Washington Center for Equitable Growth

1500 K Street NW, Suite 850
Washington, DC 20005

Working paper series

Credible research designs for minimum wage studies: A response to Neumark, Salas, and Wascher

Sylvia Allegretto
Arindrajit Dube
Michael Reich
Ben Zipperer

Working paper URL

http://equitablegrowth.org/working-papers/credible-research-designs-for-minimum-wage-studies/

Credible Research Designs for Minimum Wage Studies: A Response to Neumark, Salas and Wascher
Sylvia Allegretto, Arindrajit Dube, Michael Reich, Ben Zipperer
August 2016

Abstract

We assess the NSW (Neumark, Salas and Wascher 2014) critiques of our minimum wage studies that found small effects on teen employment. Data from 1979-2014 contradict NSW, and show that the disemployment suggested by a model assuming parallel trends across U.S. states mostly reflects 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, contradicting NSW's claim that we 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 we show to be good controls. NSW's preferred matching estimates mix treatment and control units, obtain poor matches, and find employment declines the most where the relative minimum wage falls. These findings refute NSW's key claims.

Sylvia Allegretto
Institute for Research on Labor and Employment
University of California, Berkeley
allegretto@berkeley.edu

Arindrajit Dube
Department of Economics
University of Massachusetts, Amherst
IZA
adube@econs.umass.edu

Michael Reich
Institute for Research on Labor and Employment
Department of Economics
University of California, Berkeley
mreich@econ.berkeley.edu

Ben Zipperer
Washington Center for Equitable Growth
bzipperer@equitablegrowth.org

We are grateful to Doruk Cengiz, Zachary Goldman, Carl Nadler, Thomas Peake and Luke Reidenbach for excellent research assistance. Financial support for this paper came entirely from the authors' respective institutions.

Forthcoming, ILR Review

Credible Research Designs for Minimum Wage Studies: A Response to Neumark, Salas and Wascher

Sylvia Allegretto, Arindrajit Dube, Michael Reich, Ben Zipperer

July 30, 2016

Abstract

We assess the NSW (Neumark, Salas and Wascher 2014) critiques of our minimum wage studies that found small effects on teen employment. Data from 1979-2014 contradict NSW, and show that the disemployment suggested by a model assuming parallel trends across U.S. states mostly reflects 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, contradicting NSW's claim that we 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 we show to be good controls. NSW's preferred matching estimates mix treatment and control units, obtain poor matches, and find employment declines the most where the relative minimum wage falls. These findings refute NSW's key claims.

Allegretto: Institute for Research on Labor and Employment, University of California, Berkeley; Dube: Department of Economics, University of Massachusetts Amherst and IZA; Reich: Department of Economics and Institute for Research on Labor and Employment, University of California, Berkeley; Zipperer: Washington Center for Equitable Growth. We are grateful to Doruk Cengiz, Zachary Goldman, Carl Nadler, Thomas Peake and Luke Reidenbach for excellent research assistance. Financial support for this paper came entirely from the authors' respective institutions.

Credible Research Designs for Minimum Wage Studies: A Response to Neumark, Salas and Wascher

Introduction

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 U.S. A key concern here is that the distribution of minimum wage policies among states has been far from random. If we divide the states into two equally sized groups—"high" versus "low" groups based on their average real minimum wages over the 1979-2014 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 non-random 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

[1]We classified states into high and low minimum wage groups using state-level annual minimum wages adjusted for inflation by the CPI-U-RS. 2 Other minimum wage researchers—e.g., Aaronson, French and Sorkin (2015), Magruder (2013) and Huang, Loungani and Wang (2014)—have subsequently used the border discontinuity design to estimate causal effects of minimum wage policies in both U.S. and international contexts.

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 there are sizable earnings effects 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 papers, Neumark, Salas and Wascher (NSW 2014a, 2014b) critique the use of local area controls in DLR and ADR. They make **three** important claims.

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

Second, they argue that using local area controls throws away too much useful information. In the same vein, they also claim that the small magnitudes of the employment estimates in ADR from specifications with state-specific linear trends are driven by an "endpoint bias" generated by the presence of recessions in the beginning and 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 propose a new "matching" estimator loosely based on the synthetic control approach. They argue that this matching estimator suggests substantial employment effects, at

2 Other minimum wage researchers—e.g., Aaronson, French and Sorkin (2015), Magruder (2013) and Huang, Loungani and Wang (2014)—have subsequently used the border discontinuity design to estimate causal effects of minimum wage policies in both U.S. and international contexts.

least for teens. They claim that this 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 here on this group.

We begin by presenting recent evidence on teen employment using a border discontinuity design. We also 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 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 1990-2009 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 "throwing away too much information," 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 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.

While we mostly focus on teens, we also present new evidence on restaurant employment using updated 1990-2014 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 1979-2014 period, 40.2 percent of working teens earned within 10 percent of the statutory minimum wage (higher of state or federal), as compared to 7.7 percent 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 percent of the statutory minimum wage, the teen share has fallen over time, from 32.2 percent in 1979 to 22.7 percent in 2014.[3] Finally, labor-labor substitution may imply

[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.

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, Lester and Reich (2010) developed this approach by comparing contiguous counties straddling state borders, building upon the insights of comparing nearby areas in Card and Krueger (1996, 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 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).

[4]Clemens and Wither (2014) study a different population of affected low-wage workers and find large, negative employment effects, but Zipperer (2016) finds that these results are not robust to controls for regional heterogeneity.

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

NSW (2014a) challenge 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 state (p. 632): "the cross-border county is a poor match—no better than a county chosen at random from the list of all potential comparison counties."

DLR (2016) uses the county-level Quarterly Workforce Indicators (QWI) dataset to assess whether adjacent county pairs are indeed more alike in terms of covariates than are non-adjacent county pairs. DLR (2016) consider 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 non-contiguous pairs; and in all cases but one, the gaps are statistically significant at the 1 percent level. Many of the gaps, including changes in EPOP, are substantial and exceed 25 percent. These results contradict the NSW claim that contiguous counties are not more similar to each other than two counties chosen at random.

Slichter (2016) corroborates these findings with a refinement of the contiguous county methodology by comparing counties to their neighbors, neighbors-of-neighbors, neighbors-of-neighbors-of-neighbors, etc. He shows that immediate neighbors are, indeed, more likely to have experienced similar employment changes just prior to minimum wage increases. He also shows 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 estimates 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 percent level. The estimated teen employment minimum wage elasticity from the two-way fixed effects model is -0.173 and statistically significant at the 1 percent level. In 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 finds a sizable reduction in turnover following a minimum wage increase: the turnover elasticity is -0.204 when county-pair period effects are included. Importantly, 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 employs a neighboring county discontinuity design, reinforces these conclusions. Slichter relaxes 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, etc., 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 produces small employment elasticities for teens that are similar to our findings here, ranging from -0.006 to -0.041 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 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) argue against these findings. First, they argue that major recessions near the endpoints of the ADR sample (1990-2009) lead to unreliable estimates of state-specific trends. Second, they argue that the use of third or higher order polynomial trends restores the findings of a large disemployment effect. Third, they argue that the data do not warrant using geographic controls (division-period effects). Fourth, they argue that there is little evidence of

[5]Liu, Hyclak, and Regmi (2016) uses a particular definition of a local area (BEA-based Economic Areas), QWI data from 2000-2009, 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. find 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).

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:

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

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 $\mathbf{X}_{\mathrm{it}}$ includes dummies for gender, race, Hispanic origin, age, and marital status; the teen share of the population in the state; and the non-seasonally

[6]For teen employment, we use individual-level records of 16-19 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 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 fifty states and the District of Columbia for 1979-2014 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 1979-1988 and September 1995-2014. For 1989-1993, we define imputations as observations with missing or zero “unedited” earnings but positive “edited” earnings (which we also do for hours worked and hourly wages). We do not label any observations as having imputed wages during 1994-August 1995, when there are no BLS allocation values for earnings or wages.

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 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 also allow the time effects to vary by each of the nine census divisions, denoted by d:

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

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 twelve specifications—with common or division-period fixed effects, and with polynomial trends of degree $\mathrm{k} = 0, \ldots, 5$— include the four key specifications used in ADR, which only used 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 1979-2014. The outcome variable here is the natural log of the hourly wage. All regressions include state fixed effects. The first row includes common time

[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 16 and over 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/).

effects, while 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 twelve 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, p.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. Importantly, 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 were driven by "endpoint bias" in the estimation of state trends owing to the presence of recessionary years.[10]

[10] Online Appendix B provides additional evidence that the "endpoint bias" explanation is incorrect. To summarize those findings, Online Appendix Figure B1 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

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 coefficients are always smaller than -0.09 in magnitude 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 upon the data, especially when the panel is short. By using a substantially longer panel, we estimate these trends more reliably. The estimates from including 3rd and 5th order polynomials, -0.061 and -0.088, respectively are virtually identical to the estimate with just a linear trend (-0.062). The estimate from the 2nd order trend is slightly smaller in magnitude (-0.040) while the estimate from the 4th order trend is slightly larger in magnitude (-0.088). However, in all cases, the estimates are under -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 due to lack of precision, but rather due to the small size of the coefficients.

---

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 disemployment effects in models with state trends.

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) conducts a similar analysis with 1979-2012 CPS data, and also finds that that the two-way fixed effects estimate is unique in producing a large, statistically significant disemployment estimate, and that inclusion of linear and higher order trends, as well as division-period controls, produces 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, Chernozhukov and Hansen (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. suggest 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, while all of the other 11 coefficients are under -0.09 in magnitude 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 1 through 5. 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

[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 due to the penalization term.

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 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 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 only picks the state unemployment rate as a control beyond the manually-specified two-way fixed effects. (The numerical estimates are in Online Appendix Table A1.)

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 5 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 5 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 of NSW that ADR's findings were driven by "throwing away too much information."

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; importantly, as before, no nonlinear trends are picked. Therefore, while the shorter sample produces more varied estimates using OLS and alternative trend specifications—likely due 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-

[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.

varying heterogeneity and also 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) we used a distributed lag model to demonstrate that pre-existing trends contaminate the estimates of the conventional two-way fixed effects model, which often exhibits sizable and statistically significant leading effects. Nonetheless, NSW (2014b) raise questions about our findings on pre-existing trends for teen employment. First, they argue that pre-existing trends are not clearly indicated in the two-way fixed-effects model. Second, they argue that even after differencing out the leading effects, the subsequent cumulative effects remain negative, sizable and comparable to the static estimates. Third, they argue that the inclusion of controls for spatial heterogeneity does 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 to our prior static specifications in equations (1) and (2) twelve quarters of leading and twelve quarters of lagged minimum wages. We estimate these regressions using the individual-level CPS data and control sets we used before for teens in the 1979-2014 period using four specifications. Beginning with the two-way fixed effects model

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

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

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

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 out leading coefficients from the cumulative responses.

Performance of the two-way fixed effects model

Column1 of Table 3 shows four-quarter averages of these quarterly cumulative response elasticities: $\overline{\rho}_{[\tau,\tau+3]} = \frac{1}{4} \sum_{m=0}^{3} \rho_{\tau+m}$, along with standard errors. Online Appendix C, Figure C1 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 $\overline{\rho}_{[-12,-9]}$ is -0.144, and is statistically significant at the five percent 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 ($\overline{\rho}_{[-8,-5]}$ and $\overline{\rho}_{[-4,-1]}$) are even more negative, -0.199 and -0.190, respectively; both are statistically significant at the 5 percent level. In sum, using the full 1979-2014 sample, we find unmistakable evidence that the two-way fixed-effects model fails the falsification test that leading coefficients during 1, 2 or 3 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.

Second, we 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 effect 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, $\sum_{k=0}^{\tau} \eta_k = \rho_\tau - \rho_{-1}$; i.e., this approach takes ρ_{-1} as the baseline.) However, since individual leading coefficients exhibit considerable noise, the choice of the baseline quarter can matter (e.g., see Online Appendix Figure C1). We therefore use estimates using alternative baselines averaging over quarters.

Table 3 calculates estimates for 3 and 4+ year effects from the policy. For the "medium term" or 3 year estimates, we begin by calculating the average cumulative response elasticity in the third year following the minimum wage increase $\overline{\rho}_{[8,11]}$, 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, i.e., $\overline{\rho}_{[-4,-1]}$, $\overline{\rho}_{[-8,-5]}$, or $\overline{\rho}_{[-12,-9]}$, respectively. For example, using the first year before treatment as the baseline, the 3-year estimate is: $\overline{\rho}_{[8,11]} - \overline{\rho}_{[-4,-1]}$. We also construct 4+ year or "long term" estimates as $\rho_{12} - \overline{\rho}_{baseline}$, where the baseline can again be $\overline{\rho}_{[-4,-1]}$, $\overline{\rho}_{[-8,-5]}$, or $\overline{\rho}_{[-12,-9]}$.[13]

The 3 and 4+ 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 3 and 4+ year estimates are substantially smaller than the estimate from the static specification.

[13] We say "4+ year" because ρ_{12} reflects the cumulative response at or after the 12th quarter following a minimum wage increase.

While the static estimate from Table 1 is -0.214, the 3 year and 4+ 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, there is a 40-55 percent reduction in the effect size, as compared to the static estimate, which implicitly uses a mixture baselines $\tau < 0$. Using an earlier baseline ($\tau \in [-12, -9]$) produces 3 and 4 year estimates of -0.175 and -0.152 (rows F-A and G-A), while using an even earlier baseline of $\tau < -12$ (i.e., the average 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 2014(b), who deny that there is evidence of pre-existing trends in the two-way fixed-effects model. They also argue that netting out the leading coefficients does not alter the estimates very much. To reconcile our two sets of results, in Online Appendix C, we estimate analogous regressions using their data and specification (i.e., state-by-quarter level data from 1990q1-2011q1).[15] Online Appendix C, Table C1 reports estimates similar to Table 3 but with the NSW data. In Online Appendix C, we also show the cumulative responses at quarterly frequency using the full 1979-2014 sample (Online Appendix Figure C1) as well as the NSW data (Online Appendix Figure C2).

[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 (like 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 produce NSW 2014b Figure 6: they include controls for unemployment rate, state and period fixed effects.

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. In 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 3 or 4+ year out effects exceed -0.1 in magnitude regardless of controls for state-specific trends or division-period effects. While 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 1979-2014 sample used in this paper.

Performance of models with controls for spatial heterogeneity

Table 3, columns 2, 3 and 4 show the 4-quarter averaged coefficients $\overline{\rho}_{[\tau,\tau+k]}$ 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 percent level ($\overline{\rho}_{[-8,-5]}$ in column 2 with just division-period controls), in contrast to the two-way fixed effects model where all three of the averaged leads are significant. Both the model with state linear trends (column 2) and additionally division-period effects (column 3) perform well in terms of the leading effects falsification test.

What do these models with controls for state-specific trends and division-period effects imply about medium (3 year) and longer run (4+ years) effects from the policy? In our full sample, when using either four quarters just prior to treatment ($\overline{\rho}_{[-4,-1]}$), or the four preceding quarters ($\overline{\rho}_{[-8,-4]}$) 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 Appendix Table C1, 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 2015). 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 4 quarters prior to treatment as 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 4+ year effects affect 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. While 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.

[16]This conclusion is qualitatively similar in the NSW sample (Online Appendix C, Table C1, columns 2, 3 and 4) where the equivalent range is (-0.033, 0.395).

Therefore, as an alternative, we estimate the model in first differences using state (j) by year (t) aggregated data, while including up to 3 annual lags in the average minimum wage. The baseline first-difference specification is:

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

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 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 (2015), who argue that the delayed effects of minimum wages on total employment mostly occur within 2-3 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 3-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

[17] We have chosen to weight the state-aggregated regressions by teen population weights in most parts of the paper, 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 there is less clear rationale for using weights. 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.).

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 percent level. In 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. This 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 5-8 that additionally control for leading minimum wages tend to suggest smaller (or no) disemployment effects; and 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 (2015) criticize the inclusion of state-specific trends and argue that they produce spuriously small disemployment estimates because trends soak up lagged effects. However, 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 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 find 3-year cumulative elasticities for total private-sector employment of -0.074 (column 1 of their Table 4). In 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 whether the findings in Meer and West (2015) that minimum wages reduce aggregate employment are likely to reflect causal effects.[18]

Controlling for endogeneity using factor models and synthetic controls

Existing estimates

NSW (2014a) propose 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 et al. (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 where 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.

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

Dube and Zipperer (2015) use the synthetic control approach to estimate minimum wage effects on teen wages and employment for 29 state minimum wage-increasing events during 1979-2013, and then pool the results from these individual case studies. The minimum wage is clearly binding in their sample of 29 events: 25 of 29 wage elasticities are positive and the mean and median wage elasticities are 0.237 and 0.368, respectively. In contrast, 12 of the employment elasticities are positive and the mean and median employment elasticities are relatively small: -0.051 and -0.058, respectively. Dube and Zipperer (2015) also extend the donor-based randomization inference procedure suggested by Abadie et al. (2010) to multiple events. They calculate a 95 percent 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 in their event study NSW problematically assign many minimum-wage raising states to the potential donor group, Dube and Zipperer (2015) keep the treatment-control distinction clear, as required by the case study approach of the synthetic control estimator. In order to obtain better matches, Dube and Zipperer (2015) impose a pre-treatment window of at least two years and up to four years, but NSW only use 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 percent 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 percent. NSW instead use 493 federal and state-level minimum wage increases where many treated states actually receive negative treatment relative to donor states, and the average minimum wage increase is about 2.7 percent. In addition, Dube and Zipperer (2015) provide 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) use a standard implementation of the synthetic control approach, show that the method is picking reliable controls, and find 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) estimates 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 finds unmistakable evidence that accounting for time-varying heterogeneity using factor models substantially reduces the size of the minimum wage employment estimates, consistent with the evidence in this paper. In his 1990-2010 sample, the two-way fixed effects estimate for the minimum wage elasticity of teen employment is -0.178 (statistically significant at the 5 percent level). In contrast, the estimates from the three factor models range between -0.040 and -0.065 and are not statistically significantly different from zero.[19]

NSW Matching Estimator

NSW (2014a) proposes a matching estimator based on synthetic control weights that produces different estimates from Dube and Zipperer (2015), and Totty (2015). Their sample

[19]Powell (2016) uses a "generalized synthetic control" approach and finds more sizable negative effects for teen employment. Assessing his method is beyond the scope of this paper. However, it is not clear how well his approach matches the treated and control groups. 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 popular Bai approach implemented by Totty (2015).

includes 493 federal and state minimum wage increases between 1990 and 2011 that have a four quarter 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 estimate synthetic control donor weights for each of the treatment events using a sample of donors that includes every other state—including states that had minimum wage increases during dates (t=-4,...,-1, 1,..., 3). For each event, then, they have a matched synthetic control unit for their period. Stacking this matched data and subsequently estimating standard two-way fixed effects panel regression, NSW find 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 comprise what they call a "clean sample," in which there are no minimum wage changes in the control units during 4 quarters prior or subsequent to treatment. They do not, however, use just this "clean sample;" they add an additional 364 events in which both treatment and potential control units experience minimum wage increases during treatment periods.[21] As a result, their full 493 event unclean sample, which they use for their main estimation, contains: 1) minimum wage changes in the treated units in the pre-intervention period (t=-4,...,-1), and 2) minimum wage changes in the donor (or

[20]To estimate the donor weights for each event, NSW match on residual employment, after partialing out state and time fixed effects, as well as the minimum wage. This method is not standard, and 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]Interestingly, NSW (2014a) find a small, statistically insignificant minimum wage elasticity for teen employment of -0.06 when they apply their method only to the "clean sample." They nonetheless dismiss these results, arguing that in this sample, even the two-way fixed effects estimate is 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.

potential control) states in the pre- and post-intervention periods (t=-4,...,0,...,3). This sample construction thus renders 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 percent) 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 percent more in the treated areas as compared to the control areas.

To assess NSW's (2014a) sample further, we divide the 493 events into quartiles by the extent of treatment: $\Delta \ln(MW_{treated,j}) - \Delta \ln(MW_{SC,j})$, 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), while 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 percent). Most of NSW's events thus are ill-suited for studying the effect of minimum wage increases using the synthetic control approach. There is little point in defining events, treatment groups and synthetic controls if most of these events entail so little net variation in minimum wages.

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

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 in both groups during that period in arbitrary ways.[23] Finally, NSW use 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 et al. 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, while $E_{SC,j}$ is the same in the synthetic control. Table 5 shows the pattern construct difference-in-differences estimates for log of teen employment, $\frac{1}{J}\sum_{j}^{J}\left[\Delta\ln(E_{treated,j}) - \Delta\ln(E_{SC,j})\right]$, 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 percent level. By scaling this employment effect by the differential increase in log minimum wage (0.027), we obtain an

[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 where 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 acknowledge that their logic has a "potential circularity," but argue that their results are similar whether the synthetic control weights are constructed from residual employment after partialing out minimum wages, or not. However, 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 control groups.

elasticity of $\frac{\sum_j[\Delta\ln(E_{treated,j})-\Delta\ln(E_{SC,j})]}{\sum_j[\Delta\ln(MW_{treated,j})-\Delta\ln(MW_{SC,j})]} = -0.247$.[24] This difference-in-differences 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 sized employment losses.[25] However, several pieces of evidence 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 find. 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 some suggest sizable minimum wage

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

[25]NSW do 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. 2009). Instead, they create a sample that stacks the synthetic controls and treated units and then regress 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. In contrast, the NSW panel regression additionally uses 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 is done in NSW (2014a). For this reason, we should be cautious about the statistical significance of findings from the NSW matched estimator.

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 label 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 performed 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 a valid counterfactual, then using this earlier pre-intervention data should 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.

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

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 use produces positively signed employment estimates, indicating that the treatment and 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 present are unreliable.

Relationship between local and synthetic controls

NSW (2014a, 2014b) argue 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. However, given the problems with match quality documented in the previous section, these results are not very informative. In contrast, Dube and Zipperer (2015) estimate 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 D1).

Effects on restaurant employment

NSW (2014a) devote 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 1990-2014 time period and find headline results similar to DLR. While the employment elasticity is a statistically significant -0.240 using the two-way fixed effects specification, it falls in magnitude to 0.023 (and 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 percent 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 percent 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 1990-2014 period. In 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 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 (3 year) estimates using border discontinuity design: they range between -0.026 and -0.036 depending on the baseline. Longer run (4+ years) 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 E1 reports a total of 17 employment elasticities from five key papers that include additional controls beyond the two-way fixed effects model (DLR 2010, 2016; Addison, Blackburn and Cotti 2014; Totty 2015; and NSW 2014a). 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,

[27]Aaronson, French and Sorkin (2015), who study restaurant employment using a border discontinuity design for a small number of states, obtain an overall short-run elasticity of -0.1. They describe this estimate as "very imprecise" (they do not report a standard error). They also find increased entry and exit,

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 E1, they find a zero effect even for the two-way fixed effects model.) While there may be disagreement about the merits of specific estimators, 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.

There are some remaining disagreements on the details of the restaurant findings. For instance, NSW (2014a, 2014b) criticize 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 affected by minimum wage policies. A wide variety of recent restaurant studies using different datasets, 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

---

which they interpret 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 restaurant sector in the "medium run," i.e., after 12 or 16 quarters following the minimum wage.

small effect of the policy on employment, even as the earnings effects are substantial. The main substantive disagreement—and most of this paper—thus centers upon 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) criticize the advances made in our previous papers to account for these heterogeneities. The findings in this paper 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.

In contrast, the NSW (2014a) matching estimator is riddled with internal inconsistencies, most importantly 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. (DLR 2016; Dube and Zipperer 2015; Gittings and Schmutte 2015; Manning 2016; Slichter 2016; and Totty 2015.)

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. And these effects are fairly precisely estimated for the “medium run,” including three to four years after minimum wage increases.

These findings are based upon 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 paper.

References

Aaronson, Daniel, Eric French and Isaac Sorkin 2015. “Firm Dynamics and the Minimum Wage: a Putty-Clay Approach.” Working Paper 2013-26. Research Department, Federal Reserve Board of Chicago. https://www.chicagofed.org/digital_assets/publications/working_papers/2013/wp2013_26.pdf

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.

Abowd, John, Bryce E. Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer and Simon Woodcock 2009. ”The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators.” Pp. 149-230 in Timothy Dunne, Bradford Jensen and Mark Roberts eds. *Producer Dynamics: New Evidence from Micro Data*. Chicago: University of Chicago and NBER.

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 8420, Institute for the Study of Labor (IZA). http://ftp.iza.org/dp8420.pdf

Allegretto, Sylvia and Carl Nadler 2015. “Tipped Wage Effects on Earnings and Employment in Full-Service Restaurants.” *Industrial Relations* 54, 4, October 2015.

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.

Allegretto, Sylvia, Arindrajit Dube, Michael Reich and Ben Zipperer 2013. “Credible Research Designs for Minimum Wage Studies.” Working Paper 148-13. Institute for Research on Labor and Employment, UC Berkeley. http://www.irle.berkeley.edu/workingpapers/148-13.pdf

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

Autor, David, Alan Manning and Christopher Smith 2016. “The Contribution of the Minimum Wage to U.S. Wage Inequality over Three Decades: A Reassessment.” *American Economic Journal*

Bai, Jushan 2009. “Panel Data Models With Interactive Fixed Effects.” *Econometrica* 77: 1229-1279.

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.

Brown, Charles 1982. “The Effect of the Minimum Wage on Employment and Unemployment: A Survey.” *Journal of Economic Literature* 20, 2: 487-528.

Brown, Charles 1999. “Minimum Wages, Employment, and the Distribution of Income.” In Orley C. Ashenfelter and David Card, eds., Vol. 3, Part B of *Handbook of Labor Economics*, Elsevier, pp. 2101 2163.

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: 772-793.

Card, David and Alan Krueger 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.

Dube, Arindrajit. 2014. “Designing Thoughtful Minimum Wage Policy at the State and Local Levels.” The Hamilton Project. http://www.hamiltonproject.org/papers/designing_minimum_wage_policy_at_state_and_local_levels/

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.

Dube, Arindrajit, T. William Lester and Michael Reich 2016. “Minimum Wage Shocks, Employment Flows and Labor Market Frictions.” *Journal of Labor Economics*, forthcoming. https://arindube.files.wordpress.com/2014/11/dlr2r3_fullpaper.pdf

Dube, Arindrajit and Ben Zipperer 2015. “Pooling Multiple Case Studies using Synthetic Controls: An Application to Minimum Wage Case Studies.” IZA Discussion Paper 8944. http://ftp.iza.org/dp8944.pdf

Gittings, Kaj and Ian Schmutte 2015. “Getting Handcuffs on an Octopus: Minimum Wages, Employment, and Turnover.” *Industrial and Labor Relations Review*, forthcoming. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2234118

Hansen, Christian 2013. “Econometric Methods for High-Dimensional Data.” NBER Summer Institute. http://www.nber.org/econometrics_minicourse_2013/

Hirsch, Barry T. and Schumacher, Edward J. 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. 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. http://cep.lse.ac.uk/pubs/download/dp1428.pdf.

Meer, Jonathan and Jeremy West 2015. “Effects of the Minimum Wage on Employment Dynamics.” *Journal of Human Resources,* forthcoming. http://econweb.tamu.edu/jmeer/Meer_West_Minimum_Wage.pdf

Neumark, David, J. M. Ian Salas and William Wascher 2014a. “Revisiting the Minimum Wage and Employment Debate: Throwing out the Baby with the Bathwater?” *Industrial and Labor Relations Review*, 67 (Supplement): 608-48.

Neumark, David, J. M. Ian Salas and William Wascher 2014b. “More on Recent Evidence on the Effects of Minimum Wages in the United States.” NBER Working Paper 20619. 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: 9671012.

Powell, David 2016. “Synthetic Control Estimation Beyond Case Studies: Does the Minimum Wage Reduce Employment? “RAND Working Paper WR-1142. 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.

Totty, Evan 2015. “The Effect of Minimum Wages on Employment: A Factor Model Approach.” IRLE Working Paper 110-15. Institute for Research on Labor and Employment. UC Berkeley. http://www.irle.berkeley.edu/workingpapers/110-15.pdf

Vaghul, Kavya and Ben Zipperer. 2016. “State-level historical minimum wages.” 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: 1802-1820.

Zipprer, Ben. 2016 “Did the minimum wage or the Great Recession lower low-wage employment?” Forthcoming.

Figure 1. Double-selection post-LASSO estimates for minimum wage elasticity for teen em-ployment, for alternative values of the LASSO penalization parameter, state-quarter aggre-gated CPS data 1979-2014

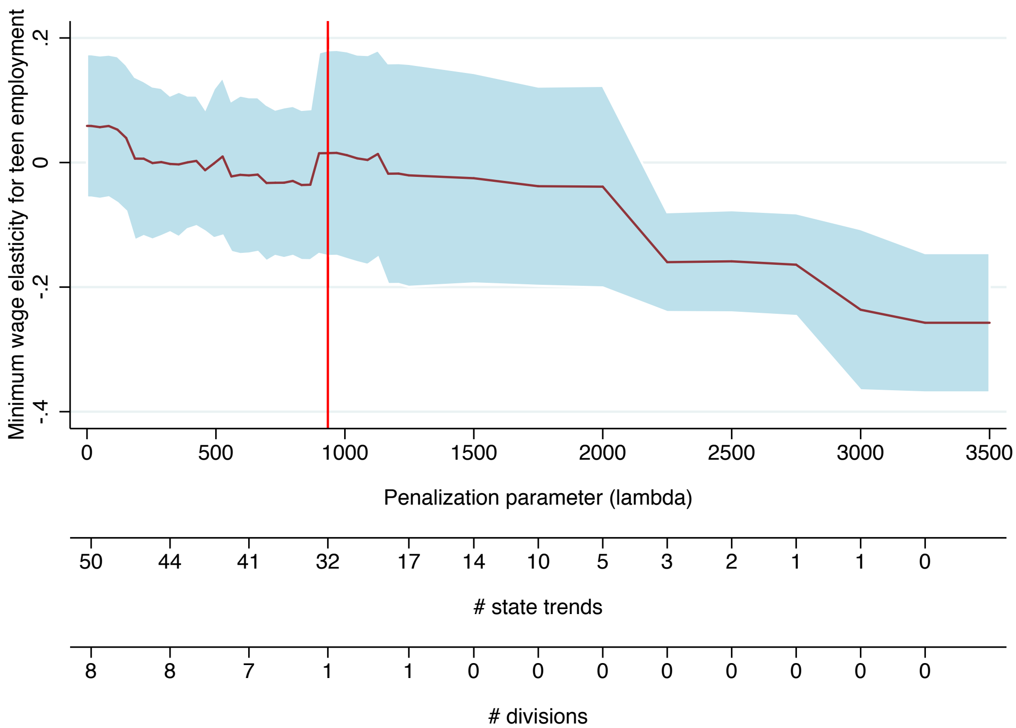

Notes: The figure reports double-selection post-LASSO estimates of minimum wage elasticity for teen employment, and associated 95 percent confidence intervals for alternative values of LASSO penalization parameter, λ, as described in the text. For each value of λ, two LASSO regressions (on log minimum 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. The 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 that are clustered at the state level. The estimates for this graph are also reported in Appendix Table A1.

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 Section 2, and state fixed effects. As reported in the table, specifications either include 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, robust standard errors (in parentheses) are clustered at the state level and significance levels are indicated by *** 1%, ** 5%, * 10%.

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 where 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%, * 10%.

Table 3. Dynamic minimum wage elasticities for teen employment, individual-level CPS data 1979-2014

		(1)	(2)	(3)	(4)
Panel A: 4-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 (3 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 (4+ 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 1979-2014. 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 Section 2.3. Panel B reports the cumulative effect in year 3, 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, and as described in Section 2, 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, robust standard errors are clustered at the state level and significance levels are indicated by *** 1%, ** 5%, * 10%

Table 4. Minimum wage elasticities for teen employment: deviations-from-means versus first-difference estimates, state-year aggregated CPS data 1979-2014

	(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 Min. 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)
3year 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)
3year 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 Min. Wage					Y	Y	Y	Y

Notes: The table reports contemporaneous and 3-year cumulative minimum wage elasticities for teen employment using state-year aggregated CPS basic monthly data from1990-2011q1. 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 additionally 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, and significance levels are indicated by *** 1%, ** 5%, * 10%.

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)		
		Δ log MW	Δ log teen emp	MW Elasticity	Δ log MW	Δ log teen emp	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)**

Table 5 (Continued). Re-analysis of results from NSW matching estimator: difference-in-differences estimates

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, … , 3) and pre-treatment period (*t*=-4 , … , -1) using the NSW (2014a) sample of 493 events, as well as between post-treatment period and earlier pre-treatment period (*t*=-8, … ,-5), using the available sub-sample of 442 events. “Treatment – Controls” rows are difference-in-difference (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%, * 10%

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 4-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 (3 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 (4+ 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

Table 6 (Continued). Minimum wage elasticities for restaurant earnings and employment, QCEW data 1990-2014

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, whose 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 4-quarter averages of the cumulative response elasticities starting at t=-12 in quarterly event time. Panel C reports the cumulative effect in year 3, after subtracting alternative baseline levels at 1, 2 or 3 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 column3 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%, * 10%.

Online Appendix A. Double-selection post-LASSO estimates

We adapted the STATA code for the post-LASSO regressions from Christian Hansen's web page: http://faculty.chicagobooth.edu/christian.hansen/research/JEPStata.zip, including the `lassoshooting.ado` file which estimates the LASSO regressions. To account for the fact that our OLS regressions using aggregated data weight the regressions by teen population, we pre-multiplied the data by the square root of teen population prior to estimating the LASSO regressions. Results using unweighted version of the double-selection post-LASSO were quantitatively similar. In `lassoshooting.ado`, we include state and time fixed effects in the "controls( )" option, which partials out these variables prior to estimating the LASSO regressions.

For the primary results in Table 2 of the main paper, we use the default level for the penalization parameter λ in the Belloni et al. program lassoshooting.ado, which is set $\lambda = 2.2 \times \sqrt{N} \times \sqrt{2ln\left(\frac{2p}{0.1/ln(N)}\right)}$, where p is the number of covariates and N is the sample size.

Online Appendix Table A1 reports the double-selection post-LASSO minimum wage elasticities for teen employment using the 1979-2014 basic monthly CPS, varying the penalization parameter, λ, from the most saturated specification to the least saturated one. All estimates include state and period fixed effects, which are partialed out prior to the LASSO based covariate selection. The covariates that LASSO chooses from include demographic controls, unemployment rate, state-specific linear trends, and division-period effects. Appendix Table A1 shows the point estimates and the confidence intervals associated with varying λ between 0 (the most saturated model) and 3500 (which only picks the state unemployment rate as a control beyond the manually-specified two-way fixed effects). The point estimate quickly falls under -0.045 in magnitude as λ is lowered to 2,000 or below. For $\lambda = 2000$, the LASSO double-selection procedure includes just 5 state-specific linear trends lowers the elasticity in magnitude to -0.040. In other words, merely adding state-specific linear trends for these 5 states (which happen to be CA, SD, OR, WA and VT) to the fixed effects model produces an estimate that is close to zero, and not statistically significant.

Finally, we note that Christian Hansen's 2013 NBER Econometric Lecture reports 5 possible asymptotically equivalent calculations for λ, which, in our case of $p = 1207, \mathrm{N} = 7344$, range between 12.562 and 1161.99. As shown in Online Appendix Table A1, this range of λ implies a range of double-selection post-LASSO estimates for the minimum wage elasticity between -0.018 and 0.059.

Online Appendix B: Sample windows, business cycles and parametric trends

In this Online Appendix, we systematically evaluate the claim in NSW (2014a, 2014b) that the results in ADR are driven by the choice of sample years (1990-2009) because of "endpoint bias." In particular, NSW claim that the presence of recessions in the beginning and end points affects the estimation of the state-specific trends, and biases the minimum wage elasticities for teen employment towards zero when such trends are included.[28] They also claim that inclusion of higher-order (cubic or greater) state-specific trends in that sample restores the more negative estimates by correcting for this "endpoint bias."

Online Appendix Figure B1 compares the four key specifications in ADR across 72 samples by varying the starting years between 1979 and 1990, and the ending years between 2009 and 2014. The two-way fixed-effects model stands out in the figure as having more negative estimates in each of the 72 samples than any of the other models. Indeed, the estimates from the two-way fixed effects model are more negative in the longer samples, consistent with the argument that long-run trends may be contaminating the estimates. The estimates from the intermediate specifications (with either state-specific linear trends, or division-period fixed effects) vary somewhat, depending on the sample. In particular, the model with just linear trends produces estimates that are somewhat negative in samples that begin with 1990, but become smaller in magnitude for estimates in samples beginning in 1989 or earlier. Moreover, extending

[28] Specifically, NSW (2014a, p. 616) states: "In models that include state-specific trends, the recessions at the beginning and end of ADRs sample period could have a large influence on the estimated state-specific trends—a so-called endpoint bias. If the recessions have a purely aggregate influence that is common across all states, this will not happen, as the year effects will absorb this common influence. But if the recessions led to cross-state deviations between teen employment rates and aggregate labor market conditions, then the estimated longer-term trends in teen employment could be biased. This, in turn, could lead to mis-classification of periods in which teen employment was high or low relative to the predicted values net of the minimum wage and hence influence the estimated minimum wage effects for reasons having nothing to do with the longer-run trends for which the specification is trying to control."

the sample forward also reduces the magnitude of the estimate. In contrast, the most saturated specification is quite robust with respect to the choice of the sample period. While stability of estimates across samples need not indicate accuracy, these results are consistent with the idea that using both state-specific trends and division-period effects guards against estimation errors when either set of control is included on its own. We made a similar observation in ADR, where we specifically warned against the reliability of estimating parametric trends in short samples and suggested the usefulness of including multiple types of controls.[29]

Online Appendix Figure B1 is also informative about “endpoint bias.” The two figures at the bottom of Figure B1 does not provide any indication that the 1990-2009 sample used in ADR produced more positive estimates (both include controls for linear trends). Indeed, the opposite appears to be the case: the estimates are more negative in the 1990-2009 sample than in the other samples. The estimates with the 1990-2009 sample using the CPS basic monthly data are somewhat more negative than the estimates reported in ADR, which used the CPS ORG data. However, the conclusion from the most saturated model remains qualitatively the same as in the original sample. Moreover, even small expansions of the sample produce estimates closer to the ones for our full 1979-2014 sample.

As another assessment of the role of business cycles in affecting estimation of trends, Online Appendix Table B1 shows how the estimates vary when we exclude recessionary periods from the sample. The table includes two definitions of recessions. One consists of the standard NBER-defined recessionary periods. The second expands the NBER concept to include quarters

[29]We wrote in ADR, p. 237: “Generally speaking, our preferred specification 4 [with division-period effects] tends to be more stable across time periods than does specification 3 with just state linear trends...While linear trends do a good job of eliminating long-term trend differences across states in longer panels, they are a less valuable means of controlling for spatially correlated shocks, and they are estimated poorly in shorter panels.”

until the national employment reaches the pre-recession peak. This expansive definition excludes the following periods from the sample: 1980q1-1980q4, 1981q3-1983q3, 1990q3-1992q4, 2001q1-2004q4, and 2007q4-2014q1. Overall, we find little indication that excluding recessionary quarters produces sizably negative estimates in models with state trends (column 3 and 4).

The exclusion of NBER recessions makes little difference to the point estimates in any of the four models. Exclusion using the expanded definition produces a point estimate of -0.033 for the specification with state specific trends only (column 3); and -0.078 for the most saturated specification (column 4); neither are statistically significant at conventional levels.

We noted in section describing the LASSO results that the models including cubic or higher order polynomial time trends by state produce estimates that were more negative in the shorter (post 1990) sample, but not in the expanded sample. Online Appendix Figure B2 shows how the results vary when higher order trends are introduced across the 72 samples with start dates varying between 1979 and 1990, and end dates varying between 2009 and 2014. (These specifications use common time fixed effects and do not additionally control for division-period effects.)

We find that the estimates using quadratic trends are similar to those with linear trends, are fairly robust to sample choice, and almost never exceed -0.1 in magnitude. However, when we extend the sample by including earlier start dates, we produce estimates that are generally smaller in magnitude. Starting the sample even a few years earlier than 1990 greatly shrinks the estimates from models with trends towards zero, even when higher order trends are included. Recall, however, that the results from LASSO-based double-selection procedure reported in the text suggest that the data do not warrant higher order trends. This result holds in both the full sample as well as the sample beginning in 1990. The combination of these two facts casts serious doubt on the relevance of the finding in NSW (2014a) that inclusion of third or higher order trends in the 1990-2011 sample reproduces more negative employment effect of minimum wages on teen employment.

Online Appendix C: Timing of teen employment effects of the minimum wage

In the first part of Online Appendix C, we report the quarterly cumulative teen employment elasticities from the 1979-2014 individual-level data (hereafter "full sample"). We averaged these underlying quarterly estimates to produce estimates in Table 3. In the second part, we provide analogous estimates using the 1990-2011 state-aggregated data used in NSW (2014b, hereafter "NSW data"). We also report with 3 and 4+ year employment elasticities netting out leading coefficients using the NSW data, analogous to the estimates provided in Table C1 using the full sample. We are therefore able to reconcile our two sets of results: the discrepancy is not due to the samples used, but rather due to the peculiar, and fragile, choice of a baseline used by NSW (2014b) when netting out leading coefficients.

Cumulative response of teen employment to a minimum wage increase in the full sample

Online Appendix Figure C1 plots the cumulative teen employment elasticities from the 1979-2014 individual-level sample using the two-way fixed effects model, along with 95 percent confidence intervals for the two-way fixed effects model

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

along with increasingly saturated models including up to state-specific linear time trends and division-period fixed effects

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

The figures show the cumulative elasticity as the sum of quarterly coefficients β_k or annualized versions ρ_τ, as described in the main text. The annualized estimates are also reported in Table 3.

The leftmost panel ("Quarterly: two-way FE") in Figure C1 plots these cumulative responses from the two-way fixed effects model, along with 95 percent confidence bands for the full sample. We can observe a clear visual pattern: every pre-treatment point estimate for the two-way fixed-effects model is negative and 5 of the 12 coefficients are statistically significant at the 5 percent level.

To reduce noise and more easily extract a signal from the data, the panel entitled "Annualized: two-way FE") in Figure C1 show four-quarter averages of these quarterly cumulative response elasticities: $\overline{\rho}_{[\tau,\tau+3]} = \frac{1}{4}\sum_{m=0}^{3} \rho_{\tau+m}$, along with the 95 percent confidence bands. (These averaged cumulative response elasticities and standard errors are also reported in the first column of Appendix Table C1) We find unmistakable evidence that the two-way fixed-effects model fails the falsification test that leading coefficients during 1, 2 or 3 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.

Cumulative response of teen employment to a minimum wage increase in the NSW sample

The results in Online Appendix Figure C1 as well as Table 3 differ from those in NSW 2014(b), who deny that there is evidence of pre-existing trends in the two-way fixed-effects model. They also argue that netting out the leading coefficients does not alter the estimates very much.

To assess their conclusions, we first estimate analogous regressions using their data and specification (i.e., state-by-quarter level data from 1990q1-2011q1).[30] Online Appendix Figure C2 shows the cumulative teen employment elasticities using the shorter NSW sample and specification and reproduces their estimates. (see their Figure 6).

[30] We use the replication data on Ian Salas' website: https://sites.google.com/site/jmisalas/data-and-code. This model is estimated using exactly the same data, sample, and specification that produce NSW 2014b Figure 6: they include controls for unemployment rate, state and period fixed effects.

When NSW (2014b) analyze estimates from this regression and compare cumulative responses using ρ_{-2} as their baseline, they conclude that "the contemporaneous elasticities are close to 0.2, building to a maximum of about 0.4 five quarters after the increase—a period around which the estimates are significantly different from zero" (p. 13). Online Appendix Figure C2 shows that their conclusion arises entirely from their peculiar choice of ρ_{-2} as the baseline, which was unusually positive (highlighted by the red circles in the figure). Had they chosen an arguably more "standard" ρ_{-1} as the baseline, they would have reached the conclusion that the 3 year or 4+ year effects are very close to zero (compare $\overline{\rho}_{[8,11]}$ or ρ_{12} to ρ_{-1}).

Since the leading coefficients appear to exhibit a seasonal pattern—or at least considerable variability—one can reach very different conclusions by picking particular quarters. To avoid cherry picking, in Online Appendix Table C1 we construct the baseline as an average of the cumulative response during four quarters just prior to treatment ($\overline{\rho}_{[-4,-1]}$), or the four preceding quarters ($\overline{\rho}_{[-8,-5]}$).We find a clear signal that netting out the leading coefficients substantially reduces the medium and long term estimates from the two-way fixed-effects model.[31]

Column 1 of Table C1 first reports the four-quarter averages of all cumulative teen employment elasticities. The four-quarter averaged cumulative response elasticities $\overline{\rho}_{[-4,-1]}$ and $\overline{\rho}_{[-8,-5]}$ are sizable, and are -0.118 and -0.126, respectively, although they are not statistically significant at conventional levels. However, as shown in Table C1 (column 1), the estimated 3

[31]Because quarterly leads and lags can be noisy, ADR and DLR (2010) included leads and lags at 4-quarter or 2-quarter frequency for the purpose of smoothing. However, we do acknowledge that there is some arbitrariness in any choice of smoothing. In this paper, for the purpose of comparability with NSW (2014b), we have used their quarterly lead/lag structure in the regression, and simply take four quarter averages of these coefficients to construct $\overline{\rho}_{[\tau,\tau+3]}$. This procedure has the added advantage of reducing arbitrariness in choosing the lead/lag frequency.

year and 4+ year effects range between -0.040 and -0.074 when using baselines that are 1 or 2 year prior to treatment, and are much smaller than their estimate from the static specification (-0.165, reported in Table 1, column 1 of NSW (2014a).[32] In other words, the estimates from the two-way fixed effects model in the sample used in NSW (2014b) show that between 50 and 75 percent of the reduction in employment implicit in the static minimum wage employment estimate occurs prior to the minimum wage increase—whether the baseline is one or two years prior to the minimum wage change.

Finally, Online Appendix Figure C2 shows using the NSW (2014b) data that both the model with state-specific linear trends, and the model with division-specific fixed effects and state-specific linear trends, pass the leading effects falsification test and generally do not obtain large, negative employment effects over the post-treatment period. Online Appendix Table C1 reports the 3 and 4+ year elasticities using the NSW (2014b) data and specifications. NSW (2014b) argues that models with additional controls for spatial heterogeneity do not outperform the two-way fixed- effects model on the leading effects falsification test. But this conclusion is based only on their evaluation of the model with just division-period effects. As in the full 1979-2014 sample, the leading coefficients in this specification exhibit a substantial negative bias.. However, including state-specific trends produces much smaller leading coefficients—with or without the inclusion of division-period effects.

Overall, our reanalysis of the 1990-2011 sample used in NSW (2014b) shows that much of the employment reduction that the two-way fixed-effects model estimates occurs substantially prior to minimum wage increase. In contrast, models with controls for state-specific trends tend to have smaller leading coefficients. And in all cases, after netting out the leading coefficient the

[32]These estimates are from rows labeled “F-C” “F-B” “G-C” and “G-B”.

employment estimates are substantially smaller. None of the 3 or 4+ year out effects exceeds -0.1 in magnitude, regardless of baselines (one, two or three years before) or specifications. As expected, the precision of some of the estimates is lower in the smaller sample; but the overall conclusion is qualitatively similar when we use the 199-2011 NSW sample, as it is for the full 1979-2014 sample used in this paper.

Online Appendix D: Donor-distance relationship based on synthetic controls

We find a clear negative relationship between relative donor weights and the geographic distance between donor and treated states for the set of minimum wage increases analyzed by Dube and Zipperer (2015).[33] For each donor j from a given treatment event, we define the relative donor weight equal to the synthetic control weight $0 \leq w_j \leq 1$ divided by the average donor weight for that event (equal to one divided by the number of donors). Figure 5 non-parametrically plots (using lowess) the mean relative donor weight as a function of the distance between the geographic centroids of the donor and treated states. The confidence bands are based on standard errors clustered by event. Since the measure of distance is less meaningful when dealing with Hawaii and Alaska, we drop these two non-contiguous states from treatment and donor samples for this exercise. (The key findings are similar when they are included.)

For the resulting 25 events, the relationship between distance and donor weights is clearly negative, especially for the first 500 miles. For example, a donor state whose centroid was 100 miles away from the treated state receives, on average, about 2.5 times the weight of a donor state that was 500 miles away from the treated state, and nearly 3.8 times the weight of a donor that was 1000 miles away. Overall, the greater weight for nearby donors constitutes evidence for the similarity of factor loadings between states that are nearer geographically, providing added support for leveraging proximity when constructing controls. This variation provides information

[33]In an earlier (2013) version of this paper we performed an analogous exercise but used synthetic control weights from a set of randomized placebo laws and found the average weights declined sharply by distance between the donor and the (placebo) treated state. We used placebo laws to assess whether nearby states are indeed more similar, which is the key contention of NSW. NSW (2014b, p.26) then criticized our use of placebo and not actual minimum wage increases, stating that this approach strikes us as uninformative about the question at hand—whether a particular subset of states provides a more valid set of controls for states where the minimum wage actually does increase. The analysis in this section addresses their concerns and shows that, on average, donor weights do indeed decline with distance when looking specifically at actual minimum wage treatments.

on the extent of spatial correlation among the loadings from the time-varying factors. With strong spatial correlations in loadings, nearby areas are likely to receive higher donor weights.

Online Appendix E: Restaurant employment effects in recent studies

We examine five recent studies of restaurant employment that incorporate some method of controlling for time-varying heterogeneity. These include the county border discontinuity papers of DLR and DLR (2016), the study by Addison, Blackburn and Cotti (2014) using parametric trend controls, the factor model approach of Totty (2015), and the "matching estimator" in NSW (2014a). Online Appendix Table E1 reports a total of 17 employment elasticities from these key five papers that include additional controls beyond the two-way fixed effects model. Altogether, the restaurant employment elasticities for models that include additional controls for time-vary heterogeneity range from -0.063 to 0.039.[34]

[34]Aaronson, French and Sorkin (2015), who study restaurant employment using a border discontinuity design for a small number of states, obtain an overall short-run elasticity of -0.1. They describe this estimate as "very imprecise" (they do not report a standard error). They also find increased entry and exit, which they interpret 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 restaurant sector in the "medium run," i.e., after 12 or 16 quarters following the minimum wage.

Online Appendix F: Spatially correlated placebos

Appendix B of DLR used a placebo-based falsification exercise to provide additional evidence on the bias of the canonical two-way fixed-effects model arising from spatial heterogeneity. The basic idea is to assess whether minimum wages in nearby areas are correlated with own employment—even when variation in the true minimum wage could not possibly affect employment across different counties in the sample. DLR ensure this by taking a sample of counties for which the minimum wage evolved in exactly the same way—they were always bound by the federal minimum wage. Yet, the results indicated that restaurant employment in these counties was correlated with nearby minimum wages—which we described as a "placebo" since there was no true minimum wage variation across counties in the estimation sample. DLR interpreted this finding to be consistent with the presence of an omitted variable that is spatial in nature.[35] NSW (2014a, 2014b) argue that this test is invalid, because the effect of the true minimum in these areas is not fully accounted for using time fixed effects. In this Appendix, we describe the original exercise in greater detail and explain why the results do, indeed, demonstrate that spatial heterogeneity contaminates the two-way fixed-effect model's employment estimates. We also discuss why the counter-argument in NSW (2014a, 2014b) is incorrect.

For this exercise, DLR started with the stacked border county pair sample, and kept only the subset of counties in which the prevailing minimum wage was always equal to the federal minimum wage: $MW_{st}^{S} = MW_{t}^{F}$. Call this the set S. Define also the set S' of cross-border counties adjacent to each of the counties $i \in S$—this will be used to define a placebo minimum

[35] Section I.G. and Appendix B of DLR (2010) describe the original exercise.

wage below.[36] To emphasize, while S-county minimum wage, MW_{it}^{S} is always equal to the federal minimum, MW_{t}^{F}, the same is not true for the minimum wage in their cross-border neighbors, $MW_{jt}^{S\prime}$.

Now consider the data generating process underlying the two-way fixed effects model:

$$E_{it}^{S} = \gamma MW_{it}^{S} + \mu_i + \tau_t + \nu_{it} \qquad (6)$$

where E_{it}^{S} is log of restaurant employment in county i, MW_{it}^{S} is the log of the prevailing minimum in that state, μ_i is a county fixed effect and τ_t is a time fixed effect, and ν_{it} is a mean zero disturbance term. Next, consider regressing E_{it}^{S}, on own county minimum wage, MW_{it}^{S}, as well cross-border neighbors' minimum wage, $MW_{jt}^{S\prime}$ along with time dummies (D_t) and county dummies, D_i using the sample S:

$$E_{it}^{S} = \gamma MW_{it}^{S} + \delta MW_{jt}^{S\prime} + D_i\theta + D_t\lambda + \epsilon_{it} \qquad (7)$$

Note that in our estimation sample of counties $i \in S$, the prevailing minimum wage is the federal one, and so $MW_{it}^{S} = MW_{t}^{F}$ is only a function of time. As a consequence, the effect of the prevailing minimum wage in S is fully soaked up by the time dummies D_t. To put it differently, by including D_t as covariates, we are "dummying out" the actual prevailing minimum wages in the states S. And so estimating (7) is equivalent to estimating:

$$E_{it}^{S} = \delta MW_{jt}^{S\prime} + D_i\theta + D_t\lambda + \epsilon_{it} \qquad (8)$$

which is what DLR did estimate. The purpose of choosing the set of border counties S where $MW_{it}^{S} = MW_{t}^{F}$ was to avoid having to obtain a consistent estimate of γ, since in this sample, any effect of own-county minimum wage effect is fully accounted for through the time dummies.

[36]In this Appendix, for comparability we use a similar notation as in NSW (2014b) instead of the original DLR (2010) notation and the notation elsewhere in this paper. However, we index the variables $MW_{it}^{S}, MW_{jt}^{S\prime}$ instead of $MW_{st}^{S}, MW_{st}^{S\prime}$ to clarify the county in question (i, or its neighbor, j).

What should we expect $\hat{\delta}$ to be? Barring cross-border spillover, the level of S'-county minimum wages should have no causal effect on S-county employment conditional on the time effects.[37] This is why we considered $MW_{jt}^{S'}$ as a "placebo" minimum wage: a $\hat{\delta} \neq 0$ indicates a likely bias due to an omitted variable, for which $MW^{S'}$ is acting as a proxy. This omitted variables bias is exactly what one expects in the presence of spatial heterogeneity—that minimum wage increases tend to be correlated with certain regional shocks.

As a point of comparison, we further estimated a two-way fixed-effect model regressing actual employment in the S' counties on S' minimum wages in this particular sample:

$$E_{jt}^{S'} = \beta MW_{jt}^{S'} + D_i\theta + D_t\lambda + \epsilon_{jt}$$

If the size of β is similar to δ, then according to the two-way fixed-effects model, the "effect" of $MW^{S'}$ on neighboring county employment, E^S, is comparable to the "effect" on own employment, $E^{S'}$, even after controlling for the minimum wage in S. DLR estimated that $\hat{\delta} = -0.123$, which is sizable, compared to $\hat{\beta} = -0.208$. They commented that "we find a negative effect in both samples (though imprecise), with elasticities exceeding -0.1 in magnitude, suggesting bias in the canonical [two-way fixed effects] specification."

NSW (2014a, 2014b) criticize this falsification test. Most importantly, they argue that that the S-county sample is "contaminated." Their argument has two parts. First they (correctly) point out that that even though the minimum wages in S-counties are equal to the federal minimum$MW_{it}^{S} = MW_t^F$, the minimum wages are changing over time, which can have a real effect. Second, they (incorrectly) argue that because the cross-border minimum wage $MW_{jt}^{S'}$ can

[37] Dube, Lester and Reich (2010, section VA, failed to find evidence of such spillovers in a test comparing border and interior counties. Moreover, NSW's criticism of our falsification exercise is not based on the possibility of such spillovers. For this reason, we do not discuss the spillover issue further in this Appendix.

be correlated with true minimum wage MW_{it}^{S}, δ can reflect some of the effect of MW_{it}^{S} on E_{it}^{S}, thereby "contaminating" the falsification exercise. They write: "But $MW_{st}^{S\prime}$ in equation (5) [similar to our equation (8) above] varies with the federal minimum wage in a way that is not perfectly correlated with the period fixed effects, because whether the federal minimum wage variation changes the cross-border minimum wage depends on whether the state or federal minimum wage is binding. Thus, federal minimum variation is not swept out by the period fixed effects, and therefore the cross-border minimum wage variation will be correlated with the actual state minimum wage variation." Formally, their argument is that since $Cov(MW_{jt}^{S\prime}, MW_{it}^{S}) = Cov(MW_{it}^{S\prime}, MW_{t}^{F}) > 0$, if $\gamma \neq 0$ and hence MW^{F} has a causal effect on E_{it}^{S}, this can be reflected in $Cov(E_{it}^{S}, MW_{jt}^{S\prime}) < 0$ under the data generating process represented by equation (7).

This argument is fundamentally flawed. By construction, the prevailing minimum wage in the sample of S-counties is the federal one*:* $MW_{it}^{S} = MW_{t}^{F}$. So, the true minimum wage effect—whatever it may be—is completely dummied out by the time dummies D_t, fully accounting for any causal effect of the federal minimum wage on E_{it}^{S}. Therefore, conditional on D_t, if MW_{it}^{S} is still correlated with E_{it}^{S}, it is *only* for a spurious reason, and *not* due to a causal effect of MW_{t}^{FED}. Formally, while *unconditionally* it may be that $Cov\left(MW_{jt}^{S'}, MW_{it}^{S}\right) = Cov\left(MW_{it}^{S'}, MW_{t}^{F}\right) > 0$, conditional on the fixed effects, we have $Cov\left(MW_{jt}^{S'}, MW_{it}^{S} \middle| D_t, D_j\right) = Cov\left(MW_{it}^{S'}, MW_{t}^{F} \middle| D_t, D_j\right) = 0$. Therefore, conditioning on the time dummies also conditions on the actual minimum wage in the S-county. So if $Cov\left(E_{it}^{S}, MW_{jt}^{S'} \middle| D_t, D_j\right) = Cov(E_{it}^{S}, MW_{jt}^{S\prime} | D_t, D_j, MW_{it}^{S}) < 0$, by definition it is via something *other than* the effect of the prevailing minimum wage, because the inclusion of the time dummies fully accounts for the impact of the prevailing (federal) minimum wage in the S-counties.

Surprisingly, NSW (2014a, 2014b) insist upon rejecting this argument and claim that the time dummies do not remove the causal effect of MW_t^F upon E_{jt}^S. As supposed evidence for this claim, they provide a decomposition of $MW_{it}^{S\prime}$ that aims to demonstrate that δ can still reflect true causal effects. In actuality, their decomposition provides an illustration of exactly the opposite argument. In particular, they rewrite equation 8 as:

$$E_{it}^S = \delta\left(MW_t^F \cdot I\left\{MW_{jt}^{S\prime} = MW_t^F\right\} + MW_{jt}^{S\prime} \cdot I\left\{MW_{jt}^{S\prime} > MW_t^F\right\}\right) + D_i\theta + D_t\lambda + \epsilon_{it} \qquad (9)$$

They then argue that the first term, $MW_t^F \cdot I\left\{MW_{jt}^{S\prime} = MW_t^F\right\}$, will not be not swept out by the time dummies because the federal minimum wage is multiplied by a dummy of whether the binding minimum wage in neighboring county j is the federal one. They write "[c]learly the federal variation can play a role here because the federal minimum wage is multiplied by a dummy that is sometimes one and sometimes zero, breaking the perfect collinearity with the time fixed effects." They interpret this to mean that δ can reflect the true effect of MW_t^Fon E_{it}^S.

However, NSW miss the implication of this decomposition. Of course, the interaction term $MW_t^F \cdot I\left\{MW_{jt}^{S\prime} = MW_t^F\right\}$ can be correlated with $\mathrm{E}_{\mathrm{it}}^{\mathrm{S}}$ even after controlling for the time effects; after all, it is part of $MW_{jt}^{S\prime}$, which we show is empirically correlated with E_{it}^S. The point is that there in no *plausible causal* interpretation of that correlation under the data generating process represented by equation (6). We have already laid this out above: by definition, conditioning on D_t conditions on $MW_{st}^S = MW_t^F$, so the conditional covariance between E_{it}^S and $M_{jt}^{S\prime}$ is precisely purged of the impact of the federal policy. So what would be the meaning of a negative correlation between E_{it}^S and the interaction term $MW_t^F \cdot I\left\{MW_{jt}^{S\prime} = MW_t^F\right\}$? Since MW_t^F is just a function of the time effects, the interaction term simply measures heterogeneity in the time effects by the nature of the minimum wage in the "neighborhood." So a negative correlation between E_{it}^{S} and the interaction term indicates that when the federal wage rises uniformly across all counties in the sample at date t, and if we take two counties i_1, i_2 that are both bound by the federal wage, employment falls more in i_1 than in i_2 when i_1's neighbor (j_1) is also bound by the federal minimum wage, while $i_2's$ neighbor (j_2) happens not to be. Again, this is for two counties i_1 and i_2 that have the identical (i.e., federal) minimum wage, so the true causal effect of the federal increase should be the same under the data generating process in equation (6). A non-zero coefficient on $MW_t^F \cdot I\{MW_{jt}^{S\prime} = MW_t^F\}$ indicates that employment changes in i_1 and i_2 turn out to depend on the characteristics of the neighbors j_1 and j_2 precisely reflecting the evidence of an omitted variable that is spatial in nature. Therefore, NSW's argument of "contamination" of the placebo by federal minimum wage is erroneous and the decomposition they marshal as evidence for "contamination" actually demonstrates the validity of our original exercise.[38]

NSW (2014b) also make a second point regarding this placebo test. They argue that correcting a small data error in DLR changes $\hat{\beta}$ from -0.208 to -0.114, while $\hat{\delta}$ is largely

[38]In Table 5 of NSW (2014b), the authors artificially change the federal minimum wage, and show that this change has a mechanical effect on the point estimate of $\hat{\delta}$ through changing in some cases the value of the placebo minimum wage ($M_{jt}^{S\prime}$) . This exercise sheds no light on the validity of our placebo test: transforming the placebo minimum wage and finding that the coefficient is altered is hardly a surprise. We are arguing that $\hat{\delta}$ should be zero under the data generating process of the two-way fixed-effects model, but it is not empirically—reflecting the correlation between employment with the neighboring minimum wage, $M_{jt}^{S\prime}$. Their exercise of changing the federal minimum wage artificially shows that the measured $\hat{\delta}$ can be changed by artificially changing MW_t^F and hence $MW_{jt}^{S\prime}$. This is both true and irrelevant. Subsequently, as an effort to "solve" the "contamination" problem, NSW proceed to eliminate over half of the sample by cutting out many of the years and by imposing an arbitrary restriction on cross-border minimum wage variation. One of these restrictions excludes all federal minimum wage increases from the sample; this restriction does not attain their objective since the real minimum wage is changing due to inflation. At any rate, once they eliminate over half of data using these arbitrary criteria, they find that the placebo estimate becomes close to zero. This "solution" does not shed any light on the validity of the placebo exercise, since their assertions about the "invalidity" of the original placebo test are erroneous.

unaffected (changing from -0.123 to -0.125).[39] We acknowledge the data error, but note that this correction actually appears to strengthen the conclusion we drew in DLR: the point estimate of the placebo minimum wage $\hat{\delta}$, is essentially of the same size (or slightly larger) than the (corrected) own minimum wage estimate, $\hat{\beta}$, and both exceed -0.1 in magnitude. The corrected sample suggests that unaccounted spatial heterogeneity in the two-way fixed-effects model explains nearly the entirety of the negative employment estimate.

[39] DLR (2010) incorrectly coded Maryland as having raised its wage to $6.15 in q1 and q2 of 2006, when in reality it was $5.15 during those two quarters. We thank NSW for catching this mistake. NSW note that this coding error does not influence any of the analysis in the paper other than the placebo exercise.

Appendix Figure B1. Minimum wage elasticities for teen employment, by panel start and end year

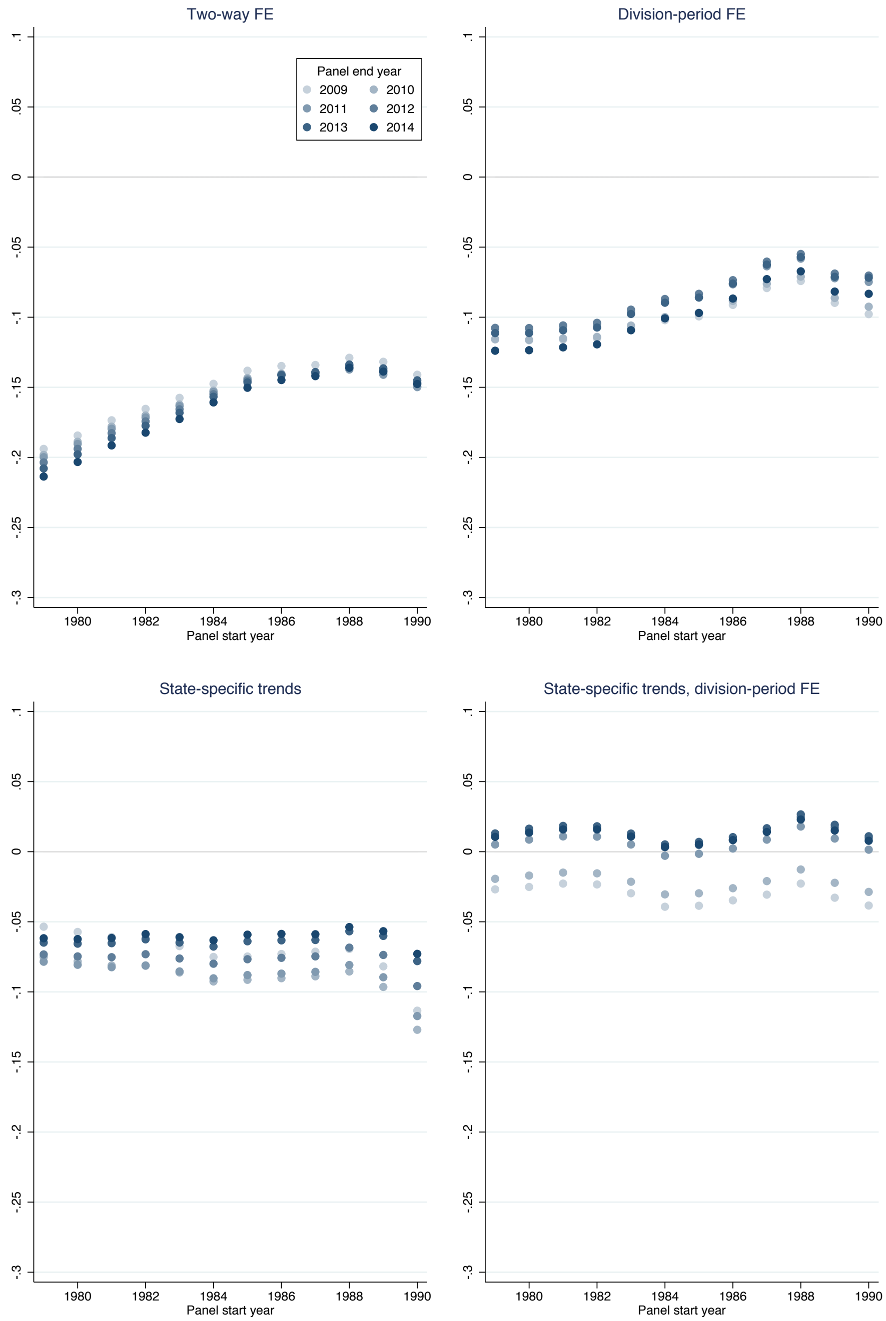

Notes: The figure reports teen employment elasticities with respect to the minimum wage for the four regression models using individual-level data described in the text. From these linear probability models, we calculate elasticities by dividing the log minimum wage coefficient by the sample mean employment.

Appendix Figure B2. Minimum wage elasticities for teen employment, by type of state-specific polynomial trend, and panel start and end year

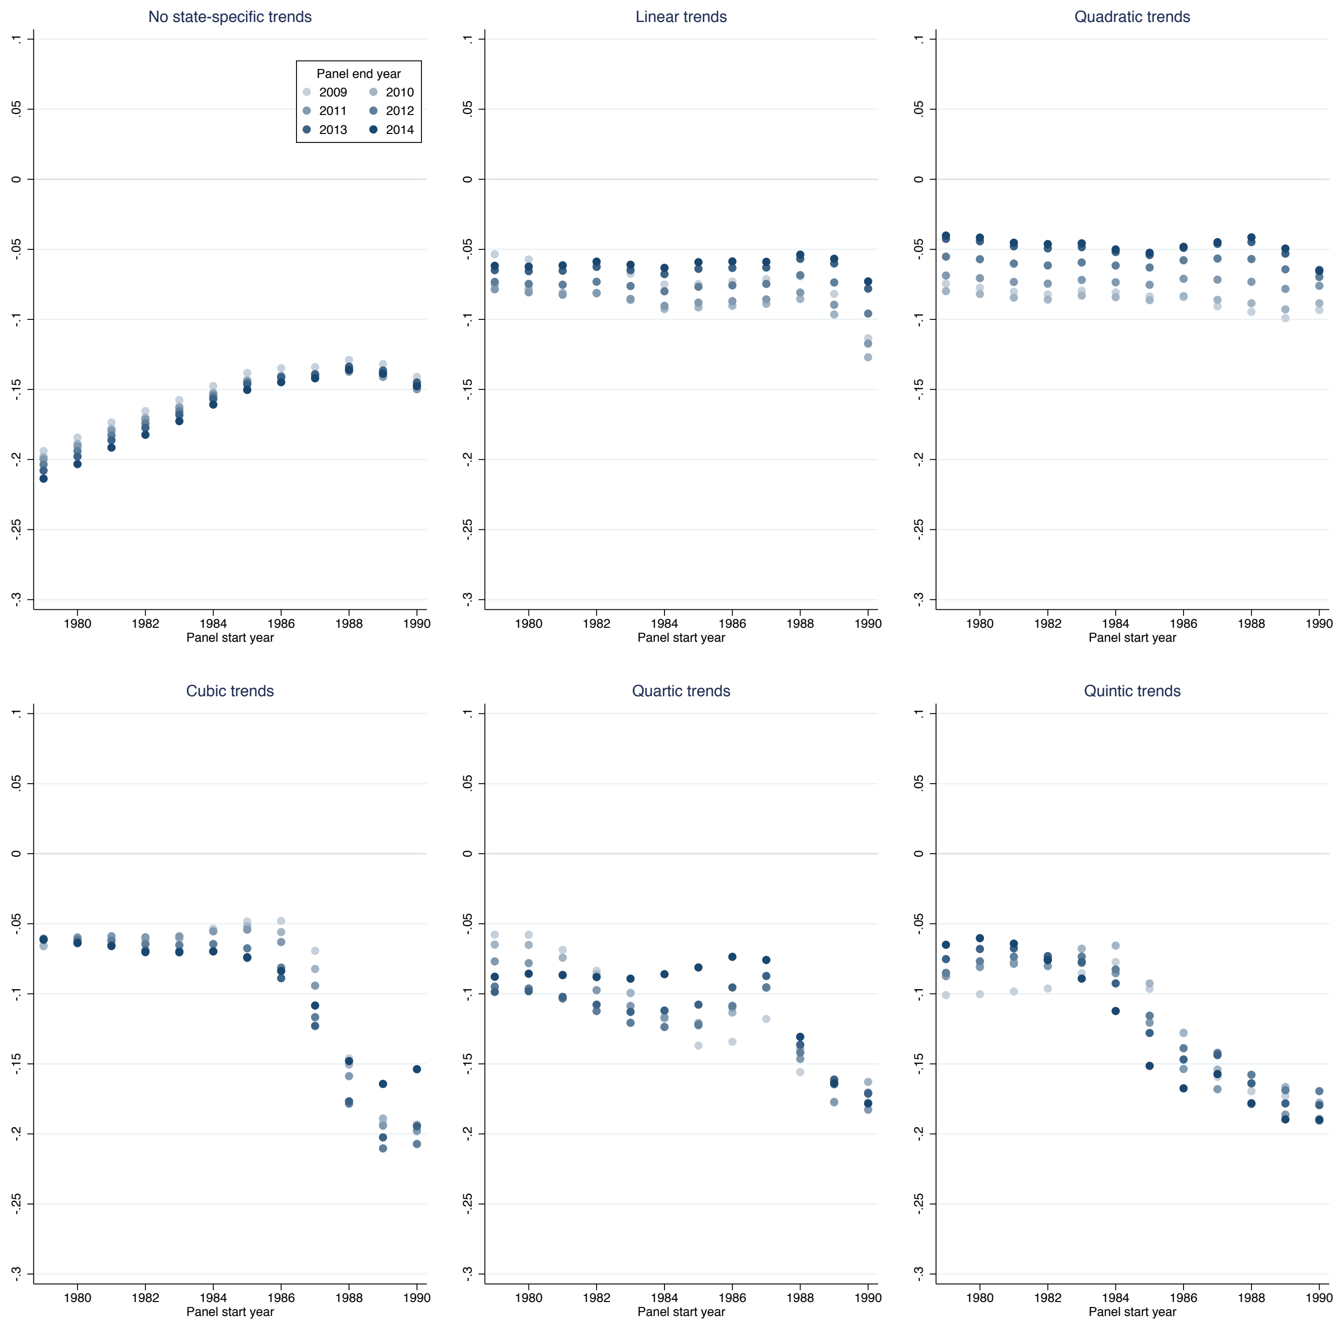

Notes: The figure reports teen employment elasticities with respect to the minimum wage for the two-way fixed-effects model, with state-specific trends of polynomials degree zero through five, using the individual-level data described in the text. From these linear probability models, we calculate elasticities by dividing the log minimum wage coefficient by the sample mean employment.

Appendix Figure C1. Cumulative response of teen employment to minimum wages, individual-level CPS data, 1979-2014

Notes: The figure shows cumulative response elasticities of teen employment to the minimum wage; 95% confidence intervals are from the individual-level distributed lag regressions for the 1979-2014 period described in the text. For each of the four regression models, the figure shows the quarterly effects and confidence intervals in blue and the 4-quarter averaged effects and confidence intervals in green. Standard errors are clustered by state.

Appendix Figure C2. Cumulative response of teen employment to minimum wages, state-quarter data from NSW (2014b), 1990-2011q1

Notes: The figure shows cumulative response elasticities of teen employment to the minimum wage; 95% confidence intervals are from the distributed lag regressions described in the text, with the state-quarter aggregated data and control set used in NSW (2014b), where the outcome is log of the teen employment-to-population ratio. For each of the four regression models, the figure shows the quarterly effects and confidence intervals in blue and the 4-quarter averaged effects and confidence intervals in green. Standard errors are clustered by state. Red colored marker for t=-2 indicates the baseline used by NSW.

Appendix Figure D1. Donor distance and relative weights

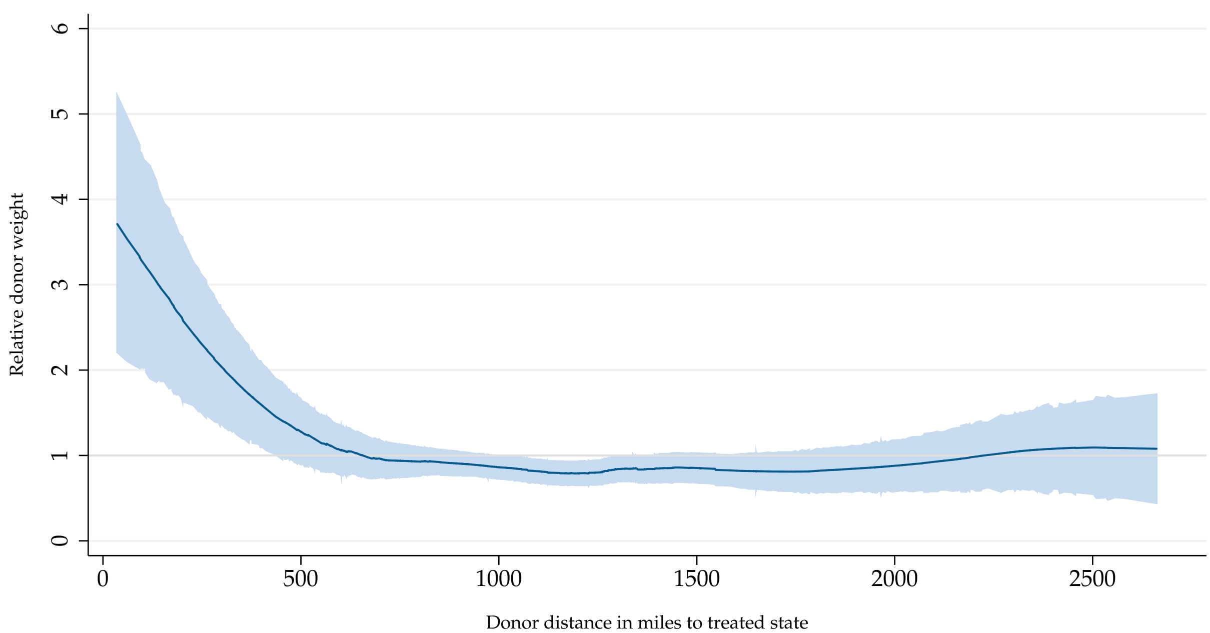

Notes: The figure shows the locally weighted regression (lowess) of the relative donor weights on donor distance to treated states for the 25 treatment events, excluding Alaska and Hawaii, as described in Section 5, where the 95 percent confidence interval is calculated from 1000 cluster bootstrap iterations at the treatment event level.

Appendix Table A1. 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

λ	# of state specific trends chosen	# of divisions selected	# division–period FE chosen	Estimate	Std. Error
0	.	.	.	0.059	0.057
16	50	8	1,139	0.059	0.057
50	48	8	1,113	0.057	0.057
84	47	8	1,091	0.059	0.057
118	48	8	1,053	0.053	0.059
152	48	8	1,012	0.039	0.059
186	47	8	948	0.006	0.065
220	47	8	860	0.006	0.062
254	46	8	779	-0.001	0.061
288	45	8	644	0.001	0.059
322	43	8	552	-0.002	0.055
356	43	8	466	-0.003	0.058
390	44	8	387	0.000	0.053
424	44	8	323	0.003	0.052
458	43	8	258	-0.012	0.049
492	44	8	200	-0.002	0.060
526	43	8	155	0.010	0.063
560	42	7	116	-0.022	0.060
594	41	7	98	-0.020	0.063
628	41	7	83	-0.021	0.063
662	40	6	73	-0.019	0.062
696	40	5	64	-0.033	0.063
730	39	3	61	-0.033	0.059
764	39	2	53	-0.032	0.060
798	37	2	48	-0.030	0.060
832	36	2	46	-0.036	0.060
866	36	2	44	-0.036	0.060
900	34	1	27	0.015	0.081
934	32	1	20	0.015	0.082
968	30	1	22	0.015	0.082
1,008	29	1	19	0.012	0.083
1,048	25	1	19	0.007	0.083
1,088	24	1	18	0.004	0.084
1,128	20	1	16	0.014	0.083
1,168	18	1	14	-0.018	0.088
1,208	18	1	14	-0.018	0.088
1,248	17	1	14	-0.021	0.089
1,500	14	0	0	-0.025	0.084
1,750	10	0	0	-0.038	0.080
2,000	5	0	0	-0.039	0.081
2,250	3	0	0	-0.160***	0.040
2,500	2	0	0	-0.159***	0.041
2,750	1	0	0	-0.164***	0.041
3,000	1	0	0	-0.236***	0.064
3,250	0	0	0	-0.257***	0.056
3,500	0	0	0	-0.257***	0.056

Notes: The table reports double-selection post-LASSO estimates regressing log teen employment on log of the quarterly minimum wage, using state-quarter aggregated CPS data. All regressions are on data after partialing out state and period fixed effects. LASSO regressions allow state-specific linear trends and division-period fixed effects, and demographic controls (see notes to Table 2 for details). λ is the penalization parameter for the LASSO regressions. Standard errors (in parentheses) are clustered at the state level, and significance levels are indicated by *** 1%, ** 5%, * 10%.

Appendix Table B1. Minimum wage elasticities for teen employment using alternative samples to exclude recessions, individual-level CPS data 1979-2014

	(1)	(2)	(3)	(4)
Full Sample	-0.214***	-0.124	-0.062	0.011
N=3,534,924	(0.044)	(0.079)	(0.041)	(0.048)
Leave out recessions	-0.204***	-0.125*	-0.061	-0.001
N=2,901,261	(0.040)	(0.077)	(0.042)	(0.053)
Leave out expanded recessions	-0.148***	-0.140*	-0.030	-0.076
N=1,924,468	(0.050)	(0.082)	(0.075)	(0.064)
Division-period FE		Y		Y
State-specific linear trends			Y	Y

Notes: The table reports minimum wage elasticities for average teen wage and employment, using individual-level Current Population Survey 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. The reported elasticities are calculated by dividing the coefficients on log minimum wage (and standard errors) 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 controls used in Table 1 and described in Section 2, and state and period fixed effects. Specifications additionally include Census division-period effects and state-specific linear trends, as indicated in the table. Samples are either the full sample, the sample excluding recessionary quarters, or the sample excluding expanded recessionary quarters. Recessionary quarters include 1980q1-1980q3, 1981q3-1982q4, 1990q3-1991q1, 2001q1-2001q4, and 2007q4-2009q2. Expanded recessions are defined to further include quarters until national employment levels reach pre-recessionary peaks: 1980q1-1980q4, 1981q3-1983q3, 1990q3-1992q4, 2001q1-2004q4, and 2007q4-2014q1. Regressions are weighted by sample weights, and robust standard errors (in parentheses) are clustered at the state level and significance levels are indicated by
*** 1%, ** 5%, * 10%.

Appendix Table C1. Dynamic minimum wage elasticities for teen employment, NSW Sample: state-quarter aggregated CPS data 1990-2011q1

		(1)	(2)	(3)	(4)
Panel A: 4-quarter averages of cumulative response elasticities					
	$\bar{\rho}_{[-12,-9]}$	0.015	-0.027	0.081	0.058
		(0.059)	(0.072)	(0.055)	(0.062)
B.	$\bar{\rho}_{[-8,-5]}$	-0.126	-0.205*	-0.038	-0.073
		(0.080)	(0.113)	(0.073)	(0.104)
C.	$\bar{\rho}_{[-4,-1]}$	-0.118	-0.143	0.005	0.056
		(0.085)	(0.152)	(0.083)	(0.133)
D.	$\bar{\rho}_{[0,3]}$	-0.169**	-0.170	-0.006	0.090
		(0.074)	(0.184)	(0.101)	(0.145)
E.	$\bar{\rho}_{[4,7]}$	-0.338***	-0.350	-0.144	-0.024
		(0.066)	(0.216)	(0.110)	(0.173)
F.	$\bar{\rho}_{[8,11]}$	-0.166**	-0.177	0.019	0.176
		(0.081)	(0.226)	(0.108)	(0.168)
G.	ρ_{12+}	-0.192**	-0.159	0.144	0.323*
		(0.092)	(0.289)	(0.144)	(0.187)
Panel B: Medium run (3 year) elasticities					
F-A	$\bar{\rho}_{[8,11]} - \bar{\rho}_{[-12,-9]}$	-0.181***	-0.149	-0.062	0.118
		(0.056)	(0.172)	(0.087)	(0.126)
F-B	$\bar{\rho}_{[8,11]} - \bar{\rho}_{[-8,-5]}$	-0.040	0.028	0.057	0.249***
		(0.062)	(0.137)	(0.093)	(0.089)
F-C	$\bar{\rho}_{[8,11]} - \bar{\rho}_{[-4,-1]}$	-0.047	-0.033	0.014	0.120
		(0.068)	(0.113)	(0.088)	(0.077)
Panel C: Longer run (4+ years) elasticities					
G-A	$\rho_{12+} - \bar{\rho}_{[-12,-9]}$	-0.207***	-0.132	0.063	0.265*
		(0.073)	(0.239)	(0.125)	(0.151)
G-B	$\rho_{12+} - \bar{\rho}_{[-8,-5]}$	-0.066	0.045	0.181	0.395***
		(0.089)	(0.208)	(0.138)	(0.119)
G-C	$\rho_{12+} - \bar{\rho}_{[-4,-1]}$	-0.074	-0.016	0.139	0.267**
		(0.074)	(0.181)	(0.124)	(0.110)
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 state-quarter aggregated CPS basic monthly data from1990-2011q1, taken from NSW (2014b) replication package. 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 an elasticity 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 Section 2.3. Panel B reports the cumulative effect in year 3, 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, 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, robust standard errors are clustered at the state level and significance levels are indicated by *** 1%, ** 5%, * 10%.

Appendix Table E1. Summary of literature: minimum wage elasticities for restaurant employment

	Sample years	Two-way FE (1)	Additional controls										
			(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Addison, Blackburn, and Cotti (2014)[c]	1990-2005	-0.101**	-0.006	-0.051***	-0.041	-0.062*	-0.046						
		(0.039)	(0.033)	(0.014)	(0.027)	(0.033)	(0.033)						
	1990-2012	(0.000)	-0.040*	-0.024	-0.035*	-0.023*	-0.010						
		(0.035)	(0.021)	(0.018)	(0.019)	(0.014)	(0.014)						
Dube, Lester, and Reich (2010)[a]	1990-2006	-0.176*						0.039	0.016				
		(0.096)						(0.050)	(0.098)				
Dube, Lester, and Reich (2016)[b]	2000-2011	-0.073*							-0.022				
		(0.042)							(0.091)				
Neumark, Salas, and Wascher (2014)[a]	1990-2006	-0.120***								-0.063***			
		(0.042)								(0.022)			
Totty (2015)[a]	1990-2010	-0.138*									-0.013	-0.013	-0.042
		[-0.297, 0.019]									[-0.042, 0.026]	[-0.046, 0.028]	[-0.085, 0.015]
Linear state trends			Y	Y	Y	Y	Y	Y					
Quadratic state trends				Y	Y	Y	Y						
Cubic state trends					Y	Y	Y						
Quartic state trends						Y	Y						
Cubic state trends							Y						
Census division-period FE								Y					
Contiguous county pair-period FE									Y				
NSW matching estimator										Y			
CCE-P											Y		
CCE-MG												Y	
IFE													Y

Sources: Addison, Blackburn, and Cotti (2014) Table 1; Allegretto, Dube, and Reich (2013) Table 3; Dube, Lester, and Reich (2010) Table 2; Dube, Lester, and Reich (2016) Table 3; Neumark, Salas, and Wascher (2014a) Table 8; Totty (2015) Table 3. *Notes:* Significance levels are indicated by *** 1%, ** 5%, * 10%. Standard errors in parentheses, except for Totty (2015) which report 95 percent confidence interval based on wild cluster bootstrap-t clustered at state level. Restaurant employment is defined as total employment in NAICS 7221 and NAICS 7222 using QCEW data[a]; employment in NAICS 722 using QWI data[b]; or employment in NAICS 722 using QCEW data[c]. The column labeled "Sample years" shows the years included in the data. CCE-P is common correlated effect, pooled; CCE-MG is common correlated effect, mean group. IFE is interactive fixed effect.