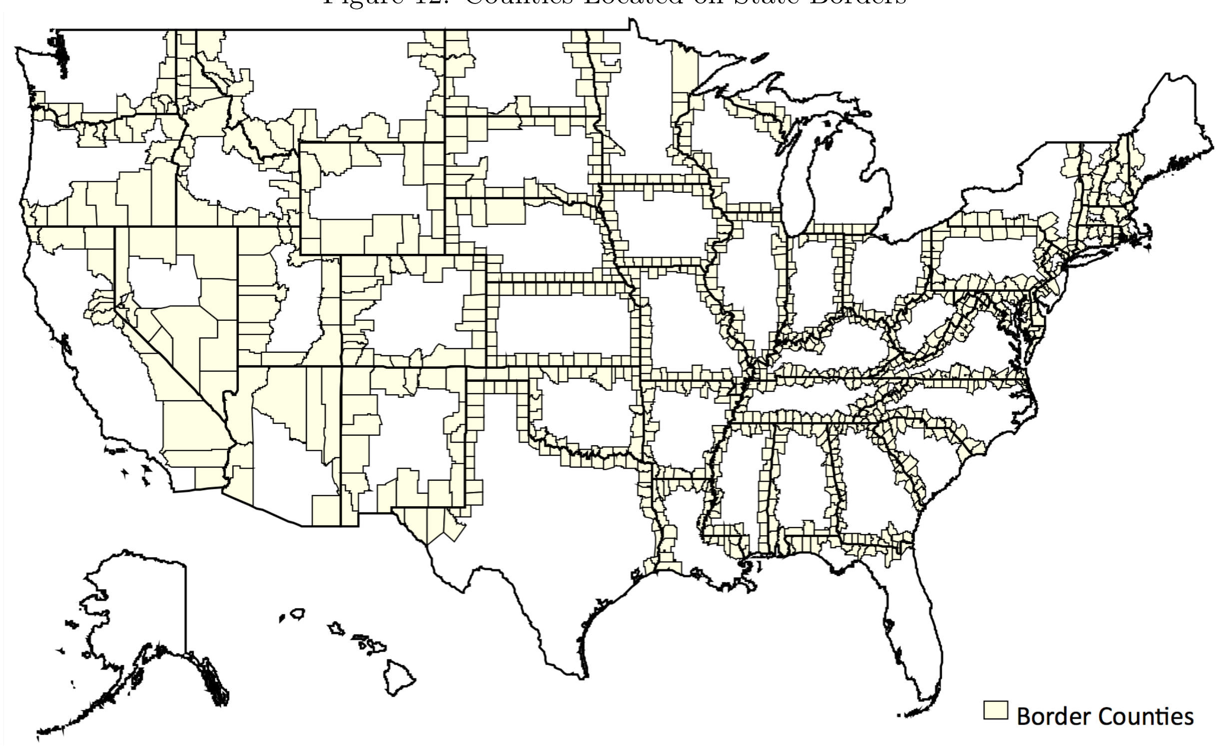

This map shows the 1,143 counties that share a land border with a county in another state.

Figure 13: County Employment Growth vs. State Employment Growth

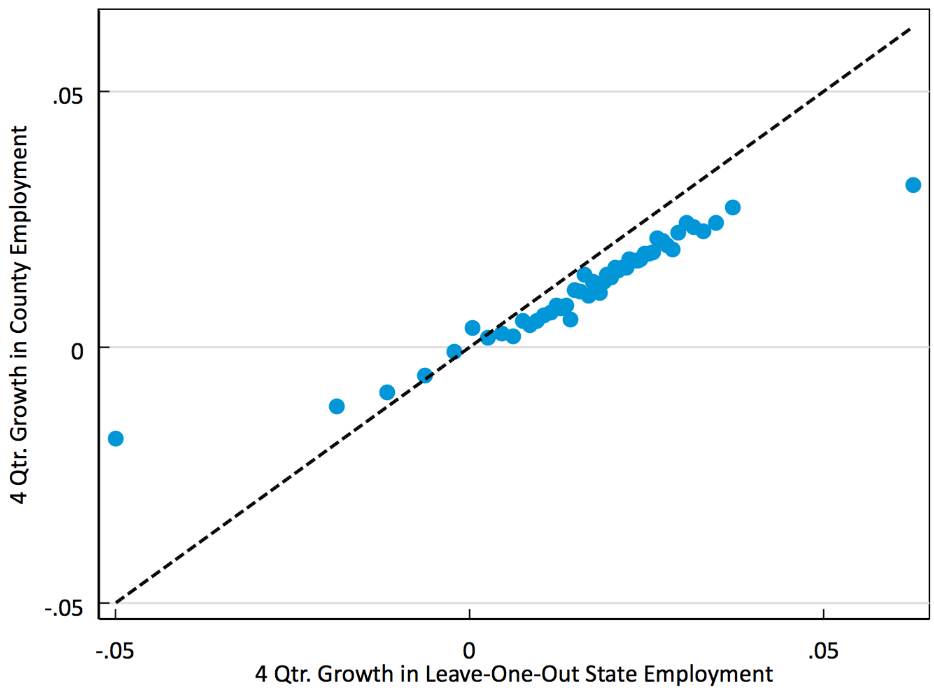

This figure shows a binned scatterplot of the yearly growth in employment at the county level against yearly employment growth in all other counties within the same state, for counties located on a state border. Both series have been residualized on yearly employment growth in the corresponding county located across the state border, along with county-border-pair fixed effects. The dotted line shows the 45 degree line. Each dot represents the mean of about 2,500 observations on both dimensions.