REPLY TO "CREDIBLE RESEARCH DESIGNS FOR MINIMUM WAGE STUDIES"

DAVID NEUMARK AND WILLIAM WASCHER*

The authors make three points in this reply to the article by Allegretto, Dube, Reich, and Zipperer (ADRZ 2017). First, ADRZ shed no new light on the sensitivity of estimated minimum wage employment effects to the treatment of trends in state-level panel data, and they make some arguments in this context that are misleading or simply wrong. Second, the key issue ADRZ emphasize—using "close controls" to account for shocks that are correlated with minimum wage changes—does not generate large differences in findings, and ADRZ do not address evidence from Neumark, Salas, and Wascher (NSW 2014a) that questions the validity of the close controls used in Allegretto, Dube, and Reich's (ADR 2011) and Dube, Lester, and Reich's (DLR 2010) work. Third, ADRZ ignore or dismiss a growing number of studies that address in various ways the same issue of potential correlations between minimum wages and shocks to low-skill labor markets that ADRZ argue generate spurious evidence of disemployment effects, yet often find rather large negative effects of minimum wages on low-skilled employment.

[G]iven the media interest in reporting novel or titillating empirical findings and politicians' desires to robe their proposals in scientific empirical cloth, however novel or inconsistent with prior research, it is crucial that as a profession we ensure that replication, or at least fear of replication, is our norm.

—Hamermesh, 2007: 725–26

The *ILR Review* (ILRR) has been at the forefront of debate on the employment effects of minimum wages in the United States, with a commitment to replication and refinement of the evidence. In addition to

publishing the original 1992 symposium credited with launching the "new minimum wage research," which focused on evidence from state-level variation in minimum wages, the ILRR published a comment and a reply on one of the papers from that symposium (Card, Katz, and Krueger 1994; Neumark and Wascher 1994), a review symposium (Brown et al. 1995) on Card and Krueger (1995), and our more recent paper (Neumark, Salas, and Wascher (NSW) 2014a) that responded to two influential papers challenging the evidence that minimum wages reduce low-skilled employment (Allegretto, Dube, and Reich (ADR) 2011; Dube, Lester, and Reich (DLR) 2010). In 2017, ILRR is publishing a response to NSW (2014a) and a related paper (NSW 2014b), along with our reply in this current article. A reader of the minimum wage literature could perhaps justifiably decide that it is too hard to wade through these articles to judge whether minimum wage increases reduce low-skilled employment; however, the exchanges promoted by the ILRR can help readers understand and evaluate the evidence more readily than reading the stand-alone articles in which researchers do not respond directly to one another. We thus agree wholeheartedly with Daniel Hamermesh: The willingness of a journal to devote considerable space to hashing out an important and complicated debate is commendable, especially regarding important policy debates in which the empirical evidence is frequently cited.

For this current reply, we were directed to provide a brief response to the most important issues raised by Allegretto, Dube, Reich, and Zipperer (ADRZ 2017), rather than a point-by-point refutation on all issues we disagree with concerning their comments. Thus, we focus on what we think are the central empirical issues on which ADRZ reach the wrong conclusions.[1]

We address three issues in particular. First, in our view the analysis in ADRZ's comment sheds no new light on the sensitivity of estimated employment effects of minimum wages to the treatment of trends in state-level panel data, an issue that was covered at length in our earlier exchanges with the authors. In this regard, some arguments that ADRZ make in defense of their work and in challenging ours are either misleading or simply wrong. Second, the key issue ADRZ emphasize—the need to use "close controls" to account for shocks that are correlated with minimum wage changes—does not generate large differences in findings. Moreover, their comment

---

[1]Our previous work primarily used Current Population Survey (CPS) data collected by the U.S. Census Bureau for the Bureau of Labor Statistics. Because we used these data to estimate the effects of minimum wage increases on employment of teens from state-level data, we focus mainly on the CPS state-level estimates in this reply. In that context, ADRZ's statement that "a substantive disagreement remains mainly for teens" (p. 561) is misleading because our past work did not focus on the effects of restaurant employment studied by DLR. We noted in NSW (2014a), however, that one might expect to see a smaller effect for restaurant employment because many restaurant workers earn more than the minimum wage so that labor–labor substitution could mask disemployment effects for the least skilled.

does not address the main evidence in our earlier work questioning the validity of the close controls used in ADR and DLR. Third, and perhaps most important, ADRZ either ignore or casually dismiss a growing number of studies that address in a variety of ways the same issue of potential correlations between minimum wages and shocks to low-skill labor markets that ADRZ argue generate spurious evidence of disemployment effects. As it turns out, most of these studies found negative effects of minimum wages on low-skilled employment that are often larger than those we have found.

ADRZ's Evidence

The main contention in ADR and DLR, which ADRZ defend, is that state minimum wages have tended to increase when labor markets for less-skilled workers were in relative decline, generating a spurious negative relationship between minimum wages and low-skilled employment, which they term *spatial heterogeneity*. Like ADR and DLR, ADRZ assert that restricting comparisons to what happens when the minimum wage increases in one state but not in another one close by solves this problem because these close controls experience the same labor market shocks, which can then be isolated from the true effects of minimum wages. ADRZ similarly defend the assertion in ADR that including state-specific trends nets out the potentially correlated shocks. They argue that combinations of these approaches yield disemployment effects of minimum wages near zero, consistent with their presumption that estimates that do not control for spatial heterogeneity are biased toward finding disemployment effects.[2]

Trends

Although the idea of using close controls to net out common shocks has some intuitive appeal, the inclusion of the state-specific linear time trends turns out to be the principal driver of the small minimum wage effects they report for teenagers, a point not emphasized by ADRZ. In their table 1, for example, the estimated minimum wage coefficient in the specification including division-period fixed effects but excluding trends is -0.124, within the range of our past estimates; adding in linear state-specific time trends then reduces the estimate to approximately zero. In addition, appendix figures C.1 and C.2 show that differences in the cumulative minimum wage effects across specifications are driven primarily by the linear trends. The cumulative minimum wage effects from the two-way fixed-effects model and from the specification that includes division-period fixed effects look

---

[2]One question they have not addressed is why the effects on wages are largely unaffected by this spatial heterogeneity. They emphasize the role of negative demand shocks to the labor market for teenagers but do not find any clear evidence that their controls for spatial heterogeneity increase the effect of minimum wages on teen wages (see ADRZ's table 1).

quite similar to each other, but both are noticeably different from the corresponding specifications with state-specific linear trends.[3]

NSW (2014a) already demonstrated that the estimated employment effects of minimum wage changes are sensitive to the inclusion of linear or nonlinear time trends in the sample period used by ADR. In their table 2, ADRZ report new evidence that including higher-order polynomial time trends also results in estimates of employment effects that are close to zero, *but only for a much longer sample period than the one we used.* Given that our earlier papers were responses to their results using post-1990 data rather than general claims about the sensitivity of the estimated effects of minimum wages to the inclusion of trends, we do not understand why it is appropriate to respond to our evidence by changing the sample period (in ADRZ) to begin in 1979.[4] Moreover, many fewer states had set their minimum wage above the federal minimum in the earlier years they added to the sample, and the longer sample period may increase the number of potential economic factors that affect teen employment but are omitted from the specification. As a result, it is debatable whether the long sample provides better identifying information to separate out the employment effects of minimum wages from trending unobservables.

Although state panel data estimates of minimum wage effects are sensitive to whether and how state-specific trends are included, we think that ADRZ (like ADR) are much too quick to view the evidence relying on linear trends as decisive. Including linear (or higher-order) state-specific trends in the model does not provide a clear sense of what the counterfactual is, and as a result it is hard to know whether the specification with trends is more convincing. Consistent with this observation, ADR described including trends as "a less valuable means of controlling for spatially correlated shocks" (2011: 237). As evidence of problems with the linear trend specification on which ADRZ rely, table 3 and appendix table C.1 show that this specification produces medium- and longer-run employment effects that are large, positive, and often statistically significant. Although ADRZ ignore this evidence, these large, positive, longer-run elasticities make us skeptical

---

[3]For example, using the 1979–2014 sample period, the correlation between the cumulative minimum wage effects from the two-way fixed-effects model and those in the specification including division-period effects is 0.84, whereas the correlation between the cumulative minimum wage effects in the latter specification and those in the specification adding in state-specific linear trends is 0.63. Indeed, ADRZ explicitly acknowledge (appendix C, p. 14) that their conclusions about negative bias in estimated minimum wage effects from prior shocks or pre-trends are driven by the specifications that include state-specific trends.

[4]In addition, their discussion of endpoint bias in appendix B misrepresents what we wrote in NSW (2014a). In particular, we did not say that all recessions at the beginning or end of the sample would produce "endpoint bias," but rather a recession might cause such bias depending on the cross-state deviations between teen and aggregate employment (because the minimum wage employment regressions include a state-level aggregate labor market control). Moreover, our suggestion was not to completely exclude recessions from the sample, as they do in table B.1, but to instead exclude recessions at the beginning or end of the sample when estimating the trend.

that these specifications provide valid shorter-run estimates.[5] Finally, given that ADRZ (like ADR) emphasize so strongly the role of close controls in capturing omitted shocks, we are puzzled as to why they think the specifications with trends provide more reliable identifying information on minimum wage effects, once division-period fixed effects are included.

In our view, given that the trends are intended to proxy for factors omitted from the specification, it would be preferable to include variables that capture these potentially omitted factors (e.g., Aaronson et al. 2007; Smith 2011)—perhaps even more preferable when extending the analysis to a period of more than three decades over which the economic content of the trends becomes even less interpretable. Absent that tactic, and as noted in NSW (2014b), a more flexible approach to controlling for trends in time-series data as well as in panel data—which imposes smoothness but not a specific functional form on the effects of the omitted factors—is to pre-filter the data using a Hodrick-Prescott (HP) filter. When we do this for ADRZ's longer sample, the estimated employment effect is –0.12 with a standard error of 0.08; using the shorter sample (1990–2014), the estimated employment effect is –0.17 with a standard error of 0.07.[6]

We see two additional problems with the more general analysis ADRZ present related to leading effects—which ADR labeled "pre-trends" and which ADR used to motivate including linear trends and to respond to our criticisms of their work more generally. First, in this analysis ADRZ again attempt to rebut our criticisms by changing the sample period. Appendix table C.1 shows the dynamic minimum wage elasticities from the sample we used in NSW (2014a), in contrast to the longer sample period in table 3. Most notably, in column (1) of appendix table C.1, the estimated employment effects prior to the increase in the minimum wage are not statistically significant, whereas the estimated employment effects following minimum wage increases are negative and statistically significant.

Second, even when using their sample period, ADRZ substantially overstate the potential bias in the two-way fixed-effects estimate due to negative leads when they suggest that the leading effects are "roughly two-thirds the size of the static employment elasticity" (pp. 574–75). In particular, in panel B of ADRZ's table 3, the medium-run elasticities based on cumulative differences that net out the negative leads remain negative and statistically significant in the two-way fixed-effects model, with elasticities ranging from –0.12

---

[5]Meer and West (2016) have emphasized the difficulty of estimating minimum wage effects when linear trends are included, especially when the effects on job growth are considered. ADRZ criticize the Meer and West analysis, although we leave the specific issues to these authors to resolve. Nonetheless, the main evidence ADRZ present (table 4) is puzzling. They report that in contrast to within-group estimates, first-difference estimates of minimum wage effects are positive in the weighted data, whereas the unweighted estimates are centered on zero. For reasons that are not clear, however, ADRZ switch to using annual data in this section of their article. Estimating the models using quarterly data, as in the other analyses, yields first-difference estimates that are negative and very similar to the within-group estimates, whether weighted or not, unless we allow for long lags.

[6]For details on how we implemented the HP filter, see Neumark, Salas, and Wascher (2013).

to –0.18. A similar calculation that subtracts the average leading effects (rows A–C of panel A, column (1)) from the average contemporaneous and lagged effects (rows D–G) yields an elasticity of –0.14. These estimates are much closer to the –0.21 elasticity estimated from the two-way fixed-effects model (table 1, panel B) than ADRZ suggest.[7]

Close Controls

In contrast to our continued skepticism about what we learn from including state-specific trends, we view the close controls approach as having some intuitive a priori appeal, as well as providing a more easily interpretable counterfactual. Still, the close controls approach need not be correct.[8] Our fundamental criticism of ADR and DLR was that neither article provided evidence that the closer controls (states in the same census division in ADR, and counties across state borders in DLR) were more valid controls for the treated states than a larger set of states (or counties), many of which are more distant.[9] Moreover, we presented evidence to the contrary, using the first stage of a synthetic control analysis to show that states inside the census division did not provide better controls for treated states (and similarly for counties). Perhaps the most succinct summary of this was figure 1 from NSW (2014b), which showed only a modest increase in the synthetic control weight for states closer to the treated state. (Curiously, this figure looks qualitatively very similar to ADRZ's appendix figure D.1, which ADRZ use to argue that within-census division donors are the appropriate controls.) Furthermore, ADRZ ignore the analysis from NSW (2014b), which pointed out that it is not valid to simply look at weight per donor state, given that

---

[7]Although ADRZ criticize our "peculiar" use of a quarter t–2 baseline for calculating minimum wage effects in NSW (2014b), we chose that quarter because quarter t–1 might pick up an anticipation effect. Moreover, figure 4 in that article showed that regardless of which leading quarter is used, a noticeable downshift occurs in teen employment four to eight quarters after the minimum wage increase, reaching a maximum estimate of about –0.2 at five quarters even relative to the quarters with the strongest negative leads.

[8]In support of the close controls approach, ADRZ argue that "when economic shocks on average vary continuously at 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" (pp. 562–63). However, other state policies also change discontinuously at the border, so even shocks that are "similar" at the border can, by way of interactions with policy, have effects that are discontinuous at the border. Moreover, economic conditions for teenagers relative to those for adults may be quite different across states within the same census division. For example, in New York the unemployment rate for teenagers in 2015 was more than four times the overall unemployment rate in that state, whereas in Pennsylvania that ratio was only two.

[9]In NSW (2014a), we also estimated a model that included bordering state pair-period interactions rather than division-period interactions, given that states in the same census division often do not share borders and can be far apart. The estimated employment effects from this specification were negative and very close to the estimates from the standard panel data estimator. Given that ADRZ embrace the border-county research design in DLR, we are surprised that in the CPS analysis they emphasize comparisons between states in the same division, rather than states that share borders.

nearby donor states are always fewer in number than donor states farther away.[10]

Which Identifying Minimum Variation to Use or Discard

The main thrust of ADRZ's response to our criticism of using same-division states as controls (in ADR) is their use of LASSO methods to ask whether the model should include division-period fixed effects. In this regard, ADRZ's comment mischaracterizes our argument. They suggest that our argument is simply that, in using only the within-division variation, ADR were "throwing away too much information" (p. 561, attributing the quote to us). In fact, we argued that across-division variation provides potentially more valid identifying information for estimating the causal effects of minimum wages than the within-division variation to which they restrict attention.[11]

This is not just a semantic point. Rather, this mischaracterization is reflected in ADRZ's use of LASSO to choose which division-period fixed effects (and state-specific trends) to include in the model. This approach reflects a fundamental misunderstanding of what LASSO can deliver. LASSO can pick controls to best predict a dependent variable, subject to a penalty for parameters to avoid over-fitting. LASSO does not tell us, however, whether the inclusion of additional variables better identifies the causal effect of interest, or, equivalently, which controls lead to a less-biased causal estimate.[12] We are particularly surprised that ADRZ would use LASSO with regard to identification, since the thrust of their earlier articles was always an a priori assertion that their approach provides more credible research designs for identification of causal parameters.

This point about valid identification versus additional controls parallels exactly the analogy used in NSW (2014b) regarding Griliches's (1979) work on estimating the returns to schooling using twin or sibling data. We know that family fixed effects increase the fit of the model and could therefore be chosen by LASSO. But that does not mean the within-family estimates are less biased as causal estimates of the returns to schooling.[13] If model-fitting techniques could resolve questions of identification, the task of labor economists would be much easier than it actually is!

---

[10]Note that their county-level analysis with the Quality Workforce Indicators data does not address the question of which counties are better controls because they examine only similarities across counties in aggregate variables. What matters in identifying valid control areas is the similarity in shocks to teen or low-skill employment conditional on the aggregate variables used as control variables.

[11]In NSW (2014a), we said they "have thrown out so much useful and potentially valid identifying information that their estimates are uninformative or invalid" (p. 644). And in NSW (2014b), we said DLR and ADR were "arbitrarily throwing away lots of *valid identifying information* and potentially *focusing on variation that generated biased estimates*" (p. 20, italics added). ADRZ do get the quote right once, on p. 566, while getting it wrong multiple other times.

[12]In contrast to LASSO, synthetic control methods (discussed below) provide a more compelling data-driven strategy for selecting controls, because these methods match on pre-treatment observations rather than on post-treatment observations of the outcome.

[13]We used this analogy to illustrate why within-division estimates are not necessarily less biased; it depends on what drives the minimum wage variation within versus across divisions (NSW 2014b).

Extending the Synthetic Control Analysis

The main thrust of our analysis was using the first stage of synthetic control methods to assess whether geographically close controls are better than potential controls that are farther away. We also experimented with an estimator that tried to use the weights this first stage had selected for potential control states. This approach led to teen employment elasticities of approximately –0.15. ADRZ spend a good deal of time criticizing the econometrics underlying this estimator. We do not respond to these criticisms directly because our primary focus clearly was on the first-stage analysis, and we acknowledged the problematic features of the estimator we were trying to use, noting that "we do not claim to have completely resolved the econometric issues involved, but we do think our approach has some intuitive, heuristic appeal" (NSW 2014a: 637).[14] Moreover, the synthetic control estimator in Dube and Zipperer (DZ 2015) also has shortcomings, most notably the restriction of the treatment sample to only 29 minimum wage increases. An apparently superior synthetic control approach to the problem of estimating minimum wage effects on employment is discussed below (Powell 2016).

The DLR Placebo Test

Although our main focus is on the CPS analysis of teen employment using state-level panel data, we must also address the placebo test in DLR, which ADRZ persist in defending. This issue is critical, given that ADRZ argue that "the county-level border discontinuity design provides one of the most compelling identification strategies for estimating minimum wage effects" (p. 562)—a conclusion that is based on this placebo test and that seems to provide their rationale for preferring the close control design more generally. We show why the placebo test they claim justifies their approach is contaminated, without appealing to spillovers.

The specification is:

$$(1) \qquad E_{st} = \gamma MW_{st}^{S} + \delta MW_{st}^{S'} + D_s\theta + D_t\lambda + \varepsilon_{st},$$

where MW_{st}^{S} is the state minimum wage and $MW_{st}^{S'}$ is the bordering state's minimum wage; MW_{t}^{F} is the federal minimum wage. ADRZ's placebo test includes only border counties for which $MW_{st}^{S} = MW_{t}^{F}$. Thus, Equation (1) becomes:

$$(2) \qquad E_{st} = \gamma MW_{t}^{F} + \delta MW_{st}^{S'} + D_s\theta + D_t\lambda + \varepsilon_{st}.$$

MW_{t}^{F} is perfectly collinear with D_t, so estimating Equation (2) is the same as estimating:

---

[14]Similarly, they write that "we should be cautious about the statistical significance of findings from the NSW matched estimator" (footnote 25), but we noted explicitly that we did not view our standard errors as correct (NSW 2014a: 640).

$$(3) \qquad \mathrm{E}_{st} = \delta \mathrm{MW}_{st}^{\mathrm{S}'} + \mathrm{D}_s \theta + \mathrm{D}_t \lambda + \varepsilon_{st}.$$

In NSW (2014b), we noted that Equation (2) is also equivalent to:

$$(4) \qquad \mathrm{E}_{st} = \gamma \mathrm{MW}_t^{\mathrm{F}} + \delta \left(\mathrm{MW}_t^{\mathrm{F}} \cdot \mathrm{I}\left\{ \mathrm{MW}_{st}^{\mathrm{S}'} = \mathrm{MW}_t^{\mathrm{F}} \right\} \right.$$
$$\left. + \mathrm{MW}_{st}^{\mathrm{S}} \cdot \mathrm{I}\{ \mathrm{MW}_{st}^{\mathrm{S}'} > \mathrm{MW}_t^{\mathrm{F}} \} \right) + \mathrm{D}_s \theta + \mathrm{D}_t \lambda + \varepsilon_{st},$$

in which case there is a built-in correlation between $\mathrm{MW}_t^{\mathrm{F}}$ and $\mathrm{MW}_{st}^{\mathrm{S}'}$ that D_t does not subsume. ($\mathrm{I}\{\cdot\}$ denotes the indicator function.) As a result, variation in the federal minimum wage affects the estimate of δ. ADRZ acknowledge this, but then argue that there is no "*plausible causal* interpretation of that correlation" (online appendix, p. 27).

A simple counter-example, however, illustrates that they, and their placebo test, are wrong, because the effect of a federal minimum wage change need not be the same across states even when it is binding, say, for example, if the underlying distribution of wages varies across states.[15] Suppose we have four states with the data in our Table 1. State 1 borders State 2, and State 3 borders State 4. In States 1 and 3 the federal minimum wage always binds; suppose it is $8 in Year 1 and $10 in Year 2. States 2 and 4 have higher state minimum wages initially: $9 in State 2 and $10 in State 4. Because State 1 borders a state in Year 1 that has a lower state minimum wage ($9 in State 2 versus $10 in State 4), we also assume that State 1 is a lower-wage state than State 3, so that employment in State 1 falls by more when the federal minimum wage increases.[16]

In Table 1, the data used in the analysis are in boldface.[17] The table shows that the federal minimum wage increase induces a larger increase in $\mathrm{MW}^{\mathrm{S}'}$—the cross-border minimum wage—in State 1 ($1) than in State 3 (zero). In addition, the federal minimum wage increase in State 1 is associated with a larger decline in employment because State 1 has lower wages. The simple difference-in-differences regression for the change in employment associated with variation in $\mathrm{MW}^{\mathrm{S}'}$ is –10, which, in contrast to ADRZ's assertion, is attributable to the federal minimum wage variation—the "contamination" of the placebo test we originally described. It can be eliminated

---

[15]By contrast, for their placebo test to be valid, ADRZ implicitly assume that the effect of an increase in the federal minimum wage on employment would be the same in every state in the absence of state-level minimum wages. Of course, this assumption is not necessary for the two-way fixed-effects model to generate an unbiased estimate of the average causal effect of minimum wage increases on employment.

[16]This assumption is supported by the data. For example, using data on average hourly earnings, for states with minimum wages above $8.45 per hour in 2015 (the average minimum wage in states above the federal level), the average wage in the leisure and hospitality industry in bordering states (with a minimum wage at or below the federal level) was $13.25, whereas for states with minimum wages above $7.25 but below $8.45, the average wage in bordering states was $12.67. (The correlation between state minimum wages and average wages in bordering states is 0.27.)

[17]The initial employment levels are set to 100 for simplicity. They could vary with initial wage levels.

Table 1. Data for Minimum Wage Example

	Year 1					Year 2			
State	MW^F	MW^S	$MW^{S'}$	Average wage	Employment	MW^F	MW^S	$MW^{S'}$	Employment
1	**8**	8	**9**	$13	**100**	**10**	10	**10**	**80**
2	8	9	—		—	10	10	—	—
3	**8**	8	**10**	$15	**100**	**10**	10	**10**	**90**
4	8	10	—		—	10	10	—	—

Notes: The "average wage" does not enter the regression as data; it simply illustrates the likelihood that states bordering lower minimum wage states (like State 1) are likely to have lower wages, which is why employment falls more in State 1 in response to the federal minimum wage increase. MW^F, federal minimum wage; MW^S, state minimum wage; $MW^{S'}$, cross-border minimum wage.

by focusing, as in NSW (2014a), on observations with no federal minimum wage variation—a restriction ADRZ characterize as "arbitrary" (online appendix, footnote 11). In this case the placebo test that DLR proposed gives an estimate very close to zero (NSW 2014a, table 6), and there is no evidence of the supposed spurious negative minimum wage effects that prompt DLR to restrict attention to cross-border county pairs.

Other Evidence Addressing Spatial Heterogeneity

In contrast to the smaller employment effects that ADRZ obtain using their close controls, a quite different conclusion emerges from many recent studies that address the problem of correlated shocks in other ways. Indeed, nearly *all* of these studies found negative effects of minimum wages on low-skilled employment—often stronger disemployment effects than we have found. Thus, ADRZ's conclusion that minimum wage effects do not reduce low-skilled employment once researchers account for correlated shocks is not supported by a wider body of evidence. This evidence is summarized in Table 2.

Studies Using Geographically Close Controls

Studies that used close controls in a fashion similar to ADR found, not surprisingly, similar results (Addison, Blackburn, and Cotti 2013; Gittings and Schmutte 2016; Slichter 2016).[18] By contrast, Liu, Hyclak, and Regmi (2016) estimated a county-level fixed-effects model that includes interactions between dummy variables for each quarter and Bureau of Economic Analysis (BEA) Economic Areas, which are intended to delineate regional markets. Because some of these economic areas cross state lines, minimum wage effects can be identified from state variation within them. Liu et al. (2016) found evidence of disemployment effects for the youngest group

---

[18]Addison et al. (2013) did find stronger evidence of disemployment effects for teens during the Great Recession.

Table 2. Recent Estimates of Minimum Wage Effects on Unskilled Employment

Authors (by order in which discussed)	Employment elasticity and groups studied	Data/approach
Geographically proximate designs		
Dube, Lester, and Reich (DLR 2010)	Near zero for teens and restaurant workers	Paired counties on opposite sides of state borders
Allegretto, Dube, and Reich (ADR 2011)	Near zero for teens	States compared only to those in same census division
Gittings and Schmutte (2016)	Near zero for teens; larger negative elasticities in markets with short non-employment durations (–0.1 to –0.98) and smaller positive elasticities in markets with long non-employment durations (0.2 to 0.46)	States compared only to those in same census division
Addison, Blackburn, and Cotti (2013)	Varying sign, more negative, generally insignificant for restaurant workers and teens; stronger negative at height of Great Recession (–0.34)	Similar methods to Dube et al. (2010) and Allegretto et al. (2011), restricted to 2005–10 period
Slichter (2016)	–0.04 (teens)	Comparisons to bordering counties and other nearby counties
Liu, Hyclak, and Regmi (2016)	–0.17 (14–18-year-olds)	Comparisons within BEA Economic Areas (EA) that cross state lines, with controls for EA-specific shocks
Other approaches		
Neumark, Salas, and Wascher (NSW 2014a, 2014b)	–0.14/–0.15 for teens, –0.05/–0.06 for restaurant workers	States compared to data-driven choice of controls (synthetic controls) and state panel data
Powell (2015)	–0.44 for teens	States compared to data-driven choice of controls (synthetic controls, estimated simultaneously with employment effect)
Totty (2015)	–0.01 to –0.04 for restaurant workers; –0.04 to –0.07 for teens	States compared to data-driven choice of controls (factor model)
Dube and Zipperer (2015)	–0.051 (mean) and –0.058 (median) for teens	States compared to data-driven choice of controls (synthetic control)
Baskaya and Rubinstein (2015)	–0.3 to –0.5 for teens	States, using federally induced variation as instrumental variable
Clemens and Wither (2016)	Approximately –0.97 for those directly affected by minimum wage increase	Targeted/affected workers versus other low-wage workers in states affected by federal increases, relative to other states
Thompson (2009)	–0.3 (for teen employment share)	Low-wage counties compared to higher-wage counties in states

Notes: The table reports our best attempts to identify the authors' preferred estimates reported in the articles. The Thompson estimate cannot be compared directly to other elasticity estimates, because the data source did not include a population count. The Clemens and Wither elasticity is based on a 6.6 percentage point decline (p. 27), divided by a 70.2% employment rate (or a 9.4% employment decline), divided by a 9.7% MW increase (50 cents, from p. 14, divided by $5.15).

covered in their data (14–18-year-olds), which are diminished only slightly—to an elasticity of –0.17—within Economic Areas. Although we have concerns about the close controls approach, the geographic controls used in Liu et al. seem more defensible than using border-control pairs that may not be in economically integrated areas.[19]

Data-Driven Choice of Controls

Other studies let the data determine the controls, rather than imposing the close control design. In this regard, Powell (2016) improved substantially on the synthetic control approach, developing an estimator that can be applied to multiple treatments with continuous variation and that simultaneously estimates the treatment effect and the weights on control states; this appears to be the most satisfactory and flexible approach, to date, of letting the data choose control states. Powell found a statistically significant estimated elasticity for teens of –0.44.[20]

Instrumental Variables Evidence

Baskaya and Rubinstein (2015) used instrumental variable (IV) techniques to confront the potential endogeneity between teen employment and minimum wages. They instrumented for state minimum wages with the federal minimum wage interacted with a measure of the propensity for each state to let the federal minimum wage bind, which purges the estimated minimum wage effect of variation from state policymakers responding to state-level economic conditions. Consistent with policymakers raising minimum wages when youth labor market conditions are strong, their IV elasticity estimates for teenagers were larger than many past estimates, in the range of –0.3 to –0.5. Their finding of *positive* bias from ignoring the endogeneity of minimum wages stands in contrast to the assertion of negative bias in ADR, DLR, and ADRZ.

Triple-Difference Estimates

The two-way fixed-effects model is a difference-in-differences (DD) estimator, comparing changes in low-skilled employment in states where the

---

[19]DLR presented evidence using only bordering counties in MSAs (table 4), which are more likely to be economically integrated areas. But their evidence was for restaurant workers rather than for teenagers.

[20]Totty (2015) used a factor model that also does not impose close controls but that is more restrictive than Powell's approach. He found small (and insignificant) negative employment effects for restaurant workers and teens. Totty (pp. 10–11) acknowledged that the factor model approach is less flexible than the synthetic control approach, because it imposes a linear structure, but noted that the latter is hard to apply to the minimum wage case. The Powell study, however, appears to solve this problem, and hence should be seen as more compelling. Similarly, it is more compelling than DZ's (2015) synthetic control analysis, which found small average elasticities for teens (around –0.055). Their analysis sharply limited the number of state minimum wage increases used, whereas this limitation is not required using Powell's approach.

minimum wage increased more to states where it increased less. A common concern with DD estimators is that the policy variation is correlated with shocks that also affect outcomes—exactly the problem that motivates ADR and DLR. To address this concern directly, some recent studies used triple-differences (DDD) estimators that isolate the effect of the policy change by introducing another group exposed to the same shock but not the policy change.[21]

Thompson (2009), which predates ADR and DLR, used a DDD approach based on differences in wage levels across counties within a state; this allowed the inclusion of state-period effects. Using variation generated by the federal increases in 1996 and 1997, he found large disemployment effects in counties where minimum wages were more binding because wages were lower and workers were lower skilled. For example, for counties in the bottom third of the teen earnings distribution within a state, a 10% federal increase reduced the teen employment share about 3%.

Clemens and Wither (2016) focused on the 2007–2009 federal minimum wage increases, comparing changes in employment for the lowest-wage workers, whose wages were differentially affected by the federal increases (because of variation in state minimum wages), to changes for workers who earned wages that were low, but high enough that the federal minimum wage had little impact on them. They estimated an employment elasticity for directly affected workers of approximately –0.97 (using SIPP data). The magnitude is likely larger than other studies because it was calculated for a more directly targeted group of workers (compared to teenagers or restaurant workers, only some of whom are directly affected by the minimum wage).[22] This larger elasticity may be more relevant to policy, however, because it measures employment effects among those most likely to experience higher wages in response to a minimum wage increase.

A similar DDD approach that compares 16–19 to 20–24-year-olds reinforces the evidence of disemployment effects for the least skilled in state-level panel data. Using standard specifications with CPS data, estimated elasticities for 16–19-year-olds range from –0.08 to –0.18 (generally statistically significant), but for 20–24-year-olds estimated elasticities range from 0.00 to 0.05 (and are never significant).[23] It seems difficult to explain the teen results as driven by shocks to the labor market for 16–19-year-olds that are correlated with minimum wage increases, absent similar shocks for 20–24-year-olds. Of course it is possible, just as it is possible in the other evidence

---

[21]This approach (as well as synthetic control matching) obviates the need to include state-specific trends.

[22]Indeed, the elasticity is smaller when using a treatment group that includes higher-wage workers and hence is "less intensively" treated.

[23]These estimates are based on March CPS data from either 1990–2011 or 1990–2016, using either the log of the minimum wage or the Kaitz index, and including controls for the share of the age group in the population, the unemployment rate for 25–54-year-olds, and state and year fixed effects. Estimates are weighted, with standard errors clustered at the state level.

using DDD approaches, that underlying trends just happen to differ between the most-affected and less-affected groups in such a way as to generate spurious evidence of disemployment effects. But this seems implausible across the different DDD research designs, and ADRZ do not make this case.

A Selective View of the Evidence?

ADRZ either dismiss this evidence, or ignore it (in the case of Baskaya and Rubinstein 2015 and Thompson 2009). Countering Clemens and Wither (2016), ADRZ cite work by Zipperer, claiming that the results are not robust to regional heterogeneity—although this seems irrelevant given that the DDD estimator controls for state-specific shocks. They criticize Liu et al. (2016) for not showing whether the results are robust to geographic grouping, ignoring the fact that the main point of the article is to use a priori sensible economic groupings. And ADRZ dismiss Powell (2016) in a brief footnote even though the issue of appropriate control states is central to the debate and their comment focuses extensively on synthetic control estimation—exactly the issues Powell addressed. Moreover, ADRZ conclude that "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" (p. 590). Table 2 demonstrates that this is an inaccurate summary of what the existing research says about the employment effects of minimum wages.

Conclusion

The studies by ADR and DLR that ADRZ defend raise a valid point: we need to explore different ways of specifying the controls in studies of the employment effects of minimum wages because minimum wages could be correlated with unmeasured shocks. Our earlier articles, however, indicated that their strategy of using geographically proximate controls to obtain more reliable estimates is not supported by the data. In our view, ADRZ fail to make a compelling case that the approaches in ADR and DLR provide more valid evidence on the effects of minimum wages on employment than the evidence that ADR and DLR critiqued. In particular, the new analyses they provide—the use of LASSO tests to choose which trends or region–time interactions to include, their critique of our synthetic control estimates, and new evidence on the validity of the border-discontinuity design—are either flawed, unconvincing, or superseded by newer research.[24]

---

[24]Moreover, the three claims that ADRZ say we made, and that they claim to have refuted regarding 1) the contamination of estimated minimum wage effects by pre-trends, 2) whether local controls throw away too much information, and 3) our "new matching estimator," are "refuted" only by modifying our claims in an erroneous, misleading, or exaggerated way, or by changing the sample period to include a much earlier period with little minimum wage variation.

This exchange further highlights that the estimates can be sensitive to the inclusion of time trends, which we (and ADR) had already demonstrated in our earlier articles. Although this evidence merits further research (including identification of other omitted influences on teen employment over the period being studied), it does not justify concluding that the effects of minimum wages on teen (or low-skill) employment are negligible. Moreover, alternative approaches can control for underlying trends that obviate the need for arbitrary parametric trend assumptions—specifically, the types of DDD or matching (synthetic control) estimators used in recent research. Most of this recent research (except for the ADR and DLR strategy of close controls) finds that higher minimum wages reduce low-skilled employment.

ADRZ ignore or dismiss the growing body of evidence indicating that higher minimum wages reduce employment of the least skilled, even when account is taken of shocks to the low-skilled labor market that are potentially correlated with minimum wages. Indeed, this evidence frequently includes elasticities that are considerably larger than the –0.1 to –0.2 range from the earlier "consensus" and many of our past estimates. Given ADRZ's conclusion that "accumulating evidence has led us . . . to the conclusion that the employment effects of U.S. minimum wage policies on low-wage employment to date have been fairly small" (p. 590),[25] they have a debate with a much larger body of minimum wage research than just ours.

References

Aaronson, Daniel, Kyung-Hong Park, and Daniel Sullivan. 2007. Explaining the decline in teen labor force participation. *Chicago Fed Letter*, No. 234 (January). Chicago: Federal Reserve Bank of Chicago.

Addison, John T., McKinley L. Blackburn, and Chad D. Cotti. 2013. Minimum wage increases in a recessionary environment. *Labour Economics* 23: 30–39.

[ADR]Allegretto, Sylvia A., 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: 205–40.

[ADRZ]Allegretto, Sylvia, Arindrajit Dube, Michael Reich, and Ben Zipperer. 2017. Credible research designs for minimum wage studies: A response to Neumark, Salas, and Wascher. *ILR Review* 70: 559–92. doi: 10.1177/0019793917692788.

Baskaya, Yusuf Soner, and Yona Rubinstein. 2015. Using federal minimum wages to identify the impact of minimum wages on employment and earnings across U.S. states. Unpublished paper. (A 2012 version can be accessed at http://cep.lse.ac.uk/seminarpapers/20-11-12-YR.pdf.)

[25]Although the word "small" is vague, elsewhere Dube has called the estimates "statistically indistinguishable from zero" in advocating for a higher minimum wage (accessed at http://opinionator.blogs.nytimes.com/2013/11/30/the-minimum-we-can-do/?_r=2&, December 8, 2016). More concretely, Reich (the R in ADRZ) and coauthors have developed a structural model they routinely apply to project the effects of proposed minimum wage increases in various cities, and their model is calibrated to "be consistent with the very small effects that researchers find for the smaller pre-2015 increases in federal and state minimum wages" (Reich et al. 2016, p. 20). They typically projected little if any job loss; for example, Reich et al. (2016) predicted that a $15 minimum wage in San Jose and Santa Clara County phased in by 2019 would cost 960 jobs in San Jose and only 80 in total over the broader region.

Brown, Charles, Richard B. Freeman, Daniel S. Hamermesh, Paul Osterman, and Finis Welch. 1995. Review Symposium on *Myth and Measurement: The New Economics of the Minimum Wage*, edited by Ronald G. Ehrenberg. *Industrial and Labor Relations Review* 48: 827–49.

Card, David, and Alan Krueger. 1995. *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton, NJ: Princeton University Press.

Card, David, Lawrence F. Katz, and Alan B. Krueger. 1994. Comment on David Neumark and William Wascher, "Employment effects of minimum and subminimum wages: Panel data on state minimum wage laws." *Industrial and Labor Relations Review* 47: 487–97.

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

[DLR]Dube, Arindrajit, T. William Lester, and Michael Reich. 2010. Minimum wage effects across state borders: Estimates using contiguous counties. *Review of Economics and Statistics* 92: 945–64.

[DZ]Dube, Arindrajit, and Ben Zipperer. 2015. Pooling multiple case studies using synthetic controls: An application to minimum wage policies. IZA Discussion Paper No. 8944. Bonn, Germany: Institute of Labor Economics.

Gittings, R. Kaj, and Ian M. Schmutte. 2016. Getting handcuffs on an octopus: Minimum wages, employment, and turnover. *ILR Review* 69: 1133–70. doi 10.1177/0019793915623519.

Griliches, Zvi. 1979. Sibling models and data in economics: Beginnings of a survey. *Journal of Political Economy* 87: S37–S64.

Hamermesh, Daniel S. 2007. Replication in economics. *Canadian Journal of Economics* 40: 715–33.

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

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

Neumark, David. 2016. Policy levers to increase jobs and increase income from work after the Great Recession. *IZA Journal of Labor Policy* 5: 6. Accessed at http://izajolp.springeropen .com/articles/10.1186/s40173-016-0064-y.

Neumark, David, and William Wascher. 1994. Employment effects of minimum and subminimum wages: Reply to Card, Katz, and Krueger. *Industrial and Labor Relations Review* 47: 497–512.

Neumark, David, J. M. Ian Salas, and William Wascher. 2013. Revisiting the minimum wage-employment debate: Throwing out the baby with the bathwater? NBER Working Paper #18681. Cambridge, MA: National Bureau of Economic Research.

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

[NSW]Neumark, David, J. M. Ian Salas, and William Wascher. 2014b. More on recent evidence on the effects of minimum wages in the United States. *IZA Journal of Labor Policy* 3: 24. Accessed at http://link.springer.com/article/10.1186/2193-9004-3-24.

Powell, David. 2016. Synthetic control estimation beyond case studies: Does the minimum wage reduce employment? RAND Labor & Population Working Paper WR-1142. Santa Monica, CA: RAND Corporation.

Reich, Michael, Claire Montialoux, Sylvia Allegretto, Ken Jacobs, Annette Bernhardt, and Sarah Thomason. 2016. The effects of a $15 minimum wage by 2019 in San Jose and Santa Clara County. Policy Brief, Institute for Research on Labor and Employment, University of California, Berkeley.

Slichter, David. 2016. The employment effects of the minimum wage: A selection ratio approach to measuring treatment effects. Unpublished paper. Accessed at http://bit.ly/1kLgpGT.

Smith, Christopher L. 2011. Polarization, immigration, education: What's behind the dramatic decline in youth employment? Finance and Economic Discussion Series 2011-41. Board of Governors of the Federal Reserve System.

Thompson, Jeffrey P. 2009. Using local labor market data to re-examine the employment effects of the minimum wage. *Industrial and Labor Relations Review* 62: 343–66.

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