IRLE
Institute for Research on Labor and Employment

IRLE WORKING PAPER
#181-09
June 2009

Spatial Heterogeneity and Minimum Wages: Employment Estimates for Teens Using Cross-State Commuting Zones

Sylvia Allegretto, Arindrajit Dube, Michael Reich

Cite as: Sylvia Allegretto, Arindrajit Dube, Michael Reich. (2009). "Spatial Heterogeneity and Minimum Wages: Employment Estimates for Teens Using Cross-State Commuting Zones." IRLE Working Paper No. 181-09.
http://irle.berkeley.edu/workingpapers/181-09.pdf

Year 2009 *Paper iirwps-181-09*

Spatial Heterogeneity and Minimum Wages: Employment Estimates for Teens Using Cross-State Commuting Zones

Sylvia Allegretto
University of California, Berkeley

Arindrajit Dube
University of California, Berkeley

Michael Reich
University of California, Berkeley

Spatial Heterogeneity and Minimum Wages: Employment Estimates for Teens Using Cross-State Commuting Zones

Abstract

Conventional approaches to estimating the effect of minimum wages on teen employment insufficiently account for heterogeneous employment patterns and selectivity of states with higher minimum wages. We overcome this problem by using policy discontinuities at state borders. Our estimates from cross-state labor markets (commuting zones) using data from the Census and the American Community Survey show that the measured negative impacts on teen employment in traditional estimates are driven by insufficient controls for spatial heterogeneity. We also replicate our key results using the Current Population Survey and show that the negative employment impact in traditional specifications is driven by pre-existing trends. Finally, by using a version of randomization inference, we devise a new test for heterogeneous effects of minimum wages across different local labor markets. We do not find evidence of such heterogeneous treatment effects using this new approach.

June 25, 2009

Spatial Heterogeneity and Minimum Wages: Employment Estimates for Teens Using Cross-State Commuting Zones

Sylvia Allegretto
IRLE, UC Berkeley
CA 94720 ,
allegretto@berkeley.edu

Arindrajit Dube
IRLE, UC Berkeley
CA 94720
adube@berkeley.edu

Michael Reich
Department of Economics
and IRLE, UC Berkeley CA 94720
mreich@econ.berkeley.edu

Institute for Research on Labor and Employment
University of California at Berkeley
Berkeley, CA 94720-5555

We thank seminar participants at the UC Berkeley Labor Lunch, UC Berkeley IGERT Seminar, University of Toronto CHRIR, University of Minnesota Carlson School of Business, and University of Massachusetts Amherst, and University of Massachusetts Boston. We are also grateful to Lisa Bell, Monica Deza, and Carol Tomas for excellent research assistance and to Eric Freeman and Oeindrila Dube for helpful suggestions.

Spatial Heterogeneity and Minimum Wages: Employment Estimates for Teens Using Cross-State Commuting Zones

ABSTRACT

Conventional approaches to estimating the effect of minimum wages on teen employment insufficiently account for heterogeneous employment patterns and selectivity of states with higher minimum wages. We overcome this problem by using policy discontinuities at state borders. Our estimates from cross-state labor markets (commuting zones) using data from the Census and the American Community Survey show that the measured negative impacts on teen employment in traditional estimates are driven by insufficient controls for spatial heterogeneity. We also replicate our key results using the Current Population Survey and show that the negative employment impact in traditional specifications is driven by pre-existing trends. Finally, by using a version of randomization inference, we devise a new test for heterogeneous effects of minimum wages across different local labor markets. We do not find evidence of such heterogeneous treatment effects using this new approach.

1. Introduction

Despite a steady stream of studies, research on the effects of minimum wage policies on employment continues to arrive at conflicting findings and implications. For example, the estimates from recent national CPS-based studies (e.g., Neumark and Wascher 2007) often imply negative net benefits for the low-wage workforce, while the results in Dube, Lester and Reich (forthcoming—henceforth DLR) suggest the opposite.

Although the conflicting findings may arise from differences in the groups being examined and/or differences in the datasets that are used, recent evidence suggests some other possibilities. Using data from establishments in the restaurant sector, and exploiting spatial policy discontinuities for contiguous counties that straddle state borders, DLR show that unobserved spatial heterogeneities in employment trends generate biases toward negative employment elasticities in traditional fixed-effects estimates. But DLR's focus is only on *jobs,* as opposed to *individuals.* While the measured effect of minimum wages on low-wage jobs may well be close to zero, the impact on low-wage workers (or potential workers) may be different. In this paper, we address this issue directly by focusing on teens.

Figure 1 displays the presence of heterogeneity in employment rates of teens across states. Although teen employment rates fell on average from 1990 to 2007 (Panel A), the changes in teen employment rates varied considerably by state (gray lines represent individual states). This cross-state variation is not captured simply by controls for business cycles or minimum wage increases. As Figure 1, Panel B shows, the cross-state variability remains largely intact when we consider changes in teen employment rates *net* of state changes in adult employment rates; whether we consider all states or just states with the same minimum wage profiles.[1] Table 1 shows that heterogeneity in teen employment trends has a regional component and that it varies considerably across the nine Census divisions.

Most existing estimates for the effects of minimum wages on teen employment do not, however, sufficiently account for this heterogeneity. If states with higher minimum wages have systematically different changes in employment demand for low-wage jobs, then insufficient controls for such heterogeneity will introduce a bias. In this paper, we document this bias in the traditional estimates that use place and time fixed-effects and overall unemployment rate as controls for underlying heterogeneity in teen employment. We show this bias using individual-level Current Population Survey (CPS) data from 1990 through 2007, as well as individual-level Census and American Community Survey (Census/ACS) data for 1990, 2000, 2005 and 2006.

There are numerous possible reasons why some states have higher minimum wages and why those states might have different employment trends for teens or other low-wage workers. For example, these states may have correlated policies (e.g., unemployment insurance, tax rates, right to work laws, zoning restrictions) that may affect labor demand or supply. Additionally, unionization, which displays considerable spatial heterogeneity (Holmes 2006) is an important determinant of minimum wage hikes (Sobel 1999, Seltzer 1995). While it is difficult to control for all possible confounding factors directly, we can utilize the strong spatially homogeneous component of many of these factors (as well of others that may codetermine both minimum wage policy and low-wage employment). Unionization, for example, exhibits strong inter-regional variation but is relatively

[1] Other evidence also shows that these variations are not simply due to school enrollment rates, relative wages of teens, or unskilled immigration. For detailed analyses that arrive at these conclusions, see Aaronson et al. 2006 and Congressional Budget Office 2004.

homogeneous across states within a region. For this reason, spatial discontinuities in minimum wage policies at state borders provide an attractive approach to estimating minimum wage effects. By considering an increase in the minimum wage in one part of a *single* labor market, we are able to control for arbitrary labor-market wide factors.

The primary contribution of this paper is to overcome the bias introduced by spatial heterogeneity by using a new research design. Similar to DLR, we use policy discontinuities at state boundaries. But while DLR uses county pairs straddling state borders we use a more economically-motivated definition of local labor markets: commuting zones. The Bureau of Labor Statistics partitions all counties into commuting zones (CZs) based on inter-county commuting flows. In our sample, 74 cross-state commuting zones had minimum wage variation, i.e. with a policy discontinuity at the state border. By allowing arbitrary time effects for each CZ, our discontinuity-based specifications use only local (within-CZ) variation to identify the minimum wage effects.[2]

Our estimates using a canonical fixed-effects specification (which exploits minimum wage variation within and across commuting zones) suggest a minimum wage employment elasticity of -0.159 (Census/ACS) and -0.153 (CPS), similar to those found in previous studies. In contrast, our preferred discontinuity-based estimates (which exploit minimum wage variation within commuting zones only) of the employment elasticity is a positive 0.129. We also provide additional evidence to demonstrate the bias in the conventional fixed-effects estimates of the minimum wage employment elasticity. First, using a dynamic specification, we show that *pre-existing* trends contaminate the traditional

[2] Throughout the paper, we use the terms "discontinuity-based" and "local" interchangeably when referring to our preferred specification.

estimates and lead to large spurious negative effects of a minimum wage increase on past employment. Second, we show that even coarse controls for spatial heterogeneity—such as inclusion of state specific linear trends, or allowing time effects to vary by Census divisions—either substantially reduce or completely eliminate the negative employment effects.

The absence of disemployment effects could result from the absence of significant minimum wage effects on average wages. However, we find strong average wage effects, with elasticities around 0.15, in all of our specifications. Our local wage elasticities are somewhat larger than those in the traditional estimates. This result suggests that employment estimates from within-CZ variation are not driven by the lack of a "bite" of the minimum wage in these areas.

We also find no relationship between the minimum wage elasticity of overall teen wages and the elasticity of employment across the 74 commuting zones. This result provides further evidence that there is no discernable disemployment effect, even when minimum wage increases lead to relatively large wage changes.

A second contribution of this paper is to test for the presence of heterogeneous wage and employment effects of minimum wages in different labor markets. While most studies in the literature have focused on the average treatment effect across labor markets, we extend the literature by considering whether the employment effects vary across the 74 particular cross-state CZs in our sample. We devise a new test motivated by the randomization inference literature. Even if the true treatment effect on employment is zero everywhere, sampling error and area-specific employment shocks (or "chance") will produce a distribution of estimated treatment effects. We examine whether the actual distribution of treatment effects across the 74 CZs can be rationalized by chance alone, or whether it also represents added variation because of heterogeneous treatment effects from minimum wage increases. To our knowledge, this represents the first analysis of this question in the minimum wage literature.

We answer this question through randomization (or permutation) based inference, in which we permute the minimum wage series from each of the 74 CZs with employment from 60 cross-state CZs *without* any minimum wage differentials. When we look at both the distribution of elasticities and t-statistics across individual commuting zones, we do not detect evidence of heterogeneous treatment effects. We therefore cannot rule out the null hypothesis that the distribution of estimated employment elasticities is produced by sampling error and area-specific shocks alone.

2. Related Literature

For the most part, minimum wage studies using the CPS state panel (or repeated cross section) data with state and year fixed-effects find modest but statistically significant negative employment effects on teens, with elasticities that range from -0.1 to -0.3. Sabia (2006) uses grouped CPS data from 1979 to 2004 to study the retail industry only. Using a canonical specification, Sabia finds a negative effect on teen workers in retail of around -0.29. Orrenius and Zavodny (2008) use a more elaborate set of business cycle controls and find negative effects for teens, although only when these added controls are included. Thompson (2009) uses county-level quarterly data on teen employment shares (not teen employment rates) and compares shares before and after two minimum wage increases in low and high wage counties. He finds very large disemployment effects. None of these studies, however, includes local controls that can convincingly overcome negative biases resulting from spatial heterogeneity.

Neumark and Wascher (2007) use individual-level repeated cross-section data from the Current Population Survey for the 1997 to 2005 period. They estimate a negative employment elasticity of -0.136 among teens, significant at the 10 percent level. Neumark and Wascher motivate their selection of the period since 1997 by arguing that welfare reform and expansions of the Earned Income Tax Credit may have changed the dynamics of the low-wage labor market. Unlike most other CPS and teen-based papers, Neumark and Wascher's (2007) estimates include a state linear trend to capture some degree of heterogeneity. As we discuss below, while such a parametric strategy works relatively well when considering longer time periods, it can be problematic for shorter panels, especially when the period straddles very different parts of the business cycle.

An alternative strategy uses border discontinuities to identify the effects of policies. This approach has been utilized previously to study a variety of state-specific laws. For example, Holmes (1998) uses policy borders to estimate the effect of right to work laws, while Huang (2008) uses cross-state border county pairs to evaluate state-level banking deregulation.[3] In the context of minimum wage research, local case studies (Card and Krueger 1994, 2000; Neumark and Wascher 2000; and Dube, Naidu and Reich 2007) make use of spatial discontinuities in minimum wage policies. In contrast to most state panel studies, these local case studies tend to find much smaller or nonexistent disemployment

[3] Using Swiss data, Lalive (2008) uses border discontinuity to estimate the effect of extended unemployment benefits on unemployment duration. Using South African data and spatial fixed effects, Magruder (2009) tests the effect of union bargaining councils on employment and business formation.

effects. However, as noted in the literature, it is difficult to make valid inference using individual case studies when there are area-specific shocks which tend to overstate the precision in these types of studies (e.g., Donald and Lang 2007).

Dube, Lester and Reich (forthcoming) generalize this local case study approach by pooling across all the spatial-discontinuity based estimates. They compare all the contiguous counties in the U.S. that lie on state borders, using sixteen years of county-level administrative data on restaurant employment. As mentioned, the authors show that previous national minimum wage studies lack adequate controls for spatial heterogeneity in employment growth. Without such controls, DLR find significant disemployment effects, within the standard -0.1 to -0.3 range of estimates. In their analysis, the economic and labor market conditions within the local area are sufficiently homogeneous to control for spatial heterogeneities in employment growth that are correlated with the minimum wage. Once they add such controls, DLR find no significant disemployment effects.

An important question is whether the DLR findings are relevant for low wage *workers* as well as low wage *jobs*. More specifically, can it help explain existing minimum wage elasticities for teens? Although this question has not been directly explored in the minimum wage literature, the importance of spatial heterogeneity in teen employment rates is evident in other research. In a study of the effect of teen population shares on teen unemployment rates, Foote (2007) finds that controlling for heterogeneous spatial trends across states generates results quite different from those using panel data with state fixed effects. Such evidence, as well as Figure 1, points to the importance of heterogeneous trends in teen employment rates. Such heterogeneity could interact with the *selectivity* of areas with greater minimum wage increases and bias estimated minimum wage elasticities. This is the key issue explored in this paper.

Our approach to testing heterogeneous effects of minimum wages across different labor markets also builds on numerous papers in the program evaluation literature. Bertrand, Duflo and Mullainathan (2001) use “placebo” laws to empirically estimate the distribution of t-statistics under the null hypothesis. Conley and Taber (2005) present a randomization or permutation-based inference in which the treatment profile from treated units is matched with outcomes from control units. But while those authors are interested in inference related to the average treatment effect, we focus on testing heterogeneity across individual cases. Our approach is inspired by Abadie, Diamond and Hainmuller (2007), who use a permutation-based test to draw inference for an individual case study with a “synthetic” control constructed from many potential control groups. In our case, we are interested not in whether the treatment effect is significant in a particular case study, but rather whether the total number of statistically significant effects (or the number of large coefficients) exceeds the expected count based on chance alone.

3. Data

We use two types of individual-level data. Our preferred identification strategy relies on county-level geographic identifiers and sufficiently large samples at local levels. For this, we use a combination of the 1990 and 2000 decennial Censuses and the American Community Survey (ACS) from 2005 and 2006. The Census provides a five percent sample of the population, while the ACS provides a one percent sample. Our sample contains 2.9 million teens aged 16 to 19. Our key outcome variables are employment, hourly wage, and usual hours worked per week. The employment variable reflects labor force status at the time of the survey. The Census/ACS hourly wage is constructed by dividing annual wage and salary income by the product of usual weekly hours and the weeks worked in the past year. Consequently, both the hourly wage and the usual hours variable is defined for a larger group of individuals than those who state that they were employed during the survey period.

Unlike most papers in the literature, the geographic unit in our paper is a Commuting Zone (CZ). To our knowledge, Autor and Dorn (2009) is the first paper in Economics to use CZs as the definition of local labor markets. We use the CZ concept for several reasons. First, analogous to the construction of a Metropolitan Statistical Area (MSA), the commuting zone definition is based on the actual degree of integration of the local labor market across counties. The BLS partitions all counties uniquely into commuting zones based on cross-county commuting flows. This is appealing because these areas are not only contiguous; they are also demonstrably linked with each other by an economically meaningful criterion. Second, unlike a MSA, a CZ is defined for *all* counties in the U.S., not just metro or urban counties. As a result, we can use a fuller range of local variation than is possible with MSA-based units.

Like Autor and Dorn, we use the most recent (1990) definition of CZs to map counties consistently on to CZs over time. The most local geographic identifier in the 1990 and 2000 Census as well as the ACS is the Public Use Microdata Area (PUMA), a sub-state area that typically comprises a population of 100,000 to 200,000 individuals. In the vast majority of cases, each PUMA can be matched to a unique county, and hence to a unique commuting zone. In some (especially rural) areas, however, PUMAs span multiple counties, although never multiple states.[4] In these instances, we assign residents of these PUMAs to several CZs. Sample weights of individuals assigned to multiple CZ's are adjusted to reflect the relative share of a CZs population in a particular PUMA.

We then merge county-level minimum wage information to the Census/ACS data. In all, we have 741 CZs in our sample. Since a CZ is an economic definition, and not based on political jurisdiction, there are 134 CZs that cross over state lines. Of these, 74 have minimum wage differences during this period; these 74 cross-state CZs provide the core of our identifying variation. To be clear, we do not drop any areas from our sample in most of our analysis, until the section in which we consider these 74 cases separately. The other CZs help identify the estimates for other control variables and help increase precision. However, for our local specification, the minimum wage elasticities will effectively be identified from the variation within these 74 CZs.

While the Census/ACS provides a dense sample allowing local-level analysis, it unfortunately is available for only four years between 1990 and 2006. For this reason, we also use the Current Population Survey to supplement our analysis. There are three additional reasons to use the CPS. First, for specifications that can be estimated using both, we want to compare the findings. Second, we want to use the higher frequency CPS data to estimate dynamic specifications and test for "pre-existing" trends. Since most of the existing literature uses the CPS, it is important to document any bias in the estimates using the same data. Third, we want to see how using the CPS and "coarser" control variables for spatial heterogeneity compares to the more local estimates from the Census/ACS.

[4] The latter point is particularly important. Even when we have to allocate individuals to multiple CZs, they are never "misallocated" across state (and hence policy) lines, which could introduce an attenuation bias in the measured effect of the policy.

We construct an individual-level repeated cross-section sample from the CPS Outgoing Rotation Groups for the years 1990 to 2007.[5] The CPS data are merged with monthly state unemployment rates and teen population shares within the states, which are used as controls. Additionally, each observation is merged with a quarterly minimum wage variable—whichever is higher of the federal or state minimum.

Tables 2 and 3 report summary statistics for the Census/ACS and CPS data, respectively. The means and standard deviations for many of the variables are quite similar across the two datasets, even though the Census/ACS data are available for only 4 of the 18 years that make up the CPS sample. In the Census/ACS data, employment rates are considerably higher for white teens (.42) than for black teens (.24) and Hispanic teens (.34), with quite similar patterns in the CPS data. Average hourly wages are somewhat higher in the Census/ACS data than in the CPS ($8.32 versus $7.93). However, differences in the definition of hourly wage in the two datasets may account for this divergence.[6] The average teen employment rate in the Census/ACS (0.38) is slightly lower than in the CPS (0.41), as is the overall unemployment rate (0.04 versus 0.05 in the CPS).

Table 2 also compares summary statistics for the full sample of 741 commuting zones in the U.S. and the sample of 134 cross-state commuting zones. In these two samples the means and standard deviations are quite similar for most variables, including

---

[5] We use a start year of 1990 to make the clearest comparison of the results with the Census-based analysis. We have also included earlier years in our CPS sample; the results are not materially affected. See Section 5.3 below.

[6] The Census/ACS hourly wage is constructed by dividing annual wage and salary income by the product of usual weekly hours and the weeks worked in the past year. Therefore, it includes all overtime, commission and tip income, which is not the case for the CPS hourly wage definition. Baum-Snow and Neal (2009) argue that usual hours worked is undercounted among part-time workers in the Census and ACS, relative to the CPS. In our Tables 2 and 3, average hours are quite close in both datasets, indicating that data quality does not affect our results.

demographics, wages, and employment (or unemployment) rates. These results suggest that there is no systematic difference between the country overall and the state border areas that are used to identify minimum wage effects in our preferred discontinuity-based approach.

4. Estimation Strategy and Main Results

4.1 Key Specifications – Census/ACS and CPS

The three outcomes we consider are (1) the natural log of hourly earnings, (2) a dichotomous employment measure that takes on the value one if the person is working, and (3) the natural log of usual weekly hours of work. The treatment variable is the natural log of minimum wage. We begin with the baseline fixed-effects specification using the Census/ACS dataset:

$$y_{icst} = \beta \ln(MW_{cst}) + X_{it}\Gamma + Z_{cst}\Omega + \phi_{cs} + \tau_t + \varepsilon_{icst} \qquad \text{(1a)}$$

Here ln(*MW*) refers to the log of the minimum wage, *i*, *c*, *s*, and *t* denote, respectively, individual, CZ, state and time indexes. *X* is a vector of individual characteristics, *Z* is a vector of area specific controls (unemployment rate, and/or average wage in the workforce), and ϕ_{cs} is a fixed effect for each commuting zone/state combination. The τ_t dummies are incremented in years and represent common (across CZs) time effects.[7] We also estimate an analogous regression using the CPS, except that the geographic unit is the state instead of a CZ and time is indexed in quarters:

$$y_{ist} = \beta \ln(MW_{st}) + X_{it}\Gamma + Z_{st}\Omega + \phi_s + \tau_t + \varepsilon_{ist} \qquad \text{(1b)}$$

[7] The individual characteristics include 2 gender categories, 4 race/ethnicity categories, 12 education categories and 4 marital status categories.

In this canonical specification, the inclusion of place and time dummies as well as the overall unemployment rate is thought to sufficiently control for local labor market conditions facing teenage workers. In contrast, for our preferred discontinuity specification using the Census/ACS, we allow the time effects to vary by each commuting zone:

$$y_{icst} = \beta \ln(MW_{cst}) + X_{it}\Gamma + Z_{cst}\Omega + \phi_{cs} + \tau_{ct} + \varepsilon_{icst} \tag{2}$$

The inclusion of CZ-specific time effects (τ_{ct}) sweeps out all the minimum wage variation between commuting zones, and uses only the within-CZ variation to identify β.

As an intermediate specification, we also estimate variants of (1a) and (1b) that allow for the time effects to vary by each of the nine census divisions *(d)*. Including relatively coarse regional controls helps us understand the scale of the spatial heterogeneity that underlies the potential bias in (1)

$$y_{icst} = \beta \ln(MW_{cst}) + X_{it}\Gamma + Z_{cst}\Omega + \phi_{cs} + \tau_{dt} + \varepsilon_{icst} \tag{3a}$$

$$y_{ist} = \beta \ln(MW_{st}) + X_{it}\Gamma + Z_{st}\Omega + \phi_{s} + \tau_{dt} + \varepsilon_{ist}. \tag{3b}$$

It is not possible to estimate the discontinuity-based specification (2) with the CPS because of the unavailability of local geographic identifiers. Instead, we estimate a specification with state-specific linear trends. Such trends represent a parametric method of controlling for heterogeneity in the underlying (long term) growth prospects of low-wage employment. For the CPS, this is the specification with the strongest controls for spatial heterogeneity.

$$y_{ist} = \beta \ln(MW_{st}) + X_{it}\Gamma + Z_{st}\Omega + \phi_{s} + \xi_s t \cdot I_s + \tau_{t} + \varepsilon_{ist} \tag{4}$$

Finally, we allow for both division-specific time effects and include state-specific linear trends. For the CPS, this is the specification with the strongest controls for spatial heterogeneity.

$$y_{ist} = \beta \ln(MW_{st}) + X_{it}\Gamma + Z_{st}\Omega + \phi_s + \xi_s t \cdot I_s + \tau_{dt} + \varepsilon_{ist} \qquad (5)$$

We report standard errors clustered at the state level to account for the lack of independence among observations within a state (our treatment unit), which might be caused by correlation in employment rates within states over time, and across individuals within a state.

4.2 Main Results

The estimated wage effects establish the clear presence of a treatment: increases in the minimum wage lead to increased average wages for the teens. Table 4 presents the estimated effects on wages from our four CPS-based specifications, and three specifications from the Census/ACS. The coefficient, which is also the wage elasticity, is positive and significant in all the specifications (in one case, only at the 10 percent level). But the magnitudes vary among the specifications—from 0.088 to 0.205. In Specification 1, the fixed-effects model using the CPS, the treatment coefficient is 0.120, while in Specification 5, the fixed effects model using the Census/ACS, the coefficient is 0.110. In our preferred discontinuity-based Specification 7, the coefficient is 0.151, while the preferred CPS based estimate using both state-specific trends and division-specific time effects (Specification 4) is 0.158. These results indicate that the effect of minimum wages on average teen wages remains similar when controls for heterogeneous spatial trends are included and that the key wage estimates from the CPS and the Census/ACS correspond well with each other.

We turn next to the employment elasticities reported in Table 4. The estimates using common time effects (Specification 1 using the CPS, and Specification 5 using the Census/ACS) are remarkably similar, -0.153 and -0.159, respectively. Both estimates are significant at conventional levels, and are consistent with the literature that uses the canonical fixed-effects model.[8]

In contrast, the employment elasticity from our discontinuity-based Specification 7 is positive (0.129) and marginally significant.[9] Moreover, we can rule out at the 5 percent level an employment elasticity more negative than -0.015. Tests of coefficient equality between Specification 7 and Specification 5 (with common time effects) can be rejected at the 1 percent level—as indicated in the p-value row of the table. These results provide strong evidence that when we account for spatial heterogeneity by using cross-border variation within commuting zones, we do not find any disemployment effects of minimum wages on teens.

We find also that intermediate levels of controls for heterogeneity produce intermediate results. For the CPS-based Specification 2, allowing for division-specific time effects reduces the elasticity to -0.105 and renders it insignificant. As Specification 3 shows, the addition of a state-specific time trend to the fixed effects model also lessens the effect of minimum wages on employment. Here the elasticity is -0.065 and it is not significant. And in Specification 4, the employment elasticity is -0.024 and remains

[8] Generally, the employment elasticity is obtained by dividing the regression coefficient (also reported in the table) by the employment-to-population ratio of the group in question.

[9] As we show later, when disaggregated by commuting zones, the average employment effect is still positive, but closer to zero when we do not weight the individual estimates by the population of the CZ—which this pooled estimate does implicitly. As a result, we do not consider the evidence to show a clear positive effect; rather, that the estimates are inconsistent with sizeable disemployment.

insignificant. Finally, in the Census/ACS data, inclusion of division-specific time effects (Specification 6) produces an elasticity estimate close to zero (0.005).

These results indicate that estimates of minimum wage employment effects using the standard fixed-effects model of Specification 1 and 5 are seriously contaminated by heterogeneous employment patterns across states. Controlling only for within-division variation substantially reduces the estimated elasticity. Allowing for long-term differential state trends makes the employment estimates indistinguishable from zero. And the cleanest discontinuity-based estimate comparing employment within commuting-zones produces a small (albeit marginally significant) positive effect.

Besides employment, we are also interested in the effects on average hours of work. Table 4 provides estimates of the effects of the minimum wage on hours worked for those who are employed. With common time effects, the elasticity on average hours is -0.096 in Specification 1 and -0.091 in Specification 5 and both are significant at the 1 percent level. In contrast, the elasticity estimate from our cross-state commuting zones (Specification 7) is close to zero: -0.031. The intermediate specifications 2, 3, 4 and 6 produce estimates between -0.038 and -0.091. These estimates imply that we can reject hours effects larger than -0.094 at the 95 percent confidence level. Similar to employment, the hours effects estimated by the canonical fixed-effects model also seem to suffer from a bias due to spatial heterogeneity.

To summarize the main results to this point, overcoming the spatial heterogeneity bias does not affect the estimated treatment effects on wages. It does, however, result in estimated employment and hours effects that are not significantly different from zero and that can rule out all but small reductions.

Some studies have found stronger negative effects for employment among minority teens (Neumark and Wascher 2007). Although not reported in the tables, using our discontinuity specification (i.e., Specification 7), we find the employment elasticities for white and minority teens are greater than zero. The estimates for individual demographic subgroups, however, are both imprecise and exhibit substantial variability, which mirrors results from other studies (e.g., Neumark and Wascher 2007).

4.3 Evidence from Dynamic Specifications

To provide further evidence on the bias of the traditional estimates, we evaluate the timing of any putative effect of minimum wages on various outcomes. We begin with the CPS sample, where we have high frequency data that enables us to estimate models spanning a wide window. We estimate dynamic versions of our key specifications 1 and 4—i.e., with common time effects, and with division-specific time effects along with state linear trends.

$$\begin{aligned} y_{ist} = & \sum_{\tau=-2}^{3} \beta_\tau \Delta_{12} \ln(MW_{cs,t+12\tau}) + \beta_4 \ln(MW_{cs,t+48}) + X_{it}\Gamma + Z_{cst}\Omega \\ & + \phi_s + \tau_t + \varepsilon_{ist} \end{aligned} \tag{6a}$$

$$\begin{aligned} y_{ist} = & \sum_{\tau=-2}^{3} \beta_\tau \Delta_{12} \ln(MW_{cs,t+12\tau}) + \beta_4 \ln(MW_{cs,t+48}) + X_{it}\Gamma + Z_{cst}\Omega \\ & + \phi_s + \tau_{dt} + \xi_s t \cdot I_s + \varepsilon_{ist} \end{aligned} \tag{6b}$$

These specifications estimate leads and lags of minimum wage changes, spanning just over 6 years (73 months) around the minimum wage change in annual increments. Δ_{12} is a 12 month difference operator, meaning $\Delta_{12} MW_{s,t+12\tau} = MW_{s,t+12\tau} - MW_{s,t+12(\tau-1)}$.

The coefficients β_{-2} to β_4 trace out the cumulative response (or time path) of the outcome variable to a log point increase in the minimum wage, starting two years prior to the increase and continuing to four years afterwards. The β_4 coefficient measures the response at year 4 or later—i.e., the "long run" effect. We divide all coefficients (and standard errors) by the employment-to-population ratios to convert them to elasticities.

We use two years of leads and four years of lags for two reasons. First, including more leads is more "costly" in terms of losing latter years in the sample, whereas lags in minimum wages are known in the early part of the sample, and hence do not lead to losing observations. Second, the reason for including leads is to capture pre-existing trends, which we believe we can capture sufficiently with two years. In contrast, the purpose of including longer lags is to investigate whether minimum wage effects occur with delay, as some have proposed in the literature.[10]

Figure 2, Panel A displays time paths of the wage effects of minimum wage increases. The left-hand column displays results for our Specification 1, while the right-hand column presents results for Specification 4, which includes both state-specific time trends and division-specific time effects. Both wage graphs show a clear increase right at the time of the minimum wage change.

Figure 2, Panel B displays the time paths of the cumulative response of employment from a log point increase in minimum wages. The timing of the negative employment response provides strong evidence against Specification 1—the canonical model without controls for time and state fixed effects. Specification 1 shows negative elasticities

[10] Using more leads produces similar results, but less precise estimates.

throughout the six-year window, including in the two years *prior* to the minimum wage increase. To emphasize, the results indicate that an increase in minimum wage at time *t* is associated with unusually low employment at *t*-2, *holding constant the minimum wage at time t-2*.[11] This dynamic evidence shows that minimum wage increases have occurred—on net—in places with lower growth (or greater reduction) in teen employment, quite apart from any causal effect of minimum wages. Moreover, there is no evidence that teen employment rates fall in the four years after any minimum wage changes.

Consistent with this interpretation, when we control for underlying heterogeneity in teen employment rates using both state-level trends and division specific time-effects in Specification 4, we find that the lead terms are close to zero and stable prior to the minimum wage increase. This result provides additional internal validity for the specification with controls for spatial heterogeneity. We also find no disemployment in the years following the increase.

For our Census/ACS sample, the lack of a full 16-year panel prevents us from jointly estimating the full set of lead and lag coefficients, as in equations (6a) and (6b). Instead, we estimate a version with a single (two year) lead, a single (two year) lag, and the contemporary minimum wage. While more limited, this version nonetheless provides valuable information about pre-existing trends that may contaminate various specifications.

$$\begin{aligned} y_{icst} = \; & \sum_{\tau=-1}^{0} \beta_{2\tau} \Delta_2 \ln(MW_{cs,t+2\tau}) + \beta_2 \ln(MW_{cs,t+2}) + X_{it}\Gamma + Z_{cst}\Omega \\ & + \phi_{cs} + \tau_t + \varepsilon_{icst} \end{aligned} \qquad (7a)$$

[11] Since we are jointly estimating the marginal impact of each lead/lag, the unusually low employment at time *t*-2 cannot be due to an unusually high minimum wage two years earlier.

$$y_{icst} = \sum_{\tau=-1}^{0} \beta_{2\tau}\Delta_2 \ln(MW_{cs,t+2\tau}) + \beta_2 \ln(MW_{cs,t+2}) + X_{it}\Gamma + Z_{cst}\Omega + \phi_{cs} + \tau_{ct} + \varepsilon_{icst} \quad (7b)$$

Here, Δ_2 is a two-year difference operator. Analogous to equation (4), the β coefficients trace the cumulative response of a log point change in the minimum wage.

The results are illustrated in Figure 3. Panel A shows that a minimum wage increase at time *t* has a positive effect on the average wage at time *t* and thereafter. This finding holds for both the common time effects (Specification 5) and commuting-zone specific time effects specifications, but appears to be stronger for the latter. This reassuring result is consistent with a clear causal effect of minimum wage on average teen wages. For employment, however, Specification 5 with common time effects produces highly counterintuitive outcomes. A minimum wage at time *t* is shown to have a negative effect two years *prior* to the minimum wage increase. Indeed, the effects are more negative on past employment than on contemporaneous employment. This anomaly is a sign of a spurious estimate driven by pre-existing trends, consistent with the evidence from the dynamic specifications using the CPS presented above. In contrast, we do not find spurious effects in our discontinuity-based Specification 7. We do see some positive effects around the times of minimum wage increases, but the effect falls to zero after a few years.

Overall, our dynamic specifications provide further evidence that in the period under consideration, failure to control for heterogeneity in employment patterns leads to a negative bias in the estimated employment response from minimum wage changes. Our preferred specifications using cross border variation provide more plausible estimates—in which minimum wage increases do not putatively affect prior employment.

We have seen that parametric controls such as state-specific linear trends may sometimes provide adequate controls for spatial heterogeneity. However, this approach is not a panacea. When considering shorter periods, or periods straddling a downturn, a linear trend may prove to be inadequate. Indeed, with our CPS data we obtain (but do not provide here) estimates from a specification that includes a state-specific trend; they are not robust for shorter windows and that they vary with the particular window under consideration.[12] This result is perhaps not surprising, since the task of estimating a linear trend parameter with a small number of years is sensitive to year-to-year movements. In contrast, our preferred discontinuity-based estimate does not rely on parametrically estimating such a long-run trend, and instead can account for any year-to-year fluctuations in a given labor market.

Clearly, low-wage or teen employment demand is growing more slowly in states with higher minimum wage increases. By construction, the pre-existing negative trends in employment in the dynamic time paths are not the result of pre-existing trends in minimum wages. Therefore, our results here indicate that the presence of systematic spatial differences that are correlated with minimum wage increases may affect low-wage employment changes for reasons that are unrelated to the specific policy of raising the minimum wage.

5. Heterogeneous Treatment Effects across Labor Markets

Most of the minimum wage literature has focused on identifying "the" effect of a minimum wage increase in an area, whether the area is a particular city (e.g., San

[12] Results available from authors upon request.

Francisco), state (e.g., New Jersey), or country (e.g., U.S., UK, Brazil). Of course, there is no *a priori* reason to believe that the employment effects from a particular increase in minimum wage would be the same across locations. One obvious source of heterogeneity is the extent of "bite" from a given increase in the minimum wage: the average wages of teens would increase more (and hence amplify the employment effect), the larger the proportion of teens at or near the minimum wage. There are other and less straightforward sources of heterogeneity as well. In a monopsonistic model, the sign and magnitude of the employment effects depends on different "regimes" of parameter values, including the distribution of firm-level labor supply elasticity (Manning 2006). The output price elasticity of demand, which in a competitive model attenuates the disemployment effect, likely varies across places or times. When looking across various labor markets across the entire United States, these (and other) factors are likely to vary.

An attractive property of our cross-state commuting zone research design is that we pool across 74 different local comparisons. Instead of estimating a pooled regression with CZ-specific time effects, in this section we estimate minimum wage elasticities separately for each of the 74 cross state CZs with minimum wage variation. We ask two specific questions: 1) Is there any relationship between the wage elasticity of the minimum wage and the employment elasticity of the minimum wage across the 74 CZs? 2) Does the distribution of individual employment elasticities and t-statistics across the 74 CZs look similar to what chance alone would generate under the sharp null hypothesis of no effect anywhere?

5.1 Relationship between Wage and Employment Effects

We begin with the first question—the relationship between wage and employment elasticities across the 74 cross state commuting zones with minimum wage discontinuities. We estimate separate regressions for each of the 74 CZs:

$$y_{icst} = \beta_c \ln(MW_{cst}) + \mathrm{X}_{icst}\Gamma + Z_{cst}\Omega + \phi_{cs} + \tau_t + \varepsilon_{icst} \quad (8)$$

For employment, we divide the coefficient by the CZ-specific teen EPOP ratio to convert it into the elasticity. Figure 4 presents the kernel density estimates of the wage and employment elasticities. In the first panel, we plot the density *without* weighting each estimate by the population weights, while we do weight in the second panel.[13] We do this both ways for two reasons. Conceptually, the population-weighted average of the elasticities across the 74 CZs is analogous to the pooled regression of Specification 7, except that the covariates are allowed to have different coefficients in different CZs. It is also instructive to consider each of these 74 "experiments" as equally informative and not to give more influence to larger CZs.

Although not reported in the figure, the weighted average of the elasticities for both wage (0.12) and employment (0.09) are indeed quite similar to results from specification 7 in Table 4. The unweighted average wage elasticity is slightly larger, at 0.18. With respect to employment, however, the unweighted average elasticity is closer to zero (0.038) and has greater dispersion. We take the unweighted distribution as evidence that the employment elasticity from our discontinuity-based estimates is close to zero (instead of a marginally

[13] To be clear, we use sampling weights when estimating each of the 74 elasticities. For the "weighted" density estimate we weight these elasticities by the population in the CZ, while we do not in the "unweighted" case. The sizes of the circles in the second panel of Figure 4 are proportional to the population of the CZ.

significant positive) if we do not weigh larger CZs more heavily. Overall, we take the evidence to indicate the lack of a substantial negative effect, as opposed to a positive one.

Are employment effects systematically different where wage effects are greater? Figure 5 provides the scatter plot and the linear projection (with 95 percent confidence bands) for the wage and employment elasticities. For both the weighted and unweighted cases, we see no evidence of systematically different effects on employment in CZs with larger wage effects. The regression coefficients are close to zero, -0.038 (0.040) and -0.002 (0.039) for the unweighted and weighted cases, respectively (standard errors in parentheses).

5.2 Randomization Inference-Based Test of Heterogeneous Treatment Effects

Figure 5 also shows considerable variation in the employment effects across the CZs. A natural question is whether this variation just reflects "chance"—from sampling noise or area-specific shocks—or whether some of this variation reflects heterogeneity in the employment effects across different commuting zones.

Consider the following data-generating process for the average teenage employment rate with a heterogeneous treatment effect of minimum wage (ignoring covariates).

$$
\begin{aligned}
y_{jst} &= \beta_j \ln(MW_{jst}) + \phi_{js} + \tau_{jt} + \varepsilon_{jst} && (9) \\
&= \left(\bar{\beta} + \eta_j\right)\ln(MW_{jst}) + \phi_{js} + \tau_{jt} + \varepsilon_{jst} \\
&= \bar{\beta}\ln(MW_{jst}) + \phi_{js} + \tau_{jt} + \left(\eta_j \ln(MW_{jst}) + \varepsilon_{jst}\right) && (10)
\end{aligned}
$$

Conditional on being within a particular commuting zone *j,* if the treatment effect is independent of treatment ($\eta_j \perp \ln(MW_{jst})$, then we can identify the average treatment effect $\bar{\beta}$ by estimating equation (10). We do not need to assume that the treatment effect is generally independent of minimum wage—only that it is conditionally so around the policy discontinuity (i.e., within commuting zone *j*). This provides a justification for our analysis using pooled data, as we did in section 4, but by necessity such analysis can only test hypotheses about the average effect $\bar{\beta}$. In our case, however, we can also estimate equation (10) separately for each commuting zone *j* and recover the estimates $\hat{\beta}_j$. If we knew the distribution of $\hat{\beta}_j$ under the null hypothesis $\beta_j = 0$ for all *j*, we can then additionally test for this heterogeneity of the treatment effect and not simply whether $\bar{\beta} = 0$.

How can we draw inference for the individual $\hat{\beta}_j$'s? One might be tempted to use t-statistics from the individual CZ regressions and test whether the number of significant t-statistics exceeds the proportion that would be expected under the (sharp) null hypothesis that the true effect is zero everywhere. This approach would be incorrect, however, since the presence of serial correlation (i.e., $Cov(e_{jxt}, e_{jst'}) \neq 0$) leads to over-rejection of the null in a fixed effects or difference-in-difference setup (Kézdi 2004; Bertrand, Duflo and Mullainathan 2004). For our pooled regression (Specification 7), the standard error for the average effect, $\bar{\beta}$, is clustered on states, which accounts for any intertemporal correlation of the error term. However, cluster-robust methods of estimating the standard error are unreliable with a small number of clusters, let alone the two clusters, as is the case for the individual CZ-level regression. What we need is the counterfactual distribution of the t-statistic under the null hypothesis of exactly zero effect everywhere. We could then determine the empirical cutoffs for t-statistics at a given confidence level, and evaluate whether the t-statistics for our 74 commuting zones exceed these empirical cutoffs more frequently than would be expected by chance alone.

Fortunately, with some added assumptions we can recover such a counterfactual distribution. We exploit the fact that there are 60 commuting zones that cross state lines and do not have any minimum wage variation within the commuting zone. We apply the treatment profile (i.e., the minimum wage series) from each of the 74 “in sample” CZs to each of these 60 “out of sample” commuting zones, and estimate the effect of this fictitious minimum wage on actual teen employment in the CZ. The minimum wage differences within the “out of sample” CZs are not only fictitious, they are also completely uncorrelated with the *actual* minimum wage differences within the commuting zone, since by definition the “out of sample” CZs have no within-CZ variation in minimum wages. The resulting distribution of the elasticities (and t-statistics) imposes the null of zero effect everywhere—with the assumption that the distribution of employment under the null in the “out of sample” CZs is the same as the “in sample” CZs. Note that this procedure preserves the time pattern of minimum wages and employment within each CZ/state in the “in sample” and “out of sample” commuting zone. Hence, we account for arbitrary serial correlation in a manner that is analogous to block-bootstrapping.

As mentioned, we need to assume that the distribution of employment under the null hypothesis is the same in the 74 “in sample” CZs as it is in the 60 “out of sample” ones. This assumption is not directly testable, since we do not observe what the counterfactual distribution would be for the “in sample” CZs in the absence of a minimum wage differential. It is instructive, however, to compare the mean and SD of the teen EPOP ratios across the CZs in the “in” versus “out” samples. The mean (SD) for the “in sample” CZs is

0.39 (0.08), while it is 0.40 (0.08) for the “out of sample” CZs.[14] For adults, the mean and SD of EPOP ratios are identical for the two samples. This result is reassuring and provides added validity for our design.

Formally, define the set J as composed of the 74 “in sample” CZs and the set K as being composed of the 60 “out of sample” CZs. Since the vast majority of cross-state CZs range across two states, for expositional simplicity denote as $s = \{1, 2\}$ each state of CZ j in J; and $s' = \{1, 2\}$ each state of CZ k in K. We merge employment data $y_{ks't}$ from side s' of the border in (the “out of sample”) commuting zone k with minimum wage data from side s in (the “in sample”) commuting zone j. In this case, there are exactly two unique ways to match the minimum wage and employment information— either s'=1 with s=1 and s'=2 with s=2; or s'=1 with s=2 and s'=2 with s=1. Each of these pairings can be uniquely indexed by j, k ,s, and for each, we estimate the following regressions:

$$y_{ks't} = \beta_{jks} \ln(MW_{jst}) + \phi_{ks'} + \tau_{kt} + \varepsilon_{ks't} \qquad (11)$$

If each commuting zone in J and K had exactly two states, we would estimate 74 x 60 x 2 = 8,880 regressions. We would draw each J, K combination, and then rotate the ordering of the “in sample” states (s) and draw each J, K combination for a second time. For the handful of “in sample” CZs with three states, we match these to analogous “out of sample” CZs with three states—for each of which there are now six ways of matching minimum wage information from the three states. Overall, the total number of regressions estimated is slightly greater, at 9,028. By definition, this process of matching each side of the border in j with each side in k creates a symmetric empirical distribution of $\hat{\beta}_{jks}$ centered

[14] SD is the standard deviation of the average EPOP in a CZ for a given year in the “in sample” or the “out sample.”

at zero. Consequently, the counterfactual distribution of minimum wage effects and the t-statistics will have zero means.

The regressions in equation (11) use data collapsed into CZ-state-year cells, which is needed for comparability when we implement the randomization inference. Since the sample sizes in CZs vary, the distribution of t-statistics from the randomization inference would not hold the sample size constant were we to use individual-level data.[15]

The regression coefficients $\hat{\beta}_{jks}$ and the t-statistics $\hat{t}_{jks}$ in each of these 9,028 regressions define the empirical Randomization Inference (RI) distribution. We then read off the critical values of the t-statistics from this empirical distribution for two-sided tests of the sharp null hypothesis that the coefficient is exactly zero—at 5, 10, 15 and 20 percent significance levels. The top row of Table 5 shows these empirical cutoffs, which are 4.38, 3.09, 2.67, and 1.81 respectively.

Next we estimate the "in sample" regressions for our 74 cross-state CZs analogous to equation (6), except now we estimate these using collapsed data at the CZ-state-year level to match equation (11):

$$y_{jst} = \beta_j \ln(MW_{jst}) + \phi_{js} + \tau_{jt} + \varepsilon_{jst} \qquad (12)$$

To test the significance of *each* estimate *j*, we compare the j'th t-statistic to the empirical cutoffs using the RI distribution.

Table 5 shows that the t-statistics from our 74 "in sample" cross-state CZs are no more likely to display significant employment effects using the empirical cutoffs from the RI distribution than what is expected by chance alone. First, when we consider just the

---

15 The collapsing of the data sacrifices the individual-level covariates, but their inclusion made almost no difference for the coefficients in all specifications in Table 4.

absolute value of the t-statistics, at all four significance levels the actual proportion of cases exceeding the cutoff is lower than the theoretical proportions resulting from sampling error and area-specific shocks. As shown in the third row, only 18.9 percent of cases in our actual sample of CZs had t-statistics greater than 1.81 in absolute value, while we would expect around 20 percent of cases from random variation alone. And only 1.4 percent of cases had t-statistics greater than 4.38 in absolute value, compared to 5 percent expected from chance alone. Second, when we tabulate the statistically significant coefficients by sign, we find that negative significant cases are not particularly prevalent in the data. The fourth and fifth rows of Table 5 show that for all significance levels beyond the 5 percent level, there are some positive and some negative estimates.

Finally, we calculate the probability (under the null) of obtaining a count of statistically significant coefficients greater than or equal to the count in our sample of 74 cross-state commuting zones. We do this by randomly sampling 74 estimated "out of sample" t-statistics from our RI distribution 10,000 times, and reporting the proportion of cases with counts of t-statistics as large or larger (in magnitude) than in our "in sample" results, and for all four significance levels. These proportions represent the probability values of obtaining a count as large as, or larger than, that in our actual sample under the null. As the last row of Table 5 shows, the probability values are all quite large for tests of individual coefficients at all four significance levels. In other words, we cannot reject the null hypothesis that our set of 74 $\hat{\beta}_j$'s from the "in sample" commuting zones are generated by a data-generating process with $\beta_j = 0$.

Next, we devise a test for heterogeneity that focuses on the magnitude and size distribution of the estimated coefficients. Besides using the distribution of t-statistics, we can also directly use the employment coefficients from our RI distribution of $\hat{\beta}_{jks}$ (from equation 11) to test whether we have a greater incidence of large or small $\hat{\beta}_j$'s in our actual "in sample" than would be expected under the null hypothesis (β_j=0). This allows us to test for the presence of heterogeneity while considering the economic significance of the estimated coefficients.

We first count the number of the estimated employment coefficients $\hat{\beta}_j$ falling in particular ranges: in particular, counts exceeding 0, 0.1, 0.3, and 0.5, or falling below 0, -0.1, -0.3, and -0.5.[16] We define the counts of the estimated employment effects exceeding (or falling below) *b* in a sample of 74 CZs as follows:

$$\hat{C}_b^+ = \sum_{j=1}^{74} \mathbf{1}(\hat{\beta}_j > b), \text{ and } \hat{C}_b^- = \sum_{j=1}^{74} \mathbf{1}(\hat{\beta}_j < b)$$

We want to test whether the sample counts of $\hat{C}_b^+$ and $\hat{C}_b^-$ fall outside the confidence intervals of C_b^+ and C_b^- that come from the RI distribution under the null hypothesis that $\beta_j = 0$. To empirically derive the confidence bounds for C_b^+ and C_b^- under the null, we repeatedly (10,000 times) sample without replacement 74 $\hat{\beta}_{jks}$'s from our RI distribution. For each of these 10,000 repetitions, we count the incidence of the employment coefficients exceeding 0, 0.1, 0.3, and 0.5, or falling below 0, -0.1, -0.3, and -0.5. By taking the appropriate (2.5, 5, 95 and 97.5) percentiles of this simulated distribution, we derive confidence bounds (at the

[16] These effects (0, ± 0.1, ± 0.3, and ± 0.5) were chosen to capture a wide enough range for plausible employment effects. None of the conclusions are sensitive to the precise cutoffs picked here.

90 and 95 percent levels) for how often one "should" expect to see the employment effect exceeding (or falling below) a certain level *b* in a sample of 74 cross-state CZs.

Formally, denote as *r* a particular simulation repetition. For each *r,* we have 74 draws of $\hat{\beta}_{jks}$, each of which we denote as $\hat{\beta}_m^r$ for $m \in \{1, \cdots, 74\}$. The empirical Cumulative Distribution Function of $\hat{C}_b^+$ under the null (β_j *=0*) is defined as:

$$\hat{F}\left(\hat{C}_b^+ < c \mid \beta_j = 0\right) = \frac{1}{10,000} \sum_{r=1}^{10,000} 1\left(\left[\sum_{m=1}^{74} \mathbf{1}(\hat{\beta}_m^r > b)\right] < c\right)$$

We use this empirical CDF $\hat{F}\left(\hat{C}_b^+ < c \mid \beta_j = 0\right)$ to construct the 90 and 95 percent confidence intervals for $\hat{C}_b^+$. We do this for different cutoff values for the employment coefficients, i.e., different values of *b*. The confidence intervals for $\hat{C}_b^-$ are computed in an analogous fashion.

The first row of Table 6 shows the actual counts of CZs with employment coefficients ($\hat{\beta}_j$) that fall within the various ranges. The subsequent rows show the cutoffs for the randomization-inference based confidence intervals at the 90 and 95 percent levels. We find that the actual counts are well within the confidence bounds. As an example, we have 12 CZs with employment coefficients smaller than -0.5, and 11 counts with coefficients exceeding 0.5; the 90 percent confidence bounds for both are (9, 20).

Formally, we cannot reject the null hypotheses that $\hat{C}_b^-$ (the count of $\hat{\beta}_j$ falling below *b* in our sample of 74 commuting zones) comes from a data-generating process with β_j=0 everywhere. This is true for all four levels of *b* we use (-0.5, -0.3, -0.1, and 0). We obtain the same finding for the incidence of $\hat{\beta}_j$ exceeding 0.5, 0.3, 0.1, or 0. We also find that the sample variance of $\hat{\beta}_j$ (0.66) is well within the RI-based confidence intervals. Together with our evidence on t-statistics, the evidence on the dispersion of $\hat{\beta}_j$ does not support the presence of heterogeneous treatment effects on employment across the various commuting zones.

6. Comparisons with restaurant studies

Similar to this paper, the evidence from DLR also showed a bias in the traditional estimates due to spatial heterogeneity. Similar to DLR, we also have results using: (1) traditional fixed effects specifications, (2) spatial discontinuity based specifications, and (3) specifications with intermediate amount of controls for spatial heterogeneity. This naturally raises the question of how the set of elasticities here compare to those in DLR. On the one hand, the elasticities are not directly comparable, since DLR focuses on jobs, while we focus on individuals. Given the possibility of labor-labor substitution, the effect on jobs may differ from effect on a particular group (teens). Moreover, DLR's focus is on a particular industry, restaurants, while ours is on teens. However, since the fraction of teens earning minimum wages is similar to the fraction of restaurant workers earning minimum wages, it is instructive to compare employment elasticities across the two studies, with the proper caveats in mind.[17]

Table 7 provides employment elasticities from this study along with those in DLR that employed similar specifications. The first column of results in Table 7 generally

[17] Moreover, the wage elasticities across the two studies are quite similar across specifications, which is consistent with the fact that a similar proportion of teens and restaurant employees are minimum wage workers.

represents the canonical fixed effects model. In the present study the elasticities are -0.153 utilizing CPS data and -0.159 using Census/ACS data. The elasticity of -0.176 in DLR is very similar. These outcomes are in the typical range of a 1 to 3 percent disemployment effect from a 10 percent increase in the minimum wage.

Moving to the second column in Table 7, we can loosely compare specifications that employ division-specific and state-specific time controls. The CPS results from this study and those from DLR each show that incorporating such controls greatly reduces the elasticities and renders each insignificant; they are -0.024 and 0.039 respectively. Lastly, Table 8 presents the discontinuity-based specifications: the Census/ACS specification from this study using commuting-zone specific year effects and the DLR specification using contiguous border county pair specific time effects. For each finding, there is no disemployment effect.

While these results are not directly comparable, they clearly demonstrate the importance of including controls for heterogeneous trends in low-wage employment. In Dube, Lester and Reich, inclusion of division-specific time effects and state-level linear time trends provide imperfect proxies for their local estimators, which also produce employment elasticities indistinguishable from zero. Including even such coarse controls in the CPS data attenuates the disemployment effect for teens in an analogous manner. When we include better local controls using the Census/ACS data (i.e., cross-state commuting zones that are comparable to the contiguous county pairs of DLR) we find no disemployment effects on teens. Omitting controls for local differences in underlying local labor market conditions induces a serious bias in the teen studies as well as in restaurant studies.

8. Discussion and conclusions

Our analysis has found that policy discontinuity-based estimates can overcome the biases we have found in national studies using time and place fixed effects. Using the canonical fixed-effects specification on the sample of teens, we estimate employment elasticities of -0.153 with the CPS, and -0.159 with the Census/ACS; both are similar to the -0.3 to -0.1 percent disemployment consensus of the estimates in other national CPS studies. In contrast, using state-based minimum wage discontinuities within commuting zones, the employment elasticity becomes positive and marginally significant. When we take the unweighted average of the elasticities separately for each CZ, we find a result closer to zero. Overall, we take the evidence to rule out a sizeable disemployment effect, as opposed to suggesting a positive one. Consistent with our interpretation, intermediate levels of controls for spatial heterogeneity also substantially attenuate the measured disemployment effects in traditional specifications.

Our dynamic plots of the time path of teen employment around the minimum wage change using only the canonical time and state controls indicate that teen employment was unusually low and falling for a substantial time period *prior* to the actual increase. These findings support the conclusion that conventional estimates are contaminated by heterogeneity bias. We also find that traditional fixed-effects estimates for specific teen demographic groups are contaminated by the omission of local controls. Overall, the evidence strongly points to the failure of the canonical fixed-effects specification to control for heterogeneity across local labor markets.

Using a Randomization Inference approach and our Census/ACS data, we also examine whether minimum wage effects vary among local labor markets. To our knowledge, this is the first study in the minimum wage literature that examines this question. We do not detect such heterogeneous treatment effects across various labor markets. We find that the distribution of estimated employment elasticities across CZs is consistent with "chance" alone—due to sampling error and area-specific shocks.

Since the proportion of teens and the proportion of restaurant workers who are paid at or near the minimum wage are very similar it is of interest to compare our estimates to those in DLR. The estimated minimum wage employment effects are similar in both studies. Moreover, the results in the two studies change in similar ways with the inclusion of controls for spatial heterogeneity. These results suggest that the effects of controlling for such heterogeneity do not result from the focus on any one demographic group or industry.

A cautionary finding about interpretations from individual case studies also emerges from both sets of studies. Figure 5 of DLR presents the kernel density for estimated employment elasticities across 64 contiguous county pairs. While the mean estimate is zero, in individual case studies there is a probability of an employment differing from zero. We obtain a similar finding in Figure 4 of the present study and we find that this pattern can be rationalized by chance alone. Together these results suggest the limitations of results that are based only on individual case studies.

In summary, inattention to spatial heterogeneity clearly compromises estimates of minimum wage effects in the U.S. Since estimates in previous national-level studies insufficiently address this issue, the interpretation of the evidence in the existing minimum wage literature must be revised accordingly. With the accumulation of annual ACS data, using cross-state commuting zones to incorporate local controls with individual-level data can provide a powerful tool for minimum wage studies and for studies of other policies as well.

References

Aaronson, Daniel, Kyung-Hong Park and Daniel Sullivan 2006. "The Decline in Teen Labor Force Participation." *Economic Perspectives,* Federal Reserve Bank of Chicago 30, 1: 2-18.

Abadie, Alberto, Alexis Diamond and Jens Hainmueller 2007. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." Work ing Paper 12831. National Bureau of Economic Research.

Autor, David and David Dorn 2009. "This Job is 'Getting Old:' Measuring Changes in Job Opportunities Using Occupational Age Structure." *American Economic Review Papers and Proceedings* 99, 2: 45-51.

Baum-Snow, Nathaniel and Derek Neal 2009. "Mismeasurement of Usual Hours Worked in the Census and ACS. *Economic Letters* 102,1: 39-41.

Bertrand, Marianne, Esther Duflo and Sendhil Mullainathan 2001. "How Much Should We Trust Differences-in-Differences Estimates?" Working Paper 01-34. MIT Department of Economics.

__________2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics* 119*,* 1: 249-75.

Card, David and Alan Krueger 1994. "Minimum Wages and Employment: a Case Study of the Fast-Food Industry in New Jersey and Pennsylvania." *American Economic Review* 84, 4: 772-93.

_________2000. "Minimum Wages and Employment: a Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Reply." *American Economic Review* 90, 5: 1397-1420.

Congressional Budget Office 2004. "What is Happening to Youth Employment Rates?" CBO Paper. Washington, D.C.: Congressional Budget Office.

Conley, Timothy and Christopher Taber 2005. "Inference with 'Difference in Differences' with a Small Number of Policy Changes." Technical working paper 312. National Bureau of Economic Research

Donald, Stephen G. and Kevin Lang 2007. "Inference with Difference-in-Differences and Other Panel Data." *Review of Economics and Statistics* 89, 2: 221-33.

Dube, Arindrajit, William Lester and Michael Reich 2010. "Minimum Wage Effects across State Borders: Estimating Using Contiguous Counties." forthcoming *Review of Economics and Statistics*. Working Paper 157-07. UC Berkeley Institute for Research on Labor and Employment. www.irle.berkeley.edu/workingpapers/157-07.pdf.

Dube, Arindrajit, Suresh Naidu and Michael Reich 2007. "The Economic Effects of a Citywide Minimum Wage." *Industrial and Labor Relations Review* 60, 4: 522-543.

Foote, Christopher 2007. "Space and Time in Macroeconomic Panel Data: Young Workers and State-Level Unemployment Revisited." Working Paper 07-10. Boston, MA: Federal Reserve Bank of Boston.

Holmes, Thomas 1998. "The Effects of State Policies on the Location of Industry: Evidence from State Borders." *Journal of Political Economy* 106, 4: 667-705.

______2006. "Geographical Spillovers of Unionism." Working Paper 12015. Cambridge, MA: National Bureau for Economic Research.

Huang, Rocco 2008. "Evaluating the Real Effect of Bank Branching Deregulation: Comparing Contiguous Counties across U.S. State Borders." *Journal of Financial Economics* 87, 3: 678-705.

Kézdi, Gábor 2004. "Robust Standard-Error Estimations in Fixed-Effect Panel Models." *Hungarian Statistical Review*. 9: 95-116.

Lalive, Rafael 2008. "How Do Extended Benefits Affect Unemployment Duration? A Regression Discontinuity Approach." *Journal of Econometrics* 142, 2: 785-806.

Magruder, Jeremy 2009. "High Unemployment Yet Few Small Firms: the Contribution of South Africa's Centralized Bargaining Agreements." Mimeo, UC Berkeley.

Manning, Alan 2006. "A Generalized Model of Monopsony." *Economic Journal* 116, 508: 84-100.

Neumark, David and William Wascher 2000. "Minimum Wages and Employment: a Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Comment." *American Economic Review* 90, 5: 1362-96.

_____ 2007. “Minimum Wages, the Earned Income Tax Credit and Employment: Evidence from the Post-Welfare Reform Era.” Working Paper 12915. Cambridge, MA: National Bureau for Economic Research.

Orrenius, Pia and Madeline Zavodny 2008. “The Effects of Minimum Wages on Immigrants.” *Industrial and Labor Relations Review* 61, 4: 544-63.

Sabia, Joseph 2006. “The Effect of Minimum Wage Increases on Retail and Small Business Employment.” Washington, D.C.: Employment Policies Institute.

Seltzer, Andrew 1995. “The Political Economy of the Fair Labor Standards Act of 1938.” *Journal of Political Economy* 103, 6: 1302-42.

Sobel, Russell 1999. “Theory and Evidence on the Political Economy of the Minimum Wage.” *Journal of Political Economy* 107, 4: 761-85.

Thompson, Jeffrey P. 2009. “Using Local Labor Market Data to Re-Examine the Employment Effects of the Minimum Wage.” *Industrial and Labor Relations Review* 62, 3: 343-66.

Figure 1 Spatial heterogeneity in teen employment, 1990 to 2007

A Evolution of state-level teen EPOP ratios

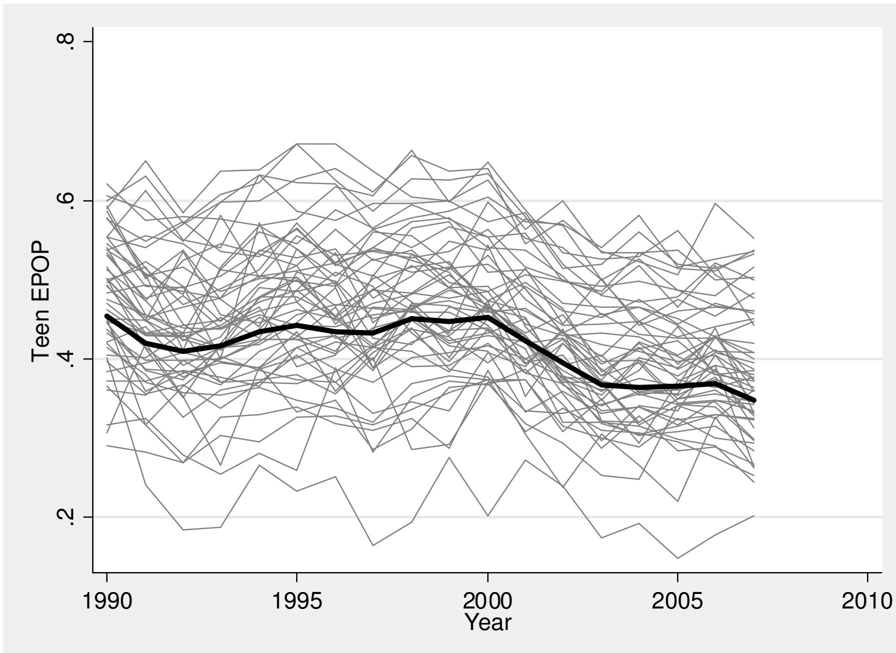

B State-level change in teen EPOP ratios: raw and net of overall EPOP ratios

All states

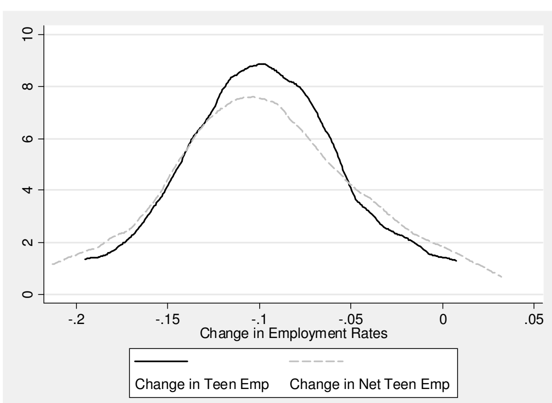

States with only federal minimum wage

Notes: CPS data 1990-2007. Panel A plots the evolution of state-level teen employment ratios using 1990-2007 CPS data. Gray lines represent individual states and the single black line is the national average. Panel B shows the kernel density estimates of state-level changes from 1990 through 2007 for (1) teen employment to population (EPOP) ratios and (2) teen EPOP ratios net of overall EPOP ratios. The first graph in Panel B is for all states, while the second is for states with no state-level minimum wages.

Figure 2 Annual time paths of wages, employment and hours in response to a minimum wage change, CPS data

Spec 1 (No additional controls) | *Spec 4 (State-linear trends and division-specific time effects)*

A Wage elasticity

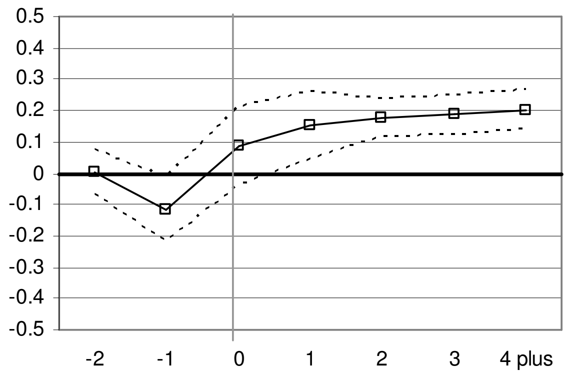

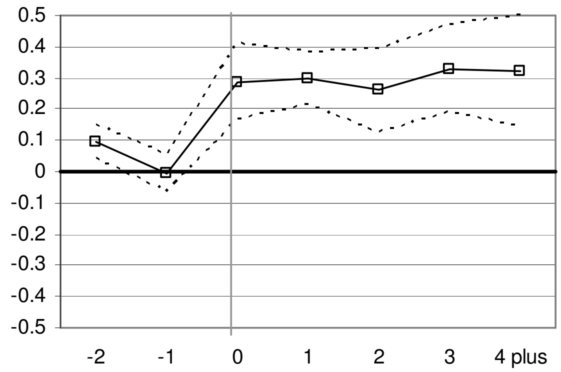

B Employment elasticity

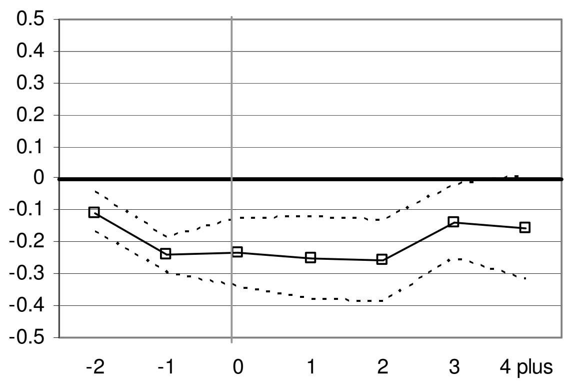

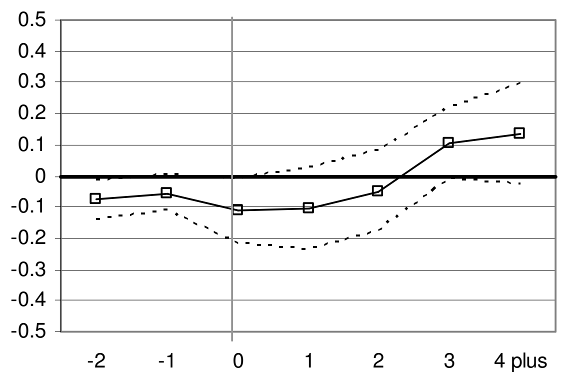

C Hours elasticity

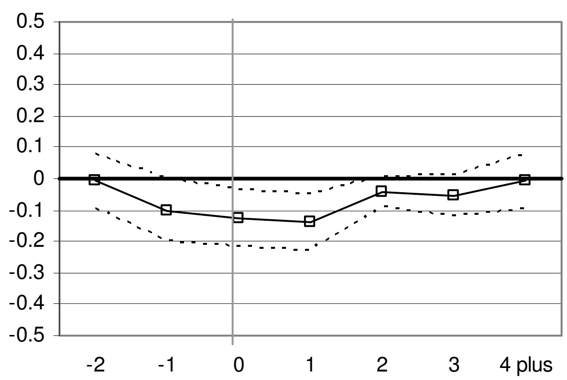

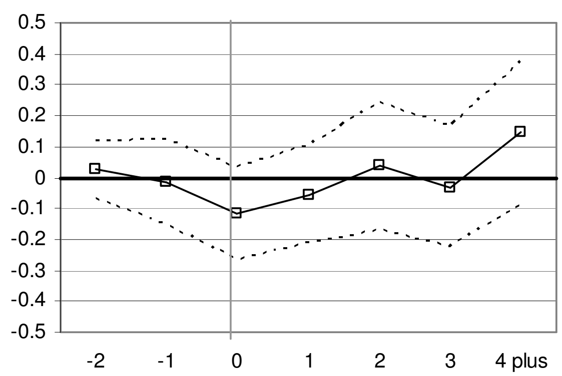

Years **Years**

Notes: The figures plot the *cumulative response* of log wage, employment and log hours using a distributed lag specification of a two-year lead, four-year lag and the contemporaneous log minimum wage considering a 25 quarter window around the minimum wage increase. For employment, coefficients are divided by average teen employment-to-population ratio, thereby representing employment elasticities. Specification 1 includes time and state fixed effects as well as the set of demographic controls reported in the text. Specification 4 additionally includes state-level linear trends and division-specific time effects (hence eliminating the variation between Census divisions). Dashed lines represent 90% confidence intervals around the estimates and were calculated using robust standard errors clustered at the state level.

Figure 3 Annual time paths of wages, employment and hours in response to a minimum wage change, Census/ACS data

Spec 5 (Common time effects) *Spec 7 (CZ-specific time effects)*

A Wage elasticity

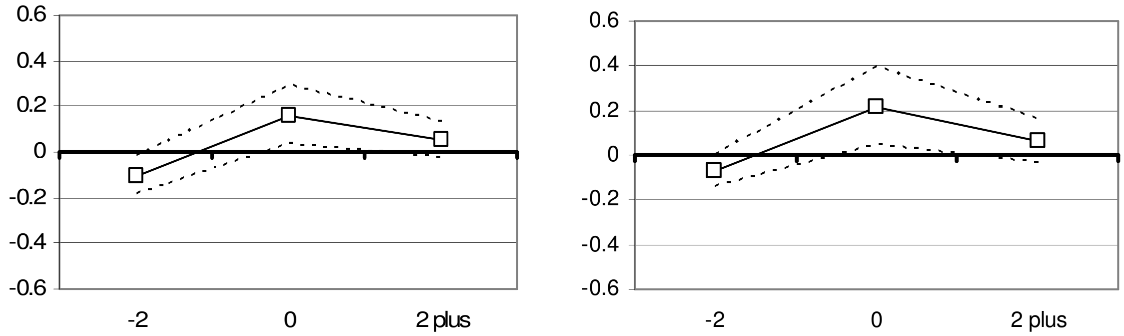

B Employment elasticity

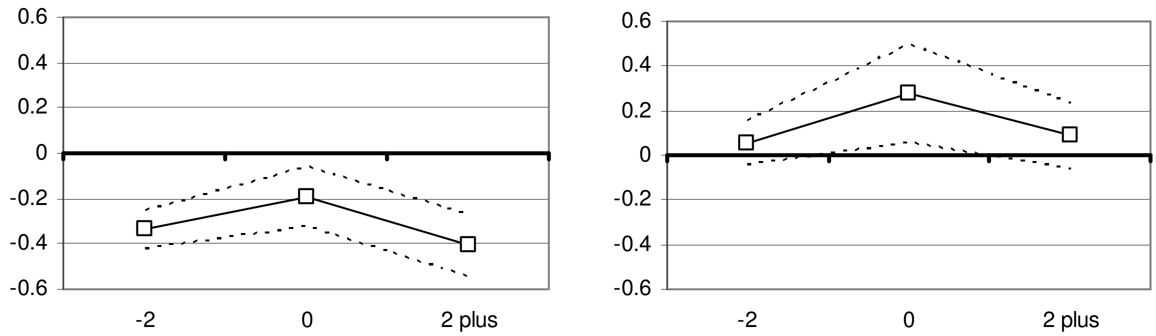

C Hours elasticity

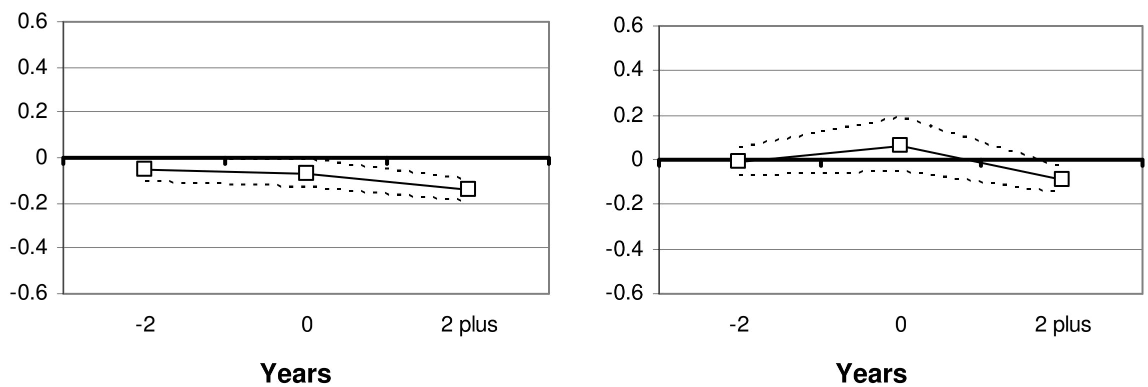

Notes: The Figures plot the *cumulative response* of log wage, employment and log hours using a distributed lag specification of a two-year lead, two-year lag and the contemporaneous log minimum wage considering a 5 year window around the minimum wage increase. For employment, coefficients are divided by average teen employment-to-population ratio, thereby representing employment elasticities. Specification 5 includes time and commuting-zone-by-state fixed effects as well as the set of demographic controls reported in the text. Specification 7 additionally includes CZ-year fixed effects (hence eliminating the variation between CZs). Dashed lines represent 90% confidence intervals around the estimates and were calculated using robust standard errors clustered at the state level.

Figure 4 Kernel density estimates of employment and wage elasticities across 74 commuting-zones (unweighted and weighted)

Unweighted

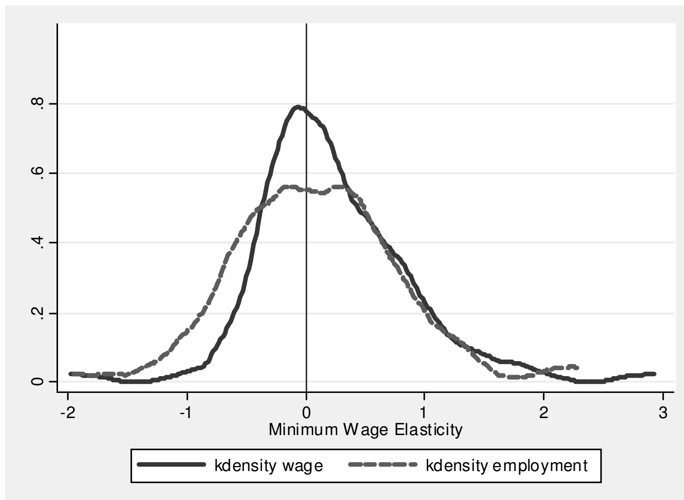

Weighted

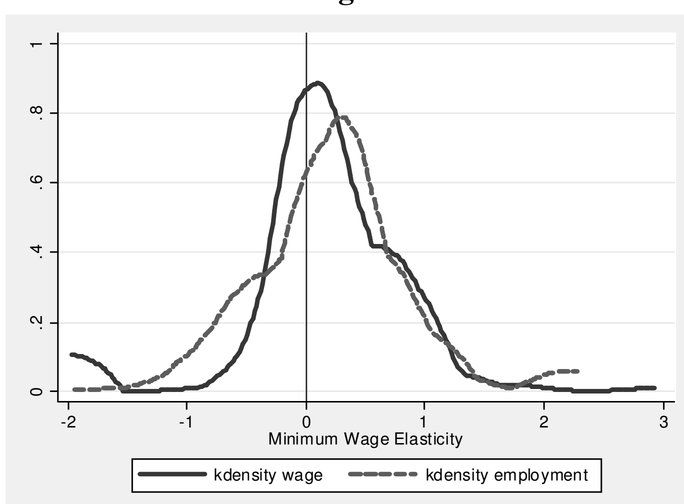

Notes: The figures plot the kernel density estimates of minimum wage elasticities (wage and employment) from 74 separate regressions for 74 cross-state commuting zones with minimum wage differences. The first figure plots the density without weighting the individual elasticities by population in that commuting zone. The second figure weights the commuting zones by population.

Figure 5 Scatter plot of employment and wage elasticities across 74 commuting zones (unweighted and weighted)

Unweighted

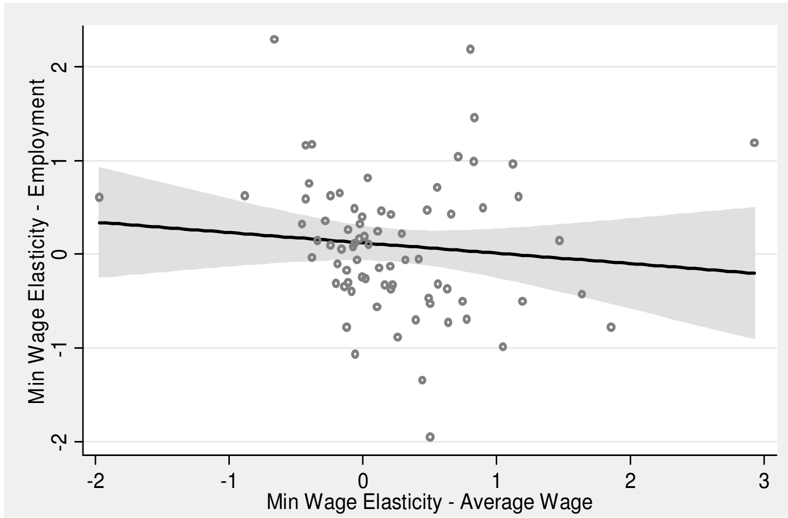

Weighted

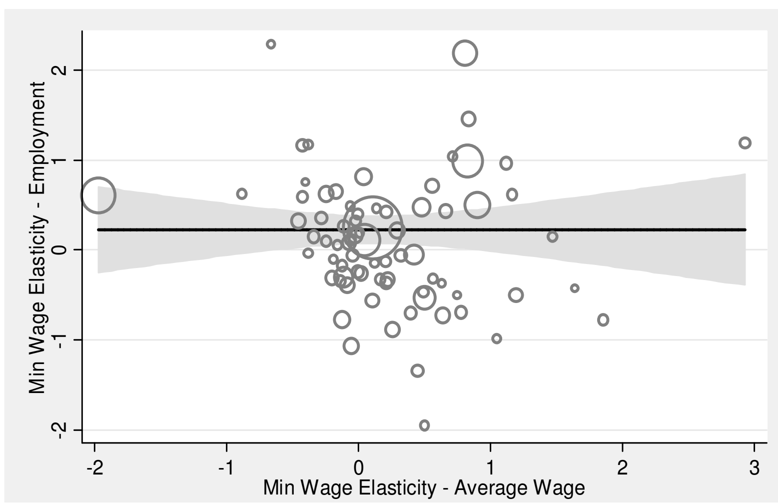

Notes: The figures show scatterplots (and fitted regression lines) of minimum wage elasticities of employment on minimum wage elasticities of average (teen) wage from 74 separate regressions for 74 cross-state commuting zones with minimum wage differences (Census/ACS data). The shaded regions show the 90% confidence intervals. The first figure shows the scatterplot and the fitted line without weighting the individual elasticities by population in that commuting zone. The second figure weights each commuting zone by population. The size of the circles in the second figure represents the population of the CZ.

Table 1 Employment to population ratios, teens 16-19, by Census division, selected years

	1990	1998	2007	Change 1990 to 1998	Change 1998 to 2007
United States	0.46	0.45	0.39	0.00	-0.07
New England	0.51	0.50	0.41	-0.01	-0.09
Maine, New Hampshire, Vermont, Massachusetts, Rhode Island, Connecticut					
Middle Atlantic	0.41	0.38	0.30	-0.02	-0.08
New York, New Jersey, Pennsylvania					
East North Central	0.51	0.52	0.39	0.01	-0.13
Ohio, Indiana, Illinois, Michigan, Wisconsin					
West North Central	0.57	0.60	0.48	0.03	-0.11
Minnesota, Iowa, Missouri, North Dakota, South Dakota, Nebraska, Kansas					
South Atlantic	0.43	0.44	0.32	0.01	-0.12
Delaware, Maryland, DC, Virginia, West Virginia, North Carolina, South Carolina, Georgia, Florida					
East South Central	0.39	0.43	0.31	0.05	-0.13
Kentucky, Tennessee, Alabama, Mississippi					
West South Central	0.39	0.39	0.33	0.01	-0.06
Arkansas, Louisiana, Oklahoma, Texas					
Mountain	0.52	0.50	0.39	-0.02	-0.11
Montana, Idaho, Wyoming, Colorado, New Mexico, Arizona, Utah, Nevada					
Pacific	0.44	0.40	0.31	-0.05	-0.09
Washington, Oregon, California, Alaska, Hawaii					

Source: Authors' calculations of Current Population Survey data

Table 2 Descriptive statistics, Census / American Community Survey

	Mean	Std Dev	N	Mean	Std Dev	N
	All Commuting Zones			Cross-State Commuting Zones		
Total			2,848,829			793,585
Male	0.51		1,468,070	0.51		408,505
Female	0.49		1,380,759	0.49		385,080
White	0.68		2,041,807	0.67		565,495
Black	0.13		319,918	0.12		83,530
Hispanic	0.13		323,464	0.14		111,101
Employed	0.38		1,101,100	0.38		308,329
Male	0.37		563,217	0.37		157,629
Female	0.39		537,883	0.39		150,700
White	0.42		872,677	0.42		242,115
Black	0.24		75,382	0.25		20,749
Hispanic	0.34		108,539	0.34		32,727
Hourly wage	$8.32	$7.44	1,664,023	$8.42	$7.47	460,338
Male	$8.60	$7.60	867,321	$8.69	$7.60	238,820
Female	$8.01	$7.26	796,702	$8.14	$7.32	221,518
White	$8.19	$7.22	1,302,866	$8.24	$7.17	359,799
Black	$8.58	$8.29	133,022	$8.84	$8.47	36,416
Hispanic	$8.64	$7.60	151,739	$8.98	$7.81	43,473
Usual Hours worked per week	25.63	13.21	1,704,142	25.45	12.94	470,685
Male	27.38	13.89	892,323	27.11	13.58	245,144
Female	23.76	12.15	811,819	23.68	11.97	225,541
White	25.02	13.08	1,336,456	24.88	12.83	368,484
Black	26.36	12.78	134,511	26.07	12.62	36,783
Hispanic	29.50	13.52	154,860	29.18	13.21	44,326
Minimum wage	$5.12	$0.95	2,848,829	$5.19	$0.98	793,585
State minimum wage (if above federal)	$5.96	$1.10	683,998	$5.90	$1.09	252,469
Unemployment rate	0.04	0.01	2,848,829	0.04	0.01	793,585

Notes: Census/ACS data from 1990, 2000, 2005 and 2006 for teens 16-19. The first three columns refer to the full sample of 741 commuting zones and the next three columns refer to the sample of 134 cross-state commuting zones. The White category excludes individuals of Hispanic origin. Standard deviations reported for continuous variables. Hourly wage is calculated as the annual wage and salary income, divided by the product of weeks worked in the year and the usual hours per week. Usual hours worked is reported for workers with positive usual hours of work. Both usual hours and wage are reported for individuals who worked at some in time during the past year, not just who were employed at the time of survey. Hourly wage is reported in 2007 dollars, while minimum wage is reported in nominal dollars.

Table 3 Descriptive statistics, CPS

	Mean	Std Dev	N
Total			401,744
Male	0.51	--	*203,320*
Female	0.49	--	*198,424*
White[1]	0.65	--	*278,931*
Black	0.15	--	*49,915*
Hispanic	0.14	--	*48,283*
Employed	0.41	--	169,783
Male	0.41	--	*85,228*
Female	0.41	--	*84,555*
White	0.46	--	*141,264*
Black	0.25	--	*12,192*
Hispanic	0.34	--	*16,327*
Hourly wage	$7.93	$8.50	165,455
Male	$8.28	$9.36	*82,356*
Female	$7.60	$7.50	*83,099*
White	$7.91	$7.71	*137,398*
Black	$7.88	$14.87	*12,039*
Hispanic	$8.06	$6.53	*16,018*
Usual hours worked per week	24.9	12.09	167,814
Male	26.4	12.61	*83,868*
Female	23.3	11.32	*83,946*
White	24.1	12.10	*139,490*
Black	25.6	11.10	*12,124*
Hispanic	29.1	11.83	*16,200*
Minimum wage	$5.06	$0.86	--
State minimum wage (above federal)	$6.17	$0.93	--
Unemployment rate	0.05	0.01	--

Notes: CPS data for 1990-2007 for teens 16-19. The White category excludes individuals of Hispanic origin. Hourly wage reported in 2007 dollars. Minimum wage reported in nominal dollars. Standard deviations reported for continuous variables. Average hourly wage is calculated for workers who reported a wage and were not self-employed or working without pay. Average hours worked is reported for workers with positive usual hours of work.

Table 4 Minimum wage effects on wages, employment and hours worked

		CPS				Census/ACS		
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log wages	η	0.120***	0.154***	0.205***	0.158***	0.110**	0.088*	0.151***
	se	(0.026)	(0.035)	(0.033)	(0.031)	(0.053)	(0.053)	(0.051)
	p-value		0.184	0.006	0.178		0.579	0.478
Employment	coeff	-0.063***	-0.043	-0.027	-0.010	-0.062*	0.002	0.050*
	se	(0.017)	(0.037)	(0.025)	(0.031)	(0.033)	(0.029)	(0.028)
	η	-0.153***	-0.105	-0.065	-0.024	-0.159*	0.005	0.129*
	p-value		0.490	0.101	0.022		0.000	0.001
Log usual hours	η	-0.096***	-0.091*	-0.038	-0.087**	-0.091***	-0.061*	-0.031
	se	(0.037)	(0.051)	(0.047)	(0.042)	(0.021)	(0.032)	(0.032)
	p-value		0.882	0.039	0.826		0.149	0.115
Division-specific time effects			Y		Y		Y	
State-specific linear time trends				Y	Y			
CZ–specific time effects								Y

Notes: Specifications 1-4 use 1990-2007 CPS data. Specifications 5-7 use 1990, 2000 Census and 2005, 2006 ACS data. All regressions include Commuting zone-by-state fixed effects. Specifications 1 and 5 include time fixed effects, quarterly for specification 1 and yearly for specification 5. Other specifications use additional time controls as indicated. *p-value* (estimated using SUR) refers to the probability value of rejecting the null hypotheses that coefficients from specifications 2, 3 and 4 are equal to those in specification 1; or whether coefficients from specifications 6 and 7 are equal to those in specification 5. Results are reported for the coefficient associated with the log of the minimum wage (i.e., the elasticity, η) for *log wage* and *log usual hours*; coefficients and elasticity η for *employment*. Each regression includes individual controls for gender, race (4 categories), age (4 categories), education (12 categories), and marital status (4 categories), as well as controls for the non-seasonally adjusted state unemployment rate. Wage regressions include workers and paid between $1 and $100 per hour in 1990 dollars and the log of hourly wage is the dependent variable. Log usual hours regressions are restricted to workers with positive hours and the log of hours is the dependent variable. Standard errors are clustered at the state level and are reported within parentheses. Significance levels are indicated by ***1%, **5%, *10%.

Table 5 Counts of individual case study t-statistics for employment that exceed cutoffs using randomization inference

	5% level	10% level	15% level	20% level
Empirical cutoffs using Randomization Inference	± 4.38	± 3.09	± 2.67	± 1.81
|t-statistic|>cutoff:				
Sample counts	1	2	6	14
Sample proportion (%)	1.4	2.7	8.1	18.9
t-statistic < cutoff (%)$^{-}$	0.0	1.4	5.4	10.8
t-statistic > cutoff (%)	1.4	1.4	2.7	8.1
Probability value: Counts>sample counts | H_0: β_j=0:	0.97	0.99	0.96	0.72

Notes: Empirical cutoffs for t-statistics are derived using Randomization Inference. The RI procedure permutes minimum wage series from 74 "in sample" commuting zones (with minimum wage differences) with employment series from 60 "out of sample" commuting zones without minimum wage differences. The t-statistics from these 9,028 regressions are used to form the empirical RI distribution of the t-statistics under the null hypothesis of zero effects everywhere. The first row of the table reports the cutoffs of the t-statistics under the null at the 5%, 10%, 15% and 20% levels. The "sample counts" row reports the number of t-statistics from the actual 74 "in sample" regressions that exceed these empirical cutoffs at various levels of significance; the "sample proportion" row reports the fraction of the 74 "in sample" regression t-statistics that exceed the empirical cutoffs. The next two rows disaggregate the proportion by the sign of the effect. Finally, the last row reports the probability that the number of counts under the null hypothesis exceeds the actual "in sample" counts observed.

Table 6 Distribution of minimum wage effects on employment: Actual counts and confidence intervals based on randomization inference

	Alternative ranges of employment effects								Variance
	(-∞, -0.5)	(-∞, -0.3)	(-∞, -0.1)	(-∞, 0)	(0, ∞)	(0.1, ∞)	(0.3, ∞)	(0.5, ∞)	
Sample count	12	19	32	37	37	31	15	11	
Sample variance									0.66
90% Confidence Interval									
Lower cutoff	9	15	24	30	30	24	15	9	0.54
Upper cutoff	20	28	38	44	44	38	28	20	1.66
95% Confidence Interval									
Lower cutoff	8	14	23	29	29	23	14	8	0.48
Upper cutoff	21	29	40	45	45	40	29	22	1.83

Notes: The first row "sample count" reports the number of minimum wage coefficients (regressed on employment) of the actual 74 "in sample" that fall in particular ranges. We then report the confidence intervals ("lower cutoff" and "upper cutoff") of the coefficients using Randomization Inference at the 90% and 95% levels. The RI procedure permutes minimum wage series from 74 "in sample" commuting zones (with minimum wage differences) with employment series from 60 "out of sample" commuting zones without minimum wage differences. The minimum wage coefficients (i.e., "employment effects") from these 9,028 regressions are used to form the empirical RI distribution of the coefficients under the null hypothesis of zero effects everywhere.

Table 7 A comparison of minimum wage employment elasticities

Study		Results		
This study, teens	η	-0.153***	-0.024	-----
CPS data	se	(0.042)	(0.075)	
1990-2007	90% CI	(-0.222, -0.083)	(-0.147, 0.099)	
This study, teens	η	-0.159*	------	0.129*
Census/ACS data	se	(0.085)		(0.072)
1990, 2000 / 2005, 2006	90% CI	(-0.299, -0.019)		(0.011, 0.247)
Dube et al. (2007), restaurants	η	-0.176*	0.039	0.016
QCEW data	se	(0.096)	(0.050)	(0.098)
1990-2006	90% CI	(-0.334, -0.018)	(-0.043, 0.121)	(-0.145, 0.177)
Division-specific time effects			Y	
State-specific linear time trends			Y	
Contiguous border county pair or commuting-zone specific time effects				Y

Notes: Elasticities are not directly comparable. They are presented to show the effects of using different model specifications and controls.
Significance levels: ***1%, **5%, *10%.