Credible Research Designs for Minimum Wage Studies

Sylvia Allegretto, Arindrajit Dube, Michael Reich and Ben Zipperer*

June 12, 2013

Abstract

Over the past two decades, the states with larger minimum wage increases have differed from other states in the depth of their business cycles, growth of upper-half wage inequality, increased job polarization, and political-economy. They have also been spatially clustered. We present estimates of minimum wage effects for teens and restaurant workers using five datasets and six different controls for spatial confounds. We show that the dis-employment results suggested by the canonical two-way fixed effects model are spurious, as these specifications generally fail falsification tests for pre-existing trends. Using policy variation within local areas (county pairs, commuting zones) or regions, as well as inclusion of state-specific trends, typically renders the employment effect small in magnitude and statistically indistinguishable from zero. Employment effects are also close to zero when we account for heterogeneity using lagged dependent variables and dynamic panel models. A synthetic control estimate that pools across state minimum wage increases between 1997 and 2007 also shows no evidence of job losses. We confirm the validity of local controls by demonstrating that synthetic control weights decline with distance: a donor state 100 miles away receives a weight seven times as large as a state 2,000 miles away. We also show that neighboring counties are more similar in terms of covariates than are other counties. These findings refute the claims made in a recent paper by Neumark, Salas and Wascher that criticizes the use of local controls.

*Allegretto: Institute for Research on Labor and Employment, University of California, Berkeley; Dube: Department of Economics, University of Massachusetts Amherst and IZA; Reich: Department of Economics and Institute for Research on Labor and Employment, University of California, Berkeley; Zipperer: Economics Department, University of Massachusetts Amherst. We are grateful to Zachary Goldman, Thomas Peake and Luke Reidenbach for excellent research assistance.

1 Introduction

The non-random distribution of minimum wage policies poses a serious threat to identification of the policy's effects. Minimum wage policies are spatially clustered, with important economic and political differences existing between states with relatively high versus low minimum wages over the past two decades. In a series of four papers (Allegretto, Dube and Reich 2009, hereafter ADR1; Allegretto, Dube and Reich 2011, hereafter ADR2); Dube, Lester and Reich 2010, hereafter DLR1; and Dube, Lester and Reich 2012, hereafter DLR2a[1]) we demonstrate the importance of accounting for spatial heterogeneity of minimum wage policies when estimating their employment effects on highly-affected groups such as teens, and employees in low-wage sectors such as restaurants. In these papers we used three different statistical approaches and five different datasets to show the importance of spatial heterogeneity and provide reliable estimates of the policy's impact. In each paper, the canonical two-way fixed effects model suggests a disemployment effect with elasticities in the -0.1 to -0.2 range. However, the inclusion of spatial controls renders the employment estimates small in magnitude and statistically indistinguishable from zero. These papers also find pre-existing negative employment trends for the canonical model, but not for the specifications with fine-grained spatial controls–indicating the latter set of results has greater internal validity.

ADR1 examines minimum wage effects on employment of teens using Census and American Community Survey (ACS) data from 1990 to 2006 and accounts for spatial heterogeneity by using variation within local labor markets ("commuting zones"). ADR2 examines minimum wage effects on teen employment using CPS panel data from 1990 to 2009 and controls for spatial heterogeneity at the level of the nine Census-defined divisions, as well as through the inclusion of state-level linear trends. DLR1 examines minimum wage effects on employment of restaurant workers using establishment-based Quarterly Census of Employment and Wages (QCEW) data from 1990-2006 and spatial heterogeneity controls at the level of contiguous county pairs across state borders. DLR2a also uses contiguous county pairs,

[1] We refer to the updated 2013 version of this paper as DLR2b.

and examines minimum wage effects on employment stocks and flows for teens and restaurant workers using establishment-based Quarterly Workforce Indicators (QWI) data from 2001-2008.

Our use of regional controls and policy discontinuities to contol for heterogeneity is well-established in the discipline. While we were one of the first sets of reserchers to apply these tools in the minimum wage literature by generalizing the case study approach of Card and Krueger (1994, 2000), they are part of the profession's standard toolkit. Holmes (1998) and Huang (2008) use the contiguous border county approach to identify the effects of right-to work legislation and bank regulation, respectively. Their approach is similar to ours in DLR1. Many other recent papers utilize a border discontinuity approach, including Magruder (2013) in the context of minimum wages in Indonesia, Black (1999) in the context of schools and housing markets, and Goldstein and Udry (2008) in the context of land rights and investment in Ghana. Lee and Lemieux (2010) discuss the methodology of spatial discontinuity estimation in their survey article. Moving to coarser forms of geographic controls, Autor (2003) uses Census division-specific time effects along with state-specific linear trends to avoid spurious correlations between state employment and outsourcing. His identifying assumptions are identical to ours in ADR. These tools, which focus on finding careful controls, are well integrated parts of the "credibility revolution" that has swept through much of labor and applied microeconomics.

In a 2013 paper, Neumark, Salas and Wascher (hereafter NSW) criticize two of our papers: ADR2 and DLR1. Their central claim is that local or even regional controls are unwarranted when estimating minimum wage effects. NSW's main evidence comes from their application of the synthetic control estimator of Abadie et al. (2010). NSW claim that this estimator frequently does not choose nearby states or counties as controls, suggesting that our specifications with local controls are faulty. They do not explain why focusing on local area variation should produce incorrect results except to say that we throw away too much variation. Yet the differences in findings between our work and theirs differ not in the degree of precision but rather in the magnitude of the estimates. Besides potentially reducing statistical power, can the inclusion of a set of controls ever exacerbate bias? Increased bias could occur only if the treatment status affects the control variables themselves (see Angrist and Pischke 2008, pp. 68). This possibility does not apply to regional controls, since a state's minimum wage policy does not affect its geographical location, nor is it likely to affect outcomes in other nearby states. For border discontinuity designs, cross-border spillover may be a threat to identification. However, as DLR1 shows, spillovers do not appear to be important at the county level in the minimum wage context. Moreover, NSW do not base their criticism of regional controls on spillover issues. NSW appear to defend the canonical two-way fixed effects model (i.e. using state panel data with state and time fixed effects), which they have used, with some variations, since the early 1990s (e.g., Neumark and Wascher 1992).

In this paper, we assess the role of spatial controls in credible research designs for minimum wage studies. In Section 2 we provide evidence on the nature of the spatial heterogeneities in minimum wage policies. We show that states experiencing greater increases in minimum wages over the past 20 years are systematically different in labor market characteristics that are unrelated to the minimum wage policy. These states have experienced more severe economic downturns; they have experienced greater job polarization in the form of sharper reduction in routine task intensive jobs; and they have seen faster growth in upper-half wage inequality. These time-varying differences suggest that the canonical fixed-effects model is likely to mis-estimate the counterfactual employment growth absent a minimum wage increase.

We then present new estimates for both teens and restaurant workers with updated data at least through 2010 from the Census, ACS, CPS, QCEW, and QWI. We utilize the three statistical approaches to accounting for spatial heterogeneity that we used in our aforementioned papers: comparing across contiguous counties, comparing within commuting zones, and using within-region variation along with state-level trend controls. Our consistent finding is that the inclusion of these spatial controls does not attenuate the treatment effect on earnings but reduces the employment effects in magnitude and renders them statistically insignificant–regardless of the data used or the approach employed. After accounting for spatial heterogeneity, the employment elasticities for teens and restaurant workers range between -0.06 and +0.13. We also find that, in contrast to the results for employment *stocks*, employment *flows* fall sharply following minimum wage increases. Along with the earnings impact, the flow results contradict claims that our local specifications throw out too much variation to detect minimum wage effects. The reduced separations and hires are also indicative of labor market search frictions, which have been hypothesized to explain the lack of minimum wage employment effects.

In Section 3 we address the validity of local-area controls. We begin by clearly showing that pre-existing trends compromise the canonical model, but not our preferred specifications. Contrary to claims in NSW, this conclusion is shown to hold independent of details, such as how many leads are included or at what frequency. We also show that contiguous counties are indeed better control groups in the sense of having more similar covariates–contradicting another claim in NSW.

The synthetic control approach, when implemented properly, can also provide a useful way to construct control groups. However, the conclusions that NSW draw from their exercise are incorrect. First, calculating synthetic control weights is uninformative about the relative merits of using local controls (as we advocate) versus using the standard two-way fixed effect model pioneered by Neumark and Wascher (1992), which NSW seem to advocate.The overlap between two sets of credible controls does not tell us about the quality of a third set of controls. While there may be different credible research designs, including using synthetic controls, the two-way fixed effect model is not one of them.

Second, we show that the synthetic control donor weights are indeed larger for nearby areas. We obtain this result for teen employment using a more general setup in which we use state-level placebo treatments and document the distance between the treated state and donor states that are picked by the synthetic control method . Our nonparametric plot shows that synthetic control weights decline rapidly with distance. In general, a donor state that is 100 miles away receives, on average, a weight seven times as large as a state 2,000 miles away. Moreover, calculations using NSW's own synthetic control weights seem to suggest the opposite of what they claim. Using their results, we show that both same-division states and neighboring counties have weights that are three to four times as large as other states and counties.

In Section 4, we provide new evidence on minimum wage elasticities for teen employment using two different approaches. First we discuss new findings from our own recent work using the synthetic control approach (Dube and Zipperer 2013). We show results that use all state-level minimum wage changes between 1997 and 2007 with at least two years of pre-intervention data and at least one year of post-intervention data without other minimum wage changes. For the nineteen resulting events, we find a clear wage effect, with a minimum wage elasticity of +0.17 that is statistically significant and within the range in our previous papers. However, we do not detect any employment losses, with the minimum wage elasticity for employment close to zero (-0.03). These results are also virtually identical to results we obtain for teens using local controls. The very method that NSW use to question our research design actually produces results that confirm our findings. Second, we apply a related approach to controlling for heterogeneity across states: estimating specifications that include lagged dependent variables (LDV). While the synthetic control is a semi-parametric generalization of the LDV method, LDVs are especially useful when the treatment continually varies, as is the case for minimum wage policies. The results from these regressions (including dynamic panel models with both lagged outcomes and panel fixed effects) clearly show that allowing for time-varying heterogeneity produces employment estimates close to zero.

In Section 5, we discuss the matching estimator presented in NSW, which is the authors' main proposed advancement over the two-way fixed effects model. They match treated units to control units (many of which actually receive treatment) based on synthetic control weights, which in turn are calculated to match residuals from a separate OLS regression. This estimator is completely *ad hoc.* The authors themselves admit that there is no actual econometric foundation for their approach and justify it instead on its supposed "intuititve" and "heuristic" appeal. We argue that the NSW matching estimator is fatally flawed for a multiplicity of reasons.

We discuss our conclusions in Section 6, where we also propose some guidelines for evaluating good research designs for minimum wage studies. An Appendix responds to some of the other claims made by NSW, including the issue of treating recessionary periods. As with the "NSW matching estimator," their proposed solutions on the topic of recessionary periods lack any solid econometric foundations, while a number of standard solutions all confirm our findings of small employment effects.[2]

2 The importance of spatial heterogeneity and controls in minimum wage studies

We begin this section by describing the nature of spatial heterogeneity as a problem for credible identification of minimum wage effects. States that have increased minimum wages above the federal level cluster regionally and differ in significant political and economic dimensions. We then present updated results using our previous approaches to control for this spatial heterogeneity.

[2]We acknowledge two small errors in DLR1 and thank NSW for discovering them during their replication of our work. First, there is a mistake in the map in DLR1 showing contiguous county pairs with minimum wage differences. Our map shows all the pairs with minimum wage differences through 2008q1, while our primary estimation sample only went through 2006q2. We did use this added minimum wage information from 2006q2-2008q1 in our dynamic specifications with leads and lags, but not for our primary specification, since our outcome data only went through 2006q2. Therefore, the map inadvertently overstated the number of border segments with policy variation. Second, there was a small error in the Maryland minimum wage for 2006 q1 and q2; we confirmed that none of our estimates were materially affected by this error. In this paper, we report updated results using data through 2010q4, and our results are all quantitatively close. We also show a map with the exact border counties used in the employment analysis here, for both balanced and unbalanced samples.

2.1 Spatial heterogeneity: the nature of the problem

States raised their minimum wages above the federal minimum primarily during the two long spells of federal inaction (1981 to 1991, and 1998 to 2007). Regional clustering of minimum wage policies has been prevalent since the early 1990s. To show this clustering succinctly, we divide the fifty states (and Washington DC) into "high" and "low" categories based on whether the average prevailing level of the minimum wage between 1990 and 2011 was above or below the median ($4.84). Because federal minimum wage increases typically erase many of the cross-state gaps in the minimum wages, the *average level* of the minimum wage in a state is also closely related to the *variance* of the minimum wage in that state over time, which is relevant for understanding the variation used in panel regressions.

Figure 1 provides two maps showing the spatial distribution of minimum wage states that have high average levels and high variance, showing a high degree of overlap. While such states are present in every region of the U.S., the maps also show significant spatial clustering. States in the Northeast, in parts of the Midwest, and in the Pacific regions are much more likely to have high state minimum wages, while states in the Southeast and the Mountain regions are much less likely.

Table 1 shows that high minimum wage states look quite different in their political economy characteristics. For example, they are much more Democratic-leaning: 88 percent of high minimum wage states voted for Barack Obama in 2008, as compared to 24 percent of low minimum wage states. High minimum wage states also have unionization rates that are nearly twice as high, and they experienced proportionately smaller reductions in these rates over the past two decades. These differences, which are all statistically significant at the 5 percent or 1 percent levels, raise the possibility that other systematic policy trends may differ between these groups of states.

The labor markets in high minimum wage states also differ substantially in dimensions that are unrelated to minimum wage increases. Table 1 displays patterns for three attributes of the labor market that differ in high and low minimum wage states: upper-half wage inequality, employment polarization, and the nature of the business cycle. This list does not exhaust all differences between the two types of states. But it does illustrate how longer run employment trends and short run fluctuations differ markedly in high versus low minimum wage states.

As Table 1 shows, during the period 1990 to 2011, high minimum wage states experienced a sharper growth in upper-half wage inequality as measured by the 90-50 ratio. In the high minimum wage states the 90-50 ratio increased from 2.07 to 2.26 between the business cycle peaks of 1990 and 2007. In low minimum wage states this ratio increased more modestly, from 2.15 to 2.22. This differential growth is statistically significant at the 1 percent level. Since these measures capture the wage distribution at or above the median, it is highly unlikely that varying levels of minimum wages could explain these differences.

This divergent pattern of inequality between high and low minimum wage states suggests different trajectories for labor demand. For example, technological change affects both inequality and labor demand for low-skill workers. A simple skill-biased technical change model predicts lower demand for low-skill workers when inequality is growing faster (Katz and Murphy 1992). In Acemoglu and Autor's (2010) three-skill-group model, technological change (or offshoring) that eases the replacement of middle-skill (e.g., routine) tasks with capital will not only increase upper-half wage inequality; it can also reduce the relative demand for low-skill workers as they compete with middle-skill workers for jobs with low-skill tasks. The greater growth in upper-half inequality in high minimum wage states could thus reflect factors that also explain why the employment rates for low-skill groups such as teens have fallen in those same states.

Relevant to this point, Smith (2011) shows that between 1980 and 2009, local labor markets with greater historical incidence of routine task intensive occupations saw (a) greater growth of adults in historically teen occupations, and (b) sharper declines in teen employment shares. We build on this evidence and show that labor market polarization patterns differ in high and low minimum wage states, generating differing trends in labor demand for teens.

Figure 2 and Table 1 show that the Routine Task Intensity (RTI) index of occupations fell more in high minimum wage states.[3] In 1990, high minimum wage states had more workers in routine task intensive occupations. This gap was statistically significant and substantial– amounting to about two-thirds of a standard deviation in the RTI index across states (results not shown). Over the next seventeen years, however, routine task intensity fell more in high minimum wage states; this trend difference too is statistically significant. As a result, the RTI gap was more than fully closed by 2007. In other words, high minimum wage states experienced greater growth in employment polarization, which likely put more downward pressure on employment for low-skill workers such as teens.

In general, the relationship between inequality, task intensity, and low-skill labor demand may be complicated. For example, in the two-skill-group model of Autor and Dorn (2013), this relationship depends on the substitutability between routine labor and computer-capital in production, substitutability of goods and services in consumption, and relative mobility costs by skill type. Nonetheless, their research indicates that areas with different trends in inequality and polarization are likely to experience different trends in low-skill labor demand. For example, Autor and Dorn show substantial geographical heterogeneity in inequality and employment trends by skill that are based on the initial occupational distribution in a local labor market (commuting zone), and hence exposure to task-biased technical change.

Consider the implications of these heterogeneities for the two-way fixed effects model, which assumes all such heterogeneity can be explicitly controlled by using common time effects and time-invariant state effects, plus a control for the overall unemployment rate. The patterns of heterogeneous trends across local labor markets directly contradict the fixed effects model. At the same time, the presence of such heterogeneity across local labor market suggests using within-local-area comparisons to identify minimum wage effects that are not contaminated by structural differences among labor markets. In ADR1, and in

[3] We use data and definitions of RTI from Autor, Levy and Murnane (2003), available from David Autor's data archive. We use the same three task measures– routine, manual and abstract–which are matched to 1990 occupation definitions. RTI is defined as $\ln(routine) - \ln(manual) - \ln(abstract)$. We calculate the employment-weighted means of this measure by state and year.

our updated analysis in this paper, we specifically consider variation in the minimum wage within commuting zones, which accounts precisely for the geographical heterogeneity in task reallocation that Autor and Dorn document.

Other evidence also points to differences in the labor market structures of high and low minimum wage states. For example, as shown in Figure 2 and Table 1, while the average unemployment rate was similar in these two groups of states, this similarity masks important differences in the nature of the business cycles. Table 1 shows that the variance (over time) in the unemployment rate was 48 percent larger in high minimum wage states, and this difference was statistically significant at the 5 percent level. Consistent with this pattern, the actual employment decline from peak to trough was 39 percent greater in high minimum wage states, when averaged over all recessions in the 1990-2011 period. The tendency of high minimum wage states to have sharper jobs recessions also raises caution flags for the two-way fixed effects model. If the patterns of demand shocks facing the two groups of states are so different, it is possible that fluctuations for low-wage workers also differ across the two groups of states. Moreover, simply including an overall unemployment rate as a control would proxy poorly for such heterogeneity. The relationship between overall and low-wage unemployment need not be the same in highly-cyclical versus moderately-cyclical states. This difference is of particular concern since minimum wage changes are much more likely to occur during economic expansions (Reich 2010).

In addition to the subtantial heterogeneity among states in the evolution of wage inequality, disappearance of routine jobs, and unemployment volatility, we also find regional clustering in these variables (Figure 2). Although the clustering patterns are not all identical, they clarify why regional controls work to reduce confoundedness. For example, the coastal states typically display a greater than average fall in routine task intensity, a greater than average increase in 90-50 wage inequality, and a greater variance in the unemployment rate. As shown in Figure 1, they are also more likely to be high minimum wage states. This geographic clustering of both the policy and of potential confounds motivates the use of spatial controls as part of a credible research design to study the effects of minimum wages. Differing labor market trends among high and low minimum wage states suggest that the two-way fixed effects model may be inappropriate, since the model implicitly assumes parallel trends in the outcomes of interest.

2.2 Updated results from previous work: four papers, five datasets, and three spatial control approaches

We have addressed the concerns of spatial heterogeneity in a variety of ways in the four key papers we have written on the topic. In part, the controls depend on the nature of the data; but generally, a more localized estimate better accounts for unobserved heterogeneity. When we have been able to use fine-grained geographic variation (such as in ADR1, DLR1, DLR2a), we have opted to use a border discontinuity approach and compared counties (or parts of commuting zones) across policy borders. When such variation has not been available, we have used coarser controls such as allowing time effects to vary by Census division and by including linear state-specific trends.

In this section, we provide new results using updated data that incorporates variation in minimum wages through 2010 or 2012 depending on the dataset.[4] We discuss the key results for the two groups we have studied in detail–teens and restaurant workers–that have also been studied extensively in the wider literature. While the details vary across datasets, in general we estimate two types of specifications. The first is the canonical model with time (t) and place (j) fixed effects. Here i denotes the unit of observation, which can be a place

[4]For the ACS, CPS, and the QWI, our samples extend through 2011, 2012q4 and 2011q4, respectively; these were the most recent data available at the time we accessed them. For the QCEW, our sample ends at 2010q4 because this appears to be the latest period for which the quarterly employment count was broken down by the 4-digit level NAICS restaurant subsectors (full and limited service restaurants, 7221 and 7222, resepectively). QCEW data beginning in 2011q1 uses 6-digit NAICS subsectors to identify full and limited service restaurants, and in a future version of this paper we hope to use those sectors if they are indeed comparable to data from earlier years, and are disclosed as frequently as before. We use the 4-digit level information for subsectors 7221 and 7222, as opposed to the somewhat broader 3-digit sector 722, because DLR1 defined the restautarant sector as the union of those two subsectors, and here we wish to retain the same definition for comparability.

j or an individual in that place depending on the data.

$$Y_{it} = \alpha + \beta MW_{jt} + \mathbf{X}_{it}\Lambda + \gamma_j + \tau_t + \nu_{it} \tag{1}$$

The key independent variable is the log of minimum wage (MW), while $\mathbf{X}$ is a vector of controls. We report all the results as elasticities.

Next, we estimate using local controls that account for spatial heterogeneity. We allow the time-effects to vary at geographic level g, where g can be (1) county pair, (2) commuting zone, or (3) Census division, depending on the data and specification.

$$Y_{it} = \alpha + \beta MW_{jt} + \mathbf{X}_{jt}\Lambda + \gamma_j + \tau_{gt} + \nu_{it} \tag{2}$$

When using the CPS, the geographic level g is the Census division–which is coarse. There we additionally include a state-specific linear time trend. Below we provide some more detail for each case and discuss the results, which are presented in Table 2. (For comparison, the results from the original papers are reproduced in Appendix Table A2.)

2.2.1 Results for teens

The relatively large proportion of minimum wage workers among teens make them an attractive group to study the effects of minimum wage policies. We are more likley to detect an impact–whatever it may be–for this demographic group. The first of our approaches uses data from the decennial Census and the American Community Survey (ACS), along with variation within local labor markets to examine minimum wage effects on teen employment. ADR1 originally used the 1990 and 2000 decennial Census and the more recent ACS from 2005-2006. Here we update these results by adding ACS data from 2007-2011, giving us substantially more local variation than in ADR1. Following Autor and Dorn (2012), we use commuting zones (CZs) as a measure of the local labor market and focus on the ones that straddle state boundaries–thus offering local variation in minimum wages. Since the CZ is a relatively unfamiliar construct to many economists, we describe here the construction of CZs in greater detail.

The Bureau of Labor Statistics partitions all U.S. counties into 741 CZs, based on inter-county commuting flows. Of these, 135 straddle state boundaries, covering 48 states (including DC). These cross-state CZs form our estimation sample. In turn, 110 of the 135 cross-state CZs had some within-CZ minimum wage variation in our sample.[5] Panel B of Figure 3 shows a map of the counties constituting the cross-state commuting zone sample and identifies those that have minimum wage variation.

The individual level regressions include fixed effects for each CZ-by-state group.[6] The "canonical" specification includes (common) year effects, which assumes that different CZs have parallel trends conditional on covariates. The "local" specification includes CZ-by-year fixed effects, which means it only uses within-CZ variation to identify minimum wage effects. Since the data is at the annual level, we use annual average minimum wages for each state. These estimates are shown in columns 1 and 5 of the upper panel of Table 2.

The second of our approaches uses Census-division specific time effects along with state-specific linear trends to account for spatial heterogeneity. In ADR2, we used CPS data between 1990-2009. Here, we have updated these results using data through 2012. These results are displayed in columns 2 and 6 of the upper panel of Table 2.

The third set of results use a relaively new BLS dataset, the Quarterly Workforce Indicators (QWI), for 2000q1 to 2011q4, and the contiguous county pair approach.[7] Like the QCEW, the QWI is also based on employer-reported official Unemployment Insurance records. The Census Bureau combines these records with with current demographic in-

[5]Similar to Autor and Dorn, we map PUMAs to CZs in the Census and ACS. In some cases, a PUMA cannot be uniquely assigned to a CZ. In such cases we assign these individuals to both CZs, with adjusted sampling weights reflecting the population shares of the PUMA in each CZ. There is never any uncertainty about which state (i.e., the policy unit) the PUMA belongs to; therefore, the probabilistic construction of CZs does not introduce any error in the treatment status of an individual. This is also true for the city-level policies we use, as both San Francisco and Washington D.C. have uniquely assiged PUMAs.

[6]The regressions also include standard demographic controls–dummies for age, sex, race/hispanic status, marital status, and educational attainment, as well as state-level unemployment rate.

[7]The QWI data for the 1990s is available for only a handful of states.

formation from other administrative censuses and population surveys. The QWI data offer employment counts and average earnings by detailed industry at the county level for specified age and gender groupings, and as well quarterly figures for hires and separations. Therefore, the data permit analyses both for teens and restaurant workers.The QWI age categories allow us to identify teens as those between 14 and 18 years of age. The employment results for teens are reported in columns 4 and 8 of Table 2; these are reproduced from an updated version of our paper, DLR2b.[8]

The results in all three of these papers indicate that adding local controls (columns 5, 6 and 8) does not substantially affect the estimated effects on teen earnings, which range between 0.12 and 0.30, depending on the data set. If anything, the estimates are somewhat larger with local controls. For employment, the canonical model suggests elasticities that are statistically significant and that range between -0.12 and -0.16. However, in all cases, the inclusion of local controls makes these estimate less negative; the elastiticies range between -0.06 and 0.13 and are always statistically indistinguishable from zero. Using commuting zone controls, the employment elasticity becomes sizably positive (0.13), although not statistically significant at conventional levels.

2.2.2 Results for restaurant workers

Minimum wage research has also focused on restaurants. Restaurants employ more minimum wage workers than any other industry, and a much higher proportion of their workers are paid near the minimum wage than workers in other sectors. DLR1, using the Quarterly Census on Employment and Wages (QCEW) and DLR2a, using the Quarterly Workforce Indicators (QWI), each examine the effects of minimum wage increases on restaurant workers. We have updated the results from the QCEW to include data from 1990q1 through 2010q4. The

[8]The previous version, DLR2a, provided results using the 2001-2008 sample period. Those results are quantitatively close and are also reproduced in this paper's Appendix Table A2. Unlike the older DLR2a (as well as DLR1), the newer DLR2b uses a maximal distance cutoff of 75 miles between county centroids. This excludes geographically large counties that are shown to be less similar to each other, and whose inclusion is less clearly justified by a border discontinuity approach. We follow the same sample restriction in the results reported here for the QWI. As shown in DLR2b, the results are not sensitive to the choice of cutoff.

QWI results are reproduced from the updated DLR2b, which uses data from 2000q1 through 2011q4. In Panel A of Figure 3 we show the minimum wage variation available in border counties with restaurant employment in the QCEW.[9]

In the lower panel of Table 2, we show the results from the QCEW and the QWI in columns 3 and 7, and 4 and 8, respectively. The results for earnings are unchanged by adding the spatial controls, and they range between 0.19 and 0.21. Again, the canonical model tends to produce more negative estimates, while the inclusion of spatial controls reduces the employment effect substantially and renders it indistinguishable from zero. The elasticities using the local specifications are -0.02 and 0.01 from the QWI and QCEW data, respectively.

Overall, the evidence summarized in Table 2 firmly establishes that for a high impact demographic group–teens–and for a high impact sector–restaurants–the earnings effects are positive and statistically significant, while the employment effects are small in magnitude and statistically indistinguishable from zero.

2.2.3 Effects on employment flows

The QWI dataset allows us to examine also minimum wage effects on employment flows, specifically hires and separations. To explain the lack of employment effects, Card and Krueger (1995) present a dynamic monopsony model, in which minimum wage increases reduce recruitment and retention costs. As Card and Krueger comment, these costs can be substantial because low-wage sectors have high turnover rates. More recently, Manning

[9]The results we report in the text follow DLR1 in using balanced panels of counties that report data for all periods. Panel A of Figure 3 illustrates, however, that roughly half of the counties in the QCEW have non-disclosed data in at least one period during 1990-2010, leading to the exclusion of roughly two-thirds of border counties due to the requirement that both counties in a pair have balanced panels. Moreover, the number of excluded counties due to non-disclosure in at least one period rises somewhat with the inclusion of additional years of data. To address concerns with selectivity of the sample, our Appendix table A3 shows the results using balanced and unbalanced samples of counties, including an intermediate specification where the full pair restriction is not imposed. The results are quantitatively similar: using the border discontinuity research design, QCEW restaurant employment elasticities are 0.008 and -0.034 for the balanced and unbalanced sample, respectively. These estimates are statistically indistinguishable from zero. The full-period reporting requirement for counties is not a major concern with the QWI since it uses "fuzzing" as opposed to outright suppression of the data to protect confidentiality.

(2003) developed a more formal version of this model, based on Burdett and Mortensen (1998), and a 2010 symposium in the *Journal of Labor Economics* presented a series of empirical studies that assessed the dynamic monopsony model by estimating firm-level separations elasticities. However, none of these studies directly test the effects of minimum wages on market-level employment flows with representative data. We do so in DLR2a.[10]

Columns 4 and 8 in Table 2 report our results for labor market flows. The upper panel presents our results for teens and the lower panel presents our results for restaurant workers. In the case of teens, the canonical model specification–column 4–shows large and statistically significant negative effects on hires and separations. When we control for spatial heterogeneity–column 8–the magnitude of the hires and separations elasticities are smaller, but they remain substantial and statistically significant. In the case of restaurant workers, the results for each specification are very similar to those for teens.

In contrast to the results for employment stocks, we find strong reductions in both separations and hires. Along with the strong earnings impact, these results contradict the claim in NSW that our spatial controls discard too much variation to find any significant effects. DLR2a also shows that these results–a substantial positive wage effect, small effect on employment stock, and a substantial negative effect on employment flows–can be explained using a job-ladder model with search frictions.[11]

3 Validity of local area controls

In this section we provide additional evidence on the desirability of using local controls. First, we expand on our previous tests for pre-existing trends and show they are present in the canonical model but not in our preferred specifications with local controls. We also

[10]For similar studies on Portugal and Canada, see Portugal and Cardoso (2006) and Brochu and Green (2012), respectively.

[11]The job-ladder model highlights the quits channel with reduced employment-to-employment transitions. A complementary explanation is that a higher minimum wage raises the cost of searching for a better match, and reduces layoffs. This channel is modeled in Brochu and Green (2012), who emphasize the reduced rate of employment-to-nonemployment transitions.

refute the suggestion in NSW that such findings are sensitive to the details of the test. Second, we demonstrate that border county pairs are statistically more similar as compared to randomly-chosen pairs. Third, we provide further evidence on the appropriateness of neighboring area controls using NSW's preferred approach: comparing donor weights picked by the synthetic control method.

3.1 Spatial heterogeneity bias: pre-existing trends

In previous work, we have used unstructured distributed lags to document that the two-way fixed effects model exhibits pre-existing negative employment trends. These spurious leading effects of minimum wages disappear after we account for spatial heterogeneity using the best available research design. Specifically, the pre-existing employment trends vanish when we use the border-county pair research design for QCEW restaurant employment in DLR1, and when we use division-period effects and state-specific linear trends for CPS teen employment in ADR1 and ADR2, and when we use commuting zone-period effects in ADR1. Depending on the dataset and paper, we have used different lag structures to demonstrate this point. Below, we first present these results using a common framework for annual QCEW and CPS data, over the period 1990-2010.[12] Then we respond to NSW's specific arguments that the pre-existing trend evidence against the canonical model is "overstated and misleading."

To estimate minimum wage employment elasticities at the point of and prior to minimum wage increases, we regress the log of state-level teen employment-to-population ratio (in the CPS) or log restaurant employment (in the QCEW) on the log of the contemporaneous minimum wage (MW) and on K leads of log minimum wage. Namely, we estimate the distributed lag model:

$$Y_{it} = \sum_{k=0}^{K} \beta_{-k} MW_{i,t-k} + \mathbf{X}_{it}\Lambda + \gamma_i + \tau_{gt} + \nu_{it} \tag{3}$$

[12] These are the most recent common samples when using three years of leading minimum wages, since the most recent minimum wage information ends in 2013.

for different values of K separately. As before, the canonical and local specifications have different types of time effects τ_{gt}. The sum of the leading coefficients, $\sum_{k=0}^{j} \beta_{-k}$, represents the cumulative response–or the timepath–of the outcome Y at j periods prior to an unit increase in MW. We report these cumulative responses for each $j \leq K$. To avoid any issues with data frequency, we show results using both annual as well as quarterly data. To address any concerns about choice of lag structure, we show how the results are affected by considering between one and three years of leads. For annual data, this means K =1, 2, or 3, while for quarterly data, $K = 4$, 8, or 12.[13]

Before discussing the results, we want to clarify the purpose of these tests. Under the hypothesis that the measured contemporaneous effect is causal, and that the causal effect occurs during the year of, or after the minimum wage increase, we should not expect to find sizeable, negative coefficients on leading values of the minimum wage for one, two, or three years prior to the change. Nor should we expect to find sizeable coefficients for sums of the leading coefficients, as these sums represent the cumulative effect of the minimum wage j periods *prior* to the minimum wage increase. Leading negative effects indicate that the underlying research design fails an important falsification test by failing to appropriately control for low or falling employment prior to minimum wage changes. If the only leading effect occurs in the quarter just prior to the minimum wage increase, it could represent a legitimate anticipation effect. However, finding such leading effects a year, or two years, or three years prior to the policy change,strongly suggests a failure of the falsification test.

We begin by summarizing the findings visually in Figure 4. The top-left panel shows cumulative employment elasticities estimated by the canonical fixed effects model with annual data: CPS state-level teen employment, QCEW county-level restaurant employment, and QCEW county-level restaurant employment with overall private sector employment as an

[13]There is a tradeoff between smoothness and precision in choosing the frequency of leads. Both ADR1 and DLR1 use semi-annual leads while using quarterly data as a way to balance these concerns. NSW assert that our demonstration of pre-existing trends in the canonical specifications was supposedly sensitive to our choice of frequency of leads, which was different from the frequency of the data. Here we use annual and quarterly data, with leads at the same frequency as the data, to demonstrate that this assertion is false.

additional control.[14] For both of these data sources, we show all annual cumulative responses for all lead lengths K =1, 2, or 3, through the contemporaneous effect at time $j = 0$. Unambiguously, the canonical model produces leading effects well before the time of treatment. These spurious effects are particularly pronounced with the CPS data. The top right panel contrasts these estimates with those generated by models with controls for spatial heterogeneity: in the case of the CPS data, these are division-period fixed effects and state-specific linear trends; with the QCEW data, these are estimates from the paired border-county research design. Local controls greatly diminish the leading effects. In fact, all leading effects using local controls are smaller in magnitude than their corresponding estimates with the canonical model. Although noisy, especially with CPS data, the quarterly estimates paint the same picture as the annual ones: with the canonical model, we see a strong presence of leading effects one year or more prior to treatment, in direct contrast to estimates using local controls. For all lead lengths and for both datasets, the cumulative response for the canonical model is clearly in the negative territory prior to the time of treatment, while the cumulative response for the local-controls specification is centered around zero. Using this falsification test, the graphical evidence demonstrates that specifications with local controls outperform the canonical model.

NSW raise three objections to the interpretation of the leading effects in DLR1 and ADR2. First, they argue that in contrast to estimating the 2-quarter or 4-quarter leading minimum wage terms reported by DLR1 and ADR2, estimates using all quarterly leading terms "do not give a clear indication" of pre-existing trends. Although the full quarterly estimates are indeed more noisy, the point estimates are similar to annual ones. These can be checked directly by comparing point estimates for our annual data in our Table 3 with those for quarterly data in Appendix Table A4. Using quarterly data with 12 leads (three years), we estimate a fourth quarter cumulative effect with QCEW data to be about -0.08 with

[14] As before, in the CPS data the dependent variable is log of the teen employment-to-population ratio and demographic controls include race, gender, and teen population shares, as well as the overall unemployment rate. In the QCEW data, the dependent variable is log of restaurant employment, with a control for log of county population (and log of total private sector employment, if specified).

overall private-sector employment control, and -0.13 without. These are virtually identical to the coefficients for one-year leads in the annual data, -0.08 and -0.12. More concisely, this conclusion is obtained from inspecting Figure 4, which shows similar patterns for annual and quarterly data. The pictures NSW present using semiannual and all quarterly leads (their Figures 7 and 8) are similar, and provide evidence of pre-existing trends using the canonical specification with either CPS teen employment or QCEW restaurant employment. Using 12 quarters of leads with QCEW restaurant employment and the canonical model, NSW's own estimates of 12-quarter and 4-quarter leading cumulative effects are about -0.07 and -0.19, the latter significant at the 10 percent level (see their Table 7). We regard these estimates as additional evidence in favor of the existence of spurious effects using the canonical model.

Second, NSW confuse the *size* of cumulative leading effects with the *change over time* in the cumulative leading effects. Evidence that the change is not always significant, as NSW present for the canonical model, is not sufficient to demonstrate that the canonical model passes the falsification test of zero leading effects. Sizeable albeit relatively *constant* cumulative leading effects are present throughout the pre-treatment period, as can be seen in the top left of our Figure 4 with QCEW data. And these leading effects indicate a violation of the parallel-trends assumptions in the canonical model.

Third, NSW object to local control specifications with the CPS that show sizeable positive employment effects three years or more after the minimum wage increase; they claim these are "unlikely" and "could be construed as evidence against using division-period effects and state-specific linear trends." However, and contrary to NSW's approach, considering the sign and magnitude of estimated *post*-treatment effects is a fundamentally unacceptable way of evaluating a research design. Estimates of *pre*-treatment effects are an important specification check precisely because the falsification test is a matter of logic and not of particular economic theories.[15]

[15]NSW also argue that controls for spatial heterogeneity yield biased estimates because using QCEW restaurant data with division-period effects and state-specific linear trends obtains estimates with statistically significant positive leading effects. First, we do not dispute that these leading effects exist with this particular design and data source. In fact, these leading effects are precisely why a better research design is needed–the

3.2 Covariate balancing

In DLR2b we use the Quarterly Workforce Indicators (QWI) dataset to show that adjacent county pairs are more alike in terms of covariates than are non-adjacent county pairs.[16] Here we reproduce those results. To examine whether local controls are indeed more similar, we consider all the six key covariates used in that paper: log of overall private sector employment, log population, employment-to-population ratio (EPOP), log of average private sector earnings, overall turnover rate and teen share of population. We note that none of these covariates is likely to be substantially affected by the treatment status. Therefore, a finding that contiguous counties are more alike in these dimensions cannot be attributed to having more similar minimum wages.

For each of these six covariates, we calculate the mean absolute differences between (1) a county in our border sample and its contiguous cross-state-border pair, and (2) a county in our border sample and every non-contiguous pair outside of the state. For the latter, each of the 972 counties in 966 cross-border pairs is paired with every possible out-of-state county, for a total of 1,737,884 pairings.[17] For each time period, we calculate the absolute differences in levels and changes of these variables between the county and (1) its cross-border pair and (2) its non-contiguous pair, respectively. Subsequently, we collapse the dataset back to the county-pair-period level and calculate the means of the absolute differences in covariates between counties within pairs. The standard errors are calculated allowing for clustering multi-dimensionally on each of the two counties in the cross-border pair.

Table 4 shows the results for these variables in levels, as well as 3 year (36 quarter) changes. In all cases, the mean absolute differences are larger for non-contiguous pairs,

view of DLR1 and our view here is that comparing border counties is a better approach, which is what why we have estimated employment effects for teens using such a design in DLR2 as well as ADR1. Second, finding that division-period effects and state-specific linear trends yield spurious leading effects with county-level QCEW restaurant employment data does not, in itself, invalidate the same research design using state-level CPS teen employment data, since the latter explicitly fails to find significant leading effects.

[16] We had conducted a similar exercise for an early version of DLR1. Huang (2008), which used a contiguous county pair design, also showed that adjacent pairs are more similar than non-adjacent pairs.

[17] Contiguous pairs in DLR2b are restricted to those whose centroids are within 75 miles of each other. Geographically large counties are shown to be less similar to each other in terms of covariates.

and in all cases the gaps are statistically significant at the 1 percent level. The average percentage gap in absolute differences for the twelve variables is 23 percent. Many of the gaps are substantial: notably, for levels of employment and earnings, and 12-quarter changes in the EPOP and the turnover rate. We conclude that cross-border counties do offer an attractive control group that better balances observed covariates–especially as they relate to the state of the labor market. These local controls therefore reduce the scope for bias stemming from omitted confounders.[18]

3.3 Synthetic control donor weights

The core argument of NSW against the local control strategy is that local controls–be they states or counties–are not picked "frequently" when they implement the synthetic control strategy of Abadie et al. (2010).[19] The synthetic control approach picks non-negative weights for donor units such that a weighted average of donors is "close" to the treated unit in the pre-intervention values of the outcome variables and/or other covariates. NSW consider 129 events including state and federal minimum wage increases.

To begin with, NSW gloss over an important conceptual issue. Examining weights within Census divisions and comparing these to weights outside divisions does not tell us whether comparing local areas is better than using state panel regressions with two-way fixed effects–the apparent preferred choice of NSW. Although synthetic control weights may put greater weights on *specific* outside-division states, that is not what a pooled state panel regression does. Conditional on covariates, a state panel regression implicitly treats all donors equally and assumes that, on average, all donors and treated groups share parallel trends. So even if the ideal synthetic control does not put more weight on nearby places, it does not follow that local area controls are inappropriate vis-a-vis the two-way fixed effects strategy. Both

[18]Although not shown, we also calculated these results for 4-quarter changes. The results were similar: all the differences were in the same direction and all were significant at the 1 percent level, except for the 4-quarter difference, which was significant at the 5 percent level.

[19]NSW also use a "ranked prediction error" approach to assess the validity of the controls. This shares much of the problems of their implementation of the synthetic control estimator; moreover, it is an *ad hoc* technique. For this reason we do not separately address it here.

local and synthetic controls may be different unbiased estimators, while the two-way fixed effects estimator remains biased.

Additionally, three reasons make NSW's set-up problematic for assessing the validity of the local controls. First, relatively few states in NSW's donor pool are located within the same division, since to be eligible, donors must not have any minimum wage changes in the four quarters leading up to the policy change. Second, the four-quarter intervention period is problematic because it is quite short, and may lead to poor match quality. Finally, in most of their synthetic control specifications, NSW select econometrically unjustified matching variables, such as residuals from an OLS regression. For these reasons, we devise a more informative test of what the synthetic control approach implies about local area controls. After presenting our results, we return to NSW, clarify the problems with their synthetic control strategy, and present evidence based on their own tables suggesting that local areas are indeed good controls.

3.3.1 Evidence from placebo laws

Since our goal is to assess whether neighboring states are indeed more similar, we do not need to rely on minimum wage events *per se*. We can simply assess whether neighboring donors receive relatively greater weights when constructing a synthetic control for any arbitrary treatment. For this exercise, we randomly assign a placebo law to an individual state in a given time period, and calculate the synthetic control donor weights for all remaining states. Using the pool of random interventions, we assess the relationship between the calculated donor weights and the donor distances to treated states. We consider all possible treatment states and quarterly time periods with a window of four years of pre-intervention data and two years of post-intervention data between 1997q4 and 2007q2, a decade during which the federal minimum wage remained stagnant.[20] By focusing on these placebo treatments rather than actual minimum wage events, the randomization approach dispenses with the

[20]We use the same sample in Section 4 to show the actual impact of minimum wage changes using the synthetic control approach.

shortcomings that plague NSW's efforts: a small number of treatment events, a short pre-intervention window, and the lack of nearby donors.

The outcome of interest is quarterly teen employment-to-population ratios estimated from the CPS between 1997q4 and 2007q2. A pre-intervention window of four years and a post-intervention window of two years limit the intervention dates to the 16 quarters during the 2001q4 through 2005q3 period. We measure distance between states using Census 2000 population-weighted centroids. For the purpose of examining the effects of distance on match quality, we exclude Hawai'i and Alaska as distance outliers; the *minimum* distance from Hawai'i to the continental United States is roughly the *maximum* distance between all continental states. Our sample therefore consists of 48 continental states, for a total of 768 placebo laws.[21]

The synthetic control procedure chooses donor weights W to minimize the distance between pre-treatment characteristics X_1 and X_0 of the treated and donor states. In other words, the procedure minimizes the sum

$$\sum_{k}^{m} v_k (X_{1k} - X_{0k} W)^2$$

over m pre-treatment characteristics, where v_k measures relative importance of the k-th predictor. The choice of exactly which pre-treatment characteristics to select as predictor variables is not obvious, although some combination of pre-intervention outcomes should be included. Intuitively, if the synthetic control fits a sufficiently large set of pre-intervention outcomes, it necessarily fits the unobserved time-varying characteristics that complicate conventional regression analysis (Abadie et al. 2010). One choice for predictors is to include every pre-treatment value of the outcome, as in the Bohn, Lofstrom, and Raphael (2011) study of Arizona's immigration law. In their synthetic control study of minimum wage

[21] NBER CPS Merged Outgoing Rotation Groups data: `http://www.nber.org/data/morg.html`. Centroid data is available through MABLE/Geocorr2K: `http://mcdc2.missouri.edu/websas/geocorr2k.html`. We calculate distance with the `globdist` tool: `http://homepages.rpi.edu/~simonk/stata/index.html`.

effects, Dube and Zipperer (2013) consider several candidate models and achieve best post-intervention fit for donors by using annualized averages of pre-treatment outcomes in addition to pre-treatment averages of demographic, labor market, and industry shares. We select the same predictor variables: the four annualized pre-treatment averages of EPOPs (for the 16 pre-intervention quarters); pre-treatment average values of the shares of teens who are white, black, and female, and the pre-treatment average age of teens; pre-treatment averages of the overall state unemployment rate and unionization rate; and pre-treatment averages of ten industry shares. The algorithm dynamically optimizes over both the set of predictors, V, as well as the set of donor weights W.[22]

We present the key evidence in Figure 3, which nonparametrically plots the average donor weight by distance between the donor and the treatment state using locally weighted regression (`lowess`). This visual evidence unambiguously demonstrates that the synthetic control algorithm assigns much greater weight to nearby states when constructing the counterfactual teen employment. For each intervention there are 47 donor states, yielding an average donor weight of 0.021. A locally weighted regression predicts this weight for donors at a distance of about 1,100 miles, suggesting that donors exceeding this distance are "worse-than-average" controls. For example, donors that are 2,000 miles away have weights of 0.013, while donors only 100 miles away are predicted to obtain a weight of 0.089. Donors that are 100 miles away are thus, on average, assigned a weight that is 7 times larger than a donor that is 2,000 miles away. As indicated by the 95 percent bootstrapped confidence bands in the figure, these differences are highly statistically significant.

Table 5 summarizes the geographic characteristics of donor weights estimated by this procedure using the four Census regions, nine Census divisions and alternative distance bands. For a given treatment, some of the 47 donors lie within that treated state's Census region. The mean weight assigned to within-region donors is 0.044, while the mean weight assigned to donors outside the region is 0.014. The ratio, or relative weight, 3.150, indicates

[22]We implement the synthetic control approach in Stata using the `synth` package with `nested` optimization and `allopt` starting point checks for robustness: `http://www.mit.edu/~jhainm/synthpage.html`.

that the synthetic control procedure strongly prefers donors within the treatment region to those outside the region. The same pattern holds for Census division categories: within-division weights are roughly four times as large as the weights outside the region. Finally, for donors whose population centroids are less than 500 miles and less than 1,000 miles from the treated state, inside-donors receive weights three to four times as large as outside-donors.

3.3.2 Reinterpreting the evidence from Neumark, Salas and Wascher

NSW argue that nearby donors are *not* assigned greater synthetic control weights, in direct contrast to the results we provided above. However, underlying their synthetic control implementation are several dubious features, which we describe below. Furthermore, our analysis of their own evidence also suggests that nearby donors do receive greater weights.

Although NSW are testing whether donors from the same division receiver greater synthetic control weight, in some cases there were no donors within the same division, as they note, forcing them to eliminate more than half of their available treatment events. Problems like these are avoided by using a more general placebo-law-based strategy, such as the one we analyzed above.

Second, using NSW's sample selection, only four quarters of pre-treatment data are available to form predictor variables. One year of data may be an insufficient amount of information for constructing quality synthetic controls. Existing synthetic control studies use many years of the pre-intervention sample to find a balanced sample.[23] Even though NSW select a relatively short four-quarter pre-treatment period for predictor variables, NSW do not discuss how well their synthetic controls actually fit the treated unit before treatment. The synthetic control methodology in principle compels the researcher to assess the validity of the synthetic counterfactual in the pre-treatment period. Yet NSW provide no such

[23]For example, in their study of California's Tobacco Control legislation, Abadie et al. (2010) use nineteen years of pre-intervention data. Another recent paper (Bohn, Lofstrom and Raphael 2011) uses nine years of pre-intervention data to study the impact of Arizona's employer sanction legislation. The only paper in the minimum wage literature to use synthetic controls–Burkhauser, Hansen and Sabia (2011), uses four years of pre-intervention data.

evidence.

NSW separately estimate synthetic controls with four sets of predictor variables: the log of the teen employment-to-population ratio; one-quarter changes in that variable; four-quarter changes in that variable; and residuals from the canonical two-way fixed effects regression of teen employment on the minimum wage. Only the first set of predictors–pre-treatment levels of teen employment–is valid. Using one-quarter or four-quarter changes as predictor variables is problematic because, as a consequence, employment levels are considered from as many as eight quarters prior to the intervention, which include periods with minimum wage changes. This complication violates a basic requirement for the synthetic control approach: weights should be calculated by matching on outcomes unaffected by a treatment.

There is absolutely no econometric justification for using residuals from an OLS panel regression as the matching variable in a synthetic control study. NSW argue (p. 21) that

> "...matching on the residuals is informative about the spatial heterogeneity arguments that ADR put forward, as their contention is that the residuals for states in the same Census division share common features that are correlated with minimum wage changes. Consequently, matching on the residuals provides information on whether the residuals for states in the same region share these commonalities and hence whether these states are good controls"

Besides being econometrically unjustifiable, this approach is also intuitively wrong. First, there seems to be a confusion between estimated and true residuals. By construction, estimated OLS residuals are uncorrelated with all regressors, including the minimum wage. If the canonical model is incorrect because its *true* residuals are correlated with minimum wage changes, then the canonical model's *estimates* of the residuals are, by definition, biased and are therefore not appropriate for constructing accurate counterfactuals. By contrast, if the canonical model is correct, then its estimates of the residuals are uninformative as predictor variables because the true (and estimated) residuals are mean-zero errors that are uncorrelated with the minimum wage. In both cases, therefore, matching on OLS residuals is wrong. We return to this matter in section 5 when discussing NSW's preferred matching estimator, which also uses these synthetic control weights based on OLS residuals and is similarly erroneous.

Using this problematic set-up, NSW report in their Table 4 the proportion of weights on donor states in the same division as treated states. In their Table 6, they also report results from an analogous exercise with cross-border and other counties, reporting the proportion of weights assigned to cross-border donors. These fractions of weights are, however, largely uninformative: it is not surprising that some good controls may be found in other parts of the country; and therefore, the proportion of weights going to local areas may be small even if they are (on average) much more similar than states that are further away. This is especially the case because so few states in NSW's donor pool are within the same division. What is more informative is the average per-donor weight of same-division donors, relative to the per-donor weight of other-division donors, which is what we calculated for our Table 5.

Using calculations not shown in their paper, NSW state that the "average weight per same-division donor states" does not differ substantially from the weight for other states.[24] However, one can calculate average per-donor weights for same-division and different-division donors using the information in their Tables 4 and 6. We report these calculations in Table 6, which very closely match the results from our placebo exercise.[25] A donor state within the same division receives weights that are on average *three times as large* as weights for donors outside of the division, similar to our placebo law exercise in the section 3.3.1 above.

[24]See their pages 21-22. For their specification with residuals, which we argued is problematic, the average weights for same- and different-division average weights are stated to be about the same. When NSW average this quantity over the three specifications using employment levels and differences, the average weight is about 34 percent higher for same-division donors.

[25]In our Table 6 we report calculations for two of four of NSW's specifications. We do not report results for their two specifications that involve matching on lags because NSW do not report the exact number of same-division or cross-border donors for these specifications (see the notes to their Tables 4 and 6). We explain also that these two specifications violate synthetic control requirements. We also ignore their county estimates based on a RMSPE limiting criterion because it is ad hoc and inappropriate for this weight-distance exercise. For a given treated county, NSW calculate synthetic controls for all potential donor counties, identify the 50 donors with the lowest pre-treatment RMSPE, and then restrict the treated counties donors to this subset of 50 counties. Limiting donors in this way necessarily worsens the pre-treatment fit of the synthetic control.

The same is true for cross-border counties, which receive substantially higher weights than do other counties. A straightforward interpretation of the evidence in NSW also shows that neighboring areas are much more alike than are places farther away, in contradiction to the central thesis of that paper.

4 New findings using alternative estimators

4.1 Synthetic control minimum wage elasticities

What does a properly implemented synthetic control approach suggest regarding minimum wage effects? By pooling multiple state minimum wage treatment events, Dube and Zipperer (2013) estimate average elasticities for teen wages and employment. We reproduce and summarize this analysis below.[26]

Since the synthetic control method depends on isolated treatment events with well-defined pre- and post-treatment periods, the approach ultimately discards much of the minimum wage variation available to conventional regression techniques. For example, use of federal minimum wage events is highly problematic since very few donor states are available to construct a reasonable synthetic control group. On the other hand, by matching on pre-treatment characteristics we may form better counterfactuals: unlike the canonical two-way fixed effects model, the synthetic control approach may provide unbiased estimates even when unobserved cofounders vary with time. Of the 89 state minimum wage changes from 1997q4 through 2007q2, there are 19 usable treatment events with pre-treatment periods of between

[26]To our knowledge, the only other paper presenting synthetic control estimates of minimum wage treatment effects is Sabia, Burkhauser, and Hansen (2013). The authors unfortunately focus only on the minimum wage change in New York state beginning in 2005. They therebyignore four other candidate treatment events also beginning that year in Florida, Minnesota, New Jersey, and Wisconsin as well as omitting other events during the 1997-2007 time period of the static federal minimum wage. The authors inexplicably fail to use any pre-treatment outcomes as predictors for their synthetic control, even though the unbiasedness of the synthetic control estimator relies specifically on pre-treatment outcome balance between the treated unit and the weighted combination of donors (see Abadie et al. 2010). As a result, the authors obtain an invalid synthetic control: their Figure 3 clearly shows that employment paths for actual and synthetic New York *never* match during the pre-treatment period.

two and four years. Table 6 lists these events, along with their pre- and post-treatment period windows, the number of potential donors, and the percent change in the minimum wage. Synthetic control weights are selected using the same set of predictor variables as in the placebo exercise in Section 3: the four annualized pre-treatment averages of EPOPs (for up to 16 pre-intervention quarters); pre-treatment average values of the shares of teens who are white, black, and female, and the pre-treatment average age of teens; pre-treatment averages of the overall state unemployment rate and unionization rate; and pre-treatment averages of ten industry shares. The algorithm dynamically optimizes over both the set of predictors, V, as well as the set of donor weights W.[27]

For each event, Table 7 shows the minimum wage elasticities for teen employment and wages.[28] Additionally, we report in parentheses the percentile ranks of the treatment effects against placebo (donor) outcomes, which is the inferential procedure suggested by ADH. When the rank is either less than 0.05 or greater than 0.95, the effect is deemed to be statistically significant at the 10 percent level. As an example, for New Jersey's 2005q4 minimum wage change, the estimated teen employment elasticity is -0.004. The effect, although negative, is insignificant with a rank of 0.565, indicating that New Jersey's effect falls around the middle of the placebo distribution. New Jersey's wage elasticity of +0.296, however, is significantly positive with a rank of 0.957. For most individual events, neither employment nor wage effects are individually significant, indicating the low power of using individual events with estimated teen EPOPs. We therefore pool across these 19 events to show average and median minimum wage impacts. We assess statistical significance of the pooled estimates by comparing the mean and median ranks, which are distributed as known

[27]Specifically, Dube and Zipperer (2013) identify all state-level minimum wage changes during 1997q4-2007q2 with 8, 12, and 16 quarters of pre-treatment data, and 4, 6, and 8 quarters of post-treatment data, and then select the maximal window configuration for each treatment event. The minimum wage change is defined as the percent change between its pre-period level and its maximum post-period level, since a treated state may alter its minimum wage more than once in the post-period.

[28]We estimate the elasticities as the percentage difference in the average post-treatment outcomes between the treated state and its synthetic control, divided by the percentage change in the minimum wage in the treated unit (the control units saw no change in the minimum wage by definition). See Dube and Zipperer (2013) for details.

functions of uniform random variables.[29] As Table 7 shows, the pooled effect for teen wages is statistically significant and sizeable for both the mean (0.17) and the median (0.15) effect. The mean teen employment elasticity, reported at the bottom of the table, is about -0.03, while the median is 0.08, and the pooled effect is not statistically significant when judged by the mean or the median rank across individual events.

These average employment and wage elasticity estimates lie within the range reported for teens using controls for spatial heterogeneity, as shown in specifications 5, 6 and 8 in Panel A of our Table 2. Overall, the synthetic-control based pooled effects confirm the general findings of local control-based approaches: clear positive impact on teen wages, along with a small or no disemployment effect that is statistically indistinguishable from zero.

4.2 Lagged dependent variables and dynamic panel data models

An alternative approach for allowing time-varying heterogeneity employs lags of the dependent variable in a panel regression. The findings in section 3.1 and in our previous work of unusually low employment prior to minimum wage increases point us in this direction. In this section we estimate panel models with a series of lagged dependent variables (LDVs) with lags between one and three years to account for trends in employment. If minimum wage increases occur during periods with unusually low and/or falling employment, then the LDV approach offers a way to account for this heterogeneity (as an alternative to using geographical controls or state trends). It is worth noting that the synthetic control approach of matching pre-intervention outcomes constitutes a generalization of the LDV approach. While less general than synthetic controls, the LDV approach has the virtue of using much more of the variation in the policy. In contrast, a proper application of the synthetic control method relies on using discrete treatments and clean pre-intervention windows, seriously limiting the amount of usable variation.

It is well known that the inclusion of both fixed effects and lagged dependent variables

[29]Namely, we compare these across-event mean (or median) ranks to the 90% extrema of the mean (or median) of 19 simulated uniform random variables.

renders the estimates inconsistent (e.g., Angrist and Pischke 2009, pp. 182-4). Dynamic panel data models can accommodate both fixed effects using older lags of the outcome in levels and differences as instruments, although identification hinges upon the assumption that the idiosyncratic component of the disturbances is not too autocorrelated. Here, we show results from (1) OLS regressions with lagged dependent variables without state fixed effects and (2) a dynamic panel data (DPD) model with lagged outcomes and state fixed effects using the Arellano-Bover/Blundell-Bond system-GMM approach, which uses all available lags as instruments.[30] The fixed effects are purged using forward-orthogonalization and we treat all the variables other than the lagged outcomes as exogenous. For all of these specifications, we use annual data for ease of use with the dynamic panel data model; with quarterly data one would have to specify a large number of lags to span three years time, which is problematic for estimation.

Table 8 reports the results for both log average teen wages and teen EPOP, using state-level annual averages of CPS data. The first column reproduces the canonical two-way fixed effects results (excluding state trends and division-time effects) with state aggregated annual data, along with controls for the overall unemployment rate, race and gender composition and teen population shares. Similar to our previous findings, we obtain a wage elasticity of 0.104. For employment, the baseline elasticity in column 1 is -0.135, similar to results from comparable specifications elsewhere.

Columns 2 to 4 in Table 8 display our results with different lagged dependent variable specifications, excluding state fixed effects. The wage estimates are all positive and statistically significant in all the LDV specifications. The inclusion of a single lag of teen EPOP reduces the magnitude of the employment effect to -0.025, and renders it statistically indistinguishable from zero. This lack of statistical significance is notable because the standard errors are actually smaller in the LDV specification. In columns 3 and 4 of Table 8, we find that incuding additional lags in the dependent variable continues to produce elasticities very

[30]This calculation is implemented in Stata using the `xtabond2` command.

close to zero (-0.014 and -0.004) while the standard errors remain small. Turning next to the DPD models (columns 5-7) that include state fixed effects, we once again find positive wage effects across all models, although they tend to be slightly larger in magnitude. Importantly, the employment elasticities are quite similar to the corresponding LDV estimates and they are all close to zero. When we include one, two and three lags, the elasticities are -0.033, -0.016 and 0.001, respectively; the associated confidence sets rule out sizeable negative effects. These results provide still another direct demonstration of how time-varying heterogeneity affects the estimation of minimum wage effects.

5 The NSW Matching Estimator

In the second version of their paper, NSW introduce a matching estimator as a "data driven" way to construct comparison groups. NSW extend their synthetic control analysis through a paired treatment-counterfactual analysis: each treated state is compared to a weighted average of control states (see their section IV). They ground their methodology in the existing literature by citing Autor et al. (2006), but the NSW approach is *ad hoc* and is fundamentally erroneous. While there are many problems with this estimator, we highlight here four key issues. (1) NSW use residuals from an OLS regression as the matching variable to construct synthetic control donor weights–a procedure without any econometric basis. (2) Their preferred estimate (the one finding a negative impact) further uses a contaminated sample, where the minimum wage is rising in some of the post-intervention control states as well as in the pre-intervention periods. This further invalidates their application of the synthetic control method. (3) NSW use a very short pre-intervention window (4 quarters) to calculate synthetic control matches, which makes finding a good match difficult. (4) They use only a single quarter of post-intervention data to measure its impact–making this the shortest-run estimate in the minimum wage literature of which we are aware. In this section we provide a more detailed description of their method and its problems.

NSW attempt to construct an average estimate of minimum wage increases by pooling case studies of minimum wage treatments. They begin with what they call their "clean" sample: the treated sample consists of 129 minimum wage increases with no other increase in the four quarters prior to treatment. The authors then stack the data into treated-counterfactual pairs. For each treatment event, they construct a counterfactual state as a weighted average of control states using synthetic control weights on this donor pool, where the synthetic control matches on OLS residuals from the canonical fixed effects regression. We note again that this method is without any merit–econometric or otherwise.[31] Given this counterfactual, NSW create a pair consisting of five treated state observations and five counterfactual observations. The five observations consist of four quarters of pre-treatment period data, and a single quarter of data during the actual treatment. They include a dummy for each event, a dummy for each event-specific treatment state, and time dummies. They regress log teen employment on log minimum wage along with these controls, using the treatment state's teen population as regression weights (for both treated and counterfactual observations).[32]

Using their "clean" sample, NSW find no minimum wage effect using either (1) equal weights for all control states, or (2) the NSW matching estimator. This result leads the authors to believe this sample of treatments is "unusual," as these estimators fail to generate a large negative employment elasticity.[33] It is worth reflecting on this finding: almost the entirety of NSW's critique of local controls uses a sample and a methodology–however flawed it may be–that does not actually show any disemployment effects. It is also noteworthy that

[31] At a basic level, the estimated OLS residuals are either incorrect or uninformative. If the canonical panel regression is flawed, as we argue, then the estimated residuals are formed incorrectly and are not unbiased estimates of the true residuals. In contrast, if the canonical panel model is appropriate, as NSW argue, the the estimated (and true) residuals are simply mean zero error term, which are by definition orthogonal to minimum wages; it does not makes sense to find controls to match such residuals, as they precisely do not include confounds. Yet even if the residuals were informative in some manner, they are not appropriate predictor variables for the synthetic control procedure, simply because the unbiasedness of the synthetic control estimator relies on the assumption that the predictor variables include pre-treatment values of employment.

[32] Event-quarter dummies would help control for event-time-specific shocks, but NSW omit these.

[33] Their point estimate for the teen employment elasticity is -0.035 with all controls equally weighted, but they do not report same-division estimates.

the authors call their self-designated "clean" sample as being too "unusual" for estimating minimum wage impact, without explaining (1) why it remains valid for criticizing the plausibility of neighbors as control groups, and (2) why it is acceptable to select estimators based on the results they produce.

Next, NSW expand their sample to include events that do not conform to clear treatment-control categories. With the expanded, unclean sample, potential control states now include those with minimum wage increases, and treated states may now receive minimum wage changes during their pre-treatment period. The expanded sample is fundamentally at odds with the synthetic control framework, which requires both clear periods of pre- and post-treatment for the treated states, as well as potential donors that receive no treatment whatsoever. Using this fundamentally problematic sample, the authors obtain large, negative teen employment elasticities in the range of -0.2 to -0.3, but not for restaurant employment using county-level QCEW data, which were all smaller than -0.05 in magnitude.[34]

While we agree with NSW that synthetic control case studies offer another promising avenue to examine minimum wage effects, NSW's particular approach cannot be expected to generate plausible causal estimates.[35] Credible estimates would begin with using the actual synthetic control estimator, and not graft different methods onto each other without justification. Most critically, such an approach would use the actual set of pre-treatment outcomes–which forms the very basis of a synthetic control strategy–and not residuals from another estimation. A proper application would also limit minimum wage treatment events to those with longer pre- and post-treatment periods. The sample of events would contain well-defined pre- and post-treatment periods, and clearly defined treatment and control units. This is exactly what we do in our application of the synthetic control approach, as described

[34]Incidentally, NSW also fail to assess the quality of their synthetic controls in the pre-treatment period, even though such a demonstration would provide a major reason for the credibility of synthetic control estimates in the first place.

[35]Perhaps least important in light of the problems above, their standard errors are certainly wrong (probably too small). They cluster on treated state and counterfactual, but the counterfactuals share many of the same states, so they are often correlated mechanically. That is why the synthetic control approach in Abadie et al. (2010) uses randomization inference.

in Section 4.1; the results confirm our previous findings.

Besides this matching estimator, NSW also argue that estimates using pairs of neighboring states suggest large disemployment effects for teens. While in principle this is a useful approach–and is similar to our specification with division-time controls–we explain in the Appendix (and Appendix Figure A1) that this specification exhibits very serious negative pre-trends in this sample, making the results from this estimator unreliable.

6 Discussion

Since 1990, a large number of policy changes have shaped minimum wage levels in the U.S. The resulting richness of the data constitute both good and bad news. The good news is that the considerable amount of policy variation permits us to use sophisticated research designs using credible identification strategies. The bad news is that it is easy to generate misleading results if one is not careful about the source of identifying variation. A casual inspection of a map of minimum wages across the U. S. shows a high degree of spatial clustering in minimum wage policies since 1990. This clustering coexists with a large array of potential confounds that vary between high and low minimum wage states. Even *a priori*, the chances would be slim that all of these factors happen to balance out on net. In practice, we show that observable confounds vary substantially across high and low minimum wage states, suggesting that unobserved factors do as well. In other words, accounting for spatial heterogeneity is of first order importance in this literature.

There are only two acceptable reasons to avoid controlling for this heterogeneity. (1) The inclusion of the controls substantially reduces statistical power. (2) The treatment affects the control variables themselves, such as through spillover effects on neighboring areas. Both of these are testable propositions. Researchers who advocate throwing out plausible and relevant controls needs to demonstrate and not merely assert these propositions. In our work, we have not found either of these factors to be critical.

If including regional controls (such as those in Autor 2003, or DLR1, or ADR1) substantially affects the results from a two-way fixed effects model, researchers should worry quite a bit about their identification strategy. Pre-existing trends in the form of sizeable leading coefficients for the minimum wage should make researchers nervous about that particular research design. Similarly, if the results are found to be sensitive to parametric controls such state-specific linear or quadratic trends, the researcher should consider the probable explanation that treated and control areas have different counterfactual outcomes. The same consideration applies to finding very different results when allowing for lagged outcomes instead of time-invariant place fixed effects. These are caution flags suggesting that the conditional independence assumption of the research design is likely being violated.

One particularly powerful way to account for unobserved time-varying heterogeneity uses a border discontinuity design. This approach provides a very effective tool to account for unobserved confounders. As we show in this paper, neighboring counties are more similar in observed dimensions, in both levels and trends. To the extent that the policy is endogenous, the policy is unlikely to respond to specific factors affecting border areas as opposed to the state as a whole, again reducing the possibility of reverse causality. For these reasons, border discontinuity designs are increasingly popular with applied economists. They are not, of course, infallible. In particular, one should consider the possibility of policy interference on control units through spillovers. We did exactly this in DLR1 and found that, at least for county-level analysis of minimum wage policies, spillovers were not a major concern.

Although the border discontinuity approach represents one way of constructing credible control groups, there may be other good ones as well. In particular, it is possible to find good controls that are not geographically close. The synthetic control approach, for example, uses past outcomes and other covariates to locate similar control groups. Although we show that this algorithm is generally much more likely to pick neighboring areas as controls, there is no reason to expect that all good controls are local. All happy families may be alike–according to Tolstoy–but that principle certainly does not apply to all unbiased estimators.

Synthetic controls promise to become an important part of the applied microeconomist's toolkit. But the approach too has important limitations, especially for minimum wage policies. In particular, the recurring and continuous nature of minimum wage increases makes it difficult to find many usable events when we require a minimally acceptable length for the pre-intervention window. And it is virtually impossible to use federal minimum wage increases, as these leave only a tiny number of (unaffected) donors to use as controls. When we consider all cases with at least two years of pre-intervention data between 1998 and 2007, the nineteen resulting events produced a small minimum wage elasticity for teen employment, consistent with our findings using border discontinuity designs, coarser spatial-control approaches, and specifications controlling for lagged dependent variables.

At this point, it is instructive to compare which types of research design produce which kind of results. Here is the full list of specifications that Neumark, Salas and Wascher (2013) argue show sizeable disemployment effects: (1) the canonical two-way fixed effects model, (2) a two-way fixed effects model using third or higher order polynomial trends by state (but not first or second order polynomials), and (3) a two-way fixed effects model using Hodrick-Prescott pre-filtered data, and (4) a two-way fixed effects model with data detrended using an out-of sample fitted trend and (5) an *ad hoc* matching estimator that constructs a comparison group using a contaminated sample by using synthetic control weights based on residuals from a panel regression, and (6) an estimator comparing pairs of neigboring states (but not neighboring counties). We established beyond reasonable doubt the problems with (1). Each of the other members of the above list (except for 6) is an unusual specification, sometimes without a clear econometric foundation, that has been used seldomly–if ever–in the discipline. This is especially the case for their preferred matching estimator. While their state-pair estimator (6) has more *a priori* justification, we show that it fails to pass the falsification test of no pre-existing trends in the actual sample.

For comparison, the following comprises the full list of specifications that we argue usually show small dis-employment effects for teens and restaurant workers, where by small we mean an employment elasticity under -0.1. (1) A border discontinuity design that uses variation within local labor markets (commuting zones), (2) a border discontinuity design using pairs of adjacent counties, (3) a model with region-specific time effects and state-specific linear trends, (4) models with lagged dependent variables without state fixed effects, (5) the Arellano-Bond/Blundell-Bover dynamic panel data estimator, and (6) a proper synthetic control design. Every one of these six approaches–or something very close–have been used in the discipline both prior to and subsequent to our original work, and are considered useful parts of our modern toolkit. When applied, they all suggest that job losses from the level of minimum wage changes in the U.S. since 1990 have not been substantial. Finally, these results have been now replicated using five different datasets and validated externally using newer data.

Of course, as is the case for most policies, it is a mistake to think of a single, true minimum wage impact. The effect of minimum wages almost surely varies by the level of the statutory minimum. Similarly, there are good theoretical reasons to expect that its effects will differ in different types of labor markets–for example, with the extent of search frictions or labor market tightness. In our work, we have typically studied the average employment effects of minimum wages on two frequently studied groups–teens and restaurant workers–that are highly affected by the policy. We have found that the employment effects are small in magnitude for the range of increases that have been implemented since 1990. And often, statistical bounds rule out substantial job losses. Our results do not imply that minimum wages never reduce employment, or that no one is ever hurt by a minimum wage increase. But our finding that moderate increases in the statutory minimum lead to substantially higher average wages for these groups, and without creating substantial attendant job losses, provides useful information for policymakers when deciding on appropriate wage standards.

References

[1] Abadie, Alberto, Alexis Diamond and Hainemuller 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Associaton* 105, 490: 493-505.

[2] Acemoglu, Daron and David Autor 2010. "Skills, Tasks, and Technologies: Implications for Employment and Earnings." Pp.1043-1171 in Orley Ashenfelter and David Card eds. *Handbook of Labor Economics*, Volume 4, Part B. Elsevier Press.

[3] Addison, John T., McKinley L. Blackburn, and Chad Cotti 2011. "Minimum Wage Increases Under Straightened Circumstances." *IZA Discussion Paper* 6036, October.

[4] Allegretto, Sylvia, Arindrajit Dube and Michael Reich 2009."Spatial Heterogeneity and Minimum Wages: Employment Estimates for Teens Using Cross-State Commuting Zones." Working Paper 181-09 Institute for Research on Labor and Employment, UC Berkeley.

[5] ___________2011. "Do Minimum Wages Really Reduce Teen Employment? Accounting for Heterogeneity and Selectivity in State Panel Data." *Industrial Relations* 50, 2: 205-40.

[6] Angrist, Joshua and Jorn-Steffen Pischke 2009. *Mostly Harmless Econometrics.* Princeton, NJ: Princeton University Press.

[7] Autor, David 2003. "Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing." *Journal of Labor Economics* 21, 1: 1-42.

[8] Autor, David H., John J. Donohue III, and Stewart J. Schwab. 2006. “The Costs of Wrongful-Discharge Laws.” *Review of Economics and Statistics* 88, 2: 211 -231.

[9] ___________and David Dorn 2013. “The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market.” *American Economic Review*, forthcoming.

[10] __________, Frank Levy and Richard Murnane 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration.” *Quarterly Journal of Economics* 118, 4: 1279 - 1333.

[11] Black, Sandra 1999. ”Do Better Schools Matter? Parental Valuation of Elementary Education.” *Quarterly Journal of Economics* 114, 2: 577-99.

[12] Bohn, Sarah, Magnus Lofstrom and Steven Raphael 2011. "Did the 2007 Legal Arizona Workers Act Reduce the State’s Unauthorized Immigrant Population?" IZA Discussion Paper 5682.

[13] Brochu, Pierre, and David Green 2012. “The Impact of Minimum Wages on Quit, Layoff and Hiring Rates.” IFS Working Paper 06/11. London: Institute for Fiscal Studies.

[14] Burdett, Kenneth and Dale Mortensen 1998. "Wage Differentials, Employer Size and Unemployment." *International Economic Review* 39: 257-73.

[15] 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- 98.

[16] ___________1995. *Myth and Measurement.* Princeton NJ: Princeton University Press.

[17] ___________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.

[18] Dube, Arindrajit, T. William Lester and Michael Reich 2010. “Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties.” *Review of Economics and Statistics* 92, 4: 945-64.

[19] __________2012. ”Minimum Wage Shocks, Employment Flows and Labor Market Frictions.” Working Paper 122-12. Institute for Research on Labor and Employment. UC Berkeley.

[20] ___________2013. ”Minimum Wage Shocks, Employment Flows and Labor Market Frictions.” *Unpublished paper*. University of Massachusetts Amherst.

[21] Dube, Arindrajit and Ben Zipperer 2013. “Pooled Synthetic Control Estimates for Recurring Treatments: An Application to Minimum Wage Case Studies.” *Unpublished paper*. University of Massachusetts Amherst.

[22] Goldstein, Markus and Christopher Udry 2008. “The Profits of Power: Land Rights and Agricultural Investment in Ghana.” *Journal of Political Economy* 11, 6: 981-1022.

[23] Heckman, James, Hidehiko Ichimura and Petra Todd 1998. “Matching as an Econometric Evaluation Estimator.” *Review of Economic Studies* 65: 261-94.

[24] Holmes, Thomas 1998. “The Effects of State Policies on the Location of Manufacturing: Evidence from State Borders.” *Journal of Political Economy* 106, 4: 667-705.

[25] 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.

[26] Katz, Lawrence and Kevin Murphy 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107, 1: 35-78.

[27] Lee, David and Thomas Lemieux 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature* 48, 2: 281-355.

[28] Magruder, Jeremy 2013. "Can Minimum Wages Cause a Big Push? Evidence from Indonesia." *Journal of Development Economics* 100, 1: 48-62.

[29] Manning, Alan 2003. *Monopsony in Motion*. Princeton NJ: Princeton University Press.

[30] Neumark, David and William Wascher 1992. "Employment Effects of Minimum and Subminimum Wages: Panel Data on State Minimum Wage Laws." *Industrial and Labor Relations Review* 45. 1: 55-81.

[31] ________ 2011. "Does a Higher Minimum Wage Enhance the Effectiveness of the Earned Income Tax Credit?" *Industrial and Labor Relations Review* 64, 4: 712-46.

[32] Neumark, David, J. M. Ian Salas and William Wascher 2013. "Revisiting the Minimum Wage and Employment Debate: Throwing out the Baby with the Bathwater?" NBER Working Paper 18681.

[33] Reich, Michael 2010. "Minimum Wages: Politics, Economics and Econometrics." Pp. 353-74 in Clair Brown, Barry Eichengreen and Michael Reich eds. *Labor in the Era of Globalization*. New York: Cambridge University Press.

[34] Sabia, Joseph, Richard Burkhauser and Benjamin Hansen 2012. "Are Minimum Wage Effects Always Small? Evidence from a Case Study of New York State." *Industrial and Labor Relations Review* 65, 2: 433-44.

Appendix: Response to other claims in Neumark, Salas and Wascher

In this appendix we respond to some of NSW's other claims. First, we address their concerns about the inclusion of recessionary periods in ADR2 through methods far more conventional than the techniques they use. When we do so, we uphold our conclusions about small effects on employment. Second, we show that their criticism of the placebo exercise in DLR1 is misplaced. Third, we discuss their results from the state-pair estimates.

Trends and cycles: Including recessionary periods

NSW argue that controlling for state trends may be problematic if the sample includes economic downturns. This claim constitutes a criticism of ADR2. NSW argue that the findings in ADR2 are driven by the inclusion of recessionary periods, leading to biased estimates of state-specific time trends, and thereby a biased estimate of the minimum wage effect. Their argument implies that recessions can have heterogeneous effects on states and hence contaminate minimum wage estimates. This argument is actually consistent with our key proposition: that there are spatial heterogeneities that may be correlated with minimum wage policies. But while this point has some merit, NSW's proposed solutions are problematic.

We argue here that directly accounting for state-specific business cycle effects in panel regressions is a more fruitful approach. To fix ideas, consider the following data-generating process:

$$y_{st} = \beta_0 + \beta_1 MW_{st} + \sum_s \delta_s \cdot I_s \cdot t + \gamma_s + \tau_t + \rho_{srt} + \epsilon_{st}. \tag{4}$$

Here, ρ_{srt} is a disturbance term for a given recessionary period. The outcome y_{st} is a function of the minimum wage (*MW*) along with state and time fixed effects (γ_s and τ_t,respectively), along with state-specific time trends (δ_s). The mean zero error term ϵ_{st} is uncorrelated with covariates. But additionally, there is a state and time-specific recession effect due to ρ_{srt}. NSW claim that the omission of a control for ρ_{srt} imparts a bias to the vector of estimated state trends $\delta_s \cdot I_s$, and that this bias leads further to a spuriously low estimate of β_1, the minimum wage effect. Note that by definition, the recessionary disturbance has to be state-specific for its ommission to bias other coefficients; otherwise it would be soaked up in the common time dummies.

To overcome this problem, NSW propose a highly unconventional strategy. They first estimate equation (1) using a subsample 1994-2007. (As we argue below, this sample does not correspond to any known defintion of "non-recessionary" periods.) They generate a predicted vector $\hat{\delta}_s$ and detrend the outcome $\tilde{y}_{st} = y_{st} - \sum_s \hat{\delta}_s \cdot I_s \cdot t$, and then regress

$$\tilde{y}_{st} = \beta_0 + \beta_1 MW_{st} + \gamma_s + \tau_t + \nu_{st}.$$

This strategy, however, does not account for the omitted variable ρ_{srt}, the recessionary disturbance. If its omission is a problem for estimating state trends, which in turn might bias the estimate of β_1, it very well might also bias the MW coefficient in its own right; that is, if $Cov(\rho_{srt}, MW_{st}) \neq 0$, then $Cov(\nu_{st}, MW_{st}) \neq 0$. This is true whether one uses y_{st} or $\tilde{y}_{st}$ as the dependent variable. Indeed, one could imagine that the omitted variables bias is exacerbated when using the detrended $\tilde{y}_{st}$: in the recessionary years, the residuals from using the NSW trend control may be more off than in the original regression. Any covariance between MW_{st} and $\tilde{y}_{st}$ during the recessionary sample may therefore be particularly spurious. In other words, the "solution" to the possible problem proposed by NSW is difficult to justify based on a reasonable data-generating process, and can worsen the very problem they seek to address.

Although the nature of the problem may seem complicated, the solution is exceedingly simple. It is very easy to add a control for ρ_{rst}. One particularly simple solution estimates regression (1) but excludes the recessionary years. Another simple solution includes dummies for state-specific recession spells. It is therefore very surprising that NSW do not report any such results. Instead, first they reproduce the same model as in NW (2011)–whose 2007 working paper version ADR cited and reproduced–but using the 1994-2007 sample. This sample period does not correspond to any known "non-recessionary" definition—there was a recession in 2001, and the recession of 1990 ended in 1991. In contrast, we use the NBER definition of recessions and estimate the model to account for recessions in conventional ways. As Table A1 shows, for our preferred specification with state-specific trends and division-time controls, these transparent methods of accounting for recession heterogeneity has no impact on our findings. In all cases, inclusion of spatial controls renders the minimum wage employment elasticity small and indistinguishable from zero. In other words, the key findings from ADR2 hold when we use standard methods to account for the problem raised by NSW.

NSW also detrend the data using several other methods, such as using a Hodrick-Prescott filter, and then run a standard regression on the detrended outcome. We do not know of any applied microeconomics papers that use an HP filter to pre-detrend panel data, and for good reasons. The basic problem with its use here is conceptual. By estimating a moving average of a series, the HP filter decomposes the variation into a "stochastic trend" and a "cyclical" component. NSW's motivation is to control for *both*, and in particular to prevent the "cycle" from contaminating the estimation of the trend. But what they actually do is to use only the "cyclical" component of the variation for estimation. Given NSW's motivation that cyclical downturns may be problematic to include, it is odd to hone in on the variation that the HP filter characterizes as business cycle variation.[36]

NSW show that while taking out a linear or quadratic trend in the 1990-2011q2 sample produces small and statistically insignificant employment effects, adding third, fourth or fifth order polynomials produce more negative estimates. The main conclusion from their exercise should be that our results are robust to including a quadratic trend, and not just a linear one. Beyond that point, fitting polynomials of higher orders primarily leaves in high

[36]Taking out a stochastic trend removes much more of the variation than would a long term linear or quadratic trend, leaving one with largely business cycle related fluctuations.

frequency variation (such as business cycle or seasonality) and noise based on parametric assumptions.

Finally, insofar as NSW are concerned that controls for trends use parametric assumptions, we wholeheartedly agree. For this reason we explicitly stated in DLR1 that finding credible comparison groups is better than putting in parametric trend controls. And that is the reason we consider the border discontinuity-based evidence in DLR1 and DLR2a and ADR1 to be convincing.

We welcome NSW's focus on the heterogeneity in the business cycle as a confounding factor. Unfortunately, their "solution" to this problem uses an *ad hoc* strategy that is unlikely to address the concerns. However, standard soutions exist–when deployed, they continue to show that the impact of minimum wages on teen employment is small.

Use of Spatially Correlated Placebos in DLR1

DLR1 uses a placebo-based falsification exercise to show directly the bias arising from spatial heterogeneity, but NSW argue this test is inappropriate.[37] Here we clarify the exercise and explain why the results in fact demonstrate that spatial heterogeneity contaminates the canonical model's employment estimates. In short, DRL1 examines the counties whose minimum wage is always equal to the federal minimum wage, and also the cross-state neighbors of these counties. DLR1 uses the canonical model to estimate *neighboring-county minimum wage* impacts on *own-county employment*. Intuitively, neighboring county minimum wages should produce an estimate close to zero unless there is spatial heterogeneity correlated with minimum wages.

Formally, define the placebo set A of border counties with minimum wages always equal to the federal minimum wage. Define also the subset B of cross-state neighbor counties for the counties in placebo set A. Construct the contiguous border county pair data for the A counties: for each time period, each county in these data has p replicates, one for each of

[37] Section I.G. and Appendix B of DLR1 describe the original exercise and section IV of NSW contains their critique.

its p neighbors in B. Using this sample with the canonical fixed effects framework, one can regress A-county employment on A-county minimum wages *and* paired B-county minimum wages:

$$y_{ipt}^{A} = \alpha MW_{ipt}^{A} + \beta MW_{ipt}^{B} + \gamma_i + \tau_t + \epsilon_{it}. \tag{5}$$

B-county minimum wages have no causal effect on A-county employment and any correlation with MW^B indicates bias due to an omitted variable, for which MW^B is acting as a proxy.[38] This omitted variables bias is exactly what one expects in the presence of spatial heterogeneity–that minimum wage increases tend to be correlated with certain regional shocks. Furthermore, note that because A-county minimum wages are always equal to the federal minimum wage, they are perfectly collinear with the time fixed effects τ_t. Estimating (2) is therefore the same as estimating:

$$y_{ipt}^{A} = \beta MW_{ipt}^{B} + \gamma_i + \tau_t + \epsilon_{it}, \tag{6}$$

which is what DLR1 does estimate, and finds that the estimate β is quite substantial and negative–indicating an omitted variables bias in the fixed effects regression.

NSW's critique suggests that they misunderstand this entire exercise. They claim that our placebo sample is "contaminated" because minimum wages are changing. They are changing, however, in exactly the same way in all counties in the placebo sample, since they all pay the federal minimum and are fully correlated with time effects. In other words, there is *zero* cross-sectional variation in minimum wages in the sample. NSW's argument about "contamination" misunderstands the basic sources of statistical variation used in a fixed effects model. In this sample, replacing the actual (common) minimum wage with a fictitious one (from the neighbor) should not produce a negative result. Yet it does, suggesting that the canonical specification is biased due to spatial heterogeneity.

As a solution to a non-problem, NSW then get rid of 80 percent of the sample by cutting

[38]This should be true except in cases of cross-border spillovers, for which DLR1 fails to find evidence in a test comparing border and interior counties (see DLR1, Section VA.)

out many of the years, and then by imposing an arbitrary restriction on cross-border minimum wage variation. Once they get rid of the 80 percent of data using arbitrary criteria, they discover that the placebo estimate becomes close to zero. This "solution" does not shed any light on the validity of the placebo exercise, because there was no problem with the exercise in the first place.

Bordering state pairs

Although NSW argue against using neighboring counties as controls with county-level data, or division-period controls for state-level data, they propose a neighboring state-pair design (see their Section IV). Using this methodology, NSW find a large and significant negative minimum wage effect on teen employment. In this section, we describe this analysis and show that the methodology does not identify appropriate counterfactuals, as it clearly fails the falsification test that there are no pre-existing trends.

Analogous to the cross-state border county design in DLR1, the bordering state research design uses paired state-period-level data: for each state s, there are p replicates of state-period observations, one for each state that shares a border with state s. State fixed effects and state pair-period fixed effects identify the minimum wage effect, and standard errors are clustered on both state and state border segment. Using the border state design and CPS data from 1990 through 2011q2, NSW estimate the teen employment elasticity to be -0.236 (see their Table 12). We confirm their estimates, finding an employment elasticity of -0.227 with CPS data from 1990 through 2010. Both estimates are significant at the level of 1 percent. NSW find small and insignificant employment effects using QCEW restaurant data, which we confirm using more recent data.[39] Addison et al. (2011) also use the border-state design with ACS and CPS data for the shorter 2006-2009 and 2006-2010 periods, respectively. Depending on the data source and exact employment outcome, Addison et al. find coefficients of a variety of signs and magnitudes, but they do estimate substantially-sized

[39] We find an insignificant employment elasticity of +0.0193 using QCEW data from 1990 through 2010, compared to NSW's estimate of -.029 using data from 1990 through 2006q2.

negative employment effects for teens: with the CPS, an insignificant elasticity of -0.153 (see their Table 10); and with the ACS, an elasticity of -0.224, significant at the ten percent level (see their Table 16).[40]

Like the estimate using within-division variation, the state-pair design has a strong *a priori* appeal: the identifying assumption of the latter is that minimum wage differences between the two neighboring states are uncorrelated with their differences in residual employment. One check of this assumption is the leading effects falsification test. Failing this specification check suggests that when there is a minimum wage differential between neighboring states, these neighbors are also on different employment paths. In Figure A1, we use CPS data on teens from 1990-2010 to estimate cumulative response effects of the minimum wage beginning at three years before the minimum wage increase, similar to our figure 4, discussed in Section 3.1. We estimate (1) the canonical two-way fixed effects model, (2) the model with state fixed effects, division-period fixed effect, and state-specific linear trends, and (3) the border state design.[41] We perform this exercise when the outcome is log state teen employment-to-population ratio, as well as log average teen hourly wage, for both annual and quarterly data, using all quarterly or annual leads from three years before the increase up through the actual increase. Figure A1 clearly illustrates that, of the three models, the best-behaved specification for employment is model (2), with division-period effects and

[40]Using different data sources, models, and sectors and demographic groups, Addison et al. (2011) report a total of nineteen employment elasticities for the border state design (see their Tables 8 through 16), five of which are statistically significant at the level of 10 percent or less. One of these estimates is positive: an elasticity of +0.097, signifcant at the ten percent level, for food and drinking establishments in the ACS. For gas station employment in the CPS and ACS, they find elastiticites of -0.473 and -0.420, respectively, both significant at the five percent level. Using CPS data from 2005-2010, the authors supplement their standard regression with an interaction of the minimum wage with a post-2007 time dummy, represent the recent recession; the non-interacted minimum wage coefficient is -0.116, significant at the ten percent level, but the interaction coefficient is just as large, at +0.144. Finally, the authors estimate the teen employment elasticity in the ACS sample to be -0.224, significant at the ten percent level.

[41]NSW also estimate state teen population-weighted versions of state-border regressions with the CPS and QCEW, obtaining qualitatively similar results to their unweighted regressions (*i.e.,* large negative employment elasticities with the CPS, and small, insignificant effects with the QCEW). We do not repeat the falsification exercise with these weighted regressions because such weights are inappropriate with state-pair fixed effects. Weighting the state-period means by population in non-paired data helps to target the underlying population regression function. With paired data and state-population weights, however, the state-pair-period fixed effects demean the data by the weighted average–in the case of a large state and small state next door, that average is effectively the value for the large state.

state-specific linear trends. Focusing on annual data in the first and second years before the increase, this preferred model estimates leading effects less than -0.1 in magnitude. During that time period the other two models estimate spurious causal effects between -0.1 and -0.3, suggesting that employment is exceptionally lower than neighboring state employment *before* the minimum wage increase.[42]

Compared to the coarse counterfactual comparisons of the canonical fixed-effects model, the border-*county* (or cross-state commuting zone) research design may be thought of as offering very fine-grained comparisons, which is the approach we have advocated. In between these two extremes are several intermediate specifications. Some intermediate specifications may fail a falsification test, while others may pass the same test. And the results may vary by data set and time period. This is exactly what happens with QCEW county-level restaurant employment: the model with division-period fixed effects and state-specific linear trends fails the falsification test (as NSW discuss; see also our Section 3.1), but the border-county approach passes the test. Here, with state-level CPS data on teens, the border-state approach estimates spurious pre-treatment effects, like the canonical model, but unlike the division-period fixed effects model with state-specific linear trends. For this reason, we do not find the CPS teen employment elasticity estimates using the border-state research design to be reliable.

[42]Although we do not show these estimates, when using QCEW restaurant employment the leading effects are better behaved in the border-state design than in the canonical model. Nevertheless, the smallest pre-treatment impacts occur when using the border-county research design. Adding state-specific linear trends does not improve the pre-treatment estimates with CPS teen or QCEW restaurant data.

Figure 1 **High versus low minimum wage states during 1990 – 2012: Means and variances**

A. Minimum wage means

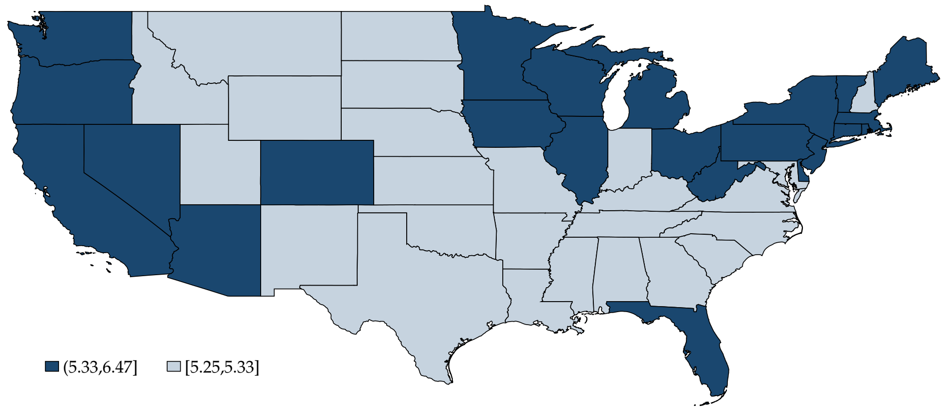

B. Minimum wage variances

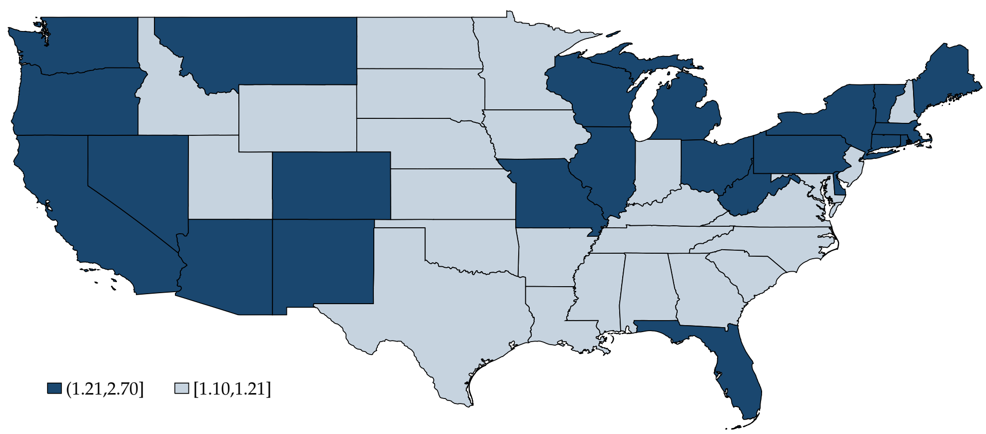

Notes. State means and variances calculated using annual state minimum wage data over 1990-2012. The shading on the maps partitions the states into above- and below-median values.

Figure 2   Regional patterns in labor market trends and cycles

A. Fall in Routine task intensity: 1990 to 2007

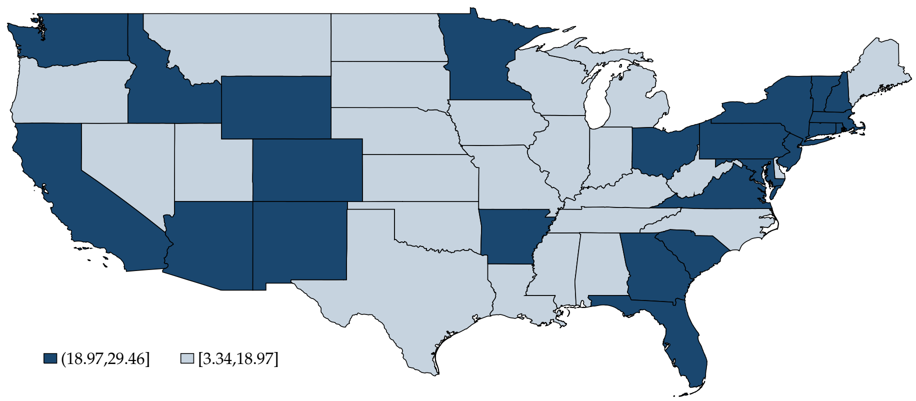

B. Increase in upper-half wage inequality (90/50 ratio)

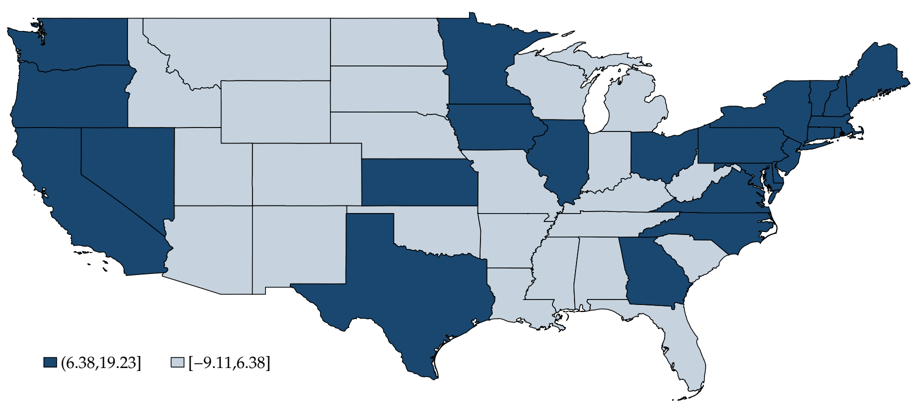

C. Variance of the annual unemployment rate, 1990-2012

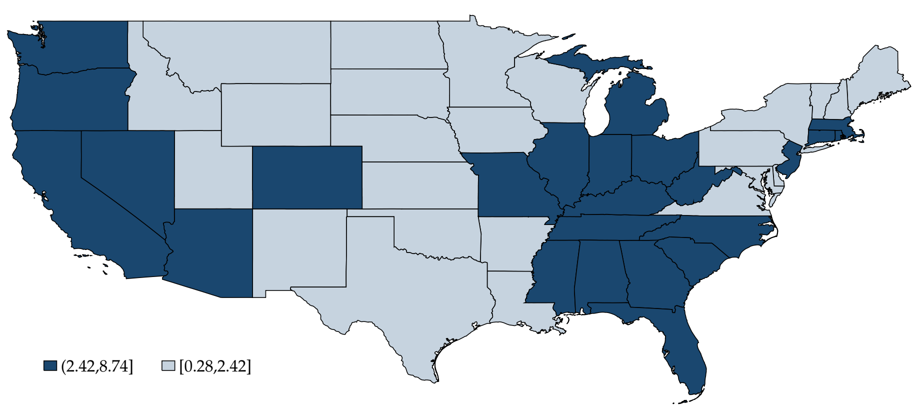

Notes. Inequality and Routine Task Intensity index use coding from Autor, Levy and Murnane (2003). Overall unemployment rate calculated from the CPS-MORG. The map shading partitions the states into above- and below-median values.

Figure 3 Maps of cross-state border county pairs and cross-state commuting zones

A. Cross-state border county pairs

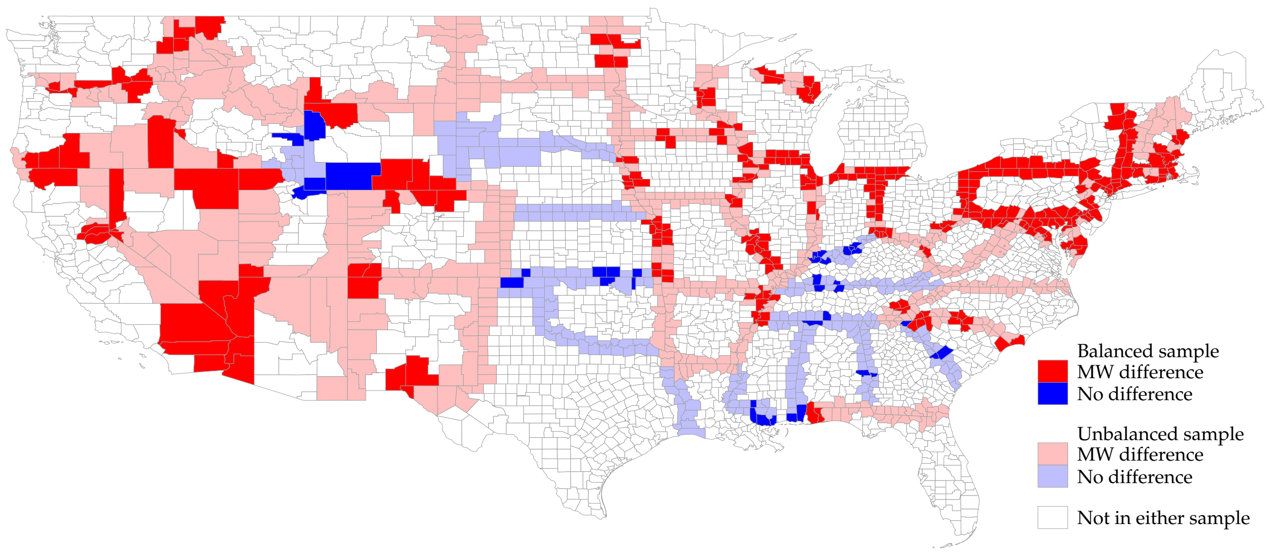

B. Cross-state commuting zones

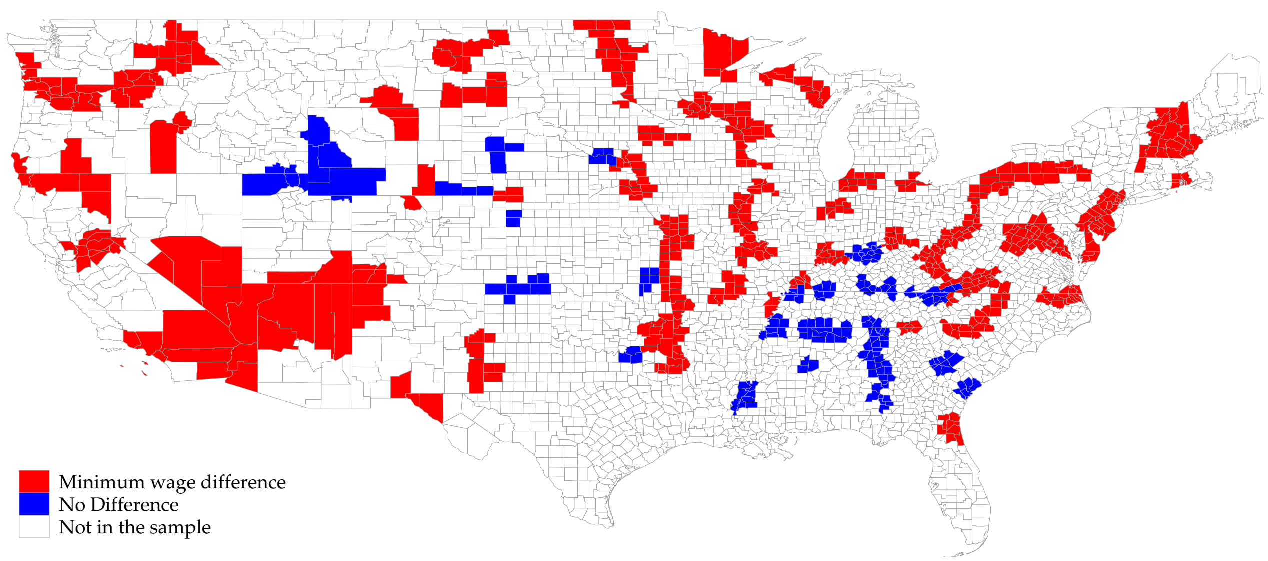

Notes. A: Red and blue colored counties indicate cross-state border county pairs. Counties colored red are part of pairs with minimum wage variation between the counties at some point in time in the sample. Darker shades indicate balanced sample. Balanced sample are those counties with employment and earnings information for all quarters, 1990-2010. Unbalanced sample includes those with limited information during that period. B: Red and blue colored counties constitute cross-state commuting zones. Counties colored red are part of commuting zones with minimum wage variation within the commuting zone at some point in time in the sample.

Figure 4 Pre-existing trends in canonical and local specifications: Annual and quarterly data

A. Annual Data

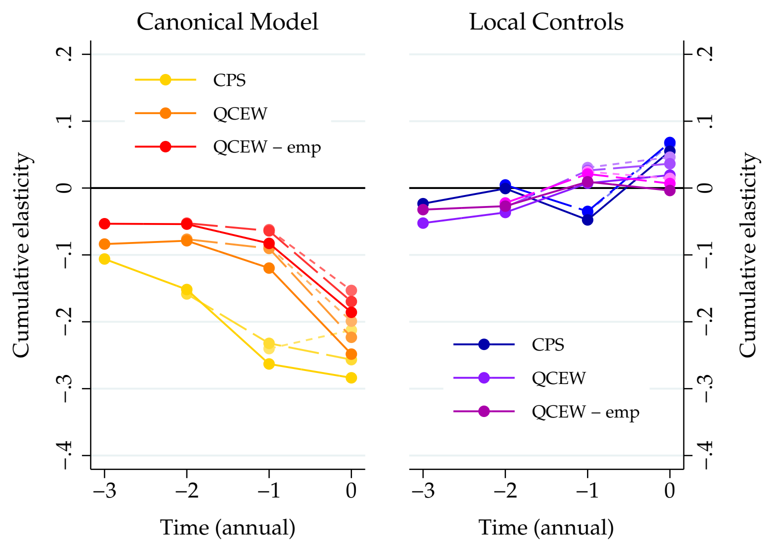

B. Quarterly Data

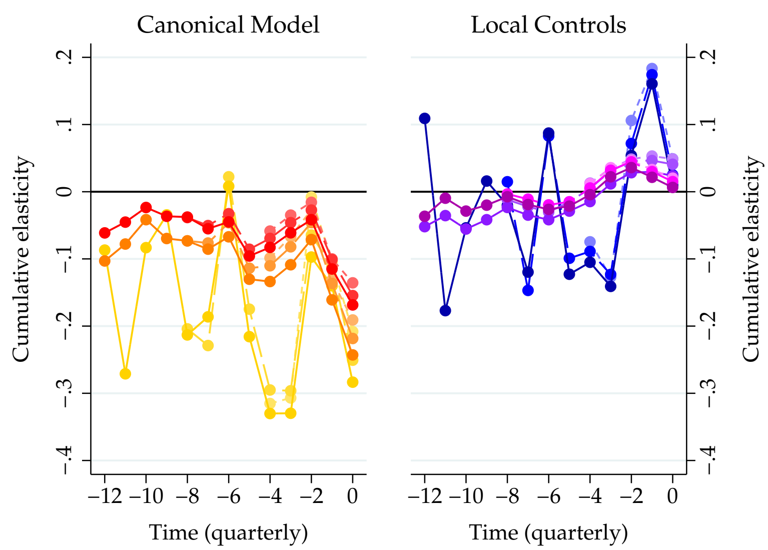

Notes. Cumulative responses for employment elasticities using from one to three years of leads of the minimum wage along with the contemporaneous minimum wage. The cumulative response at time t equals the sum of all leads up to and including time t. The actual minimum wage increase occurs at time t=0. Elasticities are calculated with the CPS, QCEW, and the QCEW with a control for overall private-sector employment. The canonical model includes period and place fixed effects (CPS - state, QCEW - county). CPS local controls are state and division×period fixed effects, as well as state-specific linear time trends. QCEW local controls are county and county pair×period fixed effects. The bottom panel repeats this exercise using quarterly data, with all four quarterly leads for a year included in the regression. CPS estimates use as dependent variable log teen employment-to-population ratio; they include controls for overall unemployment rate and racial, gender, and teen population shares. QCEW estimates use as dependent variable log restaurant-sector employment; they include log population as a control and may include log private-sector employment.

Figure 5 How synthetic control weights for donors vary by distance to treated states

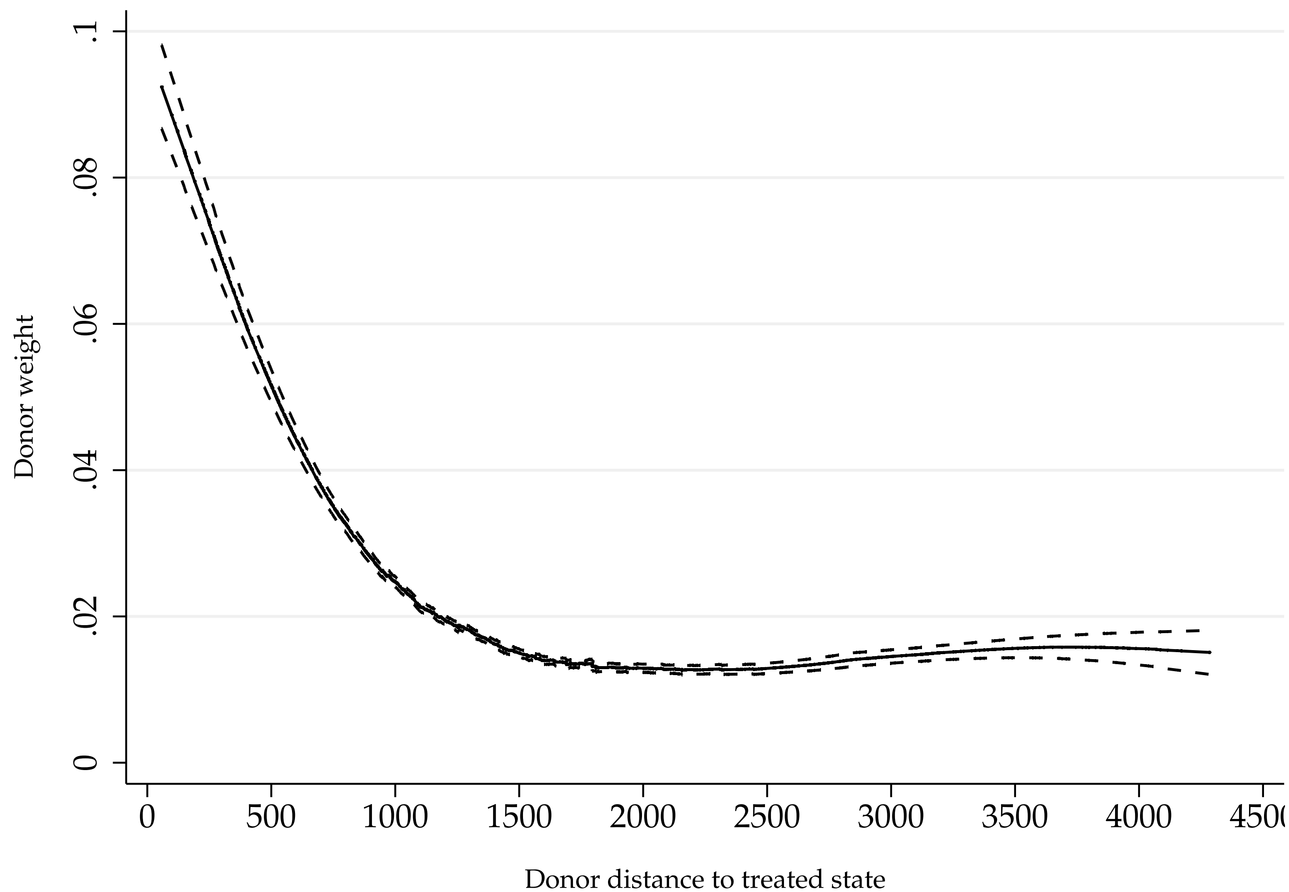

Notes. The figure shows results from a non-parametric lowess regression of synthetic control weight for a particular donor state on the distance in miles between that donor and the treated state. Donor weights were estimated based on 720 placebo treatments for the 48 contiguous states with a 4 year pre−treatment window during 1997q4−2007q2. Dashed lines indicate 95 percent confidence intervals.

Figure A1 Pre-existing employment trends using canonical, paired border-state, and local specifications

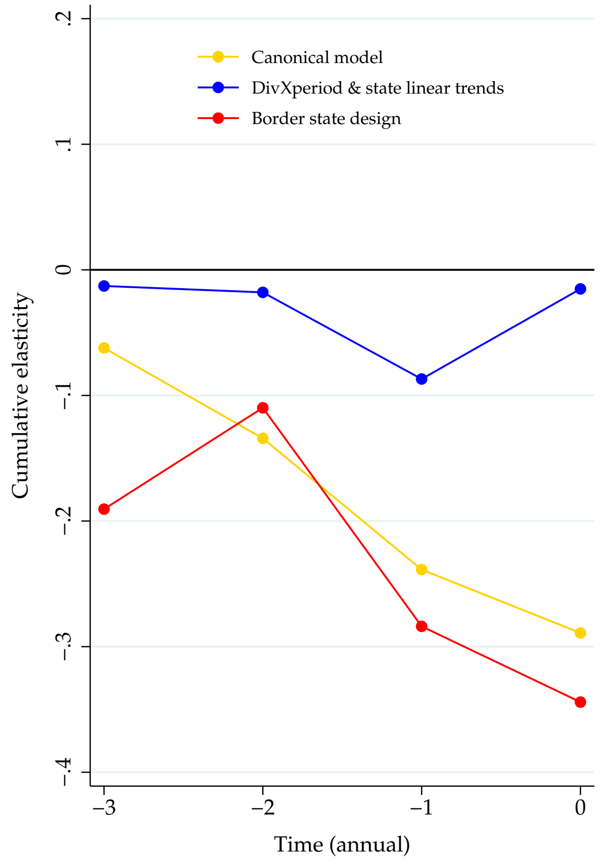

Notes. Cumulative responses are shown for employment elasticities using from one to three years of leads of the minimum wage along with the contemporaneous minimum wage. The cumulative response at time t equals the sum of all leads up to and including time t. The actual minimum wage increase occurs at time t = 0. See text for details. Elasticities are calculated with state-year-level CPS data with both the cell means and the regression weighted by teen population. The dependent variable is the log of the teen employment-to-population ratio and controls include the overall unemployment rate and racial, gender, and teen population shares. The canonical model includes year and state fixed effects. The paired border-state design includes state pairXyear and state fixed effects. Local controls with the CPS are state and divisionXyear fixed effects, as well as state-specific linear time trends.

Table 1 Heterogeneity between high and low minimum wage states

	High-MW state mean	Low-MW state mean	Difference
Political Economy			
Percent of states voting Democratic in 2008 presidential election	88.46	24.00	64.46***
			(10.81)
Percent represented by a union in 1990	20.97	12.86	8.11***
			(1.31)
Percent represented by a union in 2012	15.15	8.07	7.08***
			(1.08)
Percent change in unionization from 1990-2012	-27.64	-36.29	8.65**
			(3.85)
Inequality and polarization			
90-50 wage ratio in 1990	2.07	2.15	-0.08**
			(0.03)
90-50 wage ratio in 2007	2.26	2.22	0.040
			(0.04)
Percent change in 90-50 wage ratio	8.90	3.21	5.69***
			(1.70)
Routine Task Intensity in 1990	1.27	1.21	0.050**
			(0.02)
Routine Task Intensity in 2007	1.00	1.01	-0.007
			(0.02)
Percent change in Routine Task Intensity	-20.52	-16.69	-3.83**
			(1.53)
Business cycle variation			
Average unemployment rate over 1990-2012	5.50	5.30	0.21
			(0.33)
Average peak-trough percent change in employment over last three recessions	-2.59	-1.87	-0.72**
			(0.35)
Variance of the unemployment rate over 1990-2012	3.23	2.18	1.05**
			(0.46)

Notes. High- (low-) MW states are those with an average minimum wage over 1990-2012 that is above (below) the national median. Difference of means calculated as regression of the outcome on a high minimum wage dummy variable using robust SEs. Vote data from the Statistical Abstract: www.census.gov/compendia/statab/2012/tables/12s0406.xls. Unionization data from CPS-MORG. Inequality and RTI statistics use coding from Autor, Levy and Murnane (2003). Peak-trough comparisons use the monthly seasonally-adjusted Current Employment Statistics data from BLS. The three recessions are July 1990 - March 1991; March 2001 - November 2001; and December 2007 - June 2009.

Table 2 Minimum wage effects and spatial heterogeneity correction: Updated results from five datasets

	Canonical model with common trends				Local controls for spatial heterogeneity			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A Teens								
Wage/Earnings	0.190***	0.124***		0.182***	0.296***	0.167***		0.222***
	(0.045)	(0.025)		(0.037)	(0.064)	(0.033)		(0.048)
Employment	-0.123**	-0.162**		-0.161**	0.130	0.002		-0.059
	(0.057)	(0.042)		(0.071)	(0.126)	(0.066)		(0.086)
Hires				-0.510***				-0.219**
				(0.092)				(0.095)
Separations				-0.539***				-0.233**
				(0.099)				(0.099)
Panel B Restaurant workers								
Wage/Earnings			0.210***	0.204***			0.186***	0.207***
			(0.024)	(0.027)			(0.027)	(0.060)
Employment			-0.119	-0.079**			0.008	-0.022
			(0.076)	(0.039)			(0.051)	(0.093)
Hires				-0.468***				-0.264*
				(0.084)				(0.137)
Separations				-0.465***				-0.225*
				(0.077)				(0.129)
Time effects								
Common	Y	Y	Y	Y				
Commuting zone-period					Y			
County-pair-period							Y	Y
Division-period						Y		
State-linear-trend						Y		
Data sets	ACS/Census	CPS	QCEW	QWI	ACS/Census	CPS	QCEW	QWI
Years	1990, 2000, 2005-2011	1990-2012	1990-2010	2000-2011	1990, 2000, 2005-2011	1990-2012	1990-2010	2000-2011

Notes. In the annual ACS/Census or annual CPS microdata, the dependent variables are binary employment indicators or log of hourly wages. In the QCEW and QWI county-data, dependent variables are log of quarterly employment or quarterly average earnings. The key independent variable is log of annual (ACS/Census) or monthly (CPS) or quarterly (QCEW/QWI) minimum wage. All minimum wage effects expressed as elasticities: for the CPS and ACS/Census micro-level employment regressions, the minimum wage regression coefficients and standard errors are divided by the sample's average employment rate. Canonical regressions always include common time fixed effects and also include, with the ACS/Census, state-commuting-zone effects; with the CPS, state fixed effects; and, with the QWI and QCEW, county fixed effects. Local control specifications additionally include, with the ACS/Census, commuting-zone-year fixed effects; with the CPS, division-year fixed effects and state-specific linear trends; and, with the QWI and QCEW, county-pair-period fixed effects. All ACS/Census regressions use sample weights and include controls for gender, race, age, education, marital status, and annual state unemployment rate. All CPS regressions use sample weights and include controls for for gender, race, age, education, marital status, the teen share of the population, and the non-seasonally adjusted state unemployment rate. All QCEW and QWI regressions include controls for log of overall population and private-sector employment or earnings. Teens in ACS/Census and CPS have ages 16-19, and teens in the QWI have ages 14-18. In QCEW specifications the restaurant sector is defined by the sum of limited-service restaurants (NAICS 7221) and full-service restaurants (NAICS 7222). In QWI regressions, the restaurant sector is the overall restaurant industry (NAICS 722). Robust standard errors are clustered at the state-level for all canonical specifications 1-4 and specification 6 (Census/ACS) and 7 (CPS). Standard errors for QCEW and QWI specifications 7 ad 8 are clustered multi-dimensionally at the state-level and border segment-level. Significance levels are ***1%, **5%, *10%

Table 3 Pre-existing employment trends: Cumulative responses to minimum wage increases

	Canonical model				Local controls			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log teen EPOP (CPS)								
t=0	-0.162**	-0.212***	-0.257***	-0.284***	0.081	0.065	0.068	0.055
	(0.062)	(0.064)	(0.067)	(0.068)	(0.077)	(0.102)	(0.139)	(0.166)
t=-1		-0.240**	-0.232**	-0.263***		-0.036	-0.035	-0.048
		(0.091)	(0.088)	(0.098)		(0.112)	(0.124)	(0.148)
t=-2			-0.158*	-0.152*			0.005	-0.001
			(0.081)	(0.078)			(0.097)	(0.108)
t=-3				-0.106				-0.023
				(0.083)				(0.088)
N	1,071	1,071	1,071	1,071	1 ,071	1,071	1,071	1,071
Log restaurant employment (QCEW)								
t=0	-0.178**	-0.199**	-0.223*	-0.248**	0.037	0.046	0.036	0.019
	(0.083)	(0.095)	(0.112)	(0.123)	(0.074)	(0.086)	(0.094)	(0.096)
t=-1		-0.088	-0.090	-0.120*		0.031	0.026	0.007
		(0.058)	(0.058)	(0.071)		(0.048)	(0.053)	(0.058)
t=-2			-0.077	-0.079			-0.029	-0.037
			(0.059)	(0.058)			(0.041)	(0.045)
t=-3				-0.084*				-0.052
				(0.046)				(0.039)
N	27,132	27,132	27,132	27,132	11,886	11,886	11,886	11,886
Log restaurant employment (QCEW) - with control for overall private-sector employment								
t=0	-0.138	-0.153	-0.170	-0.186	0.007	0.015	0.007	-0.004
	(0.083)	(0.095)	(0.112)	(0.123)	(0.055)	(0.065)	(0.076)	(0.082)
t=-1		-0.062	-0.064	-0.083		0.024	0.021	0.009
		(0.055)	(0.056)	(0.068)		(0.044)	(0.049)	(0.058)
t=-2			-0.053	-0.054			-0.022	-0.027
			(0.061)	(0.061)			(0.042)	(0.046)
t=-3				-0.053				-0.032
				(0.043)				(0.037)
N	27,132	27,132	27,132	27,132	11,886	11,886	11,886	11,886

Notes. All data are annual over 1990 – 2010. The dependent variable for CPS data is log of teen employment-to-population ratio. The dependent variable for QCEW data is log restaurant employment. The main independent variables are log contemporaneous minimum wage, and up to 3 years of leads in log minimum wage. The cumulative responses are formed by successively adding the coefficients for the leading and contemporaneous log minimum wage. See text for details. For the CPS, the canonical model includes state fixed effects and time fixed effects; the local controls model includes county fixed effects, division and year fixed effects, and state-specific linear trends. Additional CPS controls are race, sex, and teen population share. For the QCEW, the canonical model includes county and time fixed effects; the local controls model include county fixed effects and county-pair-period fixed effects. Additional QCEW controls include log population and, when noted, log private-sector employment. All standard errors are clustered on state except in QCEW local control design, where they are clustered on state and border segment.

Table 4 Mean absolute differences in covariates between counties in contiguous pairs versus other pairs

	Non-contiguous pair	Contiguous pair	Gap	Percent gap
Levels				
Log employment	1.744	1.233	0.511***	41
	(0.026)	(0.027)	(0.033)	
Log population	0.042	0.039	0.003***	8
	(0.001)	(0.001)	(0.001)	
EPOP	0.229	0.170	0.060***	35
	(0.004)	(0.004)	(0.004)	
Log earnings	1.518	0.964	0.554***	57
	(0.023)	(0.023)	(0.029)	
Turnover rate	0.057	0.048	0.009***	18
	(0.001)	(0.001)	(0.001)	
Teen share	0.006	0.005	0.001***	22
	(0.0001)	(0.0001)	(0.0001)	
3-year differences				
Log employment	0.099	0.091	0.008***	8
	(0.001)	(0.002)	(0.001)	
Log population	0.069	0.066	0.004***	5
	(0.001)	(0.002)	(0.001)	
EPOP	0.037	0.027	0.001***	36
	(0.001)	(0.001)	(0.001)	
Log earnings	0.018	0.017	0.001***	8
	(0.0003)	(0.0004)	(0.0003)	
Turnover rate	0.003	0.002	0.001***	25
	(0.000)	(0.000)	(0.000)	
Teen share	0.045	0.041	0.004***	9
	(0.001)	(0.001)	(0.001)	

Notes. Reproduced from Dube, Lester and Reich 2013 (DLR2b). Each of the 972 counties in 966 cross-border pairs with centroids within 75 miles is merged with every possible out-of-state county, a total of 1,737,884 pairings. Absolute differences in levels and changes are calculated between the county, its border pair and its randomly assigned pair, respectively. Subsequently, the dataset is collapsed back to county-pair-period level and means of the absolute differences in covariates between counties within pairs are calculated subsequently, clustering standard errors multi-dimensionally on each of the two counties in the cross-border pair. "Gap" is the difference in mean absolute difference of the covariate between contiguous and other pairs. Standard errors in parentheses are clustered multi-dimensionally by each county in the pair. Significance levels are for the null hypothesis that Gap is zero and are indicated by: * for 10%, ** for 5%, and *** for 1%. "Percent Gap" divides the Gap value by the mean absolute difference for contiguous pairs.

Table 5 Donor weights and distance to treated states

	Weight per donor		
Donor relation to treatment	Inside	Outside	Relative weight
Same Census region	0.044	0.014	3.150
Same Census division	0.063	0.016	3.868
Within 0 - 500 miles	0.068	0.017	4.100
Within 0 - 1000 miles	0.040	0.013	3.022

Notes. Sample consists of 768 placebo treatments of 48 states during 1997q4-2007q2 with pre- and post-treatment windows of 4 and 2 years, respectively. Synthetic control outcome is the teen employment-to-population ratio. See text for exact model specification. Weight per donor inside (outside) a given area is equal to the total weight of all donors inside (outside) the area of their respective treatment states, divided by the number of such donors. Distance indicates distance between population-weighted state centroids.

Table 6 Relative weights for local controls group implied by NSW

States Specification	Statistic	Same-division states	Other states	Relative per donor weight
		(1)	(2)	(3)
Regression residuals	Proportion of weights	0.233	0.767	
	Average no. of donor states	2.5	21.5	
	Per donor weight	0.093	0.036	2.613
Log(EPOP)	Proportion of weights	0.323	0.677	
	Average no. of donor states	2.5	21.5	
	Per donor weight	0.129	0.032	4.103
Counties Specification	Statistic	Cross-border counties	Other counties	Relative per donor weight
Regression residuals	Proportion of weights	0.006	0.994	
	Average no. of donor states	1.7	957.9	
	Per donor weight	0.004	0.001	3.401
Log(EPOP)	Proportion of weights	0.007	0.993	
	Average no. of donor states	1.7	957.9	
	Per donor weight	0.004	0.001	3.972

Notes. Proportion of weights and average number of donor states are from NSW, Tables 4 and 6. Relative per donor weight (column 3) equals the ratio of per donor weights in column 1 to that in column 2.

Table 7 Employment and wage effects using the synthetic control approach

Event		Window length		Donors	ΔMW%	Employment elasticity		Wage elasticity	
		Pre	Post			Effect	Rank	Effect	Rank
AK	2003q1	16	8	38	0.265	0.100	(0.650)	0.378	(0.925)
CA	2001q1	8	8	41	0.174	-0.362	(0.256)	0.139	(0.698)
CT	2006q1	8	6	21	0.077	2.381	(0.957)*	0.178	(0.739)
FL	2005q2	16	8	21	0.295	0.165	(0.652)	0.127	(0.826)
HI	2002q1	16	8	39	0.190	0.114	(0.561)	0.357	(0.854)
HI	2006q1	12	6	21	0.160	0.569	(0.870)	0.574	(0.913)
IL	2004q1	16	8	33	0.262	0.182	(0.714)	0.080	(0.743)
MA	2000q1	8	8	42	0.286	-0.131	(0.273)	0.323	(0.932)
ME	2002q1	16	8	39	0.214	-0.104	(0.415)	0.063	(0.683)
MN	2005q3	16	8	20	0.194	-0.122	(0.545)	0.056	(0.636)
NJ	2005q4	16	6	21	0.388	-0.004	(0.565)	0.296	(0.957)*
NY	2005q1	16	8	31	0.311	-0.662	(0.030)*	0.149	(0.909)
OR	2003q1	16	8	38	0.085	-2.602	(0.025)*	-0.251	(0.325)
RI	2004q1	12	8	35	0.098	0.579	(0.730)	0.047	(0.568)
RI	2006q1	8	6	21	0.096	0.139	(0.565)	0.414	(0.870)
VT	1999q4	8	8	42	0.190	-0.225	(0.318)	0.422	(0.932)
VT	2004q1	12	8	35	0.120	0.137	(0.595)	0.260	(0.757)
WI	2005q2	16	8	21	0.262	0.082	(0.652)	-0.158	(0.130)
WV	2006q3	16	4	20	0.136	-0.773	(0.227)	-0.143	(0.364)
Mean				30.5	0.200	-0.028	(0.505)	0.174	(0.724)*
Median				33.0	0.190	0.082	(0.565)	0.149	(0.757)*

Notes. Reproduced from Dube and Zipperer (2013). The elasticities for wage and employment are computed using the difference in mean outcomes between the treated and the synthetic control units over the post-intervention window, and the proportional change in minimum wages. Synthetic controls are chosen based on matching on the outcome for each pre-intervention quarter in the sample. Asterisks denote significance of two-sided tests at the 90 percent significance level using randomization inference with placebo treatments. For the pooled mean and median, the 90 percent cutoffs are (0.391, 0.609) and (0.320, 0.680), respectively.

Table 8 Teen minimum wage elasticities with lagged dependent variables and dynamic panel models Annual data 1990-2012

	FE model	LDV model			Dynamic panel data model		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log average teen wage							
Log MW	0.104**	0.253***	0.213***	0.197***	0.285***	0.246***	0.217***
	(0.047)	(0.034)	(0.031)	(0.030)	(0.034)	(0.033)	(0.032)
Log Wage - Lag 1		0.290***	0.216***	0.174***	0.204***	0.170***	0.146***
		(0.055)	(0.035)	(0.031)	(0.045)	(0.033)	(0.031)
Log Wage - Lag 2			0.211***	0.157***		0.159***	0.121***
			(0.034)	(0.029)		(0.033)	(0.029)
Log Wage - Lag 3				0.184***			0.189***
				(0.029)			(0.031)
Log teen employment							
Log MW	-0.135**	-0.025	-0.014	-0.004	-0.033	-0.016	0.001
	(0.062)	(0.031)	(0.029)	(0.030)	(0.037)	(0.034)	(0.034)
Log Emp - Lag 1		0.713***	0.500***	0.448***	0.591***	0.452***	0.415***
		(0.046)	(0.029)	(0.032)	(0.075)	(0.037)	(0.032)
Log Emp - Lag 2			0.283***	0.180***		0.235***	0.144***
			(0.032)	(0.041)		(0.048)	(0.046)
Log Emp - Lag 3				0.182***			0.185***
				(0.035)			(0.037)
Controls:							
Year FE	Y	Y	Y	Y	Y	Y	Y
State FE	Y				Y	Y	Y
N	1,173	1,173	1,173	1,173	1,173	1,173	1,173

Notes. The dependent variable is log teen wage or teen EPOP and the treatment variable is log minimum wage. State-level regressions use annual CPS data from 1990 to 2012, weighted by teen population. Controls include time fixed effects and overall unemployment rate. Columns 1, 5, 6 and 7 also include state fixed effects. Columns 2 -7 include as many lags of the outcomes (log teen wage, log teen EPOP) as indicated. Columns 5, 6 and 7 use system-GMM estimate of the dynamic panel models using forward orthogonal deviations to purge fixed effects. Standard errors in parentheses are clustered by state. Significance levels are indicated by ***1%, **5%,*10%.

Table A1 Minimum wage elasticities for teen employment accounting for business cycle heterogeneity

Original Specification:	-0.162**	0.002
	-0.042	-0.066
Include state-recession dummies:	-0.176**	-0.004
	-0.043	-0.068
Leave out recessions:	-0.169**	-0.002
	-0.042	-0.068
Time effects:		
Division-time		Y
State-trends		Y

Notes. The table reports employment elasticities of the minimum wage using individual-level CPS data on teens from 1990 to 2012. The dependent variable is a binary employment indicator and the key independent variable is log of the quarterly minimum wage. The resulting semi-elasticity (and its standard error) is converted to an elasticity by dividing by the sample mean of employment. Controls include the monthly unemployment rate, the teen share of the population, and individual-level demographic controls: race, education, gender, marital status and age dummies.. Time effects are quarterly. State recession dummies are interactions of state fixed effects and a recession indicator. Recessionary periods are as designated by National Bureau of Economic Research (NBER). Standard errors are clustered at the state level.

Table A2 Minimum wage effects and spatial heterogeneity correction: Original results from five datasets

	Canonical model common trends assumption				Accounting for spatial heterogeneity			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A Teens								
Wage/Earnings	0.110**	0.123***		0.179***	0.151***	0.149***		0.202***
	(0.053)	(0.026)		(0.037)	(0.051)	(0.024)		(0.050)
Employment	-0.159*	-0.118**		-0.149**	0.129*	0.047		-0.044
	(0.033)	(0.022)		(0.073)	(0.028)	(0.024)		(0.085)
Hires				-0.474***				-0.198**
				(0.090)				(0.094)
Separations				-0.520***				-0.231**
				(0.097)				(0.099)
Panel B Restaurant workers								
Wage/Earnings			0.221***	0.193***			0.188***	0.193***
			(0.032)	(0.027)			(0.060)	(0.062)
Employment			-0.112	-0.083**			0.016	-0.028
			(0.076)	(0.041)			(0.098)	(0.095)
Hires				-0.457***				-0.265*
				(0.081)				(0.140)
Separations				-0.461***				-0.239*
				(0.073)				(0.127)
Time effects								
Common	Y	Y	Y	Y				
Commuting zone-period					Y			
County-pair-period							Y	Y
Division-period						Y		
State-linear-trend						Y		
Data sets	ACS/Census	CPS	QCEW	QWI	ACS/Census	CPS	QCEW	QWI
Years	1990, 2000, 2004-2011	1990-2009	1990-2006	2000-2011	1990, 2000, 2004-2011	1990-2009	1990-2006	2000-2011
Source	ADR1	ADR2	DLR1	DLR2a	ADR1	ADR2	DLR1	DLR2a

Notes. All specifications control for gender, race, age, education, marital status, and non-seasonally adjusted state unemployment rate. Regressions include workers paid between $1 and $100 per hour in 1990 dollars. Outcome variables are expressed in log form. Robust standard errors, in parentheses, are clustered at the state and border segment levels for all regressions. Specifications 2 and 6: Both specifications also include individual controls for the teen share of the population. Specifications 3 and 7: The employment regressions control for the log of annual county-level population and overall private sector employment or earnings (depending on the outcome). Both specifications include county fixed effects. Specification 3 includes period fixed effects. For specification 7, period fixed effects are interacted with each census division, metropolitan area, and county-pair, respectively. Specifications 4 and 8: Both regressions include controls for natural log of county population and total private sector employment. Specifications for teens provide estimates for all teens aged 14-18 regardless of industry. Specifications for restaurant workers are limited to all workers in the restaurant industry (NAICS 722). All samples and specifications include county fixed-effects. Specification 4 includes common time period fixed-effects. For specification 8, period fixed-effects are interacted with each county-pair. Significance levels are indicated by ***1%, **5%, *10%. ADR1 refers to Allegretto, Dube, and Reich (2009). ADR2 refers to Allegretto, Dube, and Reich (2011). DLR1 refers to Dube, Lester, and Reich (2011). DLR2a refers to Dube, Lester, and Reich (2012) and DLR2b refers to the updated version of this paper, Dube, Lester, and Reich (2013).

Table A3 Minimum wage effects using balanced and unbalanced QCEW samples

	All counties - canonical model		Border counties - canonical model			Border counties - border discontinuity model		
Earnings elasticity	0.210***	0.217***	0.213***	0.222***	0.225***	0.186***	0.186***	0.196***
	(0.024)	(0.026)	(0.029)	(0.030)	(0.028)	(0.039)	(0.051)	(0.049)
Employment elasticity	-0.119	-0.165**	-0.0387	-0.0742	-0.136**	0.008	0.008	-0.034
	(0.076)	(0.078)	(0.060)	(0.065)	(0.067)	(0.074)	(0.098)	(0.076)
County FE	Y	Y	Y	Y	Y	Y	Y	Y
Period FE	Y	Y	Y	Y	Y			
County pair X period FE						Y	Y	Y
Balanced Panels	Y		Y	Y		Y	Y	
Full Pairs			Y			Y		
N	108,360	161,844	28,308	38,892	58,100	47,040	82,740	122,445
Total counties	1290	2731	337	463	986	337	463	986
Average obs. per county	84	59.26	84	84	58.92	139.6	178.7	124.2
Total states	51	51	48	49	49	48	49	49
Total county pairs						280	705	1137
Total border segments						73	100	107

Notes. Data are quarterly from 1990 through 2010. The balanced sample restriction requires counties to have non-missing data for all periods. The full pairs restriction requires that any county pair×period observation contains non-missing data for both counties. All regressions have as dependent variable log restaurant employment (or earnings) and include controls for log population and log overall private sector employment (or earnings).

Table A4a Pre-existing trends in teen employment: Cumulative response elasticities using quarterly CPS data

Cumulative Response	Canonical model				Local controls			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
t=0	-0.170***	-0.208***	-0.251***	-0.283***	0.022	0.049	0.025	0.010
	(0.057)	(0.059)	(0.062)	(0.062)	(0.071)	(0.096)	(0.126)	(0.142)
t=1		-0.104	-0.111	-0.140		0.183	0.174	0.161
		(0.110)	(0.113)	(0.114)		(0.154)	(0.181)	(0.192)
t=2		-0.008	-0.065	-0.097		0.106	0.072	0.054
		(0.102)	(0.101)	(0.110)		(0.111)	(0.122)	(0.133)
t=3		-0.307***	-0.296**	-0.330***		-0.127	-0.123	-0.141
		(0.111)	(0.116)	(0.123)		(0.140)	(0.159)	(0.173)
t=4		-0.315***	-0.295***	-0.330***		-0.075	-0.089	-0.105
		(0.109)	(0.107)	(0.113)		(0.129)	(0.134)	(0.143)
t=5			-0.175	-0.216			-0.099	-0.123
			(0.127)	(0.131)			(0.148)	(0.166)
t=6			0.022	0.008			0.083	0.087
			(0.105)	(0.102)			(0.112)	(0.117)
t=7			-0.229**	-0.186**			-0.147	-0.120
			(0.088)	(0.087)			(0.132)	(0.130)
t=8			-0.204**	-0.213**			0.0148	-0.020
			(0.098)	(0.102)			(0.120)	(0.116)
t=9				-0.035				0.016
				(0.127)				(0.095)
t=10				-0.083				-0.054
				(0.096)				(0.105)
t=11				-0.271**				-0.177
				(0.121)				(0.132)
t=12				-0.087				0.109
				(0.080)				(0.130)
N	4,284	4,284	4,284	4,284	4,284	4,284	4,284	4,284

Notes. See Table 3.

Table A4b Pre-existing trends in restaurant employment: Cumulative response elasticities using quarterly QCEW data, without overall employment controls

	Canonical model - all counties				Local controls - border counties			
Cumulative Response	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
t=0	-0.166**	-0.191**	-0.218**	-0.243**	0.033	0.095	0.041	0.024
	(0.076)	(0.090)	(0.106)	(0.118)	(0.099)	(0.121)	(0.136)	(0.139)
t=1		-0.131*	-0.137*	-0.161*		0.053	0.046	0.030
		(0.071)	(0.073)	(0.085)		(0.082)	(0.091)	(0.099)
t=2		-0.027	-0.046	-0.071		0.049	0.046	0.028
		(0.050)	(0.056)	(0.068)		(0.078)	(0.085)	(0.093)
t=3		-0.068	-0.082	-0.109		0.035	0.031	0.012
		(0.050)	(0.055)	(0.068)		(0.063)	(0.078)	(0.087)
t=4		-0.099*	-0.110**	-0.134**		0.013	0.005	-0.015
		(0.050)	(0.055)	(0.066)		(0.057)	(0.074)	(0.083)
t=5			-0.114*	-0.130**			-0.015	-0.029
			(0.060)	(0.064)			(0.070)	(0.077)
t=6			-0.046	-0.067			-0.030	-0.042
			(0.051)	(0.056)			(0.065)	(0.069)
t=7			-0.076	-0.086			-0.022	-0.035
			(0.052)	(0.056)			(0.049)	(0.059)
t=8			-0.073	-0.073			-0.012	-0.023
			(0.051)	(0.054)			(0.050)	(0.065)
t=9				-0.070				-0.042
				(0.050)				(0.068)
t=10				-0.042				-0.056
				(0.041)				(0.057)
t=11				-0.078*				-0.036
				(0.042)				(0.044)
t=12				-0.103**				-0.052
				(0.044)				(0.041)
N	108,360	108,360	1083,60	108,360	47,040	47,040	47,040	47,040

Notes. See Table 3.

Table A4c Pre-existing trends in restaurant employment: Cumulative response elasticities using quarterly QCEW data, with overall employment controls

	Canonical model - all counties				**Local controls - border counties**			
Cumulative Response	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
t=0	-0.119	-0.136	-0.155	-0.168	0.008	0.020	0.015	0.006
	(0.076)	(0.090)	(0.107)	(0.119)	(0.074)	(0.091)	(0.109)	(0.118)
t=1		-0.100	-0.101	-0.115		0.033	0.030	0.022
		(0.066)	(0.070)	(0.082)		(0.069)	(0.079)	(0.092)
t=2		-0.016	-0.027	-0.042		0.046	0.045	0.036
		(0.053)	(0.059)	(0.071)		(0.066)	(0.076)	(0.087)
t=3		-0.0368	-0.041	-0.061		0.036	0.032	0.022
		(0.048)	(0.054)	(0.067)		(0.057)	(0.072)	(0.085)
t=4		-0.059	-0.069	-0.083		0.012	0.006	-0.004
		(0.044)	(0.051)	(0.063)		(0.051)	(0.068)	(0.081)
t=5			-0.086	-0.096			-0.0149	-0.022
			(0.058)	(0.063)			(0.071)	(0.079)
t=6			-0.033	-0.045			-0.020	-0.026
			(0.057)	(0.062)			(0.065)	(0.071)
t=7			-0.050	-0.055			-0.011	-0.019
			(0.055)	(0.060)			(0.050)	(0.062)
t=8			-0.038	-0.038			-0.003	-0.007
			(0.048)	(0.054)			(0.049)	(0.065)
t=9				-0.037				-0.020
				(0.047)				(0.068)
t=10				-0.023				-0.029
				(0.043)				(0.058)
t=11				-0.045				-0.010
				(0.040)				(0.045)
t=12				-0.061*				-0.037
				(0.036)				(0.037)
N	108,360	108,360	108,360	108,360	47,040	47,040	47,040	47,040

Notes. See Table 3.