

UC Berkeley

Working Paper Series

Title

Credible Research Designs for Minimum Wage Studies

Permalink

<https://escholarship.org/uc/item/3hk7s3fw>

Authors

Allegretto, Sylvia
Dube, Arindrajit
Reich, Michael
[et al.](#)

Publication Date

2013-09-23

Credible Research Designs for Minimum Wage Studies

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

September 23, 2013

Abstract

We assess alternative research designs for minimum wage studies. States in the U.S. with larger minimum wage increases differ from others in business cycle severity, increased inequality and polarization, political economy, and regional distribution. The resulting time-varying heterogeneity biases the canonical two-way fixed effects estimator. We consider alternatives including border discontinuity designs, dynamic panel data models, and the synthetic control estimator. Results from four datasets and six approaches all suggest employment effects are small. Covariates are more similar in neighboring counties, and the synthetic control estimator assigns greater weights to nearby donors. These findings also support using local area 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: Department of Economics, University of Massachusetts Amherst. We are grateful to Zachary Goldman, Thomas Peake and Luke Reidenbach for excellent research assistance. Financial support for this paper came entirely from the University of California, Berkeley and the University of Massachusetts Amherst.

1 Introduction

Recent discussions of the employment effects of minimum wage policies have centered on the issue of appropriate research design, with special attention to the desirability of local area controls. These issues are important because the non-random distribution of state minimum wage policies in the U.S. poses a serious threat to identifying the policy’s effects. Indeed, 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. Recent minimum wage research has used regional controls and policy discontinuities to control for such heterogeneity—e.g., Dube, Lester and Reich (2010), Allegretto Dube and Reich (2011), and Magruder (2013). Such designs are well-established in the discipline and constitute well-integrated parts of the “credibility revolution” that has swept through much of labor and applied microeconomics. Nonetheless, the use of local area controls in the minimum wage context is not universally accepted. In particular, a 2013 paper by Neumark, Salas and Wascher challenges this approach.

To advance the discussion of this issue, we present here a comprehensive assessment of the design of credible control groups for minimum wage studies.¹ We begin by showing that states experiencing greater increases in minimum wages over the past two decades have systematically different 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 two-way fixed effects model (with common period fixed effects) is likely to mis-estimate the counterfactual employment growth absent a minimum wage increase.

¹Appendix B of this paper presents detailed responses to Neumark, Salas, and Wascher (2013a), including critiques of their proposed estimators. We note that (2013a) refers to the publication version of their paper. Since some of the details of their argument are only present in the working paper version (2013b), we specifically refer to the latter where relevant.

Next, in sections 3 and 4, we present new estimates for both teens and restaurant workers with updated data at least through 2010 from four datasets. Here we utilize three spatial approaches to accounting for time-varying heterogeneity: comparing across contiguous counties, comparing within commuting zones, and using within-region variation along with state-level trend controls. These approaches rely on the insight that, on average, nearby areas tend to be more similar. We first show that contiguous counties are indeed better control groups in the sense of having more similar covariates. Turning to the minimum wage impact, our consistent finding is that the inclusion of 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 specific approach employed. After accounting for spatial heterogeneity, the employment elasticities for teens and restaurant workers range between -0.06 and +0.13. In contrast to the results for employment *stocks*, our results indicate that 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 were hypothesized by Card and Krueger (1995) to explain the lack of minimum wage employment effects.

We also illustrate that the canonical fixed effects model consistently exhibits substantial pre-existing employment trends. Since the minimum wage effect necessarily occurs during the year of or after the wage increase, the canonical model therefore fails an important falsification test. These pre-existing trends disappear once we include spatial controls, thereby providing additional internal validity to our research design.

While geographic controls offer one way of accounting for time-varying heterogeneity, we consider two additional approaches as well. In Section 5, we estimate specifications that include lagged dependent variables (LDV), which are estimated both with and without state fixed effects. In the latter case, the resulting dynamic panel data models are estimated using the GMM approaches of Arellano and Bond (1991), and Blundell and Bond (1998). We also

report OLS estimates from a model with lagged outcome and fixed effects using likely bounds on the autoregressive coefficient. The results from these regressions show that the inclusion of lagged outcomes as covariates renders the minimum wage employment estimate small in magnitude—much like the inclusion of local area controls. As such, they show that spatial and temporal lags offer alternative ways of accounting for the time-varying heterogeneity that contaminates the standard two-way fixed effects model. Our analysis using alternative values of the autoregressive coefficient encompasses the special case where minimum wages affect employment growth (as opposed to employment levels), as considered in Meer and West (2013). We find that for the two key low-wage groups, there is no evidence that minimum wages reduce employment growth.

The synthetic control approach of Abadie et al. (2010) provides yet another useful way to construct control groups. Using a pooled synthetic control estimator, in Section 6 we provide new evidence on minimum wage elasticities for teen employment. Our estimates 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. The pooled estimates from the nineteen resulting events show a clear wage effect, with a mean minimum wage elasticity of +0.17 that is statistically significant and within the range of estimates in our previous papers. However, we do not detect any employment losses: the mean minimum wage elasticity for employment is close to zero (-0.03). These results are also virtually identical to results we obtain for teens using local controls. In Section 6 we also compare the overlap between synthetic and local controls using state-level data. The results clearly show that synthetic control donor weights are indeed larger for nearby areas. We obtain this result for teen employment with 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 non-parametric 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 six times as large as a state 1,000 miles

away.

Our paper relates to two strands of recent research. The first concerns the use of local area controls and border discontinuity designs generally in estimating the impact of policies. The second concerns the relevance of such controls in minimum wage studies. The use of regional controls and policy discontinuities to control for heterogeneity constitutes part of the profession’s tool kit. Holmes (1998) and Huang (2008) use the contiguous border county approach to identify the effects of right-to-work legislation and bank regulation, respectively. Dube, Lester and Reich (2010, 2013) adopt a similar approach to study minimum wage effects, as do Hagedorn et al. (2013) in the context of unemployment insurance effects. Many other recent papers utilize a border discontinuity approach, including Black (1999) in the context of schools and housing markets, Goldstein and Udry (2008) in the context of land rights and investment in Ghana, and Magruder (2013) in the context of minimum wages in Indonesia. 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. Allegretto, Dube and Reich (2011) use identifying assumptions that are identical to those in Autor (2003).

In the minimum wage context, the application of local controls began with the local case study approach of Card and Krueger (1994, 1995, 2000), who compared fast-food restaurants across the New-Jersey Pennsylvania border. Dube, Naidu and Reich (2007) implemented a similar design to study a citywide minimum wage in San Francisco. Both of these local case studies find no disemployment effects of minimum wages. In contrast, multiple previous studies (for example, those reviewed in Neumark and Wascher 2008) use a two-way fixed effects model to control for state and time effects (which we will refer to here as the canonical model) and find substantial and significant negative employment effects for teens.

In previously published work—Dube, Lester and Reich (2010); Allegretto, Dube and Reich (2011)—we have used a variety of spatial controls and data on the two most affected

groups, teens and restaurant workers. In relationship to those studies, this paper makes five contributions. (1) We provide new evidence on how high and low minimum wage states are different. (2) We document the similarity of adjacent areas, which supports the validity of local controls as a way to control for time-varying heterogeneity. (3) We use expanded samples and a common framework to show results using local area controls and four datasets: the American Community Survey (ACS)/Census, the Current Population Survey (CPS), the Quarterly Census of Employment and Wages (QCEW), and the Quarterly Workforce Indicators (QWI). We show results using a new approach: cross-state commuting zones for the ACS/Census.² (4) We provide new evidence using non-spatial approaches to controlling for time-varying heterogeneity, including the dynamic panel data approach—to our knowledge the first such implementation in the minimum wage literature. We show how spatial and temporal lags offer alternative approaches to controlling for time-varying heterogeneity yet at the same time produce similar estimates of minimum wage employment elasticities. (5) We assess the comparability of local and synthetic controls, and we report estimates using a pooled synthetic control estimator.

In conclusion, we assess the economic significance of the range of employment estimates across different empirical strategies. We also propose some guidelines for evaluating what constitutes good research designs for minimum wage studies, along with some trade-offs associated with alternative approaches.

2 Minimum wages and time-varying heterogeneity

2.1 Federal and state minimum wage policy

The Fair Labor Standards Act of 1938 established a single national minimum wage floor, with individual states free to set higher minimum wages. From the 1950s through the 1970s, when

²Earlier results using commuting zones were included in a working paper version—Allegretto, Dube and Reich (2009)—but have not been used in our previously published work.

the ratio of the federal minimum wage to the median U.S. wage hovered mainly between 45 and 55 percent, states did not use this option. But during the long spell of federal inaction from 1981 to 1989, when the minimum to median wage ratio fell to 36 percent, states began to raise their minimums for the first time.³ In subsequent decades, when the minimum to median ratio hovered between 33 percent and 40 percent, an increasing number of states have set higher standards, leading to greater variation in minimum wages over time. As Figure 1 shows, increased activism at the state level has led to greater variation in minimum wages between 1984 and 2012, and especially so since the 1990s. In addition, ten states now index their minimum wage to inflation and five cities—San Francisco and San Jose, CA, Albuquerque and Santa Fe, NM and Washington, D.C.—have established citywide minimum wages.

2.2 How high and low minimum wage states differ

Along with greater variation in the statutory minimum wage, regional clustering of minimum wage policies has also increased since the early 1990s. To show the regional clustering of minimum wage policy 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 2012 was above or below the nominal median (\$5.33). Since federal minimum wage increases typically erase most 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 2 provides two maps showing the spatial distribution of minimum wage states that have high average levels and high variance, showing a considerable degree of overlap (about 85 percent of states fall into both categories). While such states are present in every

³The median wage (for full-time workers only) is available from 1960 on in the Bureau of Labor Statistics’ *Employment and Earnings Reports*. These are reported in OECD (2013). Whittaker (2003) reports the ratio of the minimum wage to the average hourly earnings of nonsupervisory production workers in manufacturing for 1941 to 2003. The two series match very closely for the 43 years included in both.

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, between the business cycle peaks of 1990 and 2007, 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 compared to a more modest increase, 2.15 to 2.22, in the low minimum wage states. 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 in-

equality and labor demand for low-skill workers. A simple skill-biased technical change model predicts lower growth in relative 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 3 and Table 1 show that the Routine Task Intensity (RTI) index of occupations fell more in high minimum wage states.⁴ 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 complex. In the two-skill-group model of Autor and Dorn (2013), the relationship

⁴We 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 $\ln(\textit{routine}) - \ln(\textit{manual}) - \ln(\textit{abstract})$. We calculate the employment-weighted means of this measure by state and year.

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 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. The two-way model 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 markets suggests using within-local-area comparisons to identify minimum wage effects that are not contaminated by structural differences among labor markets. In this paper, when we specifically consider variation in the minimum wage within commuting zones, we precisely account 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 3 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, as 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 do not occur uniformly over the course of the business cycle. Between 1990 and 2012, changes in the prevailing minimum wage in a given state (i.e., the maximum of state or federal minimum wage) were much more likely to occur in the second half of an economic expansion (16.9 percent) than during the first half (5.5 percent) or a recession (9.9 percent). Figure 4 shows these probabilities, and associated 95 percent confidence intervals. Figure 4 also shows the analogous probabilities of minimum wage increases for the pre-Great Recession sample, which excludes the federal minimum wage increase during 2008 and 2009. In this pre-2007q4 sample, which begins and ends with business cycle peaks, the probability of a minimum wage increase occurring during late expansion, early expansion, and recession were 15.0, 4.5 and 2.5 percent, respectively. In other words, minimum wage increases were consistently more likely to occur during late expansions. Overall, the cyclical nature of policy change, together with the differential nature of business cycles in high and low minimum wage states, suggest that it is important to construct control groups to account for time-varying heterogeneity in low-wage employment growth.

In addition to the substantial 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 3). 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 3, 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.

3 Econometric specifications, data and samples

In this section, we describe the econometric specification, data, and samples for both the canonical two-way fixed effects model and the models using spatial controls. 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 j or an individual in that place depending on the data:

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

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 elasticities using our preferred border discontinuity approach. We allow the time effects, τ_{gt} , to vary at the geographic level g , where g can be (1) county pair, or (2) commuting zone depending on the dataset:

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

The inclusion of the commuting zone, or county-pair, with a specific time effect sweeps out all the variation between local areas g , and only uses variation within local areas surrounding a policy border. This border discontinuity specification allows us to control for time-varying heterogeneity in the outcomes across local areas. Unbiased estimates using the canonical model (1) require the strong assumption that minimum wage differences between *any* lo-

cations j are uncorrelated with residual outcomes. In contrast, the spatial controls τ_{gt} in model (2) significantly weaken this assumption, only requiring it to hold for any locations within a given local area, g , around the state border. This approach allows us to control for the types of spatial heterogeneity in regional economic shocks we discussed in section 3.2. In the county pair case, since a single county can be a part of multiple cross-border pairs, the data is stacked by pairs; the standard errors are clustered by state and by border pair to account for multiple instances of counties in the dataset.

When the data do not permit a fine grained geographic control, the group g is set at a coarser level of the nine Census-defined divisions. In this case, we modify equation (2) to additionally include state-specific time trends:

$$Y_{it} = \alpha + \beta MW_{jt} + \mathbf{X}_{jt}\Lambda + \gamma_j + \tau_{gt} + \phi_j \times t + \nu_{it} \quad (3)$$

We use three different datasets to study minimum wage effects among teens and two different datasets to study minimum wage effects among restaurant workers. These two groups have been extensively studied in the literature because they are heavily affected by minimum wage policies. Teens, for example, are disproportionately likely to be minimum wage workers. Based on the Current Population Survey, during the 2000-2011 period 29.8 percent of working teens earned within 10 percent of the minimum wage. And teens comprised 25.2 percent of all workers earning within 10 percent of the minimum wage. The second high-impact group consists of establishments in the restaurant industry. During the same period, restaurants employed 24.3 percent of all workers paid within ten percent of the state/federal minimum wage, making restaurants the single largest employer of minimum wage workers at the 3-digit industry level. Restaurants are also the most intensive user of minimum wage workers, with 22.8 percent of restaurant workers earning within ten percent of the minimum wage (using 3-digit level industry data).⁵

⁵We use the 2000-2011 period to calculate these proportions since it represents the overlap between the teen sample (from the CPS) and the restaurant sample (from the QWI). Additionally, using the post-1999 sample allows the use of consistent industry coding in the CPS.

The specific geographic controls depend on data availability. The first of our three approaches for teens uses individual-level data from the decennial Census and its successor, the American Community Survey (ACS), along with variation within local labor markets to examine minimum wage effects on teen employment. Here we use the 1990 and 2000 decennial Census and ACS data from 2005-2011. Following Autor and Dorn (2013), 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. 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 Washington D.C.). 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.⁶ Panel A of Figure 5 shows a map of the counties constituting the cross-state commuting zone sample and identifies those that have minimum wage variation. For the Census/ACS specifications, the place fixed effects γ_j consist of state-specific commuting-zones. In the canonical model, common time effects τ_t are simply year dummies. But in the spatial controls specifications, we replace these with commuting-zone-specific year effects τ_{CZt} .

The second of our approaches uses individual-level Current Population Survey (CPS) panel data on teens between 1990 and 2012. Here we use quarterly minimum wage data and, with the canonical model, quarterly common time effects. We control for spatial heterogeneity at the level of the nine Census-defined divisions, using Census division-specific time effects τ_{dt} and also including state-specific linear trends. Sample sizes and a lack of detailed geographic information in these household survey data preclude finer spatial controls.

The third approach for studying teens uses a relatively new dataset, the Quarterly Work-

⁶Similar 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 assigned PUMAs.

force Indicators (QWI), for 2000q1 to 2011q4.⁷ Unlike the Census/ACS and the CPS, the QWI is an employer-based dataset, drawn from employer-reported administrative Unemployment Insurance records. These administrative data include 98 percent of private employment for 49 states and the District of Columbia (Massachusetts did not begin to participate until 2013). The Census Bureau uses other data—primarily from Social Security records—to either match or impute demographic information of workers. The underlying datasets consequently are much larger than the CPS. 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 and analyses of both employment stocks and employment flows. The QWI age categories identify teens as those between 14 and 18 years of age. For detailed documentation of the QWI, see Abowd et al. (2009).⁸ We limit the QWI sample to counties along state borders since our preferred specification uses the border discontinuity design, operationalized by including county pair-specific time effects τ_{gt} as in equation (2).

Our QWI sample therefore consists of the 1,130 counties that border another state. Collectively, these border counties comprise 1,181 unique county pairs. Panel B of Figure 5 provides a map of the border county sample, showing those with minimum wage variation. While most counties in the border pair sample are geographically proximate, counties in the western U.S. are much larger in size and irregular in shape. In some cases the geographic centroids of the counties in such pairs lie several hundred miles apart. For this reason, we exclude counties whose centroids are more than 75 miles apart.⁹

⁷The majority of states entered the QWI program between the late 1990s and early 2000s. Abowd and Vilhuber (2011) note that there are “differences in data quality between the 1990s and 2000s . . . due to the inclusion of 30 states beyond the original 18 included in the 2003 initial release of the QWI.” Moreover, the states were non-randomly missing: for example, large states were over-represented in early years. For these reasons, we use data from the 2000s in our analysis; by 2000, 42 states had come on line. Our 2000-2011 sample represents 77 percent of the observations in the 1990-2011 period.

⁸Abowd and Vilhuber (2011) provide an extensive comparison of the QWI to CPS and JOLTS datasets. In Abraham’s (2009) assessment of the quality of the QWI data the only major issue concerns imputed levels of education, which are not pertinent here. The QWI does not contain data on employee hours. Abraham et al. 2013 find that although the CPS data are monthly, the QWI captures many more short-term jobs.

⁹A smaller distance cutoff trades lower error variance from greater similarity for higher error variance

State minimum wage policies varied considerably during the 2000-2011 period, creating substantial variation, especially between 2004 and 2009. This period includes the three steps of the 2007-2009 federal minimum wage increases and many state-level changes. There are 196 incidents of quarter-over-quarter minimum wage increases when we pool across federal and local policy changes; and 70 percent of the sample border counties had some minimum wage variation with its contiguous pair. Limiting our attention to cross-border comparisons thus provides us with sizable policy variation that we can use for estimating minimum wage effects.

QWI data are always “fuzzed” to protect confidentiality thus a small amount of noise is added to the establishment-level data prior to aggregation at the county level. In a small number of cases, data for some cells are suppressed entirely. To avoid a sample selection bias, our estimates are for the set of counties that do not suppress the data for that relevant outcome; this reduces the sample by between 1 and 14 percent depending on the outcome.

For restaurant workers, we use both the Quarterly Census of Employment and Wages (QCEW) and the QWI datasets. The QCEW is also an employer-based dataset. Like the QWI, the QCEW provides payroll head counts and monthly earnings for detailed industries from Unemployment Insurance records. It also includes 98 percent of private sector employment, but no demographic data are included.¹⁰ Unlike the QWI, the QCEW data are available for all states during the 1990s. For this reason, our QCEW sample begins in 1990. It ends in 2010 because of an industry coding change in 2011. We construct a panel of quarterly observations of county-level employment and earnings for restaurants (NAICS 722). Some counties in the QCEW contain too few restaurants to satisfy nondisclosure requirements in every year. The QCEW does not use the fuzzing method of the QWI and it

from a smaller sample. The exact choice of cutoff was based on a data-driven randomization inference procedure that minimized the mean-squared error (MSE) of the estimator in the border sample using placebo treatments. This criterion retains about 81 percent of the sample, eliminating mostly western counties. To show that our results are not affected by the choice of cutoffs, Appendix Table A1 of Dube, Lester and Reich (2013) reports the key results with cutoffs ranging between 45 and 95 miles.

¹⁰As in the QWI, the 2 percent who are not covered are primarily certain agricultural, domestic, railroad, and religious workers.

has a higher rate of non-disclosed observations than the QWI. We exclude counties with any non-disclosure data issues because observations for these counties may be selected out of the sample if employment becomes lower because the minimum wage is high. However, including the counties with partial information did not affect their results, and that conclusion holds in this extended sample as well (results not reported). Panel C of Figure 5 provides a map of the QCEW sample, indicating pairs with minimum wage variation, and the counties that are part of a balanced panel.¹¹

The QWI data we use for restaurant workers is the same as for teens, covering the period 2000q1 to 2011q4. For the QCEW data, when estimating the canonical model we use all counties. For results using spatial controls, we limit the QCEW or QWI sample to counties along state borders, and we include county pair-specific time effects τ_{pt} .

4 Results using spatial controls

We first directly demonstrate the desirability of using local controls by showing that neighboring county pairs are more alike than are other pairs in levels and trends of key covariates. Then 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.

4.1 Similarity of local areas: Are contiguous pairs more alike?

The border discontinuity approach is predicated on the assumption that neighboring areas are good controls. Here 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.¹² To examine whether local controls are indeed more similar, we consider six key co-

¹¹For comparability with the original results in Dube, Lester and Reich (2010), we do not restrict the QCEW sample based on centroid-to-centroid distance. However, the sample with such additional restrictions produces very similar results, as reported in this paper’s footnote 16.

¹²For this exercise, we use the QWI dataset instead of the QCEW because the former provides more information, in particular the rate of employee turnover. These results are reproduced from Dube, Lester and Reich (2013).

variates: 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. 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. 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 2 shows the results for these variables in levels, as well as 4 and 12 quarter changes. In all cases, the mean absolute differences are larger for non-contiguous pairs; and in all cases but one, the gaps are statistically significant at the 1 percent level. The average percentage gap in absolute differences for the six variables in levels, and 4 quarter, and 12 quarter changes is 22.7 percent. Many of the gaps are substantial (above 25 percent): notably, for levels of employment, population and earnings; for 4 quarter change in the EPOP, and for 12 quarter changes in the EPOP and the turnover rate. We conclude that cross-border counties 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.

4.2 Main findings

In this section, we provide new results using data that incorporates the variation in minimum wages through 2010 or 2012, depending on the dataset. We begin with the results for teens and then discuss the results for restaurant workers. Both sets of results are presented in Table 3.

4.2.1 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 likely to detect an impact—whatever it may be—for this demographic group. We use three datasets to estimate earnings and employment elasticities for teens, with the results reported in columns 1, 2 and 4 of Table 3 for the “canonical” model; and in columns 5, 6 and 8 for our preferred models with local controls. Results using individual level data from the Census and the ACS are reported in columns 1 and 5. These regressions include fixed effects for each CZ-by-state group.¹³ The “canonical” specification (column 1) includes common year effects, which assumes that different CZs have parallel trends conditional on covariates. The “local” specification (column 5) includes CZ-by-year fixed effects, which means it only uses within-CZ variation to identify minimum wage effects. Since the data are annual, we use annual average minimum wages for each state.

The second set of results is based on individual-level data from the CPS. All regressions include state fixed effects. The canonical model (column 2) uses a common set of time dummies, measured in quarters. In contrast, the preferred model (column 6) includes Census-division specific time effects along with state-specific linear trends to account for spatial heterogeneity.¹⁴ The third set of results for teens uses county-level data from the QWI,

¹³The regressions also include standard demographic controls—dummies for age, sex, race/Hispanic status, marital status, and educational attainment—as well as the annualized state-level unemployment rate.

¹⁴All CPS regressions also include dummies for sex, race/Hispanic status, age, educational attainment, marital status, along with the teen share of the population in the state and the non-seasonally adjusted quarterly state unemployment rate.

restricted to border counties and stacked by county pairs. Column 4 reports the estimates from the canonical specification with common time effects, while the local specification of column 8 includes county-pair specific time effects.¹⁵

The results from all three of these datasets indicate that the earnings elasticities for teens are substantial, ranging between 0.12 and 0.30 depending on the data and the model. Comparing within datasets, adding local controls (columns 5, 6 and 8) does not reduce the estimated effects on teen earnings; if anything, the estimates are somewhat larger with local controls. This evidence clearly indicates that minimum wages increases substantially raise average teen wages—even when we compare areas right across the border. For employment, the canonical estimates are statistically significant at the 5 percent level, and range between -0.12 and -0.16 depending on the dataset. However, in each of these cases, the inclusion of local controls makes these estimate less negative; the preferred elasticities 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.

4.2.2 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. As discussed above, we present results using two data sets: the QCEW and the QWI. We present results from the QCEW that include data from 1990q1 through 2010q4. The QWI results are reproduced from (Dube, Lester and Reich, 2013), which uses data from 2000q1 through 2011q4.¹⁶

¹⁵The QWI regressions control for log of overall population and private-sector employment or earnings. These results are reproduced from Dube, Lester and Reich (2013).

¹⁶The results we report in the text follow Dube, Lester and Reich (2010) in using balanced panels of counties that report data for all periods. Panel C of Figure 5 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. To address concerns with selectivity of the sample, we estimated elasticities using balanced and unbalanced samples of counties. The results are quantitatively similar: using the border discontinuity

The lower panel of Table 3 shows the results from the QCEW and the QWI—in columns 3 and 7, and 4 and 8, respectively. All regressions include county fixed effects. Columns 3 and 4 are estimated using common time effects, while columns 7 and 8 include county-pair specific time effects. The results for earnings range between 0.19 and 0.21, are quite similar to those for teens, and they are unchanged by adding the spatial controls. Regarding employment, the canonical model tends to produce more negative estimates, while the inclusion of spatial controls reduces the 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 3 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 from our preferred specifications are small in magnitude and statistically indistinguishable from zero. We return to the issue of the size of the employment elasticities in our concluding discussion in Section 7.

4.2.3 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. For example, according

research design, QCEW restaurant employment (earnings) elasticities are 0.008 (0.186) and -0.034 (0.196) for the balanced and unbalanced sample, respectively. In both cases, the employment estimates not statistically significant at the 10 percent level, while earnings estimates are significant at the 1 percent level. Also, for comparability with Dube, Lester and Reich (2010), we do not restrict counties based on centroid-to-centroid distance using the QCEW data, as done with the QWI data in Dube, Lester and Reich (2013) and in this paper. However, the results are quite similar when we restrict the sample to counties whose centroids are within 75 miles of each other. When using the border discontinuity design, the employment (earnings) elasticity changes from 0.008 (0.186) to -0.021 (0.194) with the added centroid-distance restriction to the baseline balanced panel (this restriction reduces the number of county pairs in the sample from 280 to 265).

to QWI data, the quarterly turnover rate for restaurants in our border county sample was 42 percent. More recently, Manning (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 (Ashenfelter et al. 2010). However, none of these studies directly test the effects of minimum wages on market-level employment flows with representative data.¹⁷

Columns 4 and 8 in Table 3 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. In our local specification with county-pair specific time effects—column 8—the magnitude of the hires and separations elasticities are smaller, but they remain substantial at -0.22 and -0.23, respectively; and these estimates are statistically significant. In the case of restaurant workers, the results for each specification are very similar to those for teens; and in our preferred specification (column 8, panel B) the hires and separation elasticities are -0.26 and -0.23, respectively.

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 Neumark, Salas, and Wascher (2013a) that our spatial controls discard too much variation to find any significant effects. Dube, Lester and Reich (2013) also show that these results—a substantial positive wage effect, a small effect on employment stock, and a substantial negative effect on employment flows—can be explained using models with search friction.¹⁸

¹⁷For similar findings with Portuguese and Canadian data, see Portugal and Cardoso (2006) and Brochu and Green (2013), respectively.

¹⁸The job-ladder model highlights the quits channel with reduced employment-to-employment transitions. A complementary explanation suggests 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 (2013), who emphasize the reduced rate of employment-to-nonemployment transitions. Dube, Lester and Reich (2013) provides a detailed assessment of these two explanations.

4.3 Internal validity of alternative specifications: Pre-existing trends

Do the systematic differences documented above impart a bias in the canonical two-way fixed effects panel estimator? We use unstructured distributed lags to test whether the two-way fixed effects model exhibits pre-existing negative employment trends. We present these results using a common framework for annual QCEW and CPS data over the period 1990-2010.¹⁹

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

$$Y_{it} = \sum_{k=0}^K \beta_{-k} MW_{i,t-k} + \mathbf{X}_{it}\Lambda + \gamma_i + \tau_{gt} + \nu_{it} \quad (4)$$

for different values of K separately. As before, the canonical and local specifications have different types of time effects τ_{gt} — in the case of the canonical model, these are simply common time effects τ_t . The sum of the leading coefficients, $\sum_{k=0}^j \beta_{-k}$, represents the cumulative response—or the time path—of the outcome Y at j periods prior to a hypothetical 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 . Figure 6 displays the time paths, one for each of the three lead lengths.²⁰

¹⁹These are the most recent common samples when using three years of leading minimum wages, since the most recent minimum wage information ends in 2013.

²⁰There is a tradeoff between smoothness and precision in choosing the frequency of leads. In Dube, Lester and Reich (2010) and Allegretto, Dube and Reich (2011), we used semi-annual leads and quarterly data as a way to balance these concerns. To allay concerns about choice of frequencies, here we report results using quarterly and annual frequencies in the data and the lead structure, covering between 1 and 3 years.

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 sizable, 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 sizable coefficients for sums of the leading coefficients, as these sums represent the cumulative response of the outcome to a minimum wage increase j periods *prior* to that 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 one, two or even three years prior to the policy change strongly suggests a failure of the falsification test.

We begin by summarizing the findings visually in Figure 6. The upper-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 additional control.²¹ For both of these data sources, we show all annual cumulative responses for all lead lengths $K = 1, 2, \text{ or } 3$, through the contemporaneous effect at time $j = 0$. The canonical model unambiguously produces leading effects well before the time of treatment. These spurious effects are particularly pronounced with the CPS data. The upper 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

²¹As 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).

estimates with the canonical model. Appendix Table A1 reports the estimates underlying Figure 6.

Although noisy, especially with CPS data, the quarterly estimates in the lower panels of Figure 6 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 evidence demonstrates that specifications with local controls outperform the canonical model.

5 Lagged dependent variables

One way to relax the parallel trends assumption of the fixed effects estimator is by allowing for heterogeneous time effects by local areas, such as county pairs, commuting zones, or regions, or using parametric state-specific trend controls. We pursued those approaches in the previous section. An alternative approach directly controls for past outcomes in the panel regression to account for time-varying confounders. For example, if minimum wage increases occur during periods with unusually low and/or falling employment, then controlling for lagged employment offers a way to account for such a pattern. The findings in section 4.3 of unusually low employment prior to minimum wage increases point us in this direction. Lagged outcomes, then, offer an alternative to outcomes from local areas as a proxy for time-varying confounders.²²

In this section, we estimate state panel models for teens and restaurant employment using lagged dependent variables (LDVs). For teens, we use annual state-level CPS data during 1990-2012, where the state-year means are weighted by sample weights and the regressions

²²For more on fixed effects versus lagged dependent variables, including the bracketing property of the two estimators, see Angrist and Pischke 2009, pp. 243-6.

are weighted by teen population. For restaurants, we use annual county-level QCEW data during 1990-2010. The estimating equation with LDVs is

$$Y_{it} = \alpha + \beta MW_{i,t} + \rho Y_{it-1} + \mathbf{X}_{it}\Lambda + \tau_t + \nu_{it} \quad (5)$$

Note that equation (5) does not include state fixed effects. We can also estimate the model with LDV and fixed effects using OLS:

$$Y_{it} = \alpha + \beta MW_{i,t} + \rho Y_{it-1} + \mathbf{X}_{it}\Lambda + \gamma_i + \tau_t + \nu_{it} \quad (6)$$

However, it is well known that the inclusion of both fixed effects and lagged dependent variables renders the estimates inconsistent (e.g., Nickell 1981). In particular, the naive OLS estimate of ρ from equation (6) is biased towards zero, so it can serve as a lower bound for the true ρ . With the autoregressive structure considered here, in large samples the true ρ is also bounded from above by the estimate from equation (5) without any fixed effects (see Bond 2002). Therefore, we provide a bound on β by imposing a value of ρ in this credible range, i.e., falling between the OLS estimates from equations (5) and (6). Starting with an assumed ρ in this range, we quasi-difference the outcome $\tilde{Y}_{it} = Y_{it} - \rho Y_{it-1}$, and regress \tilde{Y}_{it} on MW_{it} along with controls \mathbf{X}_{it} and place and time fixed effects.

We also estimate equation (6) using GMM that can accommodate both fixed effects and LDVs. In particular, longer lags of the outcome in levels and differences are used as instruments, although identification requires that the idiosyncratic component of the disturbances is not too autocorrelated. We estimate two variants of the dynamic panel data (DPD) model, using both the the Arellano and Bond (1991) difference GMM approach as well as the Blundell and Bond (1998) system GMM approach. The former can suffer from a weak instrument problem, while the latter requires additional assumptions of stationarity (Roodman 2009).²³ The fixed effects are purged using forward-orthogonalization or first

²³The calculations are implemented in Stata using the `xtabond2` command. In principle can use all lags of

differencing, depending on the specification, and we treat all the variables other than the lagged outcomes as exogenous.²⁴

The first column in Table 4 reproduces the results of the FE estimate using the annual employment data.²⁵ The teen sample includes controls for the overall state-level unemployment rate, race and gender composition and teen population shares; the restaurant sample includes controls for population as well as overall private sector employment at the county level. The point estimates for both teens and restaurant workers are around -0.13, similar to the analogous results in Table 3. Next, in column 2, we report the estimates controlling for the lagged outcome (but not fixed effects). 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. In the restaurant sample, the point estimate falls to -0.05, although it is statistically significant given the small standard errors. (Generally, in the restaurant sample the inclusion of a lagged dependent variable dramatically reduces error variance, often leading to statistical significance of economically small estimates.) In column 3, we report the naive OLS estimate of equation (6); as discussed above the estimated ρ here is a lower bound for true ρ . In contrast, column 4 reports the results imposing the likely upper bound of ρ , as estimated from column 2 without the inclusion of fixed effects. The likely range of ρ for teens is therefore (0.323, 0.713), while for restaurant workers it is (0.752, 0.970). Within this likely range of ρ , the minimum wage elasticity β is small, ranging between -0.076 and -0.004

the dependent variable (in levels and differences) as instruments. In many specifications, the Arellano-Bond tests show order-2 serial correlation in $\Delta\nu_{it}$. For this reason, we limit the instrument set to lags of order 3 and greater. Imposing this limit has little effect on our estimate of the minimum wage elasticities for employment. We also collapse the instruments by lag period to mitigate the many-instruments problem, as discussed in Roodman (2009).

²⁴Forward orthogonalization entails subtracting the mean of all future observations within the state as a way to purge the state fixed effect. This technique was first proposed by Arellano and Bover (1995) to ensure that the lagged instruments are uncorrelated with the transformed data, which would not be the case when using the standard mean-differencing to purge the fixed effects. As an alternative, one can also use first-differencing to transform the data, although this approach tends to put more weight on shorter run impact on the policy. We show results using both methods.

²⁵Here we focus on employment only, and not earnings, since time-varying heterogeneity appears to mainly contaminate the employment estimates: the earnings estimates using the local and the canonical estimators in Table 3 are not substantially different and there is little evidence of pre-existing trends. For completeness, however, in Appendix Table A2 we report the earnings elasticities using analogous specifications with lagged outcomes.

for teens and -0.030 and -0.000 for restaurant workers. Figure 7 shows how the minimum wage elasticity for employment declines in magnitude with the importance of the lagged dependent variable: the greater the time-varying component of heterogeneity, the larger the downward bias in the minimum wage elasticity of employment. Within the credible range of ρ , as demarcated in grey in Figure 7, the point estimates for the employment elasticity are uniformly under -0.1 in magnitude for both teen and restaurant samples.

We also note that the results in Figure 7 subsumes the case ($\rho=1$) in which minimum wages affect net employment growth, as opposed to levels. This is the estimator used by Meer and West (2013), who find a negative association between *overall* private sector employment growth and minimum wages. However, as Figure 7 shows, for these two low-wage groups there is no indication in the $\rho = 1$ case that minimum wages lead to lower net growth in jobs. The point estimates for this case are both close to zero (0.004 for restaurants and 0.049 for teens) and not statistically significant.²⁶ These null findings from low-wage groups actually affected by the policy suggests that the negative association found by Meer and West is unlikely to be causal, perhaps the result of their lack of controls for the overall state of the labor market.

Columns 5-8 of Table 4 report the results from the dynamic panel models using GMM. Both the system GMM (columns 5, 7) and difference GMM (columns 6, 8) produce minimum wage elasticities that are no more negative than -0.04. This result holds whether we use difference or forward transforms to purge the fixed effects. We do note that some of the estimates (especially column 8) of the autoregressive coefficient ρ are well outside the credible range based on estimates in columns (2, 3). However, when they are either contained in or are close to that range, the minimum wage estimates are similarly small in magnitude, and well within the range indicated in columns (3, 4) and Figure 7.

Overall, all three methods point to small employment effects: the simple regression with only lagged employment and no fixed effects, the bounds using OLS estimation with lagged

²⁶Figure 7 also shows that the 95 percent confidence intervals for these estimates rule out effects more negative than -0.020 and -0.015 for restaurants and teens, respectively.

employment and fixed effects, and the dynamic panel models estimated using GMM. These results using lagged dependent variables provide still another direct demonstration of how time-varying heterogeneity affects the estimation of minimum wage effects.

6 Evidence using the synthetic control estimator

The synthetic control approach of Abadie et al. (2010) offers yet another way to account for time-varying heterogeneity that may contaminate the estimation of the treatment effect. Unlike a two-way fixed effects regression, the data-generating process underlying the synthetic control approach allows multiple time factors with different factor loadings across states. Denoting the intervention as D , we can write the observed outcome Y as a sum of the treatment effect and the counterfactual outcome, Y^N :

$$Y_{jt} = \alpha_{jt}D_{jt} + Y_{jt}^N = \alpha_{1t}D_{it} + \mathbf{X}_{it}\Lambda + \boldsymbol{\lambda}_t\boldsymbol{\gamma}_i + \tau_t + \epsilon_{it} \quad (7)$$

As compared to the FE model of equation 1, there is an additional $\boldsymbol{\lambda}_t\boldsymbol{\gamma}_i$ term in the equation above. Here $\boldsymbol{\lambda}_t$ is a vector of unobserved time-varying factors and $\boldsymbol{\gamma}_i$ are the unknown factor loadings. Since the factor loadings can vary across states, treatment and control states need not follow parallel trends, conditional on observables. To estimate equation (7) for a single treated state (denoted by $j = 1$), the synthetic control procedure constructs a vector of weights \mathbf{W} over J untreated donor states, such that the weighted combination of donor states closely matches the treated state in pre-intervention outcomes, thereby balancing the average factor loadings in the treated and synthetic control groups. With such a vector of donor weights, \mathbf{W} , the synthetic control estimate for the effect of the treatment can be written as

$$\hat{\alpha}_{1t} = Y_{1t} - \sum_{j=2}^{J+1} W_j Y_{jt}$$

To operationalize the estimation, the algorithm chooses \mathbf{W} to minimize the mean squared

error in pre-treatment characteristics between the treated state (\mathbf{Z}_1) and the weighted average of these characteristics among the donor states ($\mathbf{Z}_0\mathbf{W}$), summed over K predictor variables:

$$\sum_{k=1}^K V_k (Z_{1k} - \mathbf{Z}_{0k}\mathbf{W})^2$$

Here V_k is a weight measuring the relative importance of the k -th predictor. Note that the vector of predictors $\mathbf{Z} = (\mathbf{X}, \bar{\mathbf{Y}}_1^1, \dots, \bar{\mathbf{Y}}_1^M)$ includes M linear combinations of the pre-intervention outcomes $\bar{\mathbf{Y}}_1^m$ in the treated unit, besides covariates \mathbf{X} .

Abadie et al. (2010) do not specify exactly which, or how many, linear combinations of pre-intervention outcomes should be used, although some such combinations should certainly be included in \mathbf{Z} in order to ensure the balancing of factor loadings between treated and synthetic control units. One choice for predictors includes every pre-treatment value of the outcome, as in the Bohn, Lofstrom, and Raphael (2012) study of Arizona’s immigration law. Using a randomization inference procedure, and specifically considering state-level teen employment using CPS data, Dube and Zipperer (2013) consider several candidate models and achieve best out-of-sample fit for donors by using annualized averages of pre-treatment outcomes in addition to pre-treatment covariates. We select the same predictor variables here: annualized pre-treatment averages of EPOPs; 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. While the consistency of the synthetic control estimator holds for a variety of predictor weight vectors, \mathbf{V} , such weights can also be chosen optimally using a data-driven procedure (see Abadie et al. 2010 for details). Therefore, the preferred algorithm dynamically optimizes over both the vectors of predictor weights, \mathbf{V} , as well as the vector of donor weights \mathbf{W} .²⁷

²⁷We implement the synthetic control approach in Stata using the `synth` package with `nested` optimization and `allop` starting point checks for robustness: <http://www.mit.edu/~jhainm/synthpage.html>.

6.1 Pooled synthetic control estimates for teens

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. Dube and Zipperer (2013) consider all state-level minimum wage changes during the intervening period between the federal increases in 1997 and 2007.²⁸ Of the 89 state minimum wage changes from 1997q4 through 2007q2, there are 19 usable treatment events with pre-treatment periods of between two and four years. Table 5 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 four annualized pre-treatment averages of EPOPs (for up to 16 pre-intervention quarters), along with the set of predictor variables discussed above.

29

For each event, Table 5 shows the minimum wage elasticities for teen employment and wages, reproduced from Dube and Zipperer (2013).³⁰ Additionally, we report in parentheses

²⁸To our knowledge, only two papers utilize the synthetic control method when estimating minimum wage treatment effects: Sabia, Burkhauser, and Hansen (2012) and Neumark, Salas, and Wascher (2013a). Sabia et al. focus only on the minimum wage change in New York state beginning in 2005. They thereby ignore 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 unreliable synthetic control: their Figure 3 clearly shows that employment rates for actual and synthetic New York *never* match during the pre-treatment period. Neumark, Salas, and Wascher construct a “matched panel estimator” loosely based on synthetic controls. We provide a detailed discussion and critique of this estimator in Appendix B, section B2. In brief, this estimator is without an econometric foundation, as it uses residuals from an OLS panel regression to estimate synthetic control weights. In addition, the set of donors they use to construct synthetic control often themselves receive treatment, violating a key requirement of the synthetic control method.

²⁹Specifically, Dube and Zipperer 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-treatment level and its maximum post-treatment level, since a treated state may alter its minimum wage more than once in the post-treatment period.

³⁰We estimate the elasticities as the percentage difference in the average post-treatment outcomes between

the percentile ranks of the treatment effects against placebo (donor) outcomes, which is the inferential procedure suggested by Abadie et al. (2010). 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 earnings 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 examining the mean rank, which is distributed as a known function of uniform random variables.³¹ As Table 5 shows, the pooled effect for teen wages is sizeable for both the mean (0.174) and the median (0.149) effects. The mean rank test indicates that the pooled effect is statistically significant at the 90 percent level. The mean teen employment elasticity, reported at the bottom of the table, is about -0.028, while the median is 0.082, and the pooled effect is not statistically significant when judged by the mean 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 Table 3. They are also similar to the employment estimates we obtained using lagged dependent variable controls in Table 4. Overall, the synthetic-control based pooled effects confirm the general findings of local control-based approaches: clear positive impact on teen earnings, along with a small or no disemployment effect that is statistically indistinguishable from zero.

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.

³¹Namely, we compare these across-event mean ranks to the 90 percent extrema of the mean of 19 simulated uniform random variables. The pooled rank estimator is described in greater detail in Dube and Zipperer (2013).

6.2 Overlap in synthetic and local controls: Evidence from actual and placebo laws

Different control groups need not have a high degree of overlap in order to provide unbiased estimates. However, it is useful to know how often a data-driven procedure like synthetic control selects local donors to be part of the control set. As a starting point, in Table 6 we report the average weight placed on donors, based on their proximity to the treated state, for the 19 minimum wage increases analyzed in the previous section. The results show that more proximate donors are indeed more likely to be a part of the synthetic controls. On average, donors within the same (nine) Census division received weights that were four times as large, adding validity to our focus on variation within geographical regions in the United States when estimating minimum wage effects. The relative differences are still sizable, though smaller, when we consider broader geographical aggregations such as the four Census regions.

Unfortunately, these 19 treatments do not represent the ideal setting to evaluate whether neighboring areas do indeed make for better control groups because the donor pools associated with each of these treatments often exclude nearby states that have also experienced minimum wage increases within the estimation window. On average, these treatments exclude 35 percent of outside-division states but 74 percent of within-division states from the donor pool—setting up an unbalanced comparison.

Fortunately, 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.³²

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 Hawaii and Alaska as distance outliers; the minimum distance from Hawaii 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.³³ By focusing on these placebo treatments rather than actual minimum wage events, the randomization approach dispenses with the shortcomings associated with a small number of treatment events, a short pre-intervention window, and the lack of nearby potential donors.

We present the key evidence in Figure 8, which non-parametrically 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 690 miles, suggesting that donors exceeding this distance are “worse-than-average” controls. For example, donors that are 1,000 miles away have weights of 0.014, while donors only 100 miles away are predicted to obtain a weight of 0.083. Donors that are 100 miles away are thus, on average, assigned a weight that is approximately 6 times

³²This sample window is identical to the one we used in the previous section to show the impact of actual minimum wage changes using the synthetic control approach.

³³NBER 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>.

larger than a donor that is 1,000 miles away. As indicated by the 95 percent bootstrapped confidence bands in the figure, these differences are highly statistically significant.

Panel B of Table 6 also summarizes the geographic characteristics of donor weights estimated by this procedure, using the same categories we used for actual minimum wage laws: four Census regions, nine Census divisions and alternative distance bands. Similar to actual minimum wage treatments, within-division donor weights for placebo treatments are roughly four times as large as the weights outside the region. For other geographic categories, however, the preference for nearby areas is even stronger when using placebo laws. As shown in Panel B, for donors whose population centroids are less than 500 (1000) miles away from the treated state, inside-donors receive weights that are about 3.8 (2.4) times as large as outside-donors. For these two distance categories, the corresponding relative weight of more proximate donors using actual minimum wage treatments are about 2.1 and 2.0, respectively. Overall, the evidence from both actual minimum wage changes as well as placebo treatments show that neighboring areas are indeed more similar, and that in general local areas make for good control groups.³⁴

7 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

³⁴Neumark, Salas, and Wascher (2013a) use the synthetic control approach to argue that nearby areas are not more similar. In Appendix B, section B3, we show that while their calculation of synthetic control donor weights is fraught with problems, a straightforward interpretation of their results actually suggest that inside-division donors receive weights that between 2.6 and 4.1 times as large as outside-division donors.

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 each other. 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, adequately 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 need 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 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.

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 Dube, Lester and Reich (2010) 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 tool kit. But this 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 review 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) models with lagged dependent variable and fixed effects, using OLS-based bounds and GMM dynamic panel models, and (6) a synthetic control design. Every one of these six approaches—or something very close—have been used in the discipline prior to our original work, and are considered useful parts of our modern tool kit. 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 now been replicated using four different datasets and validated externally using newer data.

Figure 9 brings together and summarizes these results by visually plotting the minimum wage employment elasticities by type of estimator. We plot the canonical two-way fixed effects estimates, along with the estimates using local area controls from Table 3; estimates using lagged employment controls (with and without fixed effects) from Table 4; and the pooled synthetic control estimate from Table 5. Figure 9 shows a clear distinction between the canonical and other estimates of employment elasticities, which are completely non-overlapping sets. The *most negative* of the 21 employment elasticities using spatial, lagged employment, or synthetic controls is -0.076 , while the *least negative* of 5 canonical estimates is -0.079 .

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 Jens Hainmueller 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association* 105, 490: 493-505.
- [2] Abowd, John and Lars Vilhuber 2011. "Gross Employment, Job Flows, and the Role of Education in the Great Recession." Working paper, Labor Dynamics Institute, Cornell University.
- [3] Abowd, John et al. 2009. "The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators." Pp. 149-230 in Timothy Dunne, Bradford Jensen and Mark Roberts eds. *Producer Dynamics: New Evidence from Micro Data*. Chicago: University of Chicago and NBER.
- [4] Abraham, Katharine G. 2009. "Comment on The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators." Pp. 230-34 in Timothy Dunne, Bradford Jensen and Mark Roberts eds. *Producer Dynamics: New Evidence from Micro Data*. Chicago: University of Chicago and NBER.
- [5] Abraham, Katharine G., John Haltiwanger, Kristin Sandusky, and James R. Spletzer 2013. "Exploring Differences in Employment between Household and Establishment Data." *Journal of Labor Economics* 32, 2: S129-S172
- [6] 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.
- [7] 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.
- [8] _____ 2011. "Do Minimum Wages Really Reduce Teen Employment? Accounting for Heterogeneity and Selectivity in State Panel Data." *Industrial Relations* 50, 2: 205-40.
- [9] Allegretto, Sylvia, Arindrajit Dube, Michael Reich and Ben Zipperer 2013. "Credible Research Designs for Minimum Wage Studies." Working Paper 148-13. Institute for Research on Labor and Employment, UC Berkeley.
- [10] Angrist, Joshua and Jorn-Steffen Pischke 2009. *Mostly Harmless Econometrics*. Princeton, NJ: Princeton University Press.
- [11] Arellano, Manuel, and Stephen Bond 1991. "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations." *Review of Economic Studies* 58, 277-97.

- [12] Arellano, Manuel and Olympia Bover 1995. "Another look at the instrumental variable estimation of error-components models." *Journal of Econometrics* 68, 1: 29-51.
- [13] Ashenfelter, Orley, Henry Farber and Michael Ransom 2010. "Modern Models of Monopsony in Labor Markets: A Brief Survey." *Journal of Labor Economics* 28, 2: 203-10.
- [14] 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.
- [15] _____ and David Dorn 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market." *American Economic Review*, forthcoming.
- [16] _____, Frank Levy and Richard Murnane 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118, 4: 1279 - 1333.
- [17] Black, Sandra 1999. "Do Better Schools Matter? Parental Valuation of Elementary Education." *Quarterly Journal of Economics* 114, 2: 577-99.
- [18] Blundell, Richard and Stephen Bond 1998. "Initial Conditions and Moment Conditions in Dynamic Panel Models." *Journal of Econometrics* 87, 1: 115-43.
- [19] 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.
- [20] Bond, Stephen R. 2002. "Dynamic Panel Data Models: A Guide to Micro Data Methods and Practice." *Portuguese Economic Journal* 1: 141-62.
- [21] Brochu, Pierre, and David A. Green 2013. "The Impact of Minimum Wages on Labor Market Transitions." *The Economic Journal*, forthcoming.
- [22] Burdett, Kenneth and Dale Mortensen 1998. "Wage Differentials, Employer Size and Unemployment." *International Economic Review* 39: 257-73.
- [23] 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.
- [24] _____ 1995. *Myth and Measurement*. Princeton NJ: Princeton University Press.
- [25] _____ 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.
- [26] 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.

- [27] _____ 2013. “Minimum Wage Shocks, Employment Flows and Labor Market Frictions.” Working Paper 149-13. Institute for Research on Labor and Employment. UC Berkeley.
- [28] Dube, Arindrajit, Suresh Naidu and Michael Reich 2007. “The Economic Effects of a Citywide Minimum Wage.” *Industrial and Labor Relations Review* 60, 4: 522-43.
- [29] 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.
- [30] 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.
- [31] Hagedorn, Marcus, Fatih Karahanz, Iouri Manovskii and Kurt Mitman 2013. “Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects.” Paper presented at the Minnesota Workshop in Macroeconomic Theory.
- [32] 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.
- [33] 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.
- [34] Katz, Lawrence and Kevin Murphy 1992. “Changes in Relative Wages, 1963-1987: Supply and Demand Factors.” *Quarterly Journal of Economics* 107, 1: 35-78.
- [35] Lee, David and Thomas Lemieux 2010. “Regression Discontinuity Designs in Economics.” *Journal of Economic Literature* 48, 2: 281-355.
- [36] Magruder, Jeremy 2013. “Can Minimum Wages Cause a Big Push? Evidence from Indonesia.” *Journal of Development Economics* 100, 1: 48-62.
- [37] Manning, Alan 2003. *Monopsony in Motion*. Princeton NJ: Princeton University Press.
- [38] Meer, Jonathan and Jeremy West 2013 “Effects of the Minimum Wage on Employment Dynamics.” NBER Working Paper 19262.
- [39] Neumark, David and William Wascher 2008. *Minimum Wages*. Cambridge, MA: MIT Press.
- [40] Neumark, David, J. M. Ian Salas and William Wascher 2013a. “Revisiting the Minimum Wage and Employment Debate: Throwing out the Baby with the Bathwater?” *Industrial and Labor Relations Review*, forthcoming.
- [41] Neumark, David, J. M. Ian Salas and William Wascher 2013b. “Revisiting the Minimum Wage and Employment Debate: Throwing out the Baby with the Bathwater?” NBER Working Paper 18681.

- [42] Nickell, Stephen 1981. "Biases in Dynamic Models with Fixed Effects." *Econometrica* 49, 6: 1417-26.
- [43] OECD 2013. "Relation of Minimum Wage to Median Wage." OECD Statistics download. http://stats.oecd.org/BrandedView.aspx?oecd_bv_id=lfs-data-en&doi=data-00313-en#
- [44] Portugal, Pedro and Ana Rute Cardoso 2006. "Disentangling the Minimum Wage Puzzle: An Analysis of Worker Accessions and Separations." *Journal of the European Economic Association* 4, 5: 988-1013.
- [45] Roodman, David M. 2009. "A Note on the Theme of Too Many Instruments." *Oxford Bulletin of Economics and Statistics* 71, 1: 135-58.
- [46] 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.
- [47] Smith, Christopher 2011. "Polarization, Immigration, Education: What's Behind the Dramatic Decline in Youth Employment?" Finance and Economics Discussion Series 2011-41. Board of Governors of the Federal Reserve System.
- [48] Whittaker, William 2003. "The Federal Minimum Wage and Average Hourly Earnings of Manufacturing Production Workers." *CSR Report for Congress 98-960 E*. Washington, DC: Congressional Research Service.


Figure 1
Minimum wage variation across states and over time: 1984 - 2012




Notes. The figure plots the annual federal minimum wage, as well as the 75th, 50th, and the 90th percentile of the prevailing annual minimum wage across states over time. Annual minimum wages are constructed by averaging across months in a given year. The prevailing minimum wage is the maximum of the state and federal minimums in that year. Light grey circles represent actual prevailing annual minimum wage in each state that year.

Figure 2
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 4
Quarterly probability of a minimum wage change, by phase of the business cycle




Notes. The quarterly probability of a minimum wage change in a given business cycle phase (recession, first half of the expansion, or second half of the expansion) equals the number of quarterly minimum wage changes divided by the number of quarters in that phase. The 95 percent confidence intervals are constructed using heteroscedasticity-robust standard errors. NBER recession dates are 1990q3–1991q4, 2001q1–2001q4, and 2007q4–2009q2. Panel A contains the maximal sample in our regressions, 1990q1–2012q4; we assign all observations beginning 2009q3 to the first half of the expansion. For robustness, Panel B reports the 1990q3–2007q3 sample, which contains the two most recent complete recessions and expansions.


Figure 5

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


A. Cross-state commuting zones



B. Cross-state border county pairs (QWI)




C. Cross-state border county pairs (QCEW)




Notes. A: 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. B: Red and blue colored counties indicate cross-state border county pairs in QWI sample. Red counties are part of pairs with minimum wage variation between the counties during 2000-2011. Darker shades indicate pairs whose centroids are within 75 miles of each other. C: Red and blue colored counties indicate cross-state border county pairs in QCEW sample. Red counties are part of pairs with in-sample minimum wage variation between the counties. Darker shades indicate balanced sample for QCEW restaurant data. 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.

Figure 6
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 \times period fixed effects, as well as state-specific linear time trends. QCEW local controls are county and county pair \times 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 7
Minimum wage employment elasticities with lagged dependent variable and fixed effects, for alternative values of the autoregressive coefficient



Notes. The solid lines are minimum wage employment elasticities for restaurant and teen samples, in a fixed effects model with a lagged dependent variable for alternative values of the autoregressive coefficient, ρ . We quasi difference the outcome (EPOP for teens and log employment for restaurants) using a particular value of ρ between 0 and 1, and regress the quasi-differenced outcome on log minimum wage along with covariates and time and place fixed effects (see text). The dashed lines are 95% confidence intervals. The shaded area highlights the estimates that use a more credible region of imposed autoregressive coefficients: between those produced by the downwardly-biased OLS estimate of ρ from a model with LDV and FE, and an upwardly-biased estimate of ρ from a model with LDV only. All data are annual: restaurants, 1991-2010 (QCEW); teens 1990-2012 (CPS).


Figure 8

How synthetic control weights for donors vary by distance to treated states: Evidence from placebo laws



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 9
Minimum wage employment elasticities: Canonical model versus with alternative controls for time-varying heterogeneity



Notes: Individual employment elasticities of minimum wage for teens and restaurant workers are represented by open squares. The shaded regions denote the range of point estimates for each of the four classes of estimators (canonical two-way fixed effects (FE), local area controls, lagged dependent variable (LDV) controls, and synthetic controls). The data sources are QCEW, QWI, CPS, ACS/Census. Local area controls use contiguous counties (QCEW and QWI), commuting zones (ACS/Census), or division-specific time effects and state-specific linear trends (CPS) as reported in Table 3. LDV estimates using CPS and QCEW are as reported in Table 4: from OLS model with lagged employment and no fixed effects; and lagged employment and fixed effects estimated using OLS bounds and GMM. The pooled synthetic control estimates for teens using CPS data is as reported in Table 5.

Table 1
Heterogeneity between high and low minimum wage states

	High-MW state mean	Low-MW state mean	Difference
Panel A: 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)
Panel B: 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)
Panel C: 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. Significance levels are indicated by *** 1%, ** 5%, * 10%.

Table 2
Mean absolute differences in covariates between counties: contiguous versus other pairs

	Non-contiguous pair	Contiguous pair	Gap	Pct. gap
Panel A: Levels				
<i>Log Employment</i>	1.744*** (0.026)	1.233*** (0.027)	0.511*** (0.033)	41.4
<i>Log Population</i>	1.518*** (0.023)	0.964*** (0.023)	0.554*** (0.029)	57.5
<i>EPOP</i>	0.042*** (0.001)	0.039*** (0.001)	0.003*** (0.001)	8.0
<i>Log Earnings</i>	0.229*** (0.004)	0.1695*** (0.004)	0.060*** (0.004)	35.1
<i>Turnover Rate</i>	0.057*** (0.001)	0.048*** (0.001)	0.009*** (0.001)	18.1
<i>Teen Share</i>	0.006*** (0.0001)	0.005*** (0.0001)	0.001*** (0.0001)	21.7
Panel B: 4 Quarter difference				
<i>Log Employment</i>	0.062*** (0.001)	0.058*** (0.001)	0.003*** (0.001)	5.2
<i>Log Population</i>	0.048*** (0.001)	0.047*** (0.001)	0.002** (0.001)	3.9
<i>EPOP</i>	0.014*** (0.0002)	0.011*** (0.0002)	0.003*** (0.0002)	26.9
<i>Log Earnings</i>	0.013*** (0.0002)	0.012*** (0.0003)	0.001*** (0.0002)	7.3
<i>Turnover Rate</i>	0.002*** (0.0000)	0.001*** (0.0000)	0.0002*** (0.0000)	16.7
<i>Teen Share</i>	0.038*** (0.001)	0.036*** (0.001)	0.002*** (0.001)	5.6
Panel C: 12 Quarter difference				
<i>Log Employment</i>	0.099*** (0.001)	0.091*** (0.002)	0.008*** (0.001)	8.5
<i>Log Population</i>	0.069*** (0.001)	0.066*** (0.002)	0.004*** (0.001)	5.5
<i>EPOP</i>	0.037*** (0.001)	0.027*** (0.001)	0.010*** (0.001)	36.3
<i>Log Earnings</i>	0.018*** (0.0003)	0.017*** (0.0004)	0.001*** (0.0003)	8.5
<i>Turnover Rate</i>	0.003*** (0.0000)	0.002*** (0.0000)	0.001*** (0.0000)	25.0
<i>Teen Share</i>	0.045*** (0.001)	0.041*** (0.001)	0.004*** (0.001)	9.5

Notes. Reproduced from Dube, Lester and Reich (2013). 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, clustering standard errors multi-dimensionally on each of the two counties in the cross-border pair. "Gap" is a test of difference in mean absolute value of the covariate between contiguous and other pairs. "Pct. Gap" divides this gap value by the mean for the contiguous pairs. Significance levels are indicated by *** 1%, ** 5%, * 10%.

Table 3

Minimum wage elasticities and spatial controls for time-varying heterogeneity: 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								
<i>Wage/Earnings</i>	0.190*** (0.045)	0.124*** (0.025)		0.182*** (0.037)	0.296*** (0.064)	0.167*** (0.033)		0.222*** (0.048)
<i>Employment</i>	-0.123** (0.057)	-0.162** (0.042)		-0.161** (0.071)	0.130 (0.126)	0.002 (0.066)		-0.059 (0.086)
<i>Hires</i>				-0.510*** (0.092)				-0.219** (0.095)
<i>Separations</i>				-0.539*** (0.099)				-0.233** (0.099)
Panel B: Restaurant workers								
<i>Wage/Earnings</i>			0.210*** (0.024)	0.204*** (0.027)			0.186*** (0.027)	0.207*** (0.060)
<i>Employment</i>			-0.119 (0.076)	-0.079** (0.039)			0.008 (0.051)	-0.022 (0.093)
<i>Hires</i>				-0.468*** (0.084)				-0.264* (0.137)
<i>Separations</i>				-0.465*** (0.077)				-0.225* (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

Table 3

Minimum wage effects and spatial controls for time-varying heterogeneity: results from five datasets (contd.)

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 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 and 8 are clustered multi-dimensionally at the state-level and border segment-level. Significance levels are indicated by *** 1%, ** 5%, * 10%.

Table 4
Minimum wage employment elasticities using lagged employment controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS				GMM			
			<i>Lower bound ρ (naïve OLS)</i>	<i>Upper bound ρ (imposed at 0.71)</i>				
Panel A: Teens								
<i>Log MW</i>	-0.135** (0.062)	-0.025 (0.031)	-0.076* (0.044)	-0.004 (0.031)	-0.002 (0.021)	-0.024 (0.033)	0.012 (0.039)	0.097 (0.081)
<i>Log EPOP_{t-1}</i>		0.713*** (0.046)	0.323*** (0.044)		0.851*** (0.067)	0.599*** (0.144)	0.823*** (0.171)	0.067 (0.298)
N	1,173	1,173	1,173	1,173	1,173	1,173	1,122	1,173
Panel B: Restaurants								
<i>Log MW</i>	-0.132 (0.082)	-0.052*** (0.014)	-0.030 (0.018)	-0.000 (0.011)	-0.024*** (0.009)	-0.037*** (0.009)	0.006 (0.015)	-0.023 (0.021)
<i>Log emp_{t-1}</i>		0.970*** (0.006)	0.752*** (0.014)		1.048*** (0.023)	1.045*** (0.022)	1.012*** (0.077)	1.127*** (0.134)
N	25,840	25,840	25,840	25,840	25,840	25,840	24,548	24,548
Fixed effects:	Y		Y	Y	Y	Y	Y	Y
Transform:								
Within	Y		Y	Y				
Difference						Y		Y
Forward orthogonal					Y		Y	
GMM equations:								
Level					Y	Y		
Difference					Y	Y	Y	Y

Notes. The dependent variable is log restaurant employment or log teen EPOP and the treatment variable is log minimum wage. All regressions use annual data—restaurants, 1991-2010 (QCEW), and teens, 1990-2012 (CPS)—and include year fixed effects. Specifications include place fixed effects (FE): county for QCEW and state for CPS. Column (1) reproduces the FE estimate. Column (2) reports the estimate with the LDV but no FE. Columns (3) and (4) report minimum wage elasticities using credible bounds in the autoregressive coefficient ρ . The lower bound for ρ is from naïve OLS with FE and LDV as reported in column (3); upper bound for ρ is from column 2 (LDV only), and resulting minimum wage elasticities imposing this bound, along with fixed effects, are reported in column 4. Columns 5-8 report estimates using system or difference GMM for model with FE and LDV. Transform describes the technique used to purge FE's for the GMM estimates. Standard errors in parentheses are clustered by state. Significance levels are indicated by *** 1%, ** 5%, * 10%.

Table 5
Minimum wage elasticities of teen employment and wages using synthetic controls

Event	Window length			Δ MW%	Employment elasticity		Wage elasticity		
	Pre	Post	Donors		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		0.149	

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, the 90 percent cutoffs are (0.391, 0.609).

Table 6
Donor weights and distance to treated states

Donor relation to treatment	Average weight per donor		Relative weight
	Inside	Outside	
Panel A: Actual treatments sample			
Same Census region	0.062	0.026	2.328
Same Census division	0.113	0.029	3.97
Within 0 - 500 miles	0.061	0.029	2.112
Within 0 - 1000 miles	0.047	0.024	1.964
Panel B: Placebo treatments sample			
Same Census region	0.044	0.014	3.15
Same Census division	0.063	0.016	3.868
Within 0 - 500 miles	0.051	0.013	3.755
Within 0 - 1000 miles	0.029	0.012	2.423

Notes. Actual treatments sample consists of 19 actual minimum wage treatments of 16 states during 1997q4-2007q2 with pre- and post-treatment windows of at least 2 years and 1 year, respectively. Placebo treatments sample consists of 768 placebo treatments of 48 continental states during the same period with pre- and post-treatment windows of exactly 4 years and 2 years, respectively. Synthetic control outcome is the teen employment-to-population ratio. See text for exact sample and 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 in miles is measured between population-weighted state centroids.

Appendix A: Additional Tables

In this appendix, we provide two additional tables referred to in the text. Appendix Table A1 reports the estimates underlying the pre-existing trends shown in Figure 6. The cumulative responses are formed by successively adding the leading minimum wage coefficients. The maximum number of leads vary across columns as indicated. Panel A reports results for teens using the CPS. Panels B and C report results for restaurants using the QCEW; Panel C additionally includes overall private sector employment as a control. Columns 1-4 report results for the canonical model, while columns 5-8 are for the preferred specifications. As we note in the text when discussing Figure 6, the canonical model systematically shows negative leading terms indicating pre-existing trends for both the teen and restaurant sample. The pre-existing trends are particularly pronounced for the teen sample, and many of these leading coefficients are statistically significant at conventional levels. In contrast, for the preferred specifications with local controls, the leading terms are centered around zero and never statistically significant.

Appendix Table A2 reports wage (or earnings) elasticities using lagged dependent variable controls, analogous to the employment estimates reported in Table 4. Panel A reports results for teens using the CPS, while panel B reports them for restaurant workers using the QCEW. Columns 1-4 are estimated using OLS; Column 4 imposes a value of the autoregressive coefficient ρ deemed as an “upper bound,” as estimated in column 2 without any fixed effects. Columns 5-8 are estimated using GMM. The wage elasticities are all positive, and almost all are statistically significant at conventional levels. However, the magnitudes vary considerably across specifications: for teens (panel A) the elasticities range between 0.06 and 0.32, while for restaurant workers (panel B), they range between 0.05 and 0.23. Unlike employment, for earnings we do not find a systematic pattern from the inclusion of lagged outcomes as controls.

Appendix Table A1

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

	Canonical model				Local controls			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Log teen EPOP (CPS)								
<i>t=0</i>	-0.162** (0.062)	-0.212*** (0.064)	-0.257*** (0.067)	-0.284*** (0.068)	0.081 (0.077)	0.065 (0.102)	0.068 (0.139)	0.055 (0.166)
<i>t=-1</i>		-0.240** (0.091)	-0.232** (0.088)	-0.263*** (0.098)		-0.036 (0.112)	-0.035 (0.124)	-0.048 (0.148)
<i>t=-2</i>			-0.158* (0.081)	-0.152* (0.078)			0.005 (0.097)	-0.001 (0.108)
<i>t=-3</i>				-0.106 (0.083)				-0.023 (0.088)
N	1,071	1,071	1,071	1,071	1,071	1,071	1,071	1,071
Panel B: Log restaurant employment (QCEW)								
<i>t=0</i>	-0.178** (0.083)	-0.199** (0.095)	-0.223* (0.112)	-0.248** (0.123)	0.037 (0.074)	0.046 (0.086)	0.036 (0.094)	0.019 (0.096)
<i>t=-1</i>		-0.088 (0.058)	-0.090 (0.058)	-0.120* (0.071)		0.031 (0.048)	0.026 (0.053)	0.007 (0.058)
<i>t=-2</i>			-0.077 (0.059)	-0.079 (0.058)			-0.029 (0.041)	-0.037 (0.045)
<i>t=-3</i>				-0.084* (0.046)				-0.052 (0.039)
N	27,132	27,132	27,132	27,132	11,886	11,886	11,886	11,886
Panel C: Log restaurant employment (QCEW) - with control for overall private-sector employment								
<i>t=0</i>	-0.138 (0.083)	-0.153 (0.095)	-0.170 (0.112)	-0.186 (0.123)	0.007 (0.055)	0.015 (0.065)	0.007 (0.076)	-0.004 (0.082)
<i>t=-1</i>		-0.062 (0.055)	-0.064 (0.056)	-0.083 (0.068)		0.024 (0.044)	0.021 (0.049)	0.009 (0.058)
<i>t=-2</i>			-0.053 (0.061)	-0.054 (0.061)			-0.022 (0.042)	-0.027 (0.046)
<i>t=-3</i>				-0.053 (0.043)				-0.032 (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. Significance levels are indicated by *** 1%, ** 5%, * 10%.

Appendix Table A2
Minimum wage elasticities for average earnings using lagged wages controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS				GMM			
			<i>Lower bound ρ</i> <i>(naïve OLS)</i>	<i>Upper bound ρ</i> <i>(imposed at 0.29)</i>				
Panel A: Teens								
<i>Log MW</i>	0.104** (0.047)	0.253*** (0.034)	0.099** (0.044)	0.085** (0.036)	0.105*** (0.028)	0.321*** (0.089)	0.062** (0.028)	0.248** (0.123)
<i>Log wage_{t-1}</i>		0.290*** (0.054)	0.077* (0.038)		0.712*** (0.076)	0.096 (0.310)	0.612*** (0.103)	0.511*** (0.135)
N	1,173	1,173	1,173	1,173	1,173	1,173	1,122	1,173
Panel B: Restaurants								
<i>Log MW</i>	0.229*** (0.024)	0.052*** (0.007)	0.126*** (0.012)	0.073*** (0.010)	0.061*** (0.015)	0.085*** (0.015)	0.067*** (0.014)	0.175*** (0.016)
<i>Log earnings_{t-1}</i>		0.944*** (0.008)	0.630*** (0.024)		0.936*** (0.042)	0.910*** (0.040)	0.979*** (0.049)	0.847*** (0.048)
N	27,132	25,840	25,840	25,840	25,840	25,840	24,548	24,548
Fixed Effects:	Y		Y	Y	Y	Y	Y	Y
Transform:								
Within	Y		Y	Y				
Difference						Y		Y
Forward orthogonal					Y		Y	
GMM equations:								
Level					Y	Y		
Difference					Y	Y	Y	Y

Notes. The dependent variable is log average restaurant earnings or log teen wage and the treatment variable is log minimum wage. All regressions use annual data—restaurants, 1991-2010 (QCEW), and teens, 1990-2012 (CPS)—and include year fixed effects. Specifications include place fixed effects (FE): county for QCEW and state for CPS. Column (1) reproduces the FE estimate. Column (2) reports the estimate with the LDV but no FE. Columns (3) and (4) report minimum wage elasticities using credible bounds in the autoregressive coefficient ρ . The lower bound for ρ is from naïve OLS with FE and LDV as reported in column (3); upper bound for ρ is from column 2 (LDV only), and resulting minimum wage elasticities imposing this bound, along with fixed effects, are reported in column 4. Columns 5-8 report estimates using system or difference GMM for model with FE and LDV. Transform describes the technique used to purge FE's for the GMM estimates. Standard errors in parentheses are clustered by state. Significance levels are indicated by *** 1%, ** 5%, * 10%.

Appendix B: Responses to claims in Neumark, Salas and Wascher (2013)

In their 2013 paper, Neumark, Salas, and Wascher (hereafter NSW) criticize the use of local area controls. As we demonstrate in section 4 of this paper using multiple datasets and treatment groups, local areas are indeed more similar to each other, and use of local controls clearly eliminates spurious pre-existing trends that compromise the two-way fixed effects estimator that NSW appears to favor. As we also show in Sections 5 and 6 of this paper, non-spatial approaches to accounting for unobserved heterogeneity—using lagged outcomes as controls, as well as using the pooled synthetic control estimator—confirm that employment effects of minimum wages tend to be small. The replication of these results across four datasets and six different approaches strongly refute the central argument in NSW that there are sizable employment losses from the modest increases in the minimum wage experienced in the U.S. over the past two decades.

In this Appendix, we additionally respond to five specific claims made in NSW.³⁵ First, we show that the way they use synthetic control donor weights to assess the validity of local area controls is problematic. Moreover, their own (imperfect) evidence contradicts their claim that local areas are not better controls: a straightforward interpretation of their results shows that states within the same Census division receive much greater weight as controls, consistent with our findings in section 6.2 of this paper. Second, we discuss why their proposed matching estimator—which they offer as their main alternative to local area controls—lacks any econometric basis. This contrasts with a well-founded, properly estimated pooled synthetic control estimator that we use in Section 6.1 of this paper. Third, we demonstrate that a second estimator they propose—using state-pairs—fails a key falsification test for pre-existing trends. Fourth, we show that their concerns about the impact of

³⁵Some of these claims are explained in detail only in their working paper, Neumark, Salas and Wascher (2013b). When this is the case, we note it as such in our discussion below. However, all of the five claims we rebut here are referenced in their publication version (2013a).

the inclusion of recessionary periods in Allegretto, Dube and Reich (2011) on the estimation of state-specific trends can be easily handled using standard methods; and doing so does not alter the conclusions drawn in that paper. Fifth, we show that their criticism of the placebo exercise in Dube, Lester and Reich (2010)—used to demonstrate the presence of spatial heterogeneity in minimum wage laws—is misplaced.

As demonstrated in this paper, estimators that account for unobserved heterogeneity using local area controls, lagged dependent variable controls, and synthetic controls all suggest small employment effects from U.S. minimum wage increases during the 1990s and 2000s. These results hold across four datasets and six different approaches. In contrast, estimators favored by NSW that appear to show sizable disemployment effects either fail key falsification tests for pre-existing trends, or lack an econometric basis.

B1. 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).³⁶ 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.

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. 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

³⁶NSW 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.

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.³⁷ In a similar vein, two sets of controls may both be appropriate, without having a substantial degree of overlap. Therefore, both local and synthetic controls may be different unbiased estimators, while the two-way fixed effects estimator remains biased.

Three additional 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. Although NSW are testing whether donors from the same division receive 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 in Section 6.2 of this paper.

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.³⁸ 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 evidence.

³⁷NSW state that “in cases where most of the weight is on states outside the division, the conventional panel data estimator should provide more reliable estimates of minimum wage effects than the specification that includes division \times period controls.” This is incorrect. Just because *some* states outside of the division may be good controls does mean that *all* states are.

³⁸For 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 other paper in the minimum wage literature to use synthetic controls—Burkhauser, Hansen and Sabia (2011)—uses four years of pre-intervention data.

Third, there are important problems with the predictor variables they use to construct the synthetic control donor weights. 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. With regard to the other set of predictors, there is absolutely no econometric justification for using residuals from an OLS panel regression as the matching variable in a synthetic control study. We return to this matter in the next section, when discussing NSW’s preferred matching estimator, which also uses these synthetic control weights based on OLS residuals and is similarly erroneous.

Flawed as their evidence may be—its proper interpretation actually suggests a conclusion opposite to one they draw. NSW report in their Table 3 the proportion of weights on donor states in the same division as treated states. In their Table 5, 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. The average per-donor weight of same-division donors, relative to the per-donor weight of other-division donors is more informative; we show the results for this measure in our Appendix Table B1.

While Neumark Salas and Wascher (2013a) does not report them, one can calculate aver-

age per-donor weights for same-division and different-division donors using the information in their Tables 3 and 5.³⁹ We report these calculations in Appendix Table B1.⁴⁰ A donor state within the same division receives weights that are *2.8 to 4.1 times as large* as weights for donors outside of the division. These relative weights are quite similar to those we presented in section 6.2 (and Table 6) using the actual 19 minimum wage events, as well as the weights calculated using the placebo laws. As Appendix Table B1 shows, a similar pattern holds for cross-border counties, which receive between 4.0 to 5.1 times greater weights than do other counties. Their claims notwithstanding, a straightforward interpretation of NSW’s own evidence indicates that neighboring areas are more alike than are places farther away—contradicting their central thesis.

B2. The NSW Matching Estimator

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. They ground their methodology in the existing literature by citing Autor et al. (2006), but the NSW ap-

³⁹In their publication version (2013a), Neumark, Salas and Wascher briefly discuss average weight per donor, where they state that “[f]or the analysis in [specification] (1), the average weight per same-division donor state is higher than the random threshold of $1/(\text{number of potential donors})$ in only 18 of 50 cases. For [specifications] (2) - (4), the corresponding numbers are 24, 17, and 19.” This statement is not particularly informative about the average weight of same-division donors. In the working paper version (2013b), they state in the text using a computation not reported in their tables that for one specification “the average weight assigned to same-division states in the donor pool by the synthetic control method is no higher – and indeed is generally lower – than the average weight assigned to states in other divisions that are in the donor pool.” Averaged across their remaining specifications, they state in footnote 27 that the average weight for same-division donor is 0.199, while for different division donors it is 0.148. However, one can directly calculate average per-donor weights for same-division and different-division donors using the information in Tables 3 and 5 of Neumark, Salas and Wascher (2013a). When we do so in this paper, the results suggest a very different conclusion.

⁴⁰In our Appendix Table B1 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 3 and 5). We also explained above 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.

proach 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 employment 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. Besides being econometrically unjustifiable, this approach is also intuitively wrong. 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, then the estimated (and true) residuals are simply mean zero error terms, which by definition are orthogonal to minimum wages. It does not make sense to find controls to match such residuals, as they precisely do not include confounds. In *both* cases, therefore, matching on OLS residuals is wrong. And 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.

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).⁴¹

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 (see their Tables 7 and 8). This result leads the authors to believe this sample of treatments is “unusual,” as these estimators fail to generate a large negative employment elasticity.⁴² 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 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-

⁴¹Event-quarter dummies would help control for event-time-specific shocks, but NSW omit these.

⁴²Their point estimate for the teen employment elasticity is -0.035 with all controls equally weighted, but they do not report same-division estimates.

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.⁴³

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.⁴⁴ 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 in section 6.1; the results confirm our previous findings.

B3. 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 also propose a neighboring state-pair design (see Section IV of their working paper Neumark, Salas and Wascher (2013b)). 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

⁴³Incidentally, 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.

⁴⁴Perhaps least important in light of the problems above, their standard errors are certainly wrong (and 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.

no pre-existing trends.

Analogous to the cross-state border county design in Dube, Lester and Reich (2010), 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 Table 12 of their working paper). 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.⁴⁵ 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 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).⁴⁶

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

⁴⁵We 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.

⁴⁶Using 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, significant at the ten percent level, for food and drinking establishments in the ACS. For gas station employment in the CPS and ACS, they find elasticities 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.

specification check suggests that when there is a minimum wage differential between neighboring states, these neighbors are also on different employment paths. In Appendix Figure B1, 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 4.1. 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.⁴⁷ 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. Appendix Figure B1 clearly illustrates that, of the three models, the best-behaved specification for employment is model (2), with division-period effects and 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.⁴⁸

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

⁴⁷NSW 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.

⁴⁸Although 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.

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 4.1), 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.

B4. Trends and cycles: Including recessionary periods

NSW argue that the findings in Allegretto, Dube and Reich (2011) 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. (Some of these proposals are discussed at length only in their working paper Neumark, Salas and Wascher (2013b), but are all referenced in their publication version.)

To fix ideas, consider the following data-generating process:

$$y_{st} = \beta_0 + \beta_1 MW_{st} + \sum_s \delta_s \cdot I_s \cdot t + \gamma_s + \tau_t + \rho_{srt} + \epsilon_{st}. \quad (8)$$

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 omission 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, which they describe in detail in Section III of their working paper. They first estimate equation (1) using a subsample 1994-2007.⁴⁹ 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 the regression excluding the recessionary years. Another simple solution includes dummies

⁴⁹We note that the “non-recessionary” sample used by NSW includes the 2001 recession. Also, the early 1990s recession ended in 1991q1. See <http://www.nber.org/cycles.html>.

for state-specific recession spells. We use the NBER definition of recessions and estimate the model to account for recessions in conventional ways. As Table B2 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 Allegretto, Dube, and Reich (2011) hold when we use standard methods to account for the problem raised by NSW.⁵⁰

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 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 Dube, Lester and Reich (2010) 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 Dube, Lester and Reich (2010, 2013) and in this paper to be the most convincing evidence on the topic.

⁵⁰NSW 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. 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.

B5. Use of spatially correlated placebos

Dube, Lester and Reich (2010) uses a placebo-based falsification exercise to show directly the bias arising from spatial heterogeneity, but NSW argue that this test is inappropriate.⁵¹ 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, Dube, Lester and Reich (2010) examines the counties whose minimum wage is always equal to the federal minimum wage, and also the cross-state neighbors of these counties. Dube, Lester and Reich (2010) 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 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:

$$y_{ipt}^A = \alpha MW_{ipt}^A + \beta MW_{ipt}^B + \gamma_i + \tau_t + \epsilon_{it}. \quad (9)$$

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.⁵² 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

⁵¹Section I.G. and Appendix B of Dube, Lester and Reich (2010) describe the original exercise.

⁵²This should be true except in cases of cross-border spillovers, for which Dube, Lester and Reich (2010) fails to find evidence in a test comparing border and interior counties (see Dube, Lester and Reich (2010), Section VA.)

federal minimum wage, they are perfectly collinear with the time fixed effects τ_t . Estimating (2) is therefore the same as estimating:

$$y_{ipt}^A = \beta MW_{ipt}^B + \gamma_i + \tau_t + \epsilon_{it}, \quad (10)$$


which is what Dube, Lester and Reich (2010) 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 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.

Appendix Figure B1

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-pair-year and state fixed effects. Local controls with the CPS are state and division-year fixed effects, as well as state-specific linear time trends.

Appendix Table B1
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	0.246	0.754	
	Average no. of donor states	2.5	21.5	
	Per donor weight	0.098	0.035	2.806
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	0.009	0.991	
	Average no. of donor counties	1.7	961.7	
	Per donor weight	0.005	0.001	5.138
Log(EPOP)	Proportion of weights	0.007	0.993	
	Average no. of donor counties	1.7	961.7	
	Per donor weight	0.004	0.001	3.988

Notes. Proportion of weights and average number of donor states are from Neumark, Salas and Wascher (2013a), Tables 3 and 5. Relative per donor weight (column 3) equals the ratio of per donor weights in column 1 to that in column 2.

Appendix Table B2
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 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. Significance levels are indicated by *** 1%, ** 5%, * 10%.