Online Appendices

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

Appendix A. Double-Selection Post-LASSO Estimates

We adapted the STATA code for the post-LASSO regressions from Christian Hansen's webpage: http://faculty.chicagobooth.edu/christian.hansen/research/JEPStata.zip, including the lassoshooting.ado file that 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 an 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 article, we use the default level for the penalization parameter λ in the Belloni et al. (2014) 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 A.1 reports the double-selection post-LASSO minimum wage elasticities for teen employment using the 1979 to 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 A.1 shows the point estimates and the confidence intervals associated with varying 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 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 five state-specific linear trends and lowers the elasticity in magnitude to –0.040. In other words, merely adding state-specific linear trends for these five 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 five 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 A.1, this range of λ implies a range of double-selection post-LASSO estimates for the minimum wage elasticity between –0.018 and 0.059.

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

λ	*# of state-specific trends chosen*	*# of divisions selected*	*# division-period FE chosen*	*Estimate*	*Std. error*
0	—	—	—	0.059	0.057
16	50	8	1,140	0.059	0.057
50	48	8	1,110	0.057	0.057
84	47	8	1,092	0.059	0.057
118	48	8	1,057	0.053	0.059
152	48	8	1,012	0.039	0.059
186	47	8	941	0.006	0.065
220	47	8	857	0.006	0.062
254	46	8	779	–0.001	0.061
288	45	8	648	0.001	0.059
322	43	8	546	–0.002	0.055
356	43	8	460	–0.003	0.058
390	44	8	388	0.000	0.053
424	44	8	326	0.003	0.052
458	43	8	267	–0.012	0.049
492	44	8	209	–0.002	0.060
526	43	8	160	0.010	0.063
560	42	7	122	–0.022	0.060
594	41	7	96	–0.020	0.063
628	41	7	84	–0.021	0.063
662	40	6	72	–0.019	0.062
696	40	5	66	–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	23	0.015	0.082
968	30	1	22	0.015	0.082
1,008	29	1	20	0.012	0.083
1,048	25	1	19	0.007	0.083
1,088	24	1	17	0.004	0.084
1,128	20	1	15	0.014	0.083
1,168	18	1	14	–0.018	0.088
1,208	18	1	13	–0.018	0.088
1,248	17	1	13	–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.

Significance levels are indicated by *** 1%; ** 5%; and * 10%.

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 claimed that the presence of recessions in the beginning and end points affected the estimation of the state-specific trends and biased the minimum wage elasticities for teen employment toward zero when such trends were included.[1] They also claimed that inclusion of higher-order (cubic or greater) state-specific trends in that sample restored the more negative estimates by correcting for this endpoint bias.

Online Appendix Figure B.1 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 the sample forward also reduces the magnitude of the estimate. By contrast, the most-saturated specification is quite robust with respect to the choice of the sample period. Although stability of estimates across samples need not indicate accuracy, these results are consistent

[1]Specifically, NSW (2014a: 616) stated: "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."

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.[2]

[See Figure B.1 at end of this appendix]

Online Appendix Figure B.1 is also informative about endpoint bias. The two figures at the bottom of Figure B.1 do not provide any indication that the 1990 to 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 to 2009 sample than in the other samples. The estimates with the 1990 to 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 to 2014 sample.

As another assessment of the role of business cycles in affecting estimation of trends, online Appendix Table B.1 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 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 (columns (3) and (4)).

[2]We wrote in ADR (2011: 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."

Appendix Table B.1. 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 the text, 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.

Significance levels are indicated by *** 1%; ** 5%; and * 10%.

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 the 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 B.2 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.)

[See Figure B.2 at the end of this appendix]

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. When we extend the sample by including earlier start dates, however, 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 toward zero, even when higher-order trends are included. Recall, however, that the results from the 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 to 2011 sample reproduced more negative employment effects of minimum wages on teen employment.

Appendix Figure B.1. Minimum Wage Elasticities for Teen Employment, by Panel Start and End Year

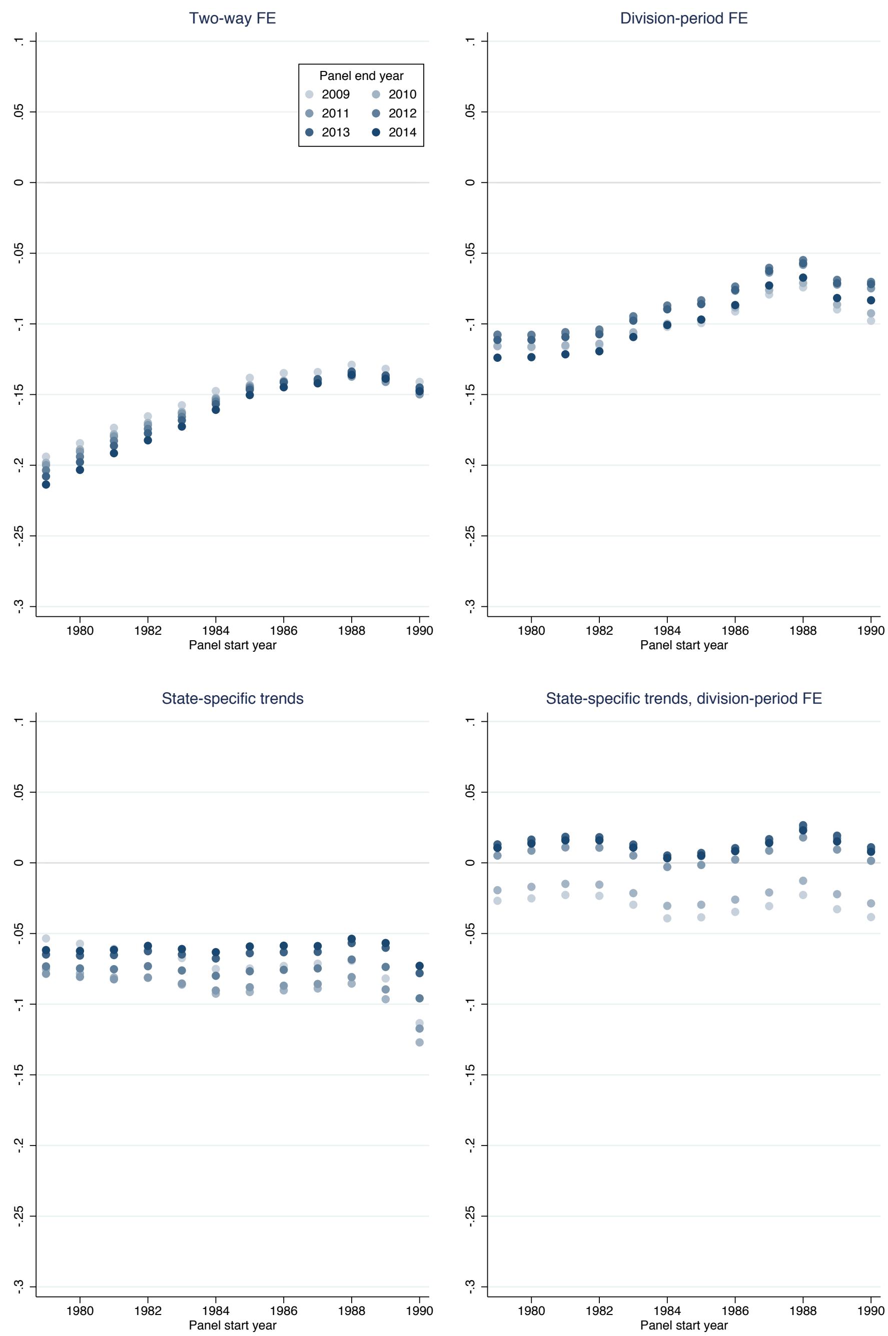

Notes: The figure reports teen employment elasticities with respect to the minimum wage for the four regression models using individual-level data described in the text. From these linear probability models, we calculate elasticities by dividing the log minimum wage coefficient by the sample mean employment.

Appendix Figure B.2. Minimum Wage Elasticities for Teen Employment, by Type of State-Specific Polynomial Trend, and Panel Start and End Year

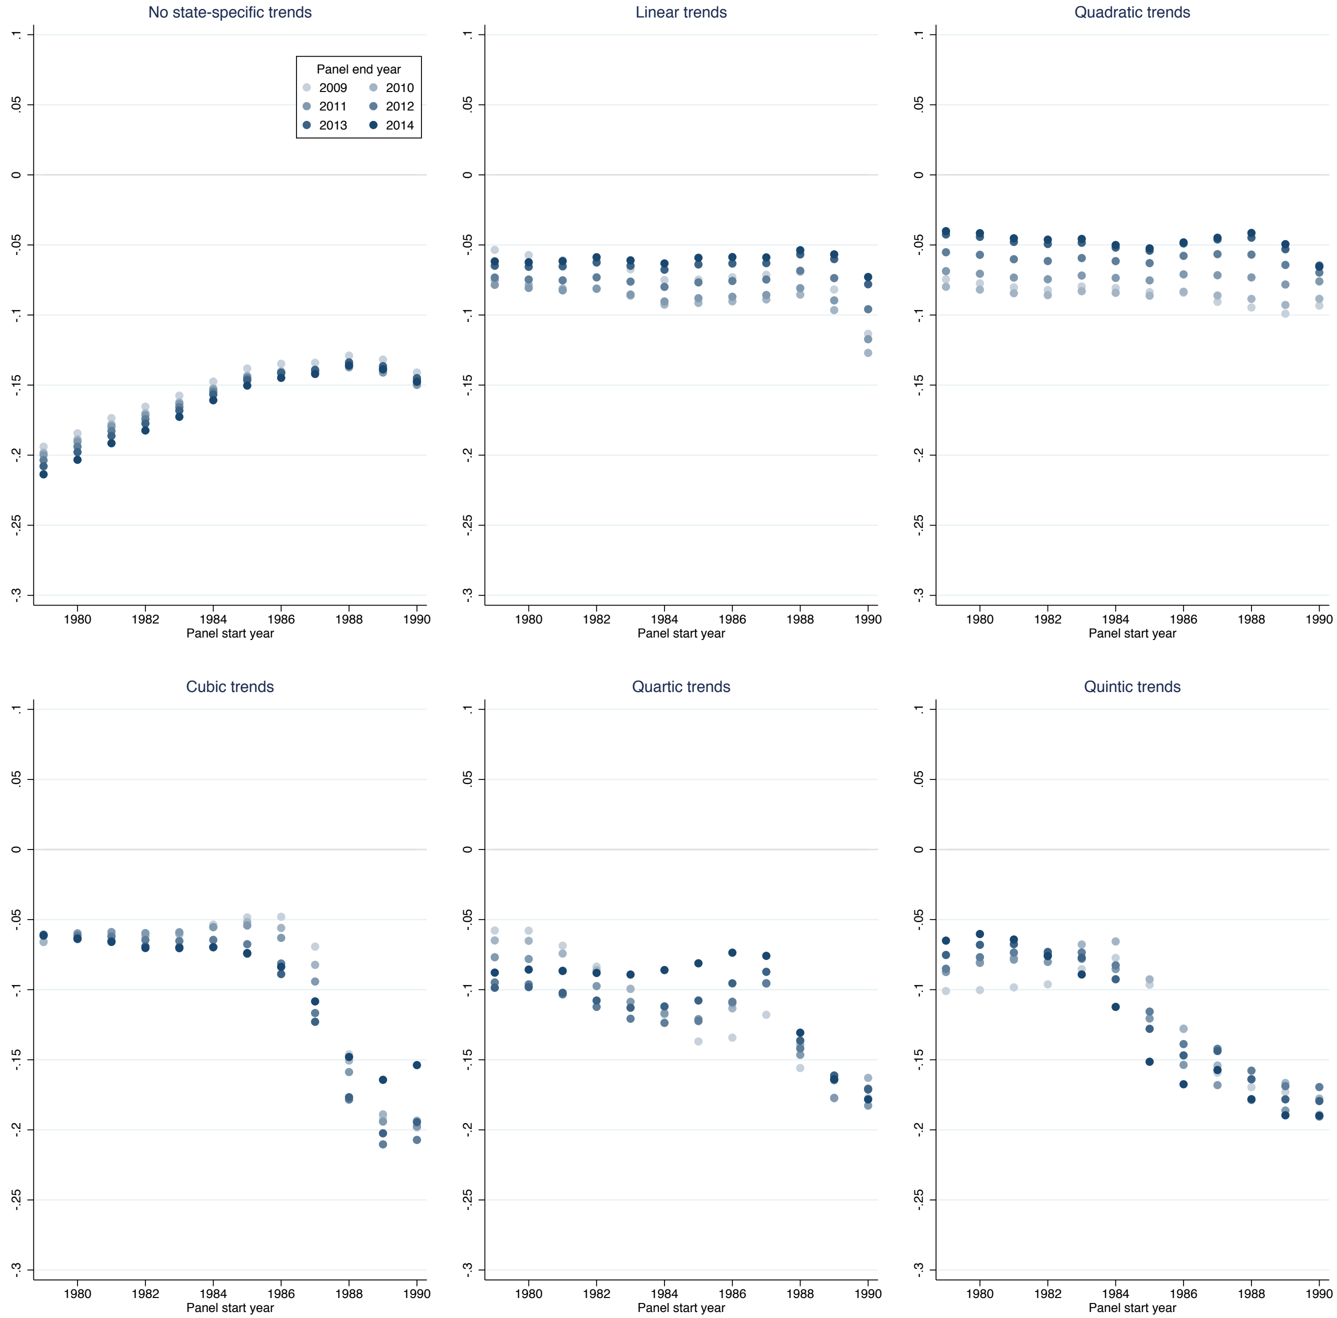

Notes: The figure reports teen employment elasticities with respect to the minimum wage for the two-way fixed-effects model, with state-specific trends of polynomials degree zero through five, using the individual-level data described in the text. From these linear probability models, we calculate elasticities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

Appendix 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 to 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 to 2011 state-aggregated data used in NSW (2014b, hereafter "NSW data"). We also report with three- and four-plus-year employment elasticities netting out leading coefficients using the NSW data, analogous to the estimates provided in Table C.1 using the full sample. We are therefore able to reconcile our two sets of results: the discrepancy is not caused by the samples used but rather by 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 C.1 plots the cumulative teen employment elasticities from the 1979 to 2014 individual-level sample using the two-way fixed-effects model, along with 95% 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.

[See Figure C.1 at the end of this appendix]

The left-most panel ("Quarterly: two-way FE") in Figure C.1 plots these cumulative responses from the two-way fixed-effects model, along with 95% 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% level.

To reduce noise and more easily extract a signal from the data, the panel entitled "Annualized: two-way FE" in Figure C.1 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 the 95% confidence bands. (These averaged cumulative response elasticities and standard errors are also reported in the first column of Appendix Table C.1) We find unmistakable evidence that the two-way fixed-effects model fails the falsification test that leading coefficients during one, two, or three years prior to treatment are zero. And since the leading effects are occurring two or three years prior to treatment, they cannot plausibly result from anticipation of the policy.

Cumulative Response of Teen Employment to a Minimum Wage Increase in the NSW Sample

The results in online Appendix Figure C.1 as well as Table 3 differ from those in NSW 2014(b), who denied evidence of pre-existing trends in the two-way fixed-effects model. They also argued 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).[3] Online Appendix Figure C.2 shows the cumulative teen employment elasticities using the shorter NSW sample and specification and reproduces their estimates (see their figure 6).

[See Figure C.2 at the end of this appendix]

[3]We use the replication data on Ian Salas's 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 included controls for unemployment rate, state, and period fixed effects.

When NSW (2014b) analyzed estimates from this regression and compared cumulative responses using ρ_{-2} as their baseline, they concluded 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 C.2 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 three-year or four-plus-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 C.1, 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.[4]

[4] Because quarterly leads and lags can be noisy, ADR and DLR (2010) included leads and lags at four-quarter or two-quarter frequency for the purpose of smoothing. However, we do acknowledge that some arbitrariness occurs in any choice of smoothing. In this article, for the purpose of comparability with NSW (2014b), we have used their quarterly lead/lag structure in the regression and simply taken 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.

Appendix Table C.1. Dynamic Minimum Wage Elasticities for Teen Employment, NSW Sample: State-Quarter Aggregated CPS Data, 1990–2011q1

		(1)	*(2)*	*(3)*	*(4)*
Panel A: Four-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 (three-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 (four-plus-year) 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 the text. Panel B reports the cumulative effect in year three, after subtracting alternative baseline levels at one, two, or three 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, and robust standard errors are clustered at the state level.

Significance levels are indicated by *** 1%; ** 5%; and * 10%.

Column (1) of Table C.1 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, –0.118 and –0.126, respectively, although they are not statistically significant at conventional levels. As shown in Table C.1 (column (1)), however, the estimated three-year and four-plus-year effects range between –0.040 and –0.074 when using baselines that are one or two years 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).[5] In other words, the estimates from the two-way fixed-effects model in the sample used in NSW (2014b) showed that between 50 and 75% 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 C.2 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 C.1 reports the three and four-plus-year elasticities using the NSW (2014b) data and specifications. NSW (2014b) argued 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 to 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 to 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. By contrast, models with controls for state-specific trends tend to have

[5] These estimates are from rows labeled F-C, F-B, G-C, and G-B.

smaller leading coefficients. And in all cases, after netting out the leading coefficient the employment estimates are substantially smaller. None of the three or four-plus-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 1990 to 2011 NSW sample, as it is for the full 1979 to 2014 sample used in this article.

Appendix Figure C.1. 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 four-quarter averaged effects and confidence intervals in green. Standard errors are clustered by state.

Appendix Figure C.2. 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 four-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 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).[6] 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). Appendix Figure D.1 non-parametrically plots (using lowess) the mean relative donor weight as a function of the distance between the geographic centroids of the donor and the 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.)

[See Figure D.1 at the end of this appendix]

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 1,000 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

[6] In an earlier (2013) version of this article, 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: 26) then criticized our use of placebo and not actual minimum wage increases, stating that this approach strikes them 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.

proximity when constructing controls. This variation provides information 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.

Appendix Figure D.1. Donor Distance and Relative Weights

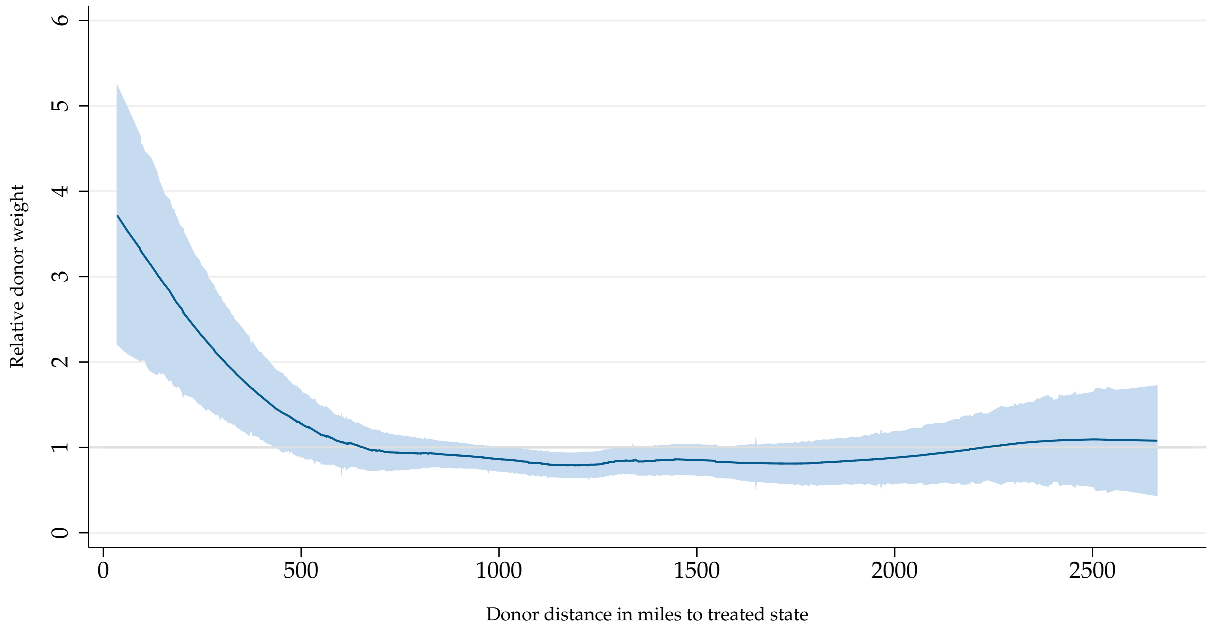

Notes: The figure shows the locally weighted regression (lowess) of the relative donor weights on donor distance to treated states for the 25 treatment events, excluding Alaska and Hawaii, as described in the text. The 95% confidence interval is calculated from 1,000 cluster bootstrap iterations at the treatment-event level.

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 (2010, 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 E.1 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-varying heterogeneity range from –0.063 to 0.039.[7]

---

[7] Aaronson, French, and Sorkin (2015), who studied restaurant employment using a border discontinuity design for a small number of states, obtained an overall short-run elasticity of –0.1. They described this estimate as "very imprecise" (they do not report a standard error). They also found increased entry and exit, which they interpreted using a calibrated putty-clay model that suggested large dis-employment effects in the longer run. However, our empirical findings here and in DLR (2010) do not suggest sizable employment losses in the restaurant sector in the "medium run," that is, after 12 or 16 quarters following the minimum wage.

Appendix Table E.1. 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: Standard errors in parentheses, except for Totty (2015), which report 95% 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, common correlated effect, pooled; CCE-MG, common correlated effect, mean group. IFE, interactive fixed effect.

Appendix F. Spatially Correlated Placebos

Appendix B of DLR (2010, 2016) 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 one's own employment—even when variation in the true minimum wage could not possibly affect employment across different counties in the sample. DLR ensured 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.[8] NSW (2014a, 2014b) argued 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

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

wage below.[9] 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, $MW_{jt}^{S'}$.

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

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

where E_{it}^{S} is the 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 as cross-border neighbors' minimum wage, $MW_{jt}^{S'}$ along with time dummies (D_t) and county dummies, D_i using the sample S:

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

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 (2) is equivalent to estimating:

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

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.

---

[9] 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 article. However, we index the variables $MW_{it}^{S}, MW_{jt}^{S'}$ instead of $MW_{st}^{S}, MW_{st}^{S'}$ 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.[10] This approach 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) criticized this falsification test. Of most importance, they argued that that the S-county sample is "contaminated." Their argument has two parts. First, they (correctly) pointed out that that even though the minimum wages in S-counties were equal to the federal minimum$MW_{it}^S = MW_t^F$, the minimum wages were changing over time, which can have a real effect. Second, they (incorrectly) argued that because the cross-border minimum wage

[10] 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.

$MW_{jt}^{S\prime}$ can 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 wrote: "But $MW_{st}^{S\prime}$ in equation (5) [similar to our equation (3) 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 (2).

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 attributable to 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) insisted upon rejecting this argument and claimed 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 provided a decomposition of $MW_{it}^{S'}$ that aimed to demonstrate that δ could still reflect true causal effects. In actuality, their decomposition provides an illustration of exactly the opposite argument. In particular, they rewrote Equation (3) as:

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

They then argued that the first term, $MW_t^F \cdot I\left\{MW_{jt}^{S'} = MW_t^F\right\}$, 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 wrote "[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 interpreted this to mean that δ can reflect the true effect of MW_t^F on E_{it}^S.

However, NSW missed the implication of this decomposition. Of course, the interaction term $MW_t^F \cdot I\left\{MW_{jt}^{S'} = MW_t^F\right\}$ can be correlated with $\mathrm{E}_{\mathrm{it}}^{\mathrm{S}}$ even after controlling for the time effects; after all, it is part of $MW_{jt}^{S'}$, 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 (1). 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 $M_{jt}^{S'}$ 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 $MW_t^F \cdot I\left\{MW_{jt}^{S'} = MW_t^F\right\}$? 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 (1). 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.[11]

[11] In table 5 of NSW (2014b), the authors artificially changed the federal minimum wage, and showed 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 shed 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 showed 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 proceeded 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 excluded all federal minimum wage increases from the sample; this restriction did not attain their objective since the real minimum wage was changing due to inflation. At any rate, once they eliminated over half of data using these arbitrary criteria, they found that the placebo estimate became 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 were erroneous.

NSW (2014b) made a second point regarding this placebo test. They argued that correcting a small data error in DLR changed $\hat{\beta}$ from –0.208 to –0.114, while $\hat{\delta}$ is largely unaffected (changing from –0.123 to –0.125).[12] 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.

[12]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 noted that this coding error does not influence any of the analysis in the article other than the placebo exercise.

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. Accessed at https://www.chicagofed.org/digital_assets/publications/working_papers/2013/wp2013_26.pdf.

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

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.

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.

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

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* 34(3): 663–704. Accessed also at https://arindube.files.wordpress.com/2014/11/dlr2r3_fullpaper.pdf.

Hansen, Christian. 2013. Econometric methods for high-dimensional data. NBER Summer Institute. Cambridge, MA: National Bureau of Economic Research. Accessed at http://www.nber.org/econometrics_minicourse_2013/.

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

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

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