

NBER WORKING PAPER SERIES

THE IMPACT OF UNEMPLOYMENT BENEFIT EXTENSIONS ON EMPLOYMENT:
THE 2014 EMPLOYMENT MIRACLE?

Marcus Hagedorn
Iouri Manovskii
Kurt Mitman

Working Paper 20884
<http://www.nber.org/papers/w20884>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2015

We thank seminar participants at the Institute for International Economic Studies and at the University of Oslo for helpful comments. Support from the National Science Foundation Grant No. SES-1357903 is gratefully acknowledged. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Marcus Hagedorn, Iouri Manovskii, and Kurt Mitman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Impact of Unemployment Benefit Extensions on Employment: The 2014 Employment Miracle?

Marcus Hagedorn, Iourii Manovskii, and Kurt Mitman

NBER Working Paper No. 20884

January 2015

JEL No. E24,J63,J64,J65

ABSTRACT

We measure the effect of unemployment benefit duration on employment. We exploit the variation induced by the decision of Congress in December 2013 not to reauthorize the unprecedented benefit extensions introduced during the Great Recession. Federal benefit extensions that ranged from 0 to 47 weeks across U.S. states at the beginning of December 2013 were abruptly cut to zero. To achieve identification we use the fact that this policy change was exogenous to cross-sectional differences across U.S. states and we exploit a policy discontinuity at state borders. We find that a 1% drop in benefit duration leads to a statistically significant increase of employment by 0.0161 log points. In levels, 1.8 million additional jobs were created in 2014 due to the benefit cut. Almost 1 million of these jobs were filled by workers from out of the labor force who would not have participated in the labor market had benefit extensions been reauthorized.

Marcus Hagedorn
Department of Economics,
University of Oslo,
Box 1095 Blindern,
0317 Oslo, Norway
marcus.hagedorn07@gmail.com

Kurt Mitman
Institute for International Economic Studies
Stockholm University
106 91 Stockholm
SWEDEN
kurt.mitman@iies.su.se

Iourii Manovskii
Department of Economics
University of Pennsylvania
160 McNeil Building
3718 Locust Walk
Philadelphia, PA 19104
and NBER
manovski@econ.upenn.edu

“We just got a jobs report today showing that we’ve now seen the fastest job growth in the United States in the first half of the year since 1999. (Applause.) So this is also the first time we’ve seen five consecutive months of job growth over 200,000 since 1999. (Applause.) And we’ve seen the quickest drop in unemployment in 30 years.”

Barack Obama
Remarks on the Economy
July 3, 2014

1 Introduction

Our objective in this paper is to assess the effect of unemployment benefit extensions on employment. Measuring the magnitude of this effect is manifestly important for understanding the economic consequences of this widely used policy instrument. Yet, the existing literature provides little information on the size, let alone the sign of this effect. In the theoretical literature the effect of benefit extensions on employment is generally ambiguous. Basic decision theory suggests that some unemployed may increase their search effort in response to a cut in benefits, while others, who were mainly searching to qualify for benefits, might drop out of the labor force once losing eligibility, leading to offsetting effects on employment. Equilibrium job search theory typically implies a positive effect of a cut in benefit duration on job creation. This makes it easier to find jobs and might induce those previously out-of-labor force to rejoin the labor force, leading to an increase in employment with an ambiguous effect on unemployment since the number of job vacancies and the number of searchers increases at the same time. The empirical micro literature has focused virtually exclusively on measuring the effects of benefit eligibility on the search effort of unemployed workers – a focus that is too narrow to infer the impact of benefit duration on employment. The estimates in the quantitative macro literature vary widely depending on the value of parameters, which are notoriously difficult to identify. Moreover, the literature generally ignores the effect of policies on participation decisions of those out-of-the-labor force, which limits their ability to measure the total effect on employment. Indeed, in the data the flow from non-participation into employment accounts for over 60% of all transitions into employment.

We propose to sidestep these difficulties by directly measuring the employment effects of a large nationwide cut in benefit duration in December 2013. The attractive property of

this quasi-natural experiment is that its effects can be measured using standard empirical techniques that do not require imposing assumptions of a particular labor market model on the data. Specifically, we measure the employment impact of the December 2013 decision by Congress to terminate the Emergency Unemployment Compensation Act of 2008 (EUC08) which abruptly lowered benefit duration in all states to their regular duration of typically 26 weeks. This decision terminated an unprecedented extension of unemployment benefit durations adopted by policymakers following the onset of the Great Recession. While benefit durations began declining in some of the states starting in 2011, even by the end of 2013, right before the reform and long after the recession ended, the average benefit duration across U.S. states stood at 53 weeks.

The decision to eliminate benefit extensions at the end of 2013 was quite controversial. Summarizing the conventional wisdom at the time, the Council of Economic Advisers and the Department of Labor (2013) predicted that 240,000 jobs would be lost in 2014 because of the negative impact on aggregate demand. Many economists voiced a concern, first articulated in Solon (1979), that without access to benefits unemployed workers will stop searching for jobs and will exit the labor force instead.

However, the U.S. labor market performance in 2014 surprised many observers. Figure A-1 in the Appendix reports some basic aggregate statistics. Average employment growth was about 25% higher in 2014 than in the best of several preceding years. The employment-to-population ratio rose. The unemployment rate declined sharply. In contrast to typical predictions, the labor force participation rate suddenly halted its steady secular decline. The number of job vacancies that employers were trying to fill increased sharply.

At this level of aggregation, however, it is difficult to ascertain the extent to which these aggregate labor market developments were induced by the elimination of unemployment benefit extensions. The fact that aggregate productivity growth was slower in 2014 than in the preceding years eliminates the most prominent alternative explanation. While that can help explain the low observed wage growth in 2014, it cannot reconcile the low wage growth with the otherwise booming labor market. However, based on aggregate data alone, it appears difficult to rule out the possibility that some other aggregate shocks (coincidental with the decline in benefit duration) suddenly spurred the decisions of firms to create job vacancies and of jobless workers to accept them.

We take a different route in this paper. In particular, we exploit the fact that, at the end of 2013, federal unemployment benefit extensions available to workers ranged from 0 to 47 weeks across U.S. states. As the decision to abruptly eliminate all federal extensions applied to all states, it was exogenous to economic conditions of individual states. In particular, states did not *choose* to cut benefits based on, e.g. their employment in 2013 or expected employment growth in 2014. This allows us to exploit the vast heterogeneity of the decline in benefit duration across states to identify the labor market implication of unemployment benefit extensions.

After describing the institutional features of the U.S. unemployment insurance system and the details of the policy change in December 2013, in Section 2 we provide a basic description of patterns in the data. We perform two simple experiments: First, we partition states into two groups based on benefit durations right before the reform in December 2013. Assuming that the pre-reform employment trends in those states would have continued into 2014 (in absence of the benefit cut), we find that the cut in unemployment benefit duration led to a 2% increase in aggregate employment, accounting for nearly all of the remarkable employment growth in the U.S. in 2014. Second, we refine the measurement of underlying employment trends by comparing only counties that border each other but belong to different states. As we explain below, the underlying economic fundamentals are expected to evolve similarly across counties bordering each other. Unemployment insurance policies, determined at the state level, however, are discontinuous at the state border. Thus, a comparison of employment growth between border counties in relation to the change in benefit durations in the states to which these border counties belong, provides another way to assess the labor market implications of unemployment benefit durations. We find that employment growth was much higher in 2014 in the border counties that experienced a larger decline in benefit durations relative to the adjacent counties. What makes this finding even more remarkable is that year after year prior to 2014 the relative employment growth was lower in the high benefit counties. Once again, the analysis based on this simple inference implies that the cut in benefits in 2014 can explain nearly all of the observed aggregate employment growth in 2014. The abrupt reversal in the relative employment growth trend of high benefit states and border counties in December 2013, right at the time when the benefit durations were cut, strongly suggests that our analysis indeed identifies the implications of this particular policy change. There were no

other policy changes at the turn of 2014 likely to have significant labor market implications. Moreover, we are not aware of any policy changes that could have differentially affected states depending on their pre-reform benefit duration.

As this discussion makes clear, the key challenge to measuring the employment growth due to the cut in benefit durations is the inference on trends in employment that various locations would have experienced without a cut in benefits. We refine our measurement of these counterfactual trends in Section 3 in which we develop an econometric methodology for formally measuring the effects of unemployment benefit extensions. Our formal measurement approach continues to rely on comparisons of counties that border each other but belong to different states. However, the effect of the benefit cut is estimated along with a flexible specification of the difference in trends between border counties in each pair using an interactive effects model developed in Bai (2009). The idea underlying our approach is that the systematic response of underlying economic conditions across counties with different benefit durations is induced by differential exposure of counties to various aggregate shocks. For example, Holmes (1998) argued that border counties may differ in the share of manufacturing industry employment, due to different state right-to-work laws. In this case, aggregate shocks affecting the relative productivity of manufacturing industry will have a different impact on the employment in the two border counties. Similarly, foreclosure laws differ across states implying that the aggregate shocks to house prices have different impact on construction industry, and, say, demand for goods and services in the two border counties. This may also induce different trends in employment in the two border counties. Thus, there are numerous aggregate shocks that potentially induce different trends across border county pairs. The interactive effects estimator accounts for these trends by identifying the important unobserved aggregate shocks and measuring their heterogeneous impacts across counties. We find that the trends implied by the factor model capture the small differences in employment growth of neighboring counties very well so that we conclude that conditional on these trends, the common trend assumption is satisfied. This allows us to obtain consistent estimates of the effects of the cut in benefit durations on employment.

The results of the empirical analysis based on this methodology are presented and discussed in Section 4. We find that changes in unemployment benefits have a large and statistically significant effect on employment: a 1 percent drop in benefit duration increases employment

by 0.0161 log points. While large, this estimate is smaller than that implied by the simple experiments described above. This happens because our estimates of the interactive effects model attribute some of the observed relative increase in employment growth in high benefit counties to the effects of economic fundamentals. In the aggregate, our estimates imply that the cut in benefit duration accounted for about 61 percent of the aggregate employment growth in 2014.

In addition, we apply our methodology to assess the effect of this policy change on the labor force. We find that more than half of the increase in employment due to the cut in benefits was due to an increase in the labor force. Our analysis thus implies that not only did the unemployed not drop out of the labor force because of losing entitlement to benefits, but instead those previously not participating in the labor market decided to enter the labor force. These effects are not unexpected in light of equilibrium labor market theory which implies an increase in job creation in response to a cut in benefit duration. The increased availability of jobs than draws non-participants into the labor market.

The only other paper to provide an estimate of the impact of unemployment benefit extensions on employment is Hagedorn et al. (2013). The objective of that paper was to measure the effects of benefits on unemployment in a way that is consistent with the standard equilibrium labor search model and to assess whether the model provides a coherent rationalization of the joint evolution of various labor market variables in response to unemployment benefit extensions. That paper exploits multiple changes in benefits over time and space which necessitates the development of a novel measurement methodology that controls for agents' expectations regarding future policy changes that is consistent with the theoretical model. Our focus in this paper is instead on the measurement of the effects of a one-time permanent change in unemployment benefit extensions on employment. We exploit the variation induced by the policy reform that lends itself to the analysis using the standard tools developed by labor economists. This allows us to conduct the measurement without imposing any theoretical restrictions of a particular labor market model. Nevertheless, we compare the results in the two papers below and find that they imply a quantitatively similar negative impact of benefit extensions on employment.

2 Data and the Unemployment Insurance Reform

2.1 Policy Environment

Prior to the onset of the Great Recession, unemployed workers in most states qualified for 26 weeks of unemployment compensation paid by the state in which the lost job was located. In response to the deterioration of labor market conditions, the federal Emergency Unemployment Compensation (EUC08) program was enacted in June 2008. The program started by allowing for an extra 13 weeks of benefits to all states and was gradually expanded to have 4 tiers, providing potentially 53 weeks of federally financed additional benefits. The availability of each tier was dependent on state unemployment rates. The EUC08 program was not originally envisioned to last for many years, but was periodically reauthorized by Congress. The last annual reauthorization took place in December 2012.

In addition, the Extended Benefits (EB) program allows for 13 or 20 weeks of extra benefits in states with elevated unemployment rates. The EB program is a joint state and federal program. The federal government pays for half of the cost, and determines a set of “triggers,” related to the state insured and total unemployment rates, that the states can adopt to qualify for extended benefits. At the onset of the recession, many states chose to opt out of the program or only adopt high triggers. The American Recovery and Reinvestment Act of 2009 turned this into a federally funded program. Following this, many states joined the program and several states adopted lower triggers to qualify for the program. Most states wrote their legislation implementing their EB program in a way that provided for their participation only as long as federal government paid for 100 percent of the cost. The provision for federal financing of the EB program was reauthorized together with reauthorizations of the EUC08 program.

An important feature of the EB program is that many triggers available to the states under the federal law contain look-back provisions. In particular, the state under those triggers qualified for federal financing only if state unemployment was 110 or 120 percent (depending on a trigger) higher than in the preceding two years. In other words, the EB program could be made available under those triggers only if unemployment is rising. As a consequence, starting in 2011 some states began losing eligibility for the EB program.¹ As total duration of available

¹To mitigate this effect, the federal government temporarily gave states an option of using a three year look-back period.

unemployment benefits began declining so did the unemployment rate resulting in some states also losing eligibility for some of the tiers of the EUC08 program.

As a result, by December 2013 there was substantial heterogeneity in the actual unemployment benefit durations across U.S. states. As Table 1 shows, 3 states had 73 weeks of benefits available, 20 states had 61-63 weeks, 9 states had 54-57 weeks, 18 states had 40-49 weeks, and one state had 19 weeks. These data on unemployment benefit durations in each state is based on trigger reports provided by the Department of Labor. These reports contain detailed information for each of the states regarding the eligibility and activation status of the EB program and different tiers of the EUC08 program.²

In December 2013 the Congress did not reauthorize the EUC08 program. As there is no “phase-out” period for EUC08 payments, all EUC08 payments ceased abruptly in all states when the program ended. Specifically, individuals who exhausted regular state unemployment compensation after December 21, 2013 (in NY, December 22, 2013) were no longer eligible for EUC08. For unemployed individuals already participating in the EUC08 program, the last payable week of EUC08 benefits was the week ending December 28, 2013 (in New York, December 29, 2013). EB program came to an end at the same time so that by January 2014 no states were offering extended benefits under this program.

From the moment the unemployment benefit extensions came to an end in December 2013, newly unemployed individuals could only qualify for the regular state unemployment compensation for a duration of 26 weeks in most states. Some states had less than 26 weeks available in 2014, including Arkansas (25), Florida (16), Georgia (18), Kansas (20), Michigan (20), Missouri (20), North Carolina (19) and South Carolina (20). Two states – Massachusetts (30) and Montana (28) – offered more generous benefit durations. Thus, the average benefit duration across states dropped from 53 to 25 weeks in December 2013.

An important property of the decision not to renew benefit extensions in December 2013 is that it applied to all states, regardless of their economic conditions. In particular, the states could not choose whether to be treated by this reform, for example, based on their employment in 2013 or expected employment growth in 2014. The fact that the policy change was exogenous from the point of view of an individual state, allows for a relatively straightforward

²See <http://ows.doleta.gov/unemploy/trigger/> for trigger reports on the EB program and <http://ows.doleta.gov/unemploy/euc.trigger/> for reports on the EUC08 program.

Table 1: Benefit Duration across States in December 2013

Weeks of Benefits	states
73 weeks	Illinois, Nevada, Rhode Island
63 weeks	Alaska, Arizona, California, Connecticut, Delaware, DC, Indiana, Kentucky, Louisiana, Maryland, Massachusetts, Mississippi, New Jersey, New York, Ohio, Oregon, Pennsylvania, Tennessee, Washington
61 weeks	Arkansas
57 weeks	Michigan
54 weeks	Alabama, Colorado, Idaho, Maine, New Mexico, Texas, West Virginia, Wisconsin
49 weeks	Missouri, South Carolina
44 weeks	Georgia
40 weeks	Florida, Hawaii, Iowa, Kansas, Minnesota, Montana, Nebraska, New Hampshire, North Dakota, Oklahoma, South Dakota, Utah, Vermont, Virginia, Wyoming
19 weeks	North Carolina

identification of the effect of the policy change on that state’s labor market. This contrasts sharply with the gradual decline in benefit durations in many states since 2011. While those declines could have had significant labor market implications, those policy changes were endogenous to a state’s labor market conditions, making the identification of the effects of policies challenging.

While from the outset, the federal unemployment benefit extension program was understood to be temporary, the decision to stop the program in December 2013 came largely as a surprise. Indeed, by December 2013 the program had been re-authorized a dozen of times. By that time it had paid benefits for a record 66 months, over two years longer than any prior discretionary benefit extension program. However, the U.S. unemployment rate was higher and the long-term unemployment rate was at least twice as high as it was at the expiration of every previous unemployment benefit extension program. Moreover, the Council of Economic Advisors, the Congressional Budget Office and others forcefully argued for the extensions on the grounds that EUC08 is among policies with “the largest effects on output and employment per dollar of budgetary cost.” In light of this, few expected the Congress to terminate the program in December 2013. Even following the Congress’ decision, there was likely some uncertainty regarding its finality throughout the first half of 2014. For example, on April 7, 2014, the Senate narrowly approved a bipartisan bill that would have restored (retroactively to December 2013) federal funding for extended unemployment benefits. The bill faced a determined opposition in the Congress, which refused to hold a vote on it. Note that, to the extent that economic agents were able to forecast the expiration of unemployment benefit extensions prior to December 2013 and adjusted their actions accordingly, and to the extent that they were uncertain about the possibility of the extensions being re-authorized at some point in 2014, our estimates will provide a lower bound on the effects of the policy change.

2.2 A First Look at the Data

As a first step in exploring whether this exogenous policy change helps account for some of the observed rise in employment, we compare the evolution of employment in states with high benefit duration to the evolution of employment in states with relatively low benefit duration in December 2013. Specifically, we split the states into two groups based on weeks of benefits

Table 2: Average State Employment Changes and Benefits

States	Employment Change 2013	Employment Change 2014	Δ Growth
High Benefit	-0.38%	0.30%	+0.68%
Low Benefit	0.37%	-0.29%	-0.66%
Δ States	- 0.75%	+0.59%	= 1.34 %

available immediately prior to the policy change in December 2013. The “high benefit” group includes all states which had strictly more than 54 weeks of benefits in December 2013 and the “low benefit” states are those with weakly less than 54 weeks of benefits in December 2013.³ The average duration of benefits was 63 weeks in the first group of states and 44 weeks in the second.

Data on employment and the labor force in each U.S. county and state are from the Local Area Unemployment Statistics (LAUS) provided by the Bureau of Labor Statistics.⁴

Table 2 describes the average employment change across states in each group in 2013 (December 2012 to December 2013) relative to the overall average for 2013, as well as the average change in employment in 2014 (December 2013 to October 2014), again relative to the overall average for 2014.⁵ Employment in high benefits states grew 0.38 percentage points less than the average in 2013 whereas employment in low benefit states increased by 0.37 percent points more than the average, a growth difference of -0.75 percent. This ranking of economic performance flipped in 2014. Employment in high benefit states grew by 0.3 percentage points more than the average, whereas the low benefit states which grew by 0.29

³The “high benefit” group includes 23 states and the District of Columbia. The “low benefit” group includes the remaining 27 states. The assignment of individual states to groups is clear from Table 1.

⁴<ftp://ftp.bls.gov/pub/time.series/la/>. Data accessed 12/29/14. Note that there is some discussion in the profession on whether LAUS unemployment data reflects genuine unemployment in the county or is to some extent imputed using state-level data. Hagedorn et al. (2013) provide direct evidence that such concerns are unfounded. Moreover, such concerns do not apply to the employment and labor force data which is the focus of the analysis in this paper. Finally, such concerns are only relevant when one is concerned with potential policy endogeneity. The policy change that is the focus of this paper was exogenous to cross-sectional differences across U.S. states, mitigating such potential concerns.

⁵At the time of writing the 2014 data at the county level is only available up to and including the month of October. For comparability of the results, we restrict our analysis of the state-level data to the same time period.

percentage points less, a difference in growth rates of 0.59 percentage points. The total change in the growth rate difference is therefore 1.34 percent.

These considerations suggest that a difference of $19 = 63 - 44$ weeks is associated with an increase in employment by 1.37 percent. The average benefit duration fell from 53 to 25 weeks in December 2013, which would imply an increase in employment by

$$\frac{53 - 25}{63 - 44} 1.34\% = 2.0\%. \quad (1)$$

Thus, this simple calculation suggests a large impact of the cut in benefits on employment. The implied employment growth due to the cut in benefits is nearly identical to the increase of 2.1 percent in U.S. employment in 2014, when the U.S. economy created 2,952,000 new jobs.

It is, however, not clear that the change in employment can be attributed to the change in benefits only. The reason is that the shocks which drive employment in high benefit states may be different from the shocks in low benefit states and these shocks may not be orthogonal to the different benefit levels in these two groups of states. In other words, it is likely that the fact that some states had high benefit durations in 2013 in part reflected worse economic fundamentals in those states relative to states with lower benefit durations. The simple calculation performed above implicitly assumed that the trends in those fundamentals among high and low benefit states would have remained the same in 2014 as they were in 2013. While this assumption appears quite plausible, it seems desirable to weaken it.

To do so, we now consider a more disaggregated approach. Specifically, we focus our analysis on a sample of county pairs that belong to different states and share a border. There are 1,178 such border county pairs for which we have complete data. Comparing the evolution of employment in counties that border each other but belong to different states overcomes the potential endogeneity problem. Neighboring locations separated by a state border share the same geography, climate, access to transportation, agglomeration benefits, access to specialized labor and supplies, etc. Indeed, Hagedorn et al. (2013) provide direct evidence that economic shocks do not stop at the state border but evolve smoothly across borders. The key feature that sets these locations apart is the difference in policies on the two sides of the border (unemployment benefit policies are set at the state level and apply to all counties within a




Figure 1: Difference in Employment Rate between High and Low Benefit Border Counties

state).

These observations imply that absent any policy differences, the employment trends induced by fundamental economic shocks are expected to be similar across border counties in the same pair. Moreover, as the median border county has only one half of one percent of its state’s employment, it seems plausible that changes in employment trends in an individual county are unlikely to induce unemployment policy changes at the state level.⁶ Thus, our next experiment is based on the assumption that the difference in benefit levels across two neighboring counties (determined at the state level) is not correlated with the difference in employment trends across the two counties. This assumption is clearly much weaker here than in previous work based on the border-county methodology as the policy change at the end of 2013 was exogenous to cross-sectional differences across U.S. states.

Figure 1 shows the average difference in employment across all border counties from 2005

⁶In the formal analysis below we will assess the sensitivity of the results to restricting the sample to counties that are small relative to the state they belong to.

to October 2014, where the county which had higher benefits in December 2013 is first.⁷ For example, Fairfax County, Virginia had 40 weeks of benefits available in December 2013, whereas its border county in Maryland, Montgomery County, had 63 weeks of benefits available. Thus, in every period the figure would reflect the employment rate in Montgomery County minus the employment rate in Fairfax County. The series represents the average of such differences among all border county pairs.

A sudden reversal of fortune experienced in 2014 by high benefit counties, evident in Figure 1, suggests that the cut in benefits led to a substantial increase in employment. After a long period of relative employment losses, the high benefit counties experienced a relative employment gain of 0.65 percent in 2014. As the average benefit duration before the policy change was 56.7 weeks in the high benefit counties and 47.6 weeks in the low benefit counties, the implied total employment gain from cutting benefits from an average level of 53 weeks to 25 weeks equals:

$$\frac{53 - 25}{56.7 - 47.6} 0.65\% = 2.0\%. \quad (2)$$

Thus, this experiment also suggests a large increase in employment due to the cut in benefit durations. Note that this experiment based on the border county pairs assumes that the average trend in underlying economic fundamentals in 2014 is the same in the set of high and low benefit counties. Thus, only the employment growth following the cut in benefits was used in this experiment. Had we assumed that the trends in fundamentals in 2014 were the same as in 2013, as we did in the experiment based on state data above, the implied effect would have been even larger (the corresponding data can be found in Appendix Table A-1).⁸

It might still be possible, however, that some unobserved trends would have led to high employment growth in border counties belonging to high benefit duration states in 2014 even in the absence of the change in benefit durations. To address this concern, in the next section we perform a more sophisticated econometric analysis that includes the estimation of a flexible specification of the difference in trends between border counties in each pair using an

⁷Data are aggregated to a quarterly frequency.

⁸Consistent, with the assumptions underlying the border county pair based inference, the average difference in 2013 employment growth between high and low benefit counties in a pair is much smaller than the difference in 2013 employment growth between high and low benefit states.

interactive effects model.⁹

3 Empirical Methodology

3.1 Identification via Border Counties

Our objective is to measure the effect of the cut of benefit durations in December 2013 on employment. The effect of this particular policy change on log employment $e_{i,t}$ in county i in calendar quarter t is captured by the coefficient α in the regression equation

$$e_{i,t} = \alpha \mathcal{I}_{t \geq 2013/Q4} \log(b_{i,t}) + \epsilon_{i,t}, \quad (3)$$

where $\mathcal{I}_{t \geq 2013/Q4}$ is one for the policy change period starting in 2013 Q4 and is zero otherwise and $b_{i,t}$ measures available benefit duration. To ensure that the identification of the effect of the change in benefit duration on employment is determined by the variation induced by this specific policy change, we do not place any assumptions on how employment and benefits are related prior to this period.

Implementing such a regression at the county level might suffer from the same endogeneity problems as a regression of state employment on benefit durations as discussed above. To avoid a bias arising from endogeneity, we first consider the difference in employment of a pair p of counties i and j which border each other but belong to different states. For each border-county pair p , we difference Equation (3) between the two border counties i and j :

$$\Delta e_{p,t} = \alpha \mathcal{I}_{t \geq 2013/Q4} \Delta b_{p,t} + \Delta \epsilon_{p,t}, \quad (4)$$

where Δ is the difference operator over counties in the same pair. More specifically, if counties i and j are in the same border-county pair p , then $\Delta e_{p,t} = e_{p,i,t} - e_{p,j,t}$ and $\Delta b_{p,t} = \log(b_{p,i,t}) - \log(b_{p,j,t})$.

The term $\Delta \epsilon_{p,t}$ contains the permanent differences in employment e across border counties caused by, e.g., permanent differences in tax policies across states they belong to. It also contains differences in employment trends. As discussed in the Introduction, these systematic

⁹In addition, the econometric analysis below corrects for a composition bias in aggregation present in the simple calculations in this section.

differences in trends arise due to the different exposure of counties to various aggregate shocks. For example, shocks to various sectors of the economy, while aggregate in their nature, are expected to have different impacts on counties depending on their sectoral composition. Clearly, shocks to the financial industry, driven in part by the evolving regulatory environment, or changes in the price of oil and gas and technological improvements in hydraulic fracking, have important but different impacts on the counties on the border of New York and New Jersey, from the counties on the border between Pennsylvania and Ohio. Thus, there are clearly numerous aggregate shocks that potentially induce heterogeneous trends across different border county pairs. To account for these trends we follow the approach in Bai (2009) who has shown that consistency and proper inference can be obtained in a panel data context, such as ours, through the use of an interactive-effects estimator.¹⁰ In particular, we decompose the error term in Equation (4) as

$$\Delta\epsilon_{p,t} = \lambda_p' F_t + \nu_{p,t}, \quad (5)$$

where λ_p ($r \times 1$) is a vector of pair-specific factor loadings and F_t ($r \times 1$) is a vector of time-specific common factors. Our baseline specification can then be written as

$$\Delta e_{p,t} = \alpha \mathcal{I}_{t \geq 2013/Q4} \Delta b_{p,t} + \lambda_p' F_t + \nu_{p,t}. \quad (6)$$

As is shown in Bai (2009), this model incorporates additive time and county pair fixed effects as special cases. It is, however, much more general and allows for a very flexible model of the heterogeneous time trends at the county pair level. The key to estimating α consistently is to treat the unobserved factors and factor loadings as parameters to be estimated. Our implementation is based on an iterative two-stage estimator described in Appendix I. Note that the factor loadings are mainly identified from the period before 2013/Q4 as our estimation sample starts in 2005/Q1.

The identification assumption is that

$$Corr(\mathcal{I}_{t \geq 2013/Q4} \Delta b_{p,t}, \nu_{p,t}) = 0. \quad (7)$$

¹⁰Note that in the presence of aggregate shocks having heterogeneous impacts on county pairs, estimating the panel regression in Equation (4), perhaps with a set of county pair and time fixed effects, is generally problematic for inference (see Andrews (2005) for the discussion of this problem in a cross-sectional regression). Gobillon and Magnac (2013) establish the superior performance of the interactive effects estimator relative to alternatives methods.

As in a simple difference-in-differences analysis above, the parameter of interest α is identified off the change in benefits when EUC08 expires. The expiration of benefits is exogenous with respect to cross-sectional differences in county employment. To understand what assumption (7) rules out, imagine for a moment that our dataset contained just two counties i and j in a pair p where county i has higher benefits at the end of 2013 than county j . In this case our estimate would not recover the true effect of benefits on employment if county i would have had higher employment growth than county j in the absence of the benefit cut. In this case we would attribute some of the differences in employment growth to the cut in benefits although not all of the employment differences are related to benefits.

However, our dataset contains not just one county pair but 1178 of such pairs. The identifying assumption then rules out that the higher benefit counties would have had on average higher employment growth in the absence of the policy reform and does not rule it out for every individual pair. Figure 1 lends support to the identifying assumption as, prior to the policy change, high benefits counties did not show on average faster employment growth than their low benefits counterparts in the border pair. Instead one clearly sees a sudden rise in employment growth just when benefits were cut at the turn of 2014. Moreover, as discussed above, with the exception of the expiration of EUC08, there were no policy changes or other developments that could have plausibly induced the co-movement between the size of the benefit cut and the subsequent employment growth across border counties. Finally, it is also important to note that to violate the identifying assumption, the higher average employment growth in the higher benefit counties in the absence of the experiment would have to be purely mechanical since counties and states could select into the experiment neither based on their employment in December 2013 nor on their expected employment growth in 2014. Thus, the exogeneity of the program rules out a version of a behavioral Ashenfelter’s “dip.”

The identifying assumption becomes even weaker once one recognizes that this correlation is conditional on using the interactive effects model to remove the trends from the data. That is our assumption even allows for a correlation of counterfactual employment growth and benefits as long as it is captured by the factor model.

Equation (6) can be estimated in the data to recover the coefficient of interest α . We then use this estimate to compute the percentage increase in U.S. employment in 2014 that

is caused by the cancellation of extended benefits as

$$\pi^E = \alpha \sum_{\text{All U.S. states } s} (\log(b_s^{2014}) - \log(b_s^{2013})) \frac{E_s^{2013}}{E_{US}^{2013}}, \quad (8)$$

where b_s^{2013} denotes the number of weeks of benefits available in state s in December 2013 (just prior to the policy change), b_s^{2014} is the number of weeks of benefits available in state s in 2014, E_s^{2013} is employment in state s in December 2013 and E_{US}^{2013} is the aggregate U.S. employment in December 2013. The corresponding gain in the total number of employed then equals

$$\Delta^E = \frac{\pi^E \times E_{US}^{2014}}{1 + \pi^E}, \quad (9)$$

where E_{US}^{2014} refers to U.S. employment in December 2014.

Estimating Equation (6) but replacing the difference in the log of the number of employed in the border county pair on the left hand side with the corresponding difference in the log of the number of labor force participants allows us to compute the effect of the cut in benefits on the labor force. Using the analogues to Equations (8) and (9), we can then measure the percentage increase in the labor force π^L and the increase in the number of labor force participants Δ^L as a consequence of the policy reform.

3.2 Estimating the Number of Factors

To implement the interactive effects estimator, we need to specify the number of factors. Bai and Ng (2002) have shown that the number of factors in pure factor models can be consistently estimated based on the information criterion approach. Bai (2009) shows that their argument can be adapted to panel data models with interactive fixed effects. Thus, we define our criterion CP as a function of the number of factors k as:

$$CP(k) = \hat{\sigma}^2(k) + \hat{\sigma}^2(\bar{k}) [k(N + T) - k^2] \frac{\log(NT)}{NT},$$

where $\bar{k} \geq r$ is the maximum number of factors, N is the number of pairs, T is the number of time observations, $\hat{\sigma}^2(k)$ is the mean squared error, defined as

$$\hat{\sigma}^2(k) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \left(\Delta e_{p,t} - \alpha \mathcal{I}_{t \geq 2013/Q4} \Delta b_{p,t} - \lambda'_i(k) F_t(k) \right)^2,$$

and $F_t(k)$ and $\lambda'_i(k)$ are the estimated factors and their loadings, respectively, when k factors are estimated. To avoid collinearity, we set \bar{k} to the minimum of seven and $T - 1$, one less than the total number of time observations. Our estimator for the number of factors is then given by

$$\hat{k} = \arg \min_{k \leq \bar{k}} CP(k).$$

3.3 Standard Errors

To properly compute standard errors, we need to take into account the potential correlation in the residuals across counties and over time. There are two possible sources of correlation. First, the employment and unemployment outcomes that we are interested in are highly serially correlated. This aspect of the data may cause serial correlation in the errors. Second, the fact that some counties appear in multiple county-pairs results in an almost mechanical correlation across county pairs. To account for these sources of correlation in the residuals, we follow Bertrand et al. (2004) and use the block-bootstrap on state border segments to compute standard errors.

4 Unemployment Benefit Extensions and Employment

4.1 Baseline Empirical Results

Column (1) of Table 3 contains the results of the estimation of the effect of unemployment benefit duration on employment using the baseline specification in Equation (6). We find that changes in unemployment benefits have a large and statistically significant effect on employment: a 1 percent drop in benefit duration increases employment by 0.0161 log point.¹¹

Our estimate implies that the drop in benefit duration led to a percentage increase in

¹¹This corresponds to an effect of -0.004 log points per quarter. This is slightly larger but comparable to the corresponding effect of -0.0035 estimated in Hagedorn et al. (2013).

Table 3: Unemployment Benefit Extensions and Employment

VARIABLES	Employment	Labor Force
Weeks of Benefits	-1.61 (0.010)	-0.75 (0.020)
Number of Factors	5	5
Observations	47,111	47,111

Note - All coefficients are multiplied by 100. Bold font denotes significance at a 95% level based on bootstrapped p-values in parentheses.

employment of¹²

$$\pi^E = \underbrace{\alpha}_{-0.0161} \underbrace{\sum_{\text{All U.S. states } s} (\log(b_s^{2014}) - \log(b_s^{2013})) \frac{E_s^{2013}}{E_{US}^{2013}}}_{-0.799} = 0.013, \quad (10)$$

that is U.S. employment increased by 1.3 percent due to the cut of benefit durations. The corresponding gain in the total number of employed then equals

$$\Delta^E = \frac{0.013 \times E_{US}^{2014}}{1 + 0.013} = 1,801,096. \quad (11)$$

The estimated trends for the difference in employment between high and low benefit duration border counties¹³ is plotted in Figure 2. Prior to 2014 the high benefits counties had been losing employment relative to the low benefit counties and these small differences in employment growth are well captured by the trend implied by the factor model. We can conclude that the common trend assumption conditional on the factor model is satisfied. Thus the employment gain counties experienced in 2014 is not a continuation of a previous trend. In fact, had that negative trend continued, we would have had to compare the employment gains relative to this negative trend, which would imply even larger employment gains than we found. The estimated trend is however not decreasing anymore in 2014 but instead is slightly increasing. This rationalizes why our estimates of the negative employment effects of unemployment

¹²Note that the effects on employment are downward biased to the extent that individual decisions in which of the border counties to live in are separated from the decision of which one to work in. Hagedorn et al. (2013) find the associated bias to be small and report only negligible amount of worker mobility across border counties in response to changes in benefit durations.

¹³As determined by the duration of benefits in December 2013.




Figure 2: Difference in Employment Rate between High and Low Benefit Border Counties and Estimated Trend

benefit extensions are smaller than the ones implied by the simple difference-in-differences analysis in Section 2.

We can also use Equation (6) with labor force on the left hand side to estimate the percentage change in the labor force attributable to the cancellation of policy. Estimating this equation, we find that a 1 percent drop in benefit duration increases the labor force by 0.0075 log points. The percentage change in the size of the labor force in the U.S. due to the cancellation of benefits then equals

$$\pi^L = \underbrace{\alpha^L}_{-0.0075} \underbrace{\sum_{\text{All U.S. states } s} (\log(b_s^{2014}) - \log(b_s^{2013})) \frac{L_s^{2013}}{L_{US}^{2013}}}_{-0.801} = 0.006, \quad (12)$$

and the corresponding increase in the size of the labor force equals

$$\Delta^L = \frac{\pi^L \times L_{US}^{2014}}{1 + \pi^L} = 931,887. \quad (13)$$

Thus, more than half of the increase in employment was due to the increase in the labor force as a result of the reduction of benefit duration. The remaining increase corresponds to a decrease in the number of unemployed by $869,209 = 1,801,096 - 931,887$. Our analysis thus shows that the dominant impact of the benefit cut on employment was not driven by a contraction in the labor force – unemployed dropping out of the labor force because they were no longer entitled to benefits – but instead by those previously not participating in the labor market deciding to enter the labor force.

It is also interesting to note that the existing empirical literature has mainly attempted to measure the “micro” effect of unemployment benefit duration on search intensity and job acceptance decisions of individual workers. Hagedorn et al. (2014) find these effects to be very small, confirming the sentiment in the literature. Clearly, this micro effect is zero for those out-of-labor force who were entitled to benefits neither in 2013 nor in 2014. Yet, it was predominantly movements from out-of-labor force that drove the rise in employment in 2014. Presumably this happened due to a large “macro” effect of the benefit cut on job creation. It is then the availability of jobs that drew non-participants back into the labor force.

When comparing the magnitude of the effects we find to the experience in the data, it is also important to keep in mind that our estimates are based on the differences across border counties. Thus, the effects of various other shocks or policies that affect these counties symmetrically are differenced out.

4.2 Robustness

4.2.1 2013 Placebo Reform

Our results imply that the turning point for employment in 2013/2014 of high relative to low benefit countries is caused to a large degree by the exogenous cut in benefits in December 2013. In particular, the turning point is not the result of an employment adjustment which would have happened anyway and with the simultaneous cut in benefit durations being a pure coincidence. To further strengthen the point that the co-movement of the benefit cuts and the employment boom is not random, we conduct a placebo analysis for the year 2013 instead of 2014 as above. To this aim, we shift the analysis back by one year and assume (counterfactually) that benefits were cut at the end of December 2012 (and not in December

Table 4: Unemployment Benefit Extensions and Employment

	(1)	(2)	(3)	(4)	(5)	(6)
Weeks of Benefits	-1.61 (0.010)	0.41 (0.315)	-1.70 (0.005)	-2.12 (0.015)	-1.83 (0.015)	-1.89 (0.005)
Number of Factors	5	5	5	5	5	5
N	47,111	44,551	42,399	30,600	22,711	37,751

Note - All coefficients are multiplied by 100. Bold font denotes significance at a 95% level based on bootstrapped p-values in parentheses.

Column(1): Baseline

Column(2): Placebo Analysis for 2013

Column(3): Sample of counties with employment share < 15%

Column(4): Sample of counties within the same CBSA

Column(5): Sample of counties with similar industrial composition

Column(6): Sample of counties excluding regular state benefit duration changes in 2013/14

2013) to the regular level of 26 weeks. We implement the same factor model on data starting in 2005 and ending in Q3/2013 (instead of Q3/2014). The result reported in Column (2) of Table 4 show that the placebo reform had no effect on employment.

Inspecting Figure 1 makes this result not very surprising as there is only a turning point in 2014 but not one in 2013. We conclude that both in 2013 and 2014 high benefit counties did not experience higher employment growth than low benefit counties due to reasons unrelated to benefits. The difference between 2013 and 2014, however, is that in 2014 benefits were cut whereas in 2013 such a cut did not happen. As a result, we find an employment boom in 2014 but not in the placebo reform. Performing the placebo experiment at points in time other than December 2012 leads to the same conclusions.

4.2.2 Dropping Large Counties

One motivation underlying our use of the empirical methodology based on comparisons between border counties was that the unemployment insurance policies are set at the state level. Thus, if individual counties are small relative to the state they belong to, changes in employment trends in an individual county would not induce unemployment policy changes at the state level. In other words, the duration of benefits prior to the reform in December 2013 (and

the resulting cut in benefits) are not determined by the economic fundamentals of individual counties in our border sample. While this is likely to be true for a median border county that has only one half of one percent of its state’s employment, some border counties are relatively large. To assess whether the presence of large counties in the sample has an important effect on the results, we now drop any pair where a county within the pair constitutes more than 15% of its states’ employment. The results reported in Column (3) of Table 4 indicate that this is not the case. The measured effect of unemployment benefit duration on employment remains virtually unaffected.

4.2.3 Border Counties within the same CBSAs

Another motivation underlying our use of the empirical methodology based on comparisons between border counties was that, absent any policy differences, the employment trends induced by fundamental economic shocks are expected to be similar across border counties in the same pair. Yet, the distance between border counties and the degree of their economic integration varies across border county pairs. To assess whether this heterogeneity has important implications for our findings, we now restrict attention to a subset of border counties with most integrated labor markets (and with population centers close to each other). To do so, we repeat the analysis on a sample of border counties that belong to the same Core Based Statistical Areas (CBSAs). CBSAs represent a geographic entity associated with at least one core of 10,000 or more population, plus adjacent counties that have a high degree of social and economic integration with the core (see Office of Management and Budget (2010) for detailed criteria). The results reported in Column (4) of Table 4 imply a slightly larger effect of unemployment benefit extensions on employment than the one found in our full sample.

4.2.4 Border Counties with Similar Industrial Composition

One reason for the presence of heterogeneous trends across county pairs and for our use of the interactive effects model was that states and counties may systematically differ in sectoral composition of employment so that aggregate changes in sectoral demand or productivity may induce heterogeneous trends in local-level employment. For example, Holmes (1998) has pointed out that the density of manufacturing industry employment varies systematically across counties within border pairs that belong to states with different right-to-work legisla-

tion. It is then possible that states with a large manufacturing sector had low employment and long benefit durations at the end of 2013. It is also possible that if 2014 witnessed a sharp rebound in demand or productivity of manufacturing industries unrelated to a change in unemployment benefits. This can potentially give rise to an endogeneity problem, and lead us to attribute this sectoral shock to the effect of the change in unemployment insurance policy. If this heterogeneity in sectoral composition across states is sufficiently empirically important, however, it will be picked up by the interactive effects estimator.

Thus, as a check on the performance of the interactive effects estimator, we now investigate whether differences in industrial composition affect our results. To this aim, we repeat the benchmark analysis on a subset of border counties with similar industrial composition. If the effects of industrial composition were not captured by the factor model and affected our inference, we would expect a different result on this subsample than on the full sample. We obtain data on county employment by industry from the Bureau of Economic Analysis, Regional Economic Information System.¹⁴ Using sample average industry employment shares within each county, we construct the l^2 -distance between border counties within each pair. The results, presented in Column (5) of Table 4, are based on the sample of 50% of county pairs with the most similar industrial composition out of all border county pairs. The effect of unemployment benefit extensions on employment on this subsample is slightly larger than the one found on our full sample.

4.2.5 Dropping Counties in States Changing Regular Benefit Durations

Kansas and North Carolina changed their benefit policies in the regular state unemployment insurance programs in 2013, and in 2012 Florida and Georgia adopted regular unemployment benefit durations which depend on the state unemployment rate. Since these changes may be endogenous to state conditions, we repeat the analysis on the sample excluding those states. The results of the estimation on the resulting sample of 944 county pairs are reported in Column (6) of Table 4. They indicate that excluding these states has little impact on the estimated effect of benefit duration on employment.

¹⁴<http://www.bea.gov/regional/>

5 Conclusion

In this paper we measure the effect of unemployment benefit extensions on employment. We exploit the variation induced by the decision of the U.S. Congress in December of 2013 to abruptly stop all federal unemployment benefit extensions. The particular usefulness of this policy change for understanding the employment effects of benefit extensions stems from the fact that the policy change at the national level was exogenous to economic conditions of individual states. Following the aftermath of the Great Recession, there was a wide heterogeneity of the federally-financed durations of benefits across U.S. states by December 2013, ranging from 0 to 47 weeks on top of the regular state-funded benefits with typical duration of 26 weeks. Averaged across all states, total benefit duration fell sharply from 53 to 25 weeks in December 2013.

A simple descriptive analysis shows a much faster employment growth in 2014 in high benefit states prior to the reform relative to their low benefit counterparts. The same finding holds if we compare the employment growth in counties that belong to high benefit states relative to their neighboring counties that belong to states with lower benefit durations prior to the reform. The implied magnitude of the negative effect of benefit duration on employment is so large that it can account for almost the entire remarkable employment growth experienced by the U.S. in 2014.

Our formal econometric analysis tackles the key challenge of precisely measuring the counterfactual employment growth that various locations would have experienced without a cut in benefits. Our formal measurement approach continues to rely on the comparisons of counties that border each other but belong to different states. However, the effect of the benefit cut is estimated alongside with a flexible specification of the difference in trends between border counties in each pair using an interactive effects model. We find that after controlling for these heterogeneous employment trends, changes in unemployment benefits continue to have a large and statistically significant effect on employment: a 1 percent drop in benefit duration increases employment by 0.0161 log point. In the aggregate, our estimates imply that the cut in benefit duration accounted for about 61 percent of the aggregate employment growth in 2014.

While we did not impose any theoretical restrictions of a particular labor market model on

our empirical analysis, the findings are consistent with the standard equilibrium labor market search model. For example, the primary labor market effect of a cut in unemployment benefit duration in the framework of Mortensen and Pissarides (1994) is the positive impact on job creation. It is this rise in job creation that leads in equilibrium to the increase in employment.

Another important finding in this paper concerns the effect of unemployment benefit duration on labor force participation. Prior to the reform, the consensus in the profession seemed to predict a negative impact of the cut in benefit durations on the size of the labor force. Instead, we found that the reform led to almost a million non-participants entering the labor market. It seems plausible that they were encouraged by the improved probability of finding jobs due to the positive effect of the reform on job creation.

It seems quite remarkable that, despite their clear importance, the aggregate labor market implications of unemployment benefit policies have been virtually unexplored in the empirical literature. This gap in knowledge seems limiting not only for our ability to develop good economic theories but also for making sound policy choices. For example, unemployment benefit extensions are routinely used for the purposes of macroeconomic stabilization. Yet, the findings in this paper imply that the negative effects of unemployment benefit extensions on employment far outweighs the potential stimulative effects often ascribed to this policy. It appears important to take these effects into account.

References

- ANDREWS, D. W. K. (2005): “Cross-Section Regression with Common Shocks,” *Econometrica*, 73, 1551–1585.
- BAI, J. (2009): “Panel Data Models with Interactive Fixed Effects,” *Econometrica*, 77, 1229–1279.
- BAI, J. AND S. NG (2002): “Determining the Number of Factors in Approximate Factor Models,” *Econometrica*, 70, 191–221.
- BERTRAND, M., E. DUFLO, AND S. MULLAINATHAN (2004): “How Much Should We Trust Differences-In-Differences Estimates?” *Quarterly Journal of Economics*, 119, 249–275.
- COUNCIL OF ECONOMIC ADVISERS AND THE DEPARTMENT OF LABOR (2013): “The Economic Benefits of Extending Unemployment Insurance,” Report, The White House.
- GOBILLON, L. AND T. MAGNAC (2013): “Regional Policy Evaluation: Interactive Fixed Effects and Synthetic Controls,” IZA Discussion Paper 7493.
- HAGEDORN, M., F. KARAHAN, I. MANOVSKII, AND K. MITMAN (2013): “Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects,” NBER Working Papers 19499, National Bureau of Economic Research, Inc.
- (2014): “Unemployment Benefits and Unemployment in the Great Recession: The Role of Micro Effects,” mimeo, University of Pennsylvania.
- HOLMES, T. J. (1998): “The Effect of State Policies on the Location of Manufacturing: Evidence from State Borders,” *Journal of Political Economy*, 106, 667–705.
- MORTENSEN, D. T. AND C. PISSARIDES (1994): “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 61, 397–415.
- OFFICE OF MANAGEMENT AND BUDGET (2010): “2010 Standards for Delineating Metropolitan and Micropolitan Statistical Areas; Notice,” *Federal Register*, 75, 37245–37252.
- SOLON, G. (1979): “Labor Supply Effects of Extended Unemployment Benefits,” *Journal of Human Resources*, 14, 247–255.

APPENDICES

I Implementation of Iterative Two-Stage Estimator

The following is a brief description of the algorithm implementing our iterative two-stage estimator.

1. Start with a guess for α , say α_1 .
2. At each iteration ξ , do the following:
 - (a) given α_ξ , for each p , construct $v_{p,t} = \Delta e_{p,t} - \alpha_j \mathcal{I}_{t \geq 2013/Q4} \Delta b_{p,t}$.
Then, $v_{p,t} = \lambda'_p F_t$ is a pure factor model and can be estimated consistently using principal components.¹⁵
 - (b) Given the estimates for λ_p and F_t , estimate equation (6) via OLS and update the guess to obtain $\alpha_{\xi+1}$.
3. Repeat 2 until α_ξ converges.¹⁶

¹⁵The exposition of the estimator assumes that there are no missing observations. We use the generalized procedure described in Bai (2009) and allow for missing observations.


¹⁶Hagedorn et al. (2013) have conducted a number of Monte Carlo simulations with sample sizes similar to our sample and found the estimator described here to converge to the true parameter.

II Appendix Tables


Table A-1: Employment Changes in Border County Pairs and Benefits

Counties	Employment Change 2013	Employment Change 2014	Δ Growth
High Benefit	-0.12%	0.33%	+0.45%
Low Benefit	0.12%	-0.32%	-0.44%
Δ States	- 0.24%	+0.65%	= 0.89%


III Appendix Figures




(a) Monthly Job Growth.




(b) Employment to Population Ratio.




(c) Unemployment Rate Change.



(d) Labor Force Participation Growth.



(e) Monthly Vacancy Growth.



(f) Productivity Growth.

Figure A-1: U.S. Labor Market Performance in 2014.

Note - Data series downloaded from the Bureau of Labor Statistics website <http://www.bls.gov/data/> on 01/09/2015 with the following series identifiers:

Panel (a) - CES0000000001, Panel (b) - LNS12300000, Panel (c) - LNS14000000, Panel (d) - LNS11300000, Panel (e) - JTS00000000JOL, Panel (f) - PRS85006093.