CREDIBLE RESEARCH DESIGNS FOR MINIMUM WAGE STUDIES: A RESPONSE TO NEUMARK, SALAS, AND WASCHER

SYLVIA ALLEGRETTO, ARINDRAJIT DUBE, MICHAEL REICH, AND BEN ZIPPERER*

The authors assess the critique by Neumark, Salas, and Wascher (2014) of minimum wage studies that found small effects on teen employment. Data from 1979 to 2014 contradict NSW; the authors show that the disemployment suggested by a model assuming parallel trends across U.S. states mostly reflects differential pre-existing trends. A data-driven LASSO procedure that optimally corrects for state trends produces a small employment elasticity (–0.01). Even a highly sparse model rules out substantial disemployment effects, contrary to NSW's claim that the authors discard too much information. Synthetic controls do place more weight on nearby states—confirming the value of regional controls—and generate an elasticity of −0.04. A similar elasticity (–0.06) obtains from a design comparing contiguous border counties, which the authors show to be good controls. NSW's preferred matching estimates mix treatment and control units, obtain poor matches, and find the highest employment declines where the relative minimum wage falls. These findings refute NSW's key claims.

Recent controversies in minimum wage research have centered on how to credibly estimate employment effects using the extensive state-level variation in minimum wages in the United States. A key concern is that the distribution of minimum wage policies among states has been far from

*Sylvia Allegretto is an Economist at the University of California, Berkeley, Center on Wage and Employment Dynamics. Arindrajit Dube is an Associate Professor at the University of Massachusetts Amherst and is a Research Fellow with the Institute for the Study of Labor (IZA). Michael Reich is a Professor at the University of California, Berkeley. Ben Zipperer is an Economist at the Economic Policy Institute. We are grateful to Doruk Cengiz, Zachary Goldman, Carl Nadler, Thomas Peake, and Luke Reidenbach for excellent research assistance. Financial support for this article came entirely from the authors' respective institutions. The replication package for the article can be found at https://dl.dropboxusercontent.com/u/15038936/adrz2016_replicationpackage_full.zip. Appendices discussed in the article can be found at http://journals.sagepub.com/doi/suppl/10.1177/0019793917692788. Additional results, appendices, and copies of the computer programs used to generate the results presented in the article are available from the corresponding author at adube@econs.umass.edu.

Keywords: minimum wage legislation, minimum wage law compliance, minimum wages, minimum wage trends

ILR Review, 70(3), May 2017, pp. 559–592
DOI: 10.1177/0019793917692788.
Journal website: journals.sagepub.com/home/ilr
Reprints and permissions: sagepub.com/journalsPermissions.nav

random. If we divide the states into two equally sized groups—"high" versus "low" groups based on their average real minimum wages over the 1979 to 2014 period—we find that minimum wage policies are highly spatially clustered. High minimum wage states are concentrated on the Pacific Coast, the Northeast, and parts of the Midwest; tend to be Democratic-leaning; and have experienced less de-unionization. These disparities raise the possibility that trends in other policies and economic fundamentals may also differ between these groups of states.[1]

The nonrandom distribution of state minimum wage policies thus poses a serious challenge to the canonical two-way fixed-effects panel approach, which assumes parallel trends across all states. To account for such heterogeneity, our past minimum wage research—Dube, Lester, and Reich (2010), hereafter DLR, and Allegretto, Dube, and Reich (2011), hereafter ADR—has used either border discontinuities or coarser regional and parametric trend controls, as nearby areas tend to experience similar shocks. When using such strategies, the estimated employment impact for highly affected groups such as restaurant workers or teens tends to be small and often statistically indistinguishable from zero, even though sizable earnings effects occur for these groups. Moreover, DLR and ADR used distributed lags and leads in minimum wages to show that the disemployment effects estimated in the two-way fixed-effects model often reflected pre-existing trends, rather than changes in employment that occurred after policy implementation.[2]

In two articles, Neumark, Salas, and Wascher, hereafter NSW (2014a, 2014b), critiqued the use of local area controls in DLR and ADR. They made three important claims.

First, they defended the results from the two-way fixed-effects estimator, arguing against the evidence that pre-existing trends contaminate those estimates. They also argued that the inclusion of controls for spatial heterogeneity does not produce smaller pre-existing trends.

Second, they argued that using local area controls throws away too much useful information. In the same vein, they claimed that the small magnitudes of the employment estimates in ADR from specifications with state-specific linear trends are driven by an endpoint bias (2014a: 616) generated by the presence of recessions in the beginning and the end of the ADR sample, and that estimates for models that include third-, fourth-, or fifth-order polynomial time trends by state suggest sizable disemployment effects.

Third, NSW proposed a new matching estimator loosely based on the synthetic control approach. They argued that this matching estimator suggests substantial employment effects, at least for teens. They claimed that this

[1]We classified states into high and low minimum wage groups using state-level annual minimum wages adjusted for inflation by the Consumer Price Index Research Series Using Current Methods (CPI-U-RS; U.S. Department of Labor, Bureau of Labor Statistics).

[2]Other minimum wage researchers—for example, Magruder (2013), Huang, Loungani, and Wang (2014), and Aaronson, French, Sorkin, and To (2017)—have subsequently used the border discontinuity design to estimate causal effects of minimum wage policies in both U.S. and international contexts.

approach provides a superior alternative to the methods we have proposed to account for time-varying confounders of minimum wage policies.

We respond here to each of these claims. We note that of the two groups discussed in this exchange (restaurant workers and teens), a substantive disagreement remains mainly for teens. We therefore focus most of our attention on this group.

We begin by presenting recent evidence on teen employment using a border discontinuity design. We review the evidence on whether neighboring counties are indeed more similar in levels and trends of covariates than are counties farther away—thereby assessing a key NSW claim about the validity of local area controls. We then turn to the evidence on teen employment from state panel studies and assess whether controls of unobserved time-varying heterogeneity beyond the two-way fixed effects are warranted. We use Current Population Survey (CPS) data between 1979 and 2014 to estimate the impact of minimum wages on teen employment. Using this expanded sample sheds light on a number of areas of contention, including any endpoint bias in the estimates in ADR's 1990 to 2009 sample, as well as providing a more precise assessment of pre-existing trends.

To provide direct evidence on NSW's contention that the small employment estimates in ADR arise from "arbitrarily throwing away lots of valid identifying information" (2014b: 18), we implement a novel, data-driven approach that adjudicates among different sets of controls: the double-selection post-LASSO estimator (Belloni, Chernozhukov, and Hansen 2014). To assess NSW's second claim that a data-driven control group does not privilege geographic proximity, we review evidence using the synthetic control approach that is presented in Dube and Zipperer (2015). This evidence explicitly shows how the donor weights chosen by synthetic controls vary by distance between the treated and the donor states.

We replicate the NSW (2014a) matching estimates and assess whether their synthetic controls are well-matched to the treated events. We also assess whether many of the events they analyze actually were subject to a clear minimum wage treatment.

Although we mostly focus on teens, we also present new evidence on restaurant employment using updated 1990 to 2014 Quarterly Census of Employment and Wages (QCEW) data and provide medium- and long-run estimates of minimum wage effects on restaurant employment using the border discontinuity design.

Our findings, using a longer sample period and new methods, as well as our re-analysis of NSW data, show clearly that none of the three key claims in NSW withstands scrutiny.

Importance of Teens in the Minimum Wage Literature

The minimum wage literature has extensively studied teens because they are heavily affected by minimum wage policies. Based on the Current

Population Survey Outgoing Rotations Group (CPS ORG) data, during the 1979 to 2014 period, 40.2% of working teens earned within 10% of the statutory minimum wage (higher of state or federal), as compared to 7.7% of workers overall. The relatively large proportion of minimum wage workers among teens makes it relatively easy to detect an effect of the policy on outcomes for this group, thus making them an attractive group to study.

At the same time, the lessons from teens may be limited. First, for an understanding of the impact of the policy more generally, teens are not representative of all minimum wage workers. Second, teens comprise a shrinking share of low-wage workers. Among workers earning within 10% of the statutory minimum wage, the teen share has fallen over time from 32.2% in 1979 to 22.7% in 2014.[3] Finally, labor–labor substitution may imply that some of the teen disemployment effects represent employment gains by other groups.[4] Therefore, estimating an overall impact of minimum wages on affected workers remains an important avenue for future research.

Nonetheless, the high incidence of minimum wage employment among teens suggests that if one is to find disemployment effects of the policy, it will likely be for teens. Therefore, the debate on teen employment still has relevance today.

Evidence for Teens Using a County-Level Border Discontinuity Design

The county-level border discontinuity design provides one of the most compelling identification strategies for estimating minimum wage effects. Dube et al. (2010) developed this approach by comparing contiguous counties straddling state borders, building on the insights of comparing nearby areas in Card and Krueger (1994, 2000). This research design can convincingly account for policy endogeneity because the identifying variation comes from treatments that are typically set at the state level. The estimates of the treatment effects are obtained by comparing adjacent border counties that tend to experience similar economic shocks, but that happen to be in states with different minimum wage policies. When economic shocks on average vary continuously across the border, but state-level policy is a function of

[3]The teen share is calculated for all workers (hourly or otherwise) with positive hourly earnings that are not imputed in the CPS ORG data.

[4]Clemens and Wither (2016) studied a different population of affected low-wage workers and found large, negative employment effects using the federal minimum wage increase during 2009. Although the use of pre-treatment earnings may be a useful way to identify workers affected by the policy, their primary findings are seriously flawed. Zipperer (2016) has shown that much—perhaps all—of the employment reductions found by Clemens and Wither reflect their failure to control for the impact of the Great Recession. In particular, low-wage employment in states bound by the minimum wage increase was much more reliant on the construction sector, which saw a big decline during the downturn. Once we account for the pre-treatment construction share—and pre-treatment sectoral shares more generally—no evidence supports a sizable fall in employment. Tellingly, Zipperer's analysis also shows that the inclusion of geographic controls (i.e., region- or division-specific time effects) largely removes the omitted variables bias arising from the Great Recession. This result provides yet another compelling piece of evidence on the validity of local area controls for identifying minimum wage effects.

shocks in all counties, this approach identifies the causal effect of the policy even if state policies are endogenous to economic conditions affecting the low-wage labor market—allaying the policy endogeneity concern raised in NSW (2014b).

Similarity of Local Areas: Are Contiguous County Pairs More Alike?

NSW (2014a) challenged the motivations behind this design, arguing that neighboring areas do not constitute good controls. Based on their synthetic control donor weights—problems with which we discuss at greater length later—they stated that "the cross-border county is a poor match—no better than a county chosen at random from the list of all potential comparison counties" (632).

DLR (2016) used the county-level Quarterly Workforce Indicators (QWI) data set to assess whether adjacent county pairs are indeed more alike in terms of covariates than are nonadjacent county pairs. DLR (2016) considered six key covariates: log of overall private-sector employment, log population, private-sector 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 minimum wage policies. Table 2 in DLR (2016) 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 noncontiguous pairs; and in all cases but one, the gaps are statistically significant at the 1% level. Many of the gaps, including changes in EPOP, are substantial and exceed 25%. These results contradict the NSW claim that contiguous counties are not more similar to each other than two counties chosen at random.

Slichter (2016) corroborated these findings with a refinement of the contiguous county methodology by comparing counties to their neighbors, neighbors-of-neighbors, neighbors-of-neighbors-of-neighbors, and so on. He showed that immediate neighbors are, indeed, more likely to have experienced similar employment changes just prior to minimum wage increases. He also showed that if unobservables behave like observables, then the border design is much better equipped to control for the unobservables than an approach that uses controls that are much farther away.

Border Discontinuity Results Using QWI Data

DLR (2016) also estimated minimum wage elasticities for teen employment using a border discontinuity approach and county-level QWI data from 2000 through 2011. The estimates on earnings are positive, sizable, and statistically significant at the 1% level. The estimated teen employment minimum wage elasticity from the two-way fixed-effects model is -0.173 and is statistically significant at the 1% level. By contrast, the estimated employment elasticity with the county-pair period effects falls in magnitude to -0.059 and is statistically indistinguishable from zero. Controlling for time-varying heterogeneity using a border discontinuity design therefore suggests employment effects for teens that are substantially smaller than the two-way fixed-effects model.

DLR (2016) also found a sizable reduction in turnover following a minimum wage increase: the turnover elasticity is -0.204 when county-pair period effects are included. Note that in conjunction with the strong earnings effects, the turnover findings undermine NSW's claim that this research design throws away too much information to detect any effects of the policy on outcomes.

Slichter (2016), who employed a neighboring county discontinuity design, reinforces these conclusions. Slichter relaxed the assumption that differences between nearby counties fully eliminate unobservable factors confounded with minimum wage differences. By using untreated neighbors of minimum wage–raising counties, along with additional control groups of neighbors-of-neighbors of treated counties, and so on, Slichter can identify minimum wage effects even when neighboring counties are imperfect controls for one another. This *selection ratio*–based refinement of the border approach produced small employment elasticities for teens that are similar to our findings here, ranging from -0.006 to -0.041 at zero to four quarters after a minimum wage increase.[5]

Effects on Teen Employment: CPS Data Using State-Level Variation

The negative bias in the two-way fixed effects estimate of the minimum wage elasticity for teen employment is also evident in state-level analysis. Using CPS data, ADR showed that the use of state-specific linear trends and division-period effects rendered the employment estimate small and statistically insignificant. This finding contrasted with the sizable, negative estimates from the two-way fixed-effects model. Using distributed lags, ADR found evidence of pre-existing trends in the form of sizable, negative coefficients associated with leading minimum wages when using the two-way fixed-effects model. Moreover, when including controls for state-specific linear trends and division period effects, there was little indication of such pre-existing trends.

NSW (2014a, 2014b) argued against these findings. First, they claimed that major recessions near the endpoints of the ADR sample (1990–2009) led to unreliable estimates of state-specific trends. Second, they contended that the use of third- or higher-order polynomial trends restores the

[5]Liu, Hyclak, and Regmi (2016) used a particular definition of a local area (U.S. Department of Commerce, Bureau of Economic Analysis, BEA-based Economic Areas), QWI data from 2000 to 2009, and the local controls to study teen employment. When controlling for spatial heterogeneity in minimum wage policies by Economic Area time-specific fixed effects, Liu et al. found more sizable negative employment estimates for teens, though not for young adults. Unfortunately, they do not provide evidence on whether their estimates are robust to the particular geographic grouping, or to their sample (which stops in 2009). Their results are at odds with the other estimates in the literature using local area controls (e.g., Dube, Lester, and Reich 2016; Slichter 2016).

findings of a large disemployment effect. Third, they disputed that the data warrant using geographic controls (division-period effects). Fourth, they argued that little evidence supports pre-existing trends in the two-way fixed-effects model, and that using additional spatial controls does not reduce the extent of such pre-existing trends.

In this section, we estimate teen employment and wage elasticities of the minimum wage using individual-level CPS data from 1979 through 2014. The use of this longer time period allows us to better assess each of the key claims in NSW.[6] We begin with estimating a canonical model with time (t) and place (j) fixed effects. Here i denotes an individual, while j denotes the state of residence of individual i:

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

The key independent variable is the log of the quarterly minimum wage (MW_{jt}), which takes on the higher of the federal minimum wage or the minimum wage in state j, while $\mathbf{X}_{\mathrm{it}}$ is a vector of controls.[7] The dependent variable Y_{it} is either the log of hourly earnings or a dummy for whether person i is currently working. For hourly workers, we use their reported hourly wage; for other workers, we construct the wage by dividing their usual weekly earnings by the usual weekly hours worked. We discard all observations with imputed wage data when estimating wage effects.[8] The vector of covariates $\mathbf{X}_{\mathbf{it}}$ includes dummies for gender, race, Hispanic origin, age, and marital status; the teen share of the population in the state; and the non-seasonally adjusted quarterly state unemployment rate.[9] All individual-level regressions are weighted by the basic monthly sample weights or earnings sample weights. We report all the results as elasticities: For earnings equations, the elasticity is simply the estimated coefficient of β, and for

[6]For teen employment, we use individual-level records of 16 to 19 year olds from the Unicon extracts of the full basic monthly sample (https://www.unicon.com/cps.html), and for wage outcomes we use the National Bureau of Economic Research (NBER) Merged Outgoing Rotation Groups (ORG) (http://www.nber.org/morg/).

[7]State-level minimum wages are quarterly means of daily state-level minimum wage levels, or federal minima when they exceed the state law, for all 50 states and the District of Columbia for 1979 to 2014 from Vaghul and Zipperer (2016).

[8]Following Hirsch and Schumacher (2004), we define wage imputations as records with positive allocation values for hourly wages (for hourly workers) and weekly earnings or hours (for other workers) during 1979 to 1988 and September 1995 to 2014. For 1989 to 1993, we define imputations as observations with missing or zero values for the NBER ORG "unedited" earnings variable but positive values for the "edited" earnings variable (which we also do for hours worked and hourly wages). We do not label any observations as having imputed wages during 1994 to August 1995, when there are no Bureau of Labor Statistics allocation values for earnings or wages.

[9]We define race as white, black, or other and interact these dummies and an indicator for Hispanic ethnicity with an indicator for period 2003 and later, as there was a large race and ethnicity classification change in the CPS after 2002. We calculate quarterly teen shares of the age 16 and older population using the full basic monthly sample. We use as the quarterly state unemployment rate the quarterly mean of the non-seasonally adjusted monthly unemployment rate from the Bureau of Labor Statistics Local Area Unemployment series (http://download.bls.gov/pub/time.series/la/).

employment equations, we divide this coefficient by the weighted sample mean of the dependent variable.

In our most saturated specification, we additionally include (up to fifth-order) state-specific time trends and allow the time effects to vary by each of the nine census divisions, denoted by *d*:

$$Y_{it} = \alpha + \beta MW_{jt} + \mathbf{X}_{jt}\Lambda + \gamma_j + \delta_{dt} + \sum_k \left(\phi_{jk} \times t^k\right) + \nu_{it} \tag{2}$$

We report the intermediate specifications with just the state-specific trends and the division-period effects as well as the most saturated specification. Altogether, these 12 specifications—with common or division-period fixed effects and with polynomial trends of degree k = 0, . . . , 5—include the four key specifications used in ADR, which used only linear and not higher-order trends. Three of these specifications—those with linear trends and/or division-period effects—are the ones criticized by NSW (2014a, 2014b).

Main Results for Teens

Panel A of Table 1 reports the wage results from the sample of teens with earnings in the individual-level CPS ORG data from 1979 to 2014. The outcome variable here is the natural log of the hourly wage. All regressions include state fixed effects. The first row includes common-time effects, whereas the second row includes time effects that vary by the nine census divisions. Column (1) contains no allowance for state-specific trends, while columns (2) through (6) add state-specific polynomial trends of successively higher orders. We find that the estimated wage effects are always economically substantial and statistically highly significant. This result holds across the 12 specifications. The wage elasticities are remarkably uniform, ranging between 0.226 and 0.271 for the common-time specification and between 0.215 and 0.256 when including division-period effects. The addition of division-period effects or higher-order trends does not substantially diminish these estimates, contrary to the claim in NSW (2014a: 644) that these more saturated models "have thrown out so much useful and potentially valid identifying information that their estimates are uninformative or invalid."

Panel B of Table 1 reports analogous results for teen employment using the full basic monthly CPS. Note that the employment elasticity is substantial and negative only in the specifications without any state-specific trend controls. Simply including state-specific linear trends reduces the common-time specification estimate in magnitude from −0.214 to −0.062 and renders it statistically insignificant. The finding in ADR that including state-specific trends diminishes the magnitude of the estimated employment effect is replicated in this expanded sample, whose end points (1979, 2014) are notably not recessionary years. The replication of the results in the expanded sample refutes NSW's key argument that the findings in ADR

Table 1. Minimum Wage Elasticities for Average Teen Wage and Employment Controlling for Time Varying Heterogeneity, Individual-Level CPS Data, 1979–2014

	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*
Panel A: Average teen wage						
Common-time FE	0.266***	0.228***	0.226***	0.271***	0.269***	0.267***
	(0.037)	(0.020)	(0.021)	(0.032)	(0.031)	(0.037)
N	295,835	295,835	295,835	295,835	295,835	295,835
Division-period FE	0.247***	0.256***	0.234***	0.230***	0.215***	0.215***
	(0.036)	(0.033)	(0.037)	(0.037)	(0.034)	(0.035)
N	295,835	295,835	295,835	295,835	295,835	295,835
Panel B: Teen employment						
Common-time FE	−0.214***	−0.062	−0.040	−0.061	−0.088	−0.065
	(0.044)	(0.041)	(0.060)	(0.065)	(0.064)	(0.071)
N	3,534,924	3,534,924	3,534,924	3,534,924	3,534,924	3,534,924
Division-period FE	−0.124	0.011	−0.009	−0.019	−0.037	−0.036
	(0.079)	(0.048)	(0.048)	(0.040)	(0.039)	(0.041)
N	3,534,924	3,534,924	3,534,924	3,534,924	3,534,924	3,534,924
State-specific trend type						
Linear		Y	Y	Y	Y	Y
Quadratic			Y	Y	Y	Y
Cubic				Y	Y	Y
Quartic					Y	Y
Quintic						Y

Notes: The table reports minimum wage elasticities for average teen wage and employment, using individual-level CPS data from 1979–2014 (basic monthly data for employment and Outgoing Rotation Groups for wage). The dependent variable is either log wage or a binary employment indicator. For the wage outcome, the table reports the coefficients on log quarterly minimum wage. For employment, the estimates are converted to elasticities by dividing the coefficients on log minimum wage (and standard error) by the sample mean employment rate. All regressions include controls for the quarterly state unemployment rate, the quarterly teen share of the working-age population, dummies for demographic variables as described in the text, and state fixed effects. As reported in the table, specifications include either common-period fixed effects or census division-period fixed effects, with up to fifth-order state-specific polynomial trends. Regressions are weighted by sample weights and robust standard errors (in parentheses) are clustered at the state level.
Significance levels are indicated by *** 1%; ** 5%; and * 10%.

were driven by endpoint bias in the estimation of state trends owing to the presence of recessionary years.[10]

Continuing with the common-time effect models in the first row of Table 1, panel B, when we include state-specific trends of higher order, the

[10]Online Appendix B provides additional evidence that the endpoint bias explanation is incorrect. To summarize those findings, online Appendix Figure B.1 shows estimates from 72 different samples with alternative starting and ending dates varying between 1979 and 1990, and 2009 and 2014, respectively, for specifications with and without state-specific linear trends. Extending the sample by considering end points away from recessionary periods does not produce more negative estimates when state trends are included. Moreover, online Appendix B shows that excluding downturns—either using the official NBER definition or a much more expansive one—does not produce evidence of substantial dis-employment effects in models with state trends.

coefficients are always less negative than −0.09 and none is statistically significant. Four out of five estimates are less than −0.07 in magnitude. These results refute the claim in NSW that inclusion of higher-order (third or greater) state-specific trends restores the finding of a sizable negative effect. Estimation of cubic, quartic, or quintic trends by state places greater demand on the data, especially when the panel is short. By using a substantially longer panel, we estimate these trends more reliably. The estimates from including third- and fifth-order polynomials, −0.061 and −0.065, respectively, are virtually identical to the estimate with just a linear trend (–0.062). The estimate from the second-order trend is slightly smaller in magnitude (–0.040), whereas the estimate from the fourth-order trend is slightly larger in magnitude (–0.088). In all cases, however, the estimates are below −0.09 in magnitude and never statistically significant. Overall, these results suggest that including higher-order trends are unlikely to change the conclusions reached in ADR.

The bottom section of panel B of Table 1 additionally includes division-period effects, isolating the identifying variation to within the nine census divisions. Including division-period effects typically produces estimates that are even less negative. For example, without any state trends (column (1)) the estimate falls from −0.214 to −0.124 in magnitude and is not statistically significant. However, inclusion of state trends renders the estimates close to zero and not statistically significant, with point estimates ranging between −0.037 and 0.011. We note that the lack of statistical significance in the more saturated models is not because of a lack of precision but rather because of the small size of the coefficients.

Overall, the evidence from the state-level CPS data is consistent with the evidence from the county-level QWI data presented above. In both cases, the two-way fixed-effects estimates are sizable and negative, −0.214 and −0.173 in the CPS and QWI, respectively. And the use of coarse controls for time-varying heterogeneity in the CPS (e.g., state trends) produces an employment estimate that is much smaller in magnitude and similar to that using a border discontinuity design (–0.062 and −0.059, respectively).

Manning (2016) conducted a similar analysis with 1979 to 2012 CPS data. He also found that the two-way fixed-effects estimate was unique in producing a large, statistically significant disemployment estimate, and that inclusion of linear and higher-order trends, as well as division-period controls, produced estimates much smaller in magnitude.

Model Selection Using LASSO

The variation in the estimates reported in Table 1 raises a fundamental question: What is the best set of controls to include in these regressions? In this section, we address this question by applying the double-selection post-LASSO approach advocated by Belloni et al. (2014). Using sparsity as a criterion for covariate selection, the LASSO regression is able to identify a small set of key predictors from a large set of potential variables, assuming such a sparse representation is feasible. The double-selection criteria apply LASSO to a program evaluation context to select the most important predictors of the outcome (in our case teen employment) or the treatment (log minimum wage). After having selected the covariates using these two LASSO regressions, Belloni et al. suggested running a simple OLS regression of the outcome on the treatment and the double-selected set of controls (hence the term "post-LASSO").[11]

As a first step, we estimate all the specifications in Table 1 using aggregated data (computational challenges in estimating LASSO with a large number of observations and variables require us to use data aggregated at the state-quarter level). These regressions are similar to those estimated in NSW (2014a, 2014b). We regress the log of the teen employment-to-population ratio on the log of the minimum wage, the state unemployment rate, and the teen share of population, while additionally controlling for state fixed effects, either common (or division-specific) period effects, and possible state-specific time trends. We also include demographic group shares analogous to covariates in the individual-level regressions: shares by gender, age groups, race categories, and marital status. We additionally weight all regressions by the state teen population.

These results, reported in Table 2, panel A, show that in most cases aggregation does not make much of a difference. The two-way fixed-effects model produces an elasticity that is substantial (–0.168) and statistically significant, whereas all of the other 11 coefficients are less negative than −0.09 and are not statistically significant.

For model selection, we estimate two LASSO regressions of the log of teen EPOP and the log minimum wage over a set of covariates: the unemployment rate, teen share of population, demographic group shares as specified above, division-period dummies, and state-specific time trends of orders one through five. The LASSO regressions partial out state and time fixed effects prior to estimation. With the superset of controls chosen by these two LASSO regressions, we estimate an OLS regression that also includes state and time fixed effects. Online Appendix A provides additional technical details of the LASSO estimation.

Column (8) of Table 2 reports the estimates from our double-selection post-LASSO regression allowing the full set of controls. Although not shown in the table, with the default recommended penalization parameter ($\lambda = 940$), the double-selection criteria for teen employment picks division-period effects from one census division (the Pacific division), 29 state-specific linear trends, and no higher-order trends. The resulting point estimate (–0.009) is numerically close to, and statistically indistinguishable

[11]This post-LASSO approach leverages the advantages of LASSO-based selection of the most important controls, while guarding against the shrinkage bias in LASSO coefficients attributable to the penalization term.

Table 2. Model Selection: Minimum Wage Elasticities for Teen Employment, State-Quarter Aggregated CPS Data, 1979–2014

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Full sample (1979–2014)								
Common-time FE	−0.168**	0.025	0.004	−0.051	−0.084	−0.069		
	(0.066)	(0.081)	(0.075)	(0.078)	(0.081)	(0.085)		
N	7,344	7,344	7,344	7,344	7,344	7,344		
Division-period FE	−0.037	0.059	0.058	0.038	0.006	0.005		
	(0.088)	(0.057)	(0.056)	(0.049)	(0.049)	(0.055)		
N	7,344	7,344	7,344	7,344	7,344	7,344		
LASSO-selected division-period FE							0.015	−0.009
							(0.082)	(0.083)
N							7,344	7,344
Panel B: Post-1990 sample (1990–2014)								
Common-time FE	−0.100	0.009	−0.053	−0.141**	−0.168**	−0.199***		
	(0.065)	(0.078)	(0.065)	(0.067)	(0.068)	(0.063)		
N	5,100	5,100	5,100	5,100	5,100	5,100		
Division-period FE	−0.021	0.076	0.051	−0.006	−0.015	−0.053		
	(0.093)	(0.063)	(0.061)	(0.057)	(0.070)	(0.062)		
N	5,100	5,100	5,100	5,100	5,100	5,100		
LASSO-selected division-period FE							−0.002	−0.024
							(0.072)	(0.069)
N							5,100	5,100
State-specific trend type								
Linear		Y	Y	Y	Y	Y		
Quadratic			Y	Y	Y	Y		
Cubic				Y	Y	Y		
Quartic					Y	Y		
Quintic						Y		
LASSO-selected trends (linear only)							Y	
LASSO-selected trends (up to quintic)								Y

Notes: The table reports minimum wage elasticities for teen employment using state-quarter aggregated CPS basic monthly data from 1979–2014. The dependent variable is the log of the state-quarter sample-weighted mean of teen employment. The reported estimates are coefficients for log quarterly minimum wage. All regressions include controls for the overall quarterly state unemployment rate, the quarterly teen share of the working-age population, and state-quarter means for demographic controls used in Table 1 and state fixed effects. Specifications include either common-period effects or census division-period effects, and up to fifth-order polynomial trends by state. Columns (7)–(8) report double-selection post-LASSO estimates in which controls (besides state and period effects) are selected using LASSO regressions predicting teen employment and minimum wage: these include demographic controls, division-period effects, and state-specific trends (linear in specification 7; up to quintic in specification 8). Regressions are weighted by teen population. Robust standard errors in parentheses are clustered by state.
Significance levels are *** 1%; ** 5%; and * 10%.

from, zero. The results from this exercise confirm that the controls for time-varying heterogeneity used in ADR—especially state trends—should be included, and that the data-driven set of controls suggests a minimum wage elasticity for teen employment that is close to zero. Comparing across

Figure 1. Double-Selection Post-LASSO Estimates for Minimum Wage Elasticity for Teen Employment, for Alternative Values of the LASSO Penalization Parameter, State-Quarter Aggregated CPS Data, 1979–2014

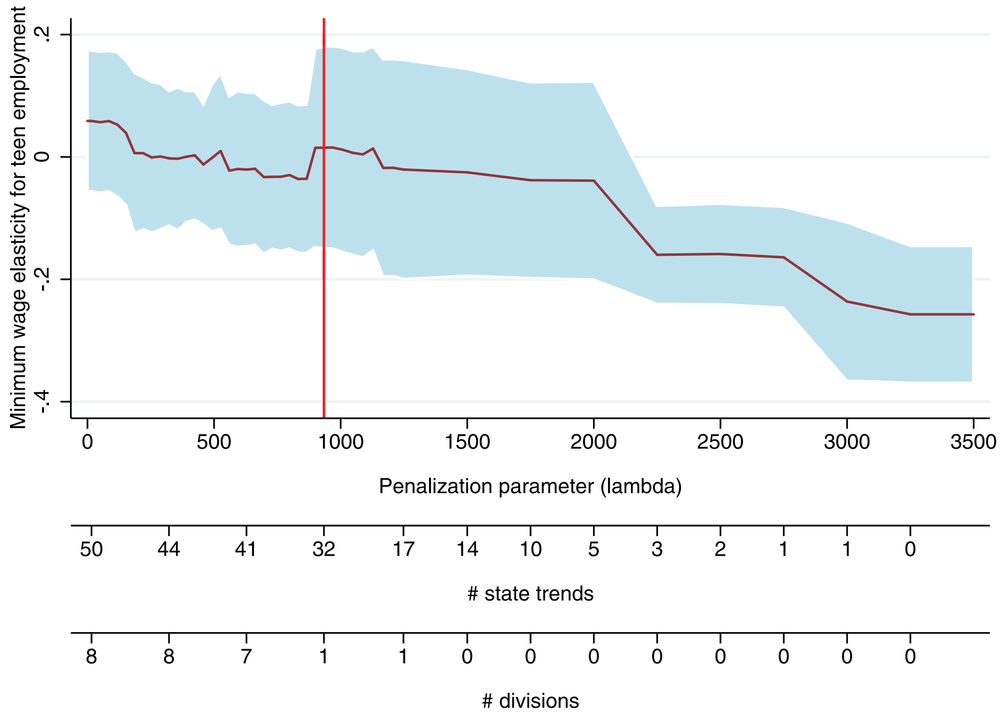

Notes: The figure reports double-selection post-LASSO estimates of minimum wage elasticity for teen employment and associated 95% confidence intervals for alternative values of LASSO penalization parameter, λ, as described in the text. For each value of λ, two LASSO regressions (on log minimum wage and log teen employment) are used to select state-specific linear trends and division-period fixed effects, and demographic controls after partialing out state and period fixed effects, using state-quarter aggregated CPS data. The subsequent post-LASSO regression of log teen employment on log of the quarterly minimum wage controls for the LASSO-selected controls, as well as state and period fixed effects. Additional horizontal axes in the figure report the number of state-specific linear trends and the number of divisions picked for division-period fixed effects picked by the double selection procedure for each value of λ. Standard errors are clustered at the state level. The estimates for this graph are also reported in online Appendix Table A.1 at http://journals.sagepub.com/doi/suppl/10.1177/0019793917692788.

columns (7) and (8), it makes no material difference if higher-order trends are allowed.

The estimates in the top panel of Table 2 are based on a penalization parameter λ that is chosen optimally, using the default plug-in method. To assess how inclusion of the most important controls (as deemed by the double-selection criteria) affects the minimum wage estimate, we also vary λ between a saturated specification with linear trends and division-period effects, and the simple two-way fixed-effects model. Figure 1 shows visually how the point estimates and the confidence intervals change as we vary λ between 0 (the most saturated model) and 3,500, which picks only the state unemployment rate as a control beyond the manually specified two-way fixed effects. (The numerical estimates are in online Appendix Table A.1.)

Starting with the canonical two-way fixed-effects estimate of -0.257, the point estimate quickly falls in magnitude to -0.039 as λ is lowered to 2,000 and never takes on a more negative value for smaller levels of λ. At $\lambda = 2,000$, the double-selection post-LASSO procedure includes just five state-specific linear trends and yet lowers the elasticity in magnitude to -0.039. In other words, merely adding state-specific linear trends for these five states (CA, SD, OR, WA, and VT) to the fixed-effects model produces an estimate that is close to zero and not statistically significant.[12] We stress that this highly sparse model, which adds only five controls for unobserved heterogeneity beyond the canonical two-way fixed-effects model, nonetheless delivers the same qualitative finding as in ADR. This result contradicts the suggestion by NSW that ADR's findings were driven by "throwing out the identifying baby along with, or worse yet instead of, the contaminated bathwater" (2014a: 611).

For comparability to the results in NSW (2014a), we also report in the bottom panel of Table 2 the double-selection post-LASSO estimates for the sample restricted to 1990 and later. The estimates across specifications in this shorter sample exhibit greater variation. Here, too, however, the double-selection post-LASSO estimate is small in magnitude (–0.024) and not statistically distinguishable from zero. The estimate for this shorter sample is based on 20 state-specific linear trends; note that, as before, no non-linear trends are picked. Therefore, although the shorter sample produces more varied estimates using OLS and alternative trend specifications—likely attributable to the imprecision of estimating many higher-order trends—a data-driven choice of predictors that considers higher-order trends produces an estimate that is close to zero in this sample as well. Online Appendix B provides additional evidence and discussion of the unreliability of estimates with higher-order trends in short panels; employment estimates are much more sensitive to the order of the polynomial for state-specific trends in samples with fewer years.

Overall, model selection techniques that make no prior assumptions about which controls should be included in a regression both confirm our approach of including controls for time-varying heterogeneity and support our original conclusion about the size of the minimum wage elasticity for teen employment.

Timing of the Employment Effects

Estimates from a given research design are less credible if the effects appear to occur substantially prior to treatment—such a pattern indicates the likelihood of contamination from pre-existing trends. In prior work (DLR 2010; ADR 2011), we used a distributed lag model to demonstrate that pre-existing trends contaminate the estimates of the conventional two-way

[12]Four of the five states are coastal, showing the importance of obtaining a valid counterfactual for the high minimum wage Pacific division. When estimating state-specific trends, the omitted state is Alabama.

fixed-effects model, which often exhibits sizable and statistically significant leading effects. Nonetheless, NSW (2014b) raised questions about our findings on pre-existing trends for teen employment. First, they argued that pre-existing trends are not clearly indicated in the two-way fixed-effects model. Second, they argued that even after differencing out the leading effects, the subsequent cumulative effects remain negative, sizable, and comparable to the static estimates. Third, they argued that the inclusion of controls for spatial heterogeneity did not produce better results, in the sense of passing the leading effects falsification test.

To shed light on this disagreement, we use exactly the same distributed lag structure as in NSW (2014b). That is, we add 12 quarters of leading and 12 quarters of lagged minimum wages to our prior static specifications in Equations (1) and (2). We estimate these regressions using the individual-level CPS data and control sets we used before for teens in the 1979 to 2014 period using four specifications. Beginning with the two-way fixed-effects model

$$Y_{it} = \alpha + \sum_{k=-12}^{12} \beta_k MW_{j,t-k} + \mathbf{X}_{it}\Lambda + \gamma_j + \delta_t + \nu_{it} \tag{3}$$

we increasingly saturate the model to include state-specific linear time trends and division-period fixed effects

$$Y_{it} = \alpha + \sum_{k=-12}^{12} \beta_k MW_{j,t-k} + \mathbf{X}_{it}\Lambda + \gamma_j + \delta_{dt} + \phi_j \times t + \nu_{it} \tag{4}$$

We also report estimates from the two intermediate specifications—with just division-time fixed effects and state-specific linear trends. We calculate the cumulative employment response from these four models by summing the coefficients for individual leads and lags and convert them to elasticities by dividing by the sample mean of teen employment rate: therefore, the cumulative response elasticity at event time τ (in quarters) is calculated as $\rho_\tau = \sum_{k=-12}^{\tau} \eta_k = \frac{1}{\bar{Y}} \sum_{k=-12}^{\tau} \beta_k$. Note that these cumulative responses are from a default baseline of $\tau < -12$; we will consider alternative baselines below by subtracting leading coefficients from the cumulative responses.

Performance of the Two-Way Fixed-Effects Model

Column (1) of Table 3 shows four-quarter averages of these quarterly cumulative response elasticities: $\bar{\rho}_{[\tau,\tau+3]} = \frac{1}{4}\sum_{m=0}^{3} \rho_{\tau+m}$, along with standard errors. Online Appendix C, Figure C.1, shows the raw cumulative responses underlying the estimates in the table.

For the two-way fixed-effects model, the four-quarter averages of the leading cumulative response elasticity $\bar{\rho}_{[-12,-9]}$ is -0.144 and is statistically

Table 3. Dynamic Minimum Wage Elasticities for Teen Employment, Individual-Level CPS Data, 1979–2011

		(1)	(2)	(3)	(4)
Panel A: Four-quarter averages of cumulative response elasticities					
A	$\bar{\rho}_{[-12,-9]}$	−0.144**	−0.094	−0.057	−0.008
		(0.072)	(0.057)	(0.050)	(0.046)
B	$\bar{\rho}_{[-8,-5]}$	−0.199**	−0.206**	−0.101	−0.098
		(0.089)	(0.080)	(0.071)	(0.067)
C	$\bar{\rho}_{[-4,-1]}$	−0.190**	−0.155	−0.058	0.005
		(0.085)	(0.113)	(0.062)	(0.094)
D	$\bar{\rho}_{[0,3]}$	−0.271***	−0.204	−0.108**	0.003
		(0.068)	(0.132)	(0.051)	(0.100)
E	$\bar{\rho}_{[4,7]}$	−0.383***	−0.300*	−0.177***	−0.039
		(0.078)	(0.165)	(0.057)	(0.131)
F	$\bar{\rho}_{[8,11]}$	−0.319***	−0.220	−0.121*	0.065
		(0.098)	(0.161)	(0.063)	(0.121)
G	ρ_{12+}	−0.296***	−0.205	0.007	0.166
		(0.112)	(0.195)	(0.065)	(0.131)
Panel B: Medium-run (three-year) elasticities					
F-A	$\bar{\rho}_{[8,11]} - \bar{\rho}_{[-12,-9]}$	−0.175***	−0.126	−0.064	0.072
		(0.049)	(0.121)	(0.048)	(0.091)
F-B	$\bar{\rho}_{[8,11]} - \bar{\rho}_{[-8,-5]}$	−0.120***	−0.014	−0.019	0.163**
		(0.040)	(0.097)	(0.050)	(0.069)
F-C	$\bar{\rho}_{[8,11]} - \bar{\rho}_{[-4,-1]}$	−0.129***	−0.065	−0.063	0.060
		(0.040)	(0.071)	(0.042)	(0.051)
Panel C: Long-run (four-plus-years) elasticities					
G-A	$\rho_{12+} - \bar{\rho}_{[-12,-9]}$	−0.152**	−0.111	0.064	0.174
		(0.067)	(0.156)	(0.063)	(0.104)
G-B	$\rho_{12+} - \bar{\rho}_{[-8,-5]}$	−0.097	0.001	0.108	0.264***
		(0.058)	(0.135)	(0.073)	(0.085)
G-C	$\rho_{12+} - \bar{\rho}_{[-4,-1]}$	−0.106*	−0.049	0.065	0.162**
		(0.060)	(0.109)	(0.062)	(0.067)
Division-period FE			Y		Y
State-specific linear trends				Y	Y

Notes: The table reports cumulative response elasticities of teen employment with respect to minimum wages using individual-level CPS basic monthly data from 1979–2011 (using minimum wage data through 2014). Regressions include the contemporaneous, 12 quarterly leads and 12 quarterly lags of log minimum wage. The dependent variable is a binary employment indicator and estimates are converted to elasticities by dividing the log minimum wage coefficients and standard errors by the sample mean employment rate. Panel A reports four quarter averages of the cumulative response elasticities starting at $t = -12$ in quarterly event time, as described in the text. Panel B reports the cumulative effect in year three, after subtracting alternative baseline levels at 1, 2, or 3 years prior to treatment, as indicated. Panel C reports the long-run cumulative response elasticity at $t = 12$ or later, after subtracting alternative baseline levels. All regressions include controls for the overall quarterly state unemployment rate, the quarterly teen share of the working-age population, dummies for demographic controls used in Table 1, as described in the text, and state and period fixed effects. Specifications may additionally include census division-period fixed effects and state-specific linear trends. Regressions are weighted by sample weights and robust standard errors are clustered at the state level.
Significance levels are indicated by *** 1%; ** 5%; and * 10%.

significant at the 5% level (row A, column (1) of Table 3). In other words, during the third year prior to the minimum wage increase, the magnitude of the average cumulative response elasticity is implausibly large and roughly two-thirds the size of the static employment elasticity of -0.214 (see Table 1). The average cumulative response elasticities during the second and the first year preceding the minimum wage increase ($\bar{\rho}_{[-8,-5]}$ and $\bar{\rho}_{[-4,-1]}$) are even more negative, -0.199 and -0.190, respectively; both are statistically significant at the 5% level. In sum, using the full 1979 to 2014 sample, we find unmistakable evidence that the two-way fixed-effects model fails the falsification test that leading coefficients during one, two, or three years prior to treatment are zero. And since the leading effects are occurring two or three years prior to treatment, they cannot plausibly result from anticipation of the policy.

We additionally find robust evidence that a sizable portion of the two-way fixed-effects estimate accrues prior to treatment. A natural approach to net out such leading effects would simply be to accumulate the contemporaneous and lagged coefficients only to form the cumulative response: $\sum_{k=0}^{\tau} \eta_k$. (In our notation, $\sum_{k=0}^{\tau} \eta_k = \rho_\tau - \rho_{-1}$; that is, this approach takes ρ_{-1} as the baseline.) Because individual leading coefficients exhibit considerable noise, however, the choice of the baseline quarter can matter (e.g., see online Appendix Figure C.1). We therefore use estimates with alternative baselines averaging over quarters.

Table 3 calculates estimates for three- and four-plus-year effects from the policy. For the medium term, or three-year estimates, we begin by calculating the average cumulative response elasticity in the third year following the minimum wage increase $\bar{\rho}_{[8,11]}$ and subtracting from this the baseline value. We use three different baselines: the average cumulative response in the first, second, or third year preceding the increase, that is, $\bar{\rho}_{[-4,-1]}, \bar{\rho}_{[-8,-5]}$, or $\bar{\rho}_{[-12,-9]}$, respectively. For example, using the first year before treatment as the baseline, the three-year estimate is $\bar{\rho}_{[8,11]} - \bar{\rho}_{[-4,-1]}$. We also construct long term, or four-plus-year estimates, as $\rho_{12} - \bar{\rho}_{baseline}$, where the baseline can again be $\bar{\rho}_{[-4,-1]}, \bar{\rho}_{[-8,-5]}$, or $\bar{\rho}_{[-12,-9]}$.[13]

The three- and four-plus-year estimates for the fixed-effects model are reported in panels B and C, column (1) of Table 3. Overall, these results show that for the two-way fixed-effects model, both three- and four-plus-year estimates are substantially smaller than the estimate from the static specification. Whereas the static estimate from Table 1 is -0.214, the three-year and the four-plus-year estimates range between -0.097 and -0.129 when using $\tau \in [-4, -1]$ or $\tau \in [-8, -5]$ averages as baselines. Although some of these estimates are statistically significant, results show a 40 to 55% reduction in the effect size, as compared to the static estimate, which implicitly uses a mixture baseline $\tau < 0$. Using an earlier baseline ($\tau \in [-12, -9]$) produces three- and four-year estimates of -0.175 and -0.152 (rows F-A and G-A), and using an even earlier baseline of $\tau < -12$ (i.e., the average

[13]We say "four-plus year" because ρ_{12} reflects the cumulative response at or after the 12th quarter following a minimum wage increase.

cumulative response elasticities in rows F and G themselves) produces estimates around -0.3 in magnitude. This pattern of more negative estimates when using earlier baselines is consistent with a bias due to pre-existing trends that are unaccounted for by the two-way fixed-effects model.[14]

These results differ from those in NSW (2014b), which denied evidence of pre-existing trends in the two-way fixed-effects model. In the article, they also argued that netting out the leading coefficients does not alter the estimates very much. To reconcile our two sets of results, we estimate analogous regressions using their data and specification (i.e., state-by-quarter level data from 1990q1–2011q1; see online Appendix C).[15] Online Appendix Table C.1 reports estimates similar to Table 3 but with the NSW data. We also show the cumulative responses at quarterly frequency using the full 1979 to 2014 sample (online Appendix Figure C.1) as well as the NSW data (online Appendix Figure C.2).

To summarize the findings in online Appendix C, the conclusion in NSW (2014b) arises entirely from their choice of ρ_{-2} as the baseline, which was unusually positive. A variety of alternative baselines shows that much of the employment reduction estimated by the two-way fixed-effects model occurs substantially prior to a minimum wage increase. By contrast, models with controls for state-specific trends tend to have smaller leading coefficients. Using a baseline of one or two years preceding the minimum wage increase produces employment estimates that are substantially smaller: none of the three- or four-plus-year-out effects exceed -0.1 in magnitude regardless of controls for state-specific trends or division-period effects. Although the precision of some of the estimates is lower in the smaller NSW sample, the conclusions from that sample are qualitatively similar to those from the full 1979 to 2014 sample we use in this article.

Performance of Models with Controls for Spatial Heterogeneity

Table 3, columns (2), (3), and (4) show the four-quarter averaged coefficients $\bar{\rho}_{[\tau, \tau+k]}$ for models with controls for spatial heterogeneity. In almost all cases the magnitudes of the leading averaged cumulative responses are smaller: Of the nine leading coefficients from the three models, only one is statistically significant at the 5% level ($\bar{\rho}_{[-8,-5]}$ in column (2) with just division-period controls), in contrast to the two-way fixed-effects model in which all three of the averaged leads are significant. Both the model with state linear trends (column (2)) and additional division-period effects (column (3)) perform well in terms of the leading effects falsification test.

[14]While netting out the leading effects should reduce bias due to pre-existing trends, the reduction may not be sufficient. If a particular model (such as the two-way fixed-effects model) produces very different estimates after netting out the leading effects, researchers should search for models that perform better in such a diagnostic test.

[15]We use the replication data on Ian Salas' website (https://sites.google.com/site/jmisalas/data-and-code) and estimate this model using exactly the same data, sample, and specification that produced NSW (2014b) figure 6. They included controls for unemployment rate, state, and period fixed effects.

What do these models with controls for state-specific trends and division-period effects imply about medium (three-year) and longer-run (four-plus-year) effects from the policy? In our full sample, when using either four quarters just prior to treatment ($\bar{\rho}_{[-4,-1]}$), or the four preceding quarters ($\bar{\rho}_{[-8,-5]}$) as the baseline, the medium- or long-run estimates range between -0.065 and 0.264 (rows F-B, F-C, G-B, G-C from Table 3, columns (2)–(4)).[16] In other words, there is scant indication of medium- or long-term disemployment effects in any of these models.

One concern with parametric trend controls is that they may incorrectly reflect delayed effects of treatment (Wolfers 2006; Meer and West 2016). However, including 12 quarters of leads and lags in our dynamic specifications means that the trends are identified using only variation outside of the 25-quarter window around minimum wage increases and are unlikely to reflect lagged or anticipation effects.

When using the four quarters prior to treatment as a baseline, the long-run estimates in Table 3 for models with some controls for time-varying heterogeneity range between -0.049 (column (2)) to 0.162 (column (4)). These estimates compare to an estimate of -0.106 from the two-way fixed-effects model (column (1)). Two limitations are important when interpreting these longer-term effects. First, the variation to estimate these effects is more limited, making them less precise. Second, different from short- and medium-term effects, the four-plus-year effects influence the estimation of state-specific trends. With those caveats in mind, we find little indication of more negative impacts in the longer run.

First-Difference versus Deviations-from-Means Estimators

When using state-aggregated data, first-differencing is an alternative to taking deviations-from-means for purging the state fixed effects. Although each approach has its advantages, the first-difference estimator is less prone to bias if the state effects are not fixed and are time-varying instead.

Therefore, as an alternative, we estimate the model in first-differences using state (j) by year (t) aggregated data, while including up to three annual lags in the average minimum wage. The baseline first-difference specification is

$$\Delta Y_{jt} = \alpha + \sum_{k=0}^{3} \eta_k \Delta MW_{j,t-k} + \Delta \mathbf{X}_{jt}\Lambda + \delta_t + \nu_{jt} \tag{5}$$

As before, we saturate this baseline model to account for division-period effects, as well as state-specific trends. In the first-differenced version, adding state fixed effects is analogous to including state-specific linear trends in the deviations-from-means version (since the first-differencing purges the

[16]This conclusion is qualitatively similar in the NSW sample (online Appendix C, Table C.1, columns (2), (3), and (4)) in which the equivalent range is (–0.033, 0.395).

state fixed effects). We also report two intermediate specifications with just state fixed effects or just division-period effects. The four specifications are very close to the specifications estimated by Meer and West (2016), who argued that the delayed effects of minimum wages on total employment mostly occur within two to three years of the implementation of the policy. We report estimates both with and without teen population weights and with and without leads in log minimum wage.[17]

Table 4 reports the cumulative three-year minimum wage elasticities for teen employment $\rho_3 = \sum_{k=0}^{3} \eta_k$, as well as the contemporaneous elasticity η_0. For comparability, panel A reports estimates from the models using the deviations-from-means estimator—as in previous sections—and broadly reproduces the results in Table 3 using annual data. In column (1), the contemporaneous and the three-year cumulative elasticity are sizable and negative, ranging between -0.220 and -0.146 depending on weights, and three out of the four estimates are statistically significant at the 5% level. By contrast, the estimates with controls for state trends and division-period effects, or when including leading minimum wage as controls, tend to be more positive; and none of the negative coefficients are statistically significant.

Panel B of Table 4 reports the first-difference estimates. Now the two-way fixed-effects model in column (1) produces estimates ranging between -0.007 and 0.143, and none of these estimates are statistically significant. To emphasize, the sizable negative estimates of the two-way fixed-effects model obtain only when the model is estimated using deviations-from-means, and not first-differences—and is true even when we account for up to three years of lags in minimum wages. This result is consistent with the idea that the first-difference estimates are less likely to be picking up time-varying heterogeneity correlated with the minimum wage.

Estimates in columns (2), (3), and (4) of Table 4 further control for state fixed effects and division-period effects, and those in columns (5) to (8) that additionally control for leading minimum wages tend to suggest smaller (or no) disemployment effects; none of the negative coefficients are statistically significant. To emphasize, none of the first-difference estimates in Table 4—whether or not they include additional controls for time-varying heterogeneity—suggest substantial employment loss, even three years after the increase in minimum wage.

We make one additional observation about the results in Table 4. Meer and West (2016) criticized the inclusion of state-specific trends and argued that they

---

[17]We have chosen to weight the state-aggregated regressions by teen population weights in most parts of the article, so they correspond more closely to estimates using individual-level data (see Angrist and Pischke 2009 for a discussion). The first-difference specification, however, does not have a corresponding individual-level representation, and the rationale for using weights is less clear. For this reason, we report weighted and unweighted variants of regressions in Table 4. For the first-difference specification, weights are defined as $\frac{pop_t \times pop_{t-1}}{pop_t + pop_{t-1}}$. (Borjas, Freeman, and Katz [1997] provide a discussion of weights in differenced specification.)

Table 4. Minimum Wage Elasticities for Teen Employment: Deviations-from-Means versus First-Difference Estimates, State-Year Aggregated CPS Data

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Deviations-from-means								
Population weighted								
Contemporaneous MW elasticity	−0.158**	−0.005	0.047	0.110*	0.005	0.079	0.036	0.114*
	(0.074)	(0.087)	(0.094)	(0.063)	(0.080)	(0.070)	(0.076)	(0.062)
3-year cumulative MW elasticity	−0.146	0.015	0.223*	0.250**	−0.075	0.060	0.140	0.243**
	(0.120)	(0.175)	(0.127)	(0.105)	(0.098)	(0.155)	(0.108)	(0.114)
Unweighted								
Contemporaneous MW elasticity	−0.160**	−0.026	0.003	0.111	−0.035	−0.005	0.002	0.040
	(0.064)	(0.084)	(0.063)	(0.071)	(0.071)	(0.080)	(0.071)	(0.079)
3-year cumulative MW elasticity	−0.220**	−0.102	0.140*	0.200*	−0.138*	−0.040	0.101	0.169*
	(0.090)	(0.132)	(0.071)	(0.089)	(0.079)	(0.123)	(0.073)	(0.095)
Division-period FE		Y		Y		Y		Y
State-specific linear trends			Y	Y			Y	Y
Controls for leads in minimum wage					Y	Y	Y	Y
Panel B: First-difference								
Population weighted								
Contemporaneous MW elasticity	0.030	0.093	0.037	0.100*	0.024	0.092	0.032	0.099*
	(0.082)	(0.058)	(0.085)	(0.058)	(0.078)	(0.058)	(0.081)	(0.058)
3-year cumulative MW elasticity	0.143	0.330**	0.158	0.343**	0.121	0.375**	0.147	0.399**
	(0.137)	(0.142)	(0.142)	(0.145)	(0.134)	(0.165)	(0.145)	(0.176)
Unweighted								
Contemporaneous MW elasticity	−0.007	0.007	−0.001	0.014	−0.027	0.009	−0.023	0.015
	(0.060)	(0.069)	(0.062)	(0.071)	(0.070)	(0.073)	(0.072)	(0.074)
3-year cumulative MW elasticity	0.020	0.033	0.035	0.051	−0.051	0.054	−0.036	0.075
	(0.091)	(0.128)	(0.093)	(0.133)	(0.099)	(0.129)	(0.106)	(0.137)
Division-period FE		Y		Y		Y		Y
State FE			Y	Y			Y	Y
Controls for leads in minimum wage					Y	Y	Y	Y

Notes: The table reports contemporaneous and three-year cumulative minimum wage elasticities for teen employment using state-year aggregated CPS basic monthly data: the sample is 1979–2014 for Panel A and 1980–2014 for Panel B. All specifications include the contemporaneous log annual minimum wage, and three years of lags of the log annual minimum wage, in levels or differences. The dependent variable is the log of the state-year sample-weighted mean of teen employment (in levels or differences). All regressions include controls for the overall quarterly state unemployment rate, the quarterly teen share of the working-age population, and state-year means for demographic controls used in Table 1 in levels or differences. The table reports the coefficient on the contemporaneous log minimum and the sum of the contemporaneous and lagged terms. Estimates in panel A are from the deviation-from-means estimator, and estimates in panel B are from the first-difference estimator. The deviation-from-means specifications always include state fixed effects, and may additionally include state linear trends as indicated. The first-difference specifications may additionally include state fixed effects as indicated. All specifications include period fixed effects and may additionally include division-period effects as indicated. Columns (5)–(8) also control for three years of leading minimum wages (in levels or differences). Regressions are unweighted or weighted by the state-year teen population size, as indicated. Robust standard errors (in parentheses) are clustered at the state level.
Significance levels are indicated by *** 1%; ** 5%; and * 10%.

produce spuriously small disemployment estimates because trends soak up lagged effects. This argument is categorically not true here. Using Meer and West's preferred distributed-lag first-difference specification also produces an employment estimate for teens that is close to zero, similar to estimates with state-specific trends but different from the two-way fixed-effects estimate in levels. Relatedly, we note that the negative employment effects for aggregate employment reported in Meer and West do not appear in analogous specifications for teen employment, at least with state-level CPS data from 1979 to 2014 (close to their sample of 1977–2011 using Business Dynamics Statistics data). For their baseline specification, they found three-year cumulative elasticities for total private-sector employment of −0.074 (column (1) of their table 4). By contrast, our closest first-difference specification (unweighted, with state fixed effects, without leads) in Table 4 (panel B, column (3)) suggests an elasticity for teen employment of around 0.035. Table 4 thus raises questions about whether the findings that minimum wages reduce aggregate employment in Meer and West (2016) are likely to reflect causal effects.[18]

Controlling for Endogeneity Using Factor Models and Synthetic Controls

Existing Estimates

NSW (2014a) proposed a matching estimator based on synthetic control weights that obtains sizable and statistically significant employment elasticities for teens of about −0.14. In this section we contrast this finding with other existing results based on synthetic controls and factor models.

The synthetic control approach of Abadie, Diamond, and Hainmueller (2010) offers one way to account for time-varying factors that may contaminate the estimation of the minimum wage effect. For a single treatment event in which a state raises its minimum wage, the procedure constructs a vector of weights over a set of untreated donor states, such that the weighted combination of donor states closely matches the treated state in pre-intervention outcomes.

Dube and Zipperer (2015) used the synthetic control approach to estimate minimum wage effects on teen wages and employment for 29 state minimum wage–increasing events during 1979 to 2013 and then pooled the results from these individual case studies. The minimum wage is clearly binding in their sample: 25 of 29 wage elasticities were positive and the mean and median wage elasticities were 0.237 and 0.368, respectively. By contrast, 12 of the employment elasticities were positive, and the mean and the median employment elasticities were relatively small: −0.051 and −0.058, respectively. Dube and Zipperer (2015) also extended the donor-based randomization inference procedure suggested by Abadie et al. (2010) to multiple events. They calculated a 95% confidence interval for the pooled employment elasticity of (–0.170, 0.087), which statistically rejects the point estimate of −0.214 that we find above for the OLS two-way fixed-effects model.

Dube and Zipperer's (2015) implementation of the synthetic control estimator contrasts sharply with that of NSW. Whereas NSW's event study problematically assigned many minimum wage–raising states to the

[18]The lack of evidence for teen disemployment using the first-difference specification holds whether we include the state-level unemployment rate as a control and whether we restrict the sample to 1990 and later (results not shown).

potential donor group, Dube and Zipperer (2015) kept the treatment–control distinction clear, as required by the case study approach of the synthetic control estimator. To obtain better matches, Dube and Zipperer (2015) imposed a pre-treatment window of at least two years and up to four years, but NSW used only a one-year pre-treatment period, the shortest pre-treatment length we are aware of in the literature using synthetic controls. These restrictions, along with requirements of at least five potential donors and a 5% nominal minimum wage increase, reduce Dube and Zipperer's (2015) sample from 215 state-level quarterly minimum wage changes to 29 events, with an average minimum wage increase of 19.3%. NSW instead use 493 federal- and state-level minimum wage increases, for which many treated states actually received negative treatment relative to donor states, and the average minimum wage increase was about 2.7%. In addition, Dube and Zipperer (2015) provided a visual demonstration (see their figure 3) that employment was unchanging prior to treatment, without much change up to three years after an initial minimum wage increase. In sum, Dube and Zipperer (2015) used a standard implementation of the synthetic control approach, showed that the method is picking reliable controls, and found little effect on teen employment up to three years following the implementation of the policy.

An alternative estimation strategy to forming synthetic controls explicitly estimates the unobserved factor and factor loadings that underlie the data-generating process. Using this approach, Totty (2015) estimated minimum wage effects on teen employment using two panel-data factor models: the Bai (2009) interactive fixed-effects estimator and two variants of the common correlated estimator of Pesaran (2006). Totty found unmistakable evidence that accounting for time-varying heterogeneity using factor models substantially reduced the size of the minimum wage employment estimates, consistent with the evidence in this article. In his 1990 to 2010 sample, the two-way fixed-effects estimate for the minimum wage elasticity of teen employment was −0.178 (statistically significant at the 5% level). By contrast, the estimates from the three factor models ranged between −0.040 and −0.065 and were not statistically significantly different from zero.[19]

NSW Matching Estimator

NSW (2014a) proposed a matching estimator based on synthetic control weights that produces estimates that differ from Dube and Zipperer (2015) and Totty (2015). Their sample included 493 federal and state minimum wage increases between 1990 and 2011 that had a four-quarter

[19]Powell (2016) used a "generalized synthetic control" approach and found more sizable negative effects for teen employment. He does not, however, provide evidence on how well his approach actually matches the treated and control groups prior to treatment. In addition, given the similarity of his approach to the panel factor models, it would be useful to show why his estimator appears to produce results that are quite different from the more standard Bai approach implemented by Totty (2015).

pre-treatment period ($t = -4, -3, -2$, and -1 in event time), along with a four-quarter treatment period ($t = 0, 1, 2, 3$). Using state-level CPS data on teens, they estimated synthetic control donor weights for each of the treatment events using a sample of donors that included every other state—including states that had minimum wage increases during dates ($t = -4, \ldots, -1, 1, \ldots, 3$). For each event, then, they had a matched synthetic control unit for their period. Stacking this matched data and subsequently estimating standard two-way fixed-effects panel regression, NSW found statistically significant employment elasticities of -0.143 and -0.145, depending on estimation details.[20]

The most fundamental shortcoming of the NSW matching estimator concerns their sample. Of the 493 events studied by NSW, 129 constitute what they call a "clean sample," in which no minimum wage changes occur in the control units during four quarters prior or subsequent to treatment. They did not, however, use just this clean sample; they added an additional 364 events in which both treatment and potential control units experienced minimum wage increases during treatment periods.[21] As a result, their full 493-event unclean sample, which they used for their main estimation, contained: 1) minimum wage changes in the treated units in the pre-intervention period ($t = -4, \ldots, -1$), and 2) minimum wage changes in the donor (or potential control) states in the pre- and post-intervention periods ($t = -4, \ldots, 0, \ldots, 3$). This sample construction thus rendered the distinction between treatment and control units nearly meaningless.

We report a re-analysis of NSW in Table 5. As column (1) shows, when using their full sample of 493 events, the treated units experienced an average 0.098 log point minimum wage increase.[22] But during the same time period, the control units experienced a 0.071 log point minimum wage increase, yielding only a 0.027 log point (approximately 2.7%) net increase in the treated versus control units. This increase is very small: for comparison, in the 29 events analyzed by Dube and Zipperer (2015), the minimum wage rose 19.3% more in the treated areas as compared to the control areas.

[20]To estimate the donor weights for each event, NSW matched on residual employment, after partialing out state and time fixed effects, as well as the minimum wage. This method is not standard and is possibly problematic because the minimum wage effect is what one is trying to estimate. Nonetheless, to keep our results comparable, in our re-analysis of their data we follow their practice and use residual employment.

[21]NSW (2014a) found a small, statistically insignificant minimum wage elasticity for teen employment of -0.06 when they applied their method only to the clean sample. They nonetheless dismissed these results, arguing that in this sample even the two-way fixed-effects estimate was not sizably negative. This argument is indefensible. The two-way fixed-effects estimate in their clean sample may simply be less biased than in the expanded (unclean) sample. In general, we see little justification in expanding the sample to include events inappropriate for the synthetic control approach, just because the two-way fixed-effects estimate in that sample matches that from the full state panel sample.

[22]We used the programs and data set posted at http://j.mp/datacodeILRR.

Table 5. Re-analysis of Results from NSW Matching Estimator: Difference-in-Differences Estimates

		NSW pre-treatment period ($t = -4, -3, -2, -1$)			*Earlier pre-treatment period ($t = -8, -7, -6, -5$)*		
		(1) Δ log MW	*(2) Δ log teen emp*	*(3) MW Elasticity*	*(4) Δ log MW*	*(5) Δ log teen emp*	*(6) MW Elasticity*
Overall	Treatment	0.098	−0.048		0.160	−0.080	
		(0.003)	(0.008)		(0.006)	(0.012)	
	Control	0.071	−0.042		0.122	−0.088	
		(0.003)	(0.009)		(0.004)	(0.013)	
	Treatment–Control	**0.027*****	**–0.007***	**–0.247***	**0.038*****	**0.008**	**0.205**
		(0.003)	**(0.004)**	**(0.128)**	**(0.005)**	**(0.006)**	**(0.156)**
Quartile 1	Treatment	0.055	−0.058		0.118	−0.119	
		(0.006)	(0.016)		(0.012)	(0.019)	
	Control	0.080	−0.046		0.146	−0.102	
		(0.004)	(0.011)		(0.007)	(0.015)	
	Treatment–Control	**–0.024*****	**–0.012**	**0.490**	**–0.027*****	**–0.018****	**0.646**
		(0.003)	**(0.009)**	**(0.382)**	**(0.009)**	**(0.009)**	**(0.437)**
Quartile 2	Treatment	0.101	−0.029		0.153	−0.050	
		(0.003)	(0.011)		(0.004)	(0.016)	
	Control	0.096	−0.028		0.146	−0.064	
		(0.003)	(0.008)		(0.004)	(0.014)	
	Treatment–Control	**0.005*****	**0.000**	**–0.106**	**0.007****	**0.014**	**1.938**
		(0.000)	**(0.010)**	**(2.043)**	**(0.003)**	**(0.011)**	**(1.458)**
Quartile 3	Treatment	0.103	−0.055		0.171	−0.092	
		(0.005)	(0.021)		(0.007)	(0.030)	
	Control	0.075	−0.049		0.129	−0.121	
		(0.003)	(0.022)		(0.006)	(0.033)	
	Treatment–Control	**0.028*****	**–0.006**	**–0.205**	**0.041*****	**0.029****	**0.695****
		(0.002)	**(0.006)**	**(0.225)**	**(0.003)**	**(0.012)**	**(0.307)**
Quartile 4	Treatment	0.133	−0.044		0.200	−0.048	
		(0.010)	(0.015)		(0.008)	(0.018)	
	Control	(0.033)	−0.037		0.068	−0.052	
		(0.005)	(0.013)		(0.007)	(0.017)	
	Treatment–Control	**0.099*****	**–0.007**	**–0.074**	**0.132*****	**0.004**	**0.029**
		(0.007)	**(0.007)**	**(0.075)**	**(0.009)**	**(0.013)**	**(0.097)**

Notes: The table reports mean differences of log minimum wage and log teen employment rate for both control and treatment groups between post-treatment period ($t = 0, \ldots, 3$) and pre-treatment period ($t = -4, \ldots, -1$) using the NSW (2014a) sample of 493 events, as well as between post-treatment period and earlier pre-treatment period ($t = -8, \ldots, -5$), using the available subsample of 442 events. "Treatment–Control" rows are difference-in-differences (DD) estimates, in boldface. The top panel reports the estimates for the overall samples. The subsequent panels report estimates from four quartiles of the extent of treatment (i.e., DD in log minimum wage). Minimum wage elasticities are obtained by dividing DD estimate for log teen employment by the DD estimate for log minimum wage. Robust standard errors (in parentheses) of elasticities are clustered at the state level and calculated using SUEST command in STATA.
Significance levels are indicated only for the DD estimates by *** 1%; ** 5%; and * 10%.

To assess NSW's (2014a) sample further, we divide the 493 events into quartiles by the extent of treatment: $\Delta\ln\left(MW_{treated,j}\right) - \Delta\ln\left(MW_{SC,j}\right)$, the differential increase of the log minimum wage in the treated versus in the synthetic control units. As shown in the first column of Table 5, the bottom quartile (quartile 1) actually received a net negative treatment: the treated units experienced a 0.024 net *decrease* in log minimum wage as compared to their synthetic controls. The second quartile received essentially no net treatment (a very small increase of 0.005), and the third quartile received a 0.028 increase in log minimum wage. Only the fourth quartile received a substantial treatment—a net minimum wage increase of around 0.099 log points (approximately 10.4%). Most of NSW's events thus are ill-suited for studying the effect of minimum wage increases using the synthetic control approach. Defining events, treatment groups, and synthetic controls have little point if most of these events entail such limited net variation in minimum wages.

In addition, it is not clear why we should expect a reliable counterfactual for the treated state by matching past outcomes in the treated and synthetic control units in the pre-intervention period, when the treatment status in NSW's sample of events was in reality changing arbitrarily in both groups during that period.[23] Finally, NSW used only four quarters of pre-treatment data, a very short window to estimate synthetic control donor weights. Other existing work using synthetic control methods use several years of pre-treatment data (Abadie et al. 2010; Bohn, Lofstrom, and Raphael 2014). Overall, the nature of NSW's sample raises serious questions about their findings.

What does NSW's sample imply about the employment effects of the minimum wage? A difference-in-differences approach provides a straightforward way to estimate an employment elasticity using these 493 events. $E_{treated,j}$ is the teen employment rate in the treated unit, and $E_{SC,j}$ is the teen employment rate in the synthetic control. Table 5 shows the pattern construct difference-in-differences estimates for log of teen employment, $\frac{1}{J}\sum_{j}^{J}\left[\Delta ln\left(E_{treated,j}\right) - \Delta ln\left(E_{SC,j}\right)\right]$, where J is the total number of events. For the full sample (top panel, Treatment–Control row), log employment changes by -0.007 in the treatment units differentially following the minimum wage increase; this decline is statistically significant at the 10% level. By scaling this employment effect by the differential increase in log minimum wage (0.027), we obtain an elasticity of $\frac{\sum_j\left[\Delta\ln\left(E_{treated,j}\right)-\Delta\ln\left(E_{SC,j}\right)\right]}{\sum_j\left[\Delta\ln\left(MW_{treated,j}\right)-\Delta\ln\left(MW_{SC,j}\right)\right]} = -0.247.$[24] This difference-in-differences

[23]Matching on residual employment, after partialing out minimum wage effects, may guard against the bias from aligning employment in the treatment and synthetic control groups in a pre-treatment period in which treatment status was in reality changing. But this approach relies on having the right estimate for the minimum wage, which is unknown, and is estimated using the very two-way fixed-effects model that is in contention. NSW acknowledged that their logic has a "potential circularity," but argued that their results are similar whether the synthetic control weights are constructed from residual employment after partialing out minimum wages or not. Nonetheless, the weights may still not be constructed correctly. For example, if the mean minimum wage effect is small but with some heterogeneity, the weights constructed from matching the treatment and control units will be wrong, and the resulting estimates may be biased. More generally, it is problematic to use an unclean sample in which treatment status is changing in the pre- and post-treatment periods in both the treatment and the control groups.

[24]Standard errors for the elasticity were computed using the SUEST command in STATA, clustering on state.

elasticity estimate is somewhat larger in magnitude than the -0.145 elasticity estimate obtained by NSW using a panel regression. Nonetheless, both results suggest at least moderate employment losses.[25] Several pieces of evidence, however, suggest that these disemployment estimates are unreliable. First, the largest fall in employment (–0.012) occurs in quartile 1, which experiences a minimum wage *decline*, implying a positive minimum wage employment elasticity of 0.490. Second, for quartile 4—the only quartile with a substantial increase in the minimum wage—the employment fall is more muted (–0.007), and it is not statistically significant. The implied minimum wage elasticity based on the fourth quartile is -0.074, substantially smaller in magnitude than what NSW found. Third, for many events (e.g., quartile 2) the change in the minimum wage is virtually the same in treated and control groups: these observations provide little usable information to identify the effect of the treatment.

To summarize to this point, when using NSW's sample of events and their synthetic controls, some events suggest sizable job loss and others suggest sizable minimum wage increases, but these typically happen to be different events; and many events are simply uninformative. Consistent with these findings, when we limit our analysis to the 129 events that NSW labeled as their clean sample, we find a minimum wage elasticity of 0.025.

Given the shortcomings of the NSW sample, we should worry about the quality of matches obtained by their procedure. To assess the impact of match quality on the estimates, we perform the following exercise. The synthetic control weights in NSW (2014a) were estimated using quarters $t = -4, \ldots, -1$ in event time, and the minimum wage estimates were constructed by taking the difference between the post- and pre-treatment periods. As a check, we use a slightly earlier pre-intervention period $t = -8, \ldots, -5$ to form the difference-in-differences estimates. Since this earlier period was not used to estimate the synthetic controls, it provides a test of internal validity: if control groups are well constructed and represent a valid counterfactual, then using this earlier pre-intervention data should

[25] NSW did not conduct this type of difference-in-differences (or perhaps a simple post-treatment difference), even though such an approach corresponds to the standard application of the synthetic control method (Abadie et al. 2010). Instead, they created a sample that stacks the synthetic controls and treated units and then regressed log employment on log minimum wage, controlling for time period dummies, event-by-state dummies, state unemployment rate, and teen population share. Their estimate of -0.15 is somewhat smaller than the -0.25 difference-in-differences elasticity we report in Table 5. The difference-in-differences estimate presented here is based on the actual variation in minimum wages induced by the treatments in 493 events. By contrast, the NSW panel regression additionally used variation in minimum wages 1) between synthetic control units, and 2) between treated units, which seems contrary to the purpose of defining treatment events. Moreover, the difference-in-differences formulation allows us to diagnose what drives the mean estimate by considering different subsets of events, as we do in Table 5. We cluster the standard errors at the level of events; this likely understates the standard errors by not accounting for estimation of synthetic controls, and also the possible correlation in the control units across events. However, clustering by events strikes us as more accurate than clustering by the narrower event-by-treatment status, as was done in NSW (2014a). For this reason, we should be cautious about the statistical significance of findings from the NSW matched estimator.

provide broadly similar results.[26] As column (4) of Table 5 shows, when using this earlier pre-treatment period, the relative increase in the minimum wage between treatment and synthetic control groups is somewhat larger, 0.038 as opposed to 0.027; this occurs in each of the four quartiles as well. In other words, using this earlier pre-treatment period does not cause any attenuation in the extent of treatment. Using this earlier pre-treatment period, however, the employment estimate (0.008) is now positive in sign and not statistically significant, and the implied overall minimum wage elasticity is 0.205. If the NSW synthetic control weights were reliable, this earlier pre-intervention period should not suggest such different estimates of the employment effect of the policy.

The sensitivity of the results to an earlier pre-intervention period suggests serious problems with match quality and indicates that the synthetic control and treatment units were not following parallel trends prior to treatment. In quartile 4, the only quartile with a sizable treatment magnitude (0.132), the earlier pre-treatment estimates are close to zero, with a minimum wage elasticity of 0.029, compared to an estimate of −0.074 using the later pre-treatment period. In other words, for arguably the most informative events in the NSW sample, neither pre-intervention period suggests a substantial disemployment effect.

Overall, our re-analysis of the NSW (2014a) data suggests serious flaws in their sample construction and their estimation of synthetic control groups. Tellingly, most of the minimum wage increases used to construct their synthetic controls do not exhibit a clear treatment. In combination with the short pre-intervention period used to estimate the synthetic control weights, the unclean nature of the sample appears to produce poor matches. Moreover, in the cases in which the treatment group actually experiences a sizable increase in the relative minimum wage as compared to their synthetic control, there is no indication of a sizable reduction in employment. Match quality is also poor: a slightly earlier pre-intervention period than NSW used produces positively signed employment estimates, indicating that the treatment and the control units did not track each other very well or follow parallel trends, prior to the intervention. The conceptual problems with the NSW matching estimator, the problems with their sample construction, and the discouraging findings from simple diagnostic tests all strongly suggest that the estimates they presented are unreliable.

[26]The sample of events shrinks from 493 to 442 when using the earlier period since the events in 1991 in the NSW data set do not have a balanced earlier period ($t = -8, \ldots, -5$). This sample restriction, however, has little impact on the baseline difference-in-differences estimates (results not shown). For example, the overall minimum wage elasticity for teen employment using the sample of 493 events is −0.247 (Table 5, column (3)), whereas the analogous elasticity for the restricted sample is −0.271 (results not shown in tables).

Relationship between Local and Synthetic Controls

NSW (2014a, 2014b) argued that synthetic control–based donor weights for states within the same census division as the treated states are not greater than weights for donors that are outside of the division. Given the problems with match quality documented in the previous section, however, these results are not very informative. By contrast, Dube and Zipperer (2015) estimated synthetic controls using a clear control–treatment distinction, longer pre-intervention periods, a data-driven set of predictors, and a formal assessment of the quality of matches—making their findings on this question more informative than those in NSW. In the teen employment specification of Dube and Zipperer (2015), the total weight per donor inside the same census division is about 3.1 times that of the weights per donor outside the division. Online Appendix D examines these weights further and finds a clear negative relationship between relative donor weights and the geographic distance between donor and treated states (see online Appendix Figure D.1).

Effects on Restaurant Employment

NSW (2014a) devoted substantial attention to critiquing the methods and details of DLR on the effects of minimum wages on restaurant employment. In this section, we show that on the leaded effects falsification test, a county border discontinuity design for restaurant employment outperforms the two-way fixed-effects specification. We also show that in recent studies attempting to account for time-varying heterogeneity, including NSW (2014a), the range of estimated restaurant employment elasticities is remarkably narrow.

We first extend the restaurant employment sample in DLR to the 1990 to 2014 time period and find headline results similar to DLR. Although the employment elasticity is a statistically significant -0.240 using the two-way fixed-effects specification, it falls in magnitude to 0.023 (and is not distinguishable from zero) when using the county border discontinuity design (columns (1) and (3) of Table 6).

We also estimate dynamic specifications, just as we did previously for teens. In the all-counties sample, for the two-way fixed-effects model, the average cumulative response elasticity in the four quarters preceding the minimum wage increase is -0.198 and is statistically significant at the 1% level (see Table 6, panel B, row C, column (1)). Between the 9th and 12th quarter preceding the minimum wage, the cumulative response averages -0.118 and is statistically significant at the 1% level. In other words, we find unambiguous evidence of pre-existing trends that contaminate the two-way fixed-effects estimate for employment in the food services and drinking places sector over the 1990 to 2014 period. By contrast, the specification with pair-specific period effects (column (3) of Table 6) shows no indication of pre-existing trends: The point estimates are all 0.023 or less in

Table 6. Minimum Wage Elasticities for Restaurant Earnings and Employment, QCEW Data, 1990–2014

		(1)	(2)	(3)
Panel A: Contemporaneous minimum wage elasticities				
Earnings		0.233***	0.245***	0.209***
		(0.026)	(0.030)	(0.033)
Employment		−0.240***	−0.184**	0.023
		(0.075)	(0.076)	(0.069)
Panel B: Four-quarter averages of cumulative response elasticities for employment				
A	$\bar{\rho}_{[-12,-9]}$	−0.118***	−0.044	0.014
		(0.046)	(0.048)	(0.054)
B	$\bar{\rho}_{[-8,-5]}$	−0.136***	−0.056	0.023
		(0.046)	(0.055)	(0.075)
C	$\bar{\rho}_{[-4,-1]}$	−0.198***	−0.120*	0.014
		(0.058)	(0.065)	(0.087)
D	$\bar{\rho}_{[0,3]}$	−0.277***	−0.164**	0.022
		(0.078)	(0.085)	(0.101)
E	$\bar{\rho}_{[4,7]}$	−0.329***	−0.201**	−0.016
		(0.088)	(0.096)	(0.115)
F	$\bar{\rho}_{[8,11]}$	−0.358***	−0.206*	−0.012
		(0.106)	(0.108)	(0.127)
G	ρ_{12+}	−0.506***	−0.348**	−0.059
		(0.147)	(0.158)	(0.164)
Panel C: Medium-run (three-year) elasticities for employment				
F-A	$\bar{\rho}_{[8,11]} - \bar{\rho}_{[-12,-9]}$	−0.240***	−0.163**	−0.026
		(0.075)	(0.079)	(0.089)
F-B	$\bar{\rho}_{[8,11]} - \bar{\rho}_{[-8,-5]}$	−0.221***	−0.150**	−0.036
		(0.071)	(0.073)	(0.068)
F-C	$\bar{\rho}_{[8,11]} - \bar{\rho}_{[-4,-1]}$	−0.160***	−0.086	−0.026
		(0.056)	(0.060)	(0.056)
Panel D: Long-run (four-plus-years) elasticities for employment				
G-A	$\rho_{12+} - \bar{\rho}_{[-12,-9]}$	−0.388***	−0.305**	−0.074
		(0.115)	(0.129)	(0.131)
G-B	$\rho_{12+} - \bar{\rho}_{[-8,-5]}$	−0.369***	−0.292**	−0.083
		(0.113)	(0.125)	(0.112)
G-C	$\rho_{12+} - \bar{\rho}_{[-4,-1]}$	−0.308***	−0.228**	−0.074
		(0.097)	(0.110)	(0.097)
Sample		All counties	Border county pairs	Border county pairs
County pair-period FE				Y

Notes: All specifications use quarterly county-level data for Food Services and Drinking Places (NAICS 722) from the 1990–2014 QCEW. The dependent variable is log of county-quarter restaurant employment or average earnings, as indicated, and right-hand-side controls include log of county-quarter population and overall private-sector employment. Panel A specifications include only the contemporaneous log quarterly minimum wage, for which coefficients are reported. The specifications in panels B through D include the contemporaneous, 12 quarterly leads and 12 quarterly lags of log minimum wage. Panel B reports four-quarter averages of the cumulative response elasticities starting at $t = -12$ in quarterly event time. Panel C reports the cumulative effect in year three, after subtracting alternative baseline levels at one, two, or three years prior to treatment, as indicated. Panel D reports long-run cumulative response elasticity at $t = 12$, after subtracting alternative baseline levels. Column (1) uses the balanced panel of counties with common-period fixed effects, column (2) uses the contiguous-border-county-pair sample with common-period effects, and column (3) uses the contiguous-border-county-pair sample with pair-period effects. Robust standard errors in parentheses are clustered at the state-level in column (1) and clustered multi-dimensionally at the state-level and border segment-level in columns (2) and (3).
Significance levels are indicated by *** 1%; ** 5%; and * 10%.

magnitude, and not one is statistically significant. This result provides strong evidence that the border discontinuity design provides more reliable estimates by using more similar comparisons. Table 6, column (3) also reports medium-run (three-year) estimates using border discontinuity design: they range between −0.026 and −0.036 depending on the baseline. Longer-run (four-plus-year) effects are more imprecise, but range between −0.074 and −0.083.

How do these border discontinuity design estimates for restaurant employment compare with other existing work? Online Appendix Table E.1 reports a total of 17 employment elasticities from five key publications that include additional controls beyond the two-way fixed-effects model (DLR 2010, 2016; Addison, Blackburn, and Cotti 2014; NSW 2014a; Totty 2015). The restaurant employment elasticities range from −0.063 to 0.039 for models that incorporate additional controls for time-varying heterogeneity, including NSW's preferred matching estimator.[27] Since the earnings elasticities in these studies are typically around 0.2, these results imply labor demand elasticities generally smaller than −0.3 in magnitude. Moreover, all of these estimators, including NSW's preferred matching estimator, suggest employment effects that are usually substantially smaller than the two-way fixed-effects model. (An exception is Addison et al.'s 1990–2012 sample; as reported in our online Appendix Table E.1, they found a zero effect even for the two-way fixed-effects model.) There may be disagreement about the merits of specific estimators, but these results comprise a highly robust set of findings. They confirm 1) at most a modest impact of minimum wages to date on restaurant employment, and 2) the violation of the parallel trends assumption in the two-way fixed-effects model and likely bias toward finding evidence of job loss.

Some disagreements remain on the details of the restaurant findings. For instance, NSW (2014a, 2014b) criticized a falsification test we performed in DLR to demonstrate the unreliability of the two-way fixed-effects estimates; we respond to these criticisms in online Appendix F. The key takeaway nevertheless remains: the research literature seems to be reaching an agreement on the medium-run effects of minimum wages on restaurant employment.

Conclusion

Much of the minimum wage research on employment effects has focused on teens and on restaurant workers because these two groups are especially

[27]Aaronson et al. (2017), who studied restaurant employment using a border discontinuity design for a small number of states, obtained an overall short-run elasticity of −0.1. They described this estimate as "very imprecise" (they do not report a standard error). They also found increased entry and exit, which they interpreted using a calibrated putty-clay model that suggests large disemployment effects in the longer run. However, our empirical findings here and in DLR (2010) do not suggest sizable employment losses in the restaurant sector in the "medium run," that is, after 12 or 16 quarters following the minimum wage.

affected by minimum wage policies. A wide variety of recent restaurant studies using different data sets, time periods, and estimators arrive at similar findings. In these studies, the preferred elasticities of employment with respect to minimum wages lie within a fairly narrow range of −0.063 and 0.039, suggesting at most a small effect of the policy on employment, even as the earnings effects are substantial. The main substantive disagreement—and most of this article—thus centers on the effects on teens.

Our key contention is that when using cross-state variation to estimate minimum wage effects, it is critically important to account for time-varying heterogeneity. This clustering coexists with an array of potential confounds that vary between high and low minimum wage states. The chances seem small that all of these factors happen to balance each other.

NSW (2014a, 2014b) criticized the advances made in our previous articles to account for these heterogeneities. The findings in this article show that the key claims made by NSW (2014a, 2014b) do not withstand scrutiny. The minimum wage employment elasticities estimated using the canonical two-way fixed-effects model have a substantial negative bias—visible in the form of pre-existing trends. These sizable negative estimates largely disappear once we control for time-varying heterogeneity with linear or higher-order state-specific trends and division-period controls. The double-selection post-LASSO approach finds that controls for time-varying heterogeneity should be included and that the resulting employment elasticity is close to zero (–0.009). Additionally, this approach establishes that including a handful of state trends reduces the employment estimate to close to zero, contradicting the assertion in NSW that our findings in ADR resulted from throwing away too much data. Large, negative teen employment effects are also absent when estimating the two-way fixed-effects model using first-differences.

By contrast, the NSW (2014a) matching estimator is riddled with internal inconsistencies, the most important of which is mixing treatment and control groups, and is sensitive to the choice of pre-intervention period, indicating treatment and control groups are likely not following parallel trends. Notably, the NSW matching estimator is one of the few in the recent minimum wage literature on teen employment that goes beyond the two-way fixed-effects strategy and still finds a substantial negative employment effect (Dube and Zipperer 2015; Totty 2015; DLR 2016; Gittings and Schmutte 2016; Manning 2016; Slichter 2016).

Accumulating evidence has led us, as well as many other economists, to the conclusion that the employment effects of U.S. minimum wage policies on low-wage employment to date have been fairly small. Also, these effects are fairly precisely estimated for the medium run, including three to four years after minimum wage increases.

These findings are based on state and federal minimum wage changes between 1979 and 2014, when the federal minimum wage was relatively low by both historical and international standards (Dube 2014). Future research will tell us whether the impacts of higher minimum wages will differ from the effects of the policies studied in this article.

References

Aaronson, Daniel, Eric French, Isaac Sorkin, and Ted To. 2017. Firm dynamics and the minimum wage: A putty-clay approach. *International Economic Review*, forthcoming.

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.

Addison, John, McKinley Blackburn, and Chad Cotti. 2014. On the robustness of minimum wage effects: Geographically-disparate trends and job growth equations. IZA Discussion Paper No. 8420. Bonn, Germany: Institute for the Study of Labor (IZA). Accessed at http://ftp.iza.org/dp8420.pdf.

[ADR] Allegretto, Sylvia, Arindrajit Dube, and Michael Reich. 2011. Do minimum wages really reduce teen employment? Accounting for heterogeneity and selectivity in state panel data. *Industrial Relations* 50(2): 205–40.

Angrist, Joshua, and Jorn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.

Bai, Jushan. 2009. Panel data models with interactive fixed effects. *Econometrica* 77: 1229–79.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014. High-dimensional methods and inference on treatment and structural effects in economics. *Journal of Economic Perspectives* 28(2): 29–50.

Bohn, Sarah, Magnus Lofstrom, and Steven Raphael. 2014. Did the 2007 Legal Arizona Workers Act reduce the state's unauthorized immigrant population? *Review of Economics and Statistics* 96(2): 258–69.

Borjas, George J., Richard B. Freeman, and Lawrence F. Katz. 1997. How much do immigration and trade affect labor market outcomes? *Brookings Papers on Economic Activity* 1: 1–90.

Card, David, and Alan Krueger. 1994. Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review* 84(4): 772–93.

———. 2000. Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania: Reply. *American Economic Review* 90(5): 1397–1420.

Clemens, Jeffrey, and Michael Wither. 2016. The minimum wage and the Great Recession: Evidence of effects on the employment and income trajectories of low-skilled workers. Accessed at http://econweb.ucsd.edu/~j1clemens/pdfs/ClemensWitherMinimumWageGreatRecession.pdf.

Dube, Arindrajit. 2014. Designing thoughtful minimum wage policy at the state and local levels. The Hamilton Project. Accessed at http://www.hamiltonproject.org/papers/designing_minimum_wage_policy_at_state_and_local_levels/.

Dube, Arindrajit, and Ben Zipperer. 2015. Pooling multiple case studies using synthetic controls: An application to minimum wage case studies. IZA Discussion Paper No. 8944. Bonn, Germany: Institute for the Study of Labor. Accessed at http://ftp.iza.org/dp8944.pdf.

[DLR] 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.

———. 2016. Minimum wage shocks, employment flows and labor market frictions. *Journal of Labor Economics* 34(3): 663–704.

Gittings, R. Kaj, and Ian M. Schmutte. 2016. Getting handcuffs on an octopus: Minimum wages, employment, and turnover. *ILR Review* 69(5): 1133–70.

Hirsch, Barry T., and Edward J. Schumacher. 2004. Match bias in wage gap estimates due to earnings imputation. *Journal of Labor Economics* 22(3): 689–722.

Huang, Yi, Prakash Loungani, and Gewei Wang. 2014. Minimum wages and firm employment: Evidence from China. IMF Working Paper 14/184. Washington, DC: International

Monetary Fund. Accessed at http://www.imf.org/external/pubs/ft/wp/2014/wp14184.pdf.

Liu, Shanshan, Thomas Hyclak, and Krishna Regmi. 2016. Impact of the minimum wage on youth labor markets. *Labour* 30(1): 18–37.

Magruder, Jeremy. 2013. Can minimum wages cause a big push? Evidence from Indonesia. *Journal of Development Economics* 100(1): 48–62.

Manning, Alan. 2016. The elusive employment effects of the minimum wage. CEP Discussion Paper No. 1428. London: Centre for Economic Performance. Accessed at http://cep.lse.ac.uk/pubs/download/dp1428.pdf.

Meer, Jonathan, and Jeremy West. 2016. Effects of the minimum wage on employment dynamics. *Journal of Human Resources* 51(2): 500–522.

[NSW] Neumark, David, J. M. Ian Salas, and William Wascher. 2014a. Revisiting the minimum wage and employment debate: Throwing out the baby with the bathwater? *ILR Review* 67(Supplement): 608–48.

———. 2014b. More on recent evidence on the effects of minimum wages in the United States. NBER Working Paper No. 20619. Cambridge, MA: National Bureau of Economic Research. Accessed at http://www.nber.org/papers/w20619.

Pesaran, M. Hashem. 2006. Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica* 74(4): 967–1012.

Powell, David. 2016. Synthetic control estimation beyond case studies: Does the minimum wage reduce employment? RAND Working Paper WR-1142. Santa Monica, CA: RAND Corporation. Accessed at https://www.rand.org/content/dam/rand/pubs/working_papers/WR1100/WR1142/RAND_WR1142.pdf.

Slichter, David. 2016. The employment effects of the minimum wage: A selection ratio approach to measuring treatment effects. Mimeo. Accessed at https://docs.google.com/viewer?a=v&pid=sites&srcid=ZGVmYXVsdGRvbWFpbnxzbGljaHRlcmRhdmlkfGd4OmNmZjkwNDgxYmUyOGZhNA.

Totty, Evan. 2015. The effect of minimum wages on employment: A factor model approach. IRLE Working Paper 110-15. UC Berkeley: Institute for Research on Labor and Employment. Accessed at http://www.irle.berkeley.edu/workingpapers/110-15.pdf.

Vaghul, Kavya, and Ben Zipperer. 2016. State-level historical minimum wages. Washington, DC: Washington Center for Equitable Growth.

Wolfers, Justin. 2006. Did unilateral divorce laws raise divorce rates? A reconciliation and new results. *American Economic Review* 96(5): 1802–20.

Zipperer, Ben. 2016. Did the minimum wage or the Great Recession reduce low-wage employment? Working Paper. Washington, DC: Washington Center for Equitable Growth. Accessed at http://cdn.equitablegrowth.org/wp-content/uploads/2016/12/02155549/120616-WP-comments-on-clemens-and-wither.pdf.