Revisiting the Minimum Wage-Employment Debate:
Throwing Out the Baby with the Bathwater?[*]

David Neumark
UCI, NBER, and IZA

J.M. Ian Salas
UCI

William Wascher
Board of Governors of the Federal Reserve System

July 2013

Abstract:

We revisit the minimum wage-employment debate, which is as old as the Department of Labor. In particular, we assess new studies claiming that estimates from the panel data approach used in much of the "new minimum wage research" are flawed because this approach fails to account for spatial heterogeneity. These new studies use research designs intended to control for this heterogeneity and conclude that minimum wages in the United States have not reduced employment. We explore the ability of these research designs to isolate reliable identifying information and test the untested assumptions in this new research about the construction of better control groups. Our evidence points to serious problems with these research designs. Moreover, new evidence based on methods that let the data identify the appropriate control groups leads to evidence of disemployment effects, with teen employment elasticities near −0.15. We conclude that the evidence still shows that minimum wages pose a tradeoff of higher wages for some against job losses for others, and that policymakers need to bear this tradeoff in mind when making decisions about increasing the minimum wage.

Debates about the economic effects and the merits of the minimum wage date back at least as far as the establishment of the Department of Labor as a cabinet-level agency in 1913. Economists and statisticians from the Department have contributed importantly to the debate through their empirical studies on the economic effects of minimum wages over the past century. Indeed, one of the first statistical analyses of minimum wages in a U.S. state was conducted in 1915 by Marie Obenauer and Bertha von der Nienburg of the Bureau of Labor Statistics (BLS), who examined the effects of a minimum wage for women that was introduced in Oregon between October 1913 and February 1914. For this study, which was a precursor to the case study approach that constitutes a key branch of the empirical literature that blossomed after 1990, the BLS collected data on employment and wages by age and sex, as well as sales, from 40 retail stores in Oregon for March and April of 1913, about five months prior to the introduction of the minimum wage in the state, and for March and April of 1914, about five months after the minimum wage took effect. The study then compared the changes in the employment of women and men over that period. Although quite cautious about the power of their difference-in-differences statistical approach (for example, the analysis was complicated by a recession in 1914 and by a legislated reduction in working hours for women), the study concluded that the minimum wage had a positive effect on wages and little or no effect on women's employment in the aggregate, but that stores substituted teenage girls (who were subject to a lower minimum wage) for adult women in lesser-skilled jobs.

Similarly, some of the first empirical analyses of the federal minimum wage law enacted in 1938 were undertaken by analysts from the Wage and Hours and Public Contracts Division of the Department of Labor. In particular, the Department conducted a series of studies examining the effects of the new minimum wage on wages and employment in certain industries, either by

comparing changes in wages and employment in plants in low-wage southern states with their counterparts in higher-wage northeastern states or by comparing employment changes in plants with different levels of average wages before the federal minimum wage took effect. These studies, which tended to find modest disemployment effects, were followed by similar efforts by the Department to assess the effects of increases in the federal minimum wage in the 1940s and 1950s. The findings of these studies were at the center of a vigorous debate between Richard Lester and John Peterson on the merits of the minimum wage in the late 1950s and early 1960s (Lester 1960; Peterson 1957, 1959, 1960).

Over time, empirical analyses, especially the time-series studies conducted in the 1960s and 1970s, increasingly found that minimum wages tended to reduce employment among teenagers, who were viewed as a proxy for low-skilled labor more generally. A famous paper by Charles Brown, Curtis Gilroy, and Andrew Kohen, published in 1982, surveyed the existing literature on minimum wages and established the "consensus" that a 10 percent increase in the minimum wage would reduce teenage employment by 1 to 3 percent (Brown, Gilroy, and Kohen 1982). Following that study, economists began to coalesce around the idea that minimum wages have adverse effects on low-skilled employment.

That consensus turned out to be relatively short-lived. After a decade of near-silence, the debate over the employment effects of the minimum wage reemerged in the early 1990s with the publication of a special issue of the *Industrial and Labor Relations Review* (ILRR). This issue featured four studies that used different analytical approaches and that took advantage of the increasing divergence of minimum wages at the state level to estimate the employment effects of minimum wages. These studies, which formed the basis for what is sometimes termed the "new minimum wage research," were diverse in their findings, ranging from disemployment effects

similar to the earlier consensus (Neumark and Wascher 1992), to no effect on employment (Card 1992a) to a positive effect of the minimum wage on employment (Card 1992b; Katz and Krueger 1992).

The ILRR symposium launched a new body of contemporary research on the minimum wage, much of which was summarized in the 2008 book *Minimum Wages* written by two of us (Neumark and Wascher 2008). In that book, our evaluation and summary of the evidence concluded that "… [M]inimum wages reduce employment opportunities for less-skilled workers, especially those who are most directly affected by the minimum wage" (Neumark and Wascher 2008, p. 6). This paper, in part, extends this evaluation and summary to the present by evaluating two recent studies that have questioned the empirical methods and conclusions in much of the recent literature (Allegretto, Dube, and Reich 2011; Dube, Lester, and Reich 2010).

The key question raised by these recent studies is how researchers can best identify the employment effects of the minimum wage – a question that is nearly as long-running as the debate over the minimum wage. In particular, the identification of minimum wage effects requires both a sufficiently sharp focus on potentially affected workers and the construction of a valid counterfactual "control group" for what would have happened absent increases in the minimum wage. The latter is critical to account for other influences on the employment of potentially affected workers that may be confounded with the effects of changes in the minimum wage. In the research of the past two decades, economists have frequently used state variation in minimum wages to generate comparisons between states with different minimum wage levels or changes at the same point in time, to avoid confounding minimum wage effects with other aggregate influences on the labor market (e.g., the national business cycle).

Dube et al. (2010, hereafter DLR) and Allegretto et al. (2011, hereafter ADR) have put

forward a severe critique of the state panel-data approach, including the work discussed at length in Neumark and Wascher (2008). The essence of the argument in DLR and ADR is summarized in a review of *Minimum Wages* by Dube (2011), which draws heavily on the findings from the two papers he co-authored:

> "…[V]ariation over the past two decades in minimum wages has been highly selective spatially, and employment trends for low-wage workers vary substantially across states … This has tended to produce a spurious negative relationship between the minimum wage and employment for low wage workers – be it for sectors such as restaurant and retail or for demographic groups such as teenagers" (Dube 2011, p. 763).

Commenting on the econometric evidence more specifically, Dube writes: "Even simple regional controls and trends produce employment effects close to zero, as do more sophisticated approaches such as comparing contiguous counties across policy boundaries – which essentially embeds the "case study" approach within panel data analysis …" (pp. 763-4). Dube defines his and his co-authors' studies as "a fourth generation of recent work that tries to make sense of the sometimes contradictory evidence" (p. 763), and argues that their work raises serious questions about the conclusions drawn by Neumark and Wascher – and much of the broader literature – regarding the employment effects of minimum wages.

Echoing Dube, ADR assert without reservation that their results overturn the conclusion that minimum wages reduce employment of low-skilled workers: "Interpretations of the quality and nature of the evidence in the existing minimum wage literature … must be revised substantially. Put simply, our findings indicate that minimum wage increases – in the range that have been implemented in the United States – do not reduce employment among teens" (ADR 2011, p. 238). Similarly, DLR conclude that there are "no detectable employment losses from the kind of minimum wage increases we have seen in the United States" (DLR 2010, p. 962).

Our principal goal in this paper is to evaluate this new research because of the strong

challenge it poses to the large body of prior research that found that minimum wages reduce employment of low-skilled workers. As the description of the work above suggests, the central element of this new research is the issue of how to construct counterfactuals for the places where minimum wages are increased. The authors of both studies argue that one must compare places that are geographically proximate to have valid controls, because, according to them, minimum wage changes are correlated with unobserved economic shocks to areas that can confound the estimation of minimum wage effects. Consequently, much of the analysis in this paper focuses on the validity of this criticism, and on the approaches these studies take to address this potential problem. The overriding concern we have with these studies is that their research designs, out of concerns about avoiding minimum wage variation that is potentially confounded with other sources of employment change, discard a great deal of valid identifying information – throwing out the identifying "baby" along with, or worse yet instead of, the contaminated "bathwater." Our findings, in a nutshell, indicate that neither the conclusions of these studies nor the methods they use are supported by the data.

Recent Research Challenging the Conclusion that Minimum Wages Reduce the Employment of Low-Skilled Workers

Of the two papers by Dube and his colleagues, the analysis in ADR is the most direct extension of the state panel-data approach used extensively in the existing research on the employment effects of minimum wages. In this study, ADR focus on state-level panel data specifications of minimum wage effects on the employment of teenagers, using information on state-level minimum wages and individual-level data from the Current Population Survey (CPS) from 1990-2009. When they estimate a model that includes state and period fixed effects along with other standard controls, they find a negative employment effect of minimum wages.

However, when they include either state-specific linear trends or Census division × period interactions (or both), the estimated employment effects of minimum wages fall to approximately zero and are statistically insignificant.

In contrast, DLR's analysis focuses primarily on restaurant employment using county-level Quarterly Census of Employment and Wages (QCEW) data from 1990-2006. Although they present some results from panel data models that include state-specific trends and Census division × period interactions (along with county fixed effects), their core analysis uses a research design based on cross-border county pairs. Their specification includes county pair × period interactions intended to control for shocks common to both counties, and thus identifies the effect of minimum wages from differences in employment changes in paired counties on either side of a state border. This narrowing of identification to within-county-pair comparisons causes the employment effects to go from negative and sometimes statistically significant to small and insignificant.[1]

To put this new evidence in context, it is useful first to assess the implications of these results for the existing state-level panel studies, especially since ADR and DLR have explicitly used their findings to cast doubt on the evidence from these studies. With regard to DLR's paper, it is worth noting that very little of the existing work is on the restaurant sector or the retail sector more broadly, so new evidence on restaurant or retail employment does not address the far more pervasive evidence on teens or other very low-skilled workers. For one thing, the evidence from the earlier research is strongest for individuals most directly affected by the minimum wage, and many workers within the restaurant or retail sector earn well more than the minimum wage. For another, the minimum wage can lead employers to substitute higher-skilled

---

[1] Closely related findings are reported in Addison, Blackburn, and Cotti (2009, 2012).

workers for lower-skilled workers without reducing net employment very much.[2]

ADR's research focuses primarily on teenagers and can therefore be viewed as posing more of a direct challenge to the findings from the state-level panel data approach. Even in this case, however, the potential for labor-labor substitution among teenagers with different skill levels means that the effects of minimum wages on overall teenage employment can be difficult to detect; for example, larger gross disemployment effects among the least-skilled teens may be masked by inflows of other teens into employment.[3] Indeed, the most recent estimates Neumark and Wascher have presented for teenagers show negative effects only for male teens when disaggregating by sex, and only for black or Hispanic male teens when disaggregating male teens into whites vs. black or Hispanic (Neumark and Wascher 2011). And other work, focused on the lowest-wage workers rather than on teenagers per se, finds negative employment effects for them as well (Neumark, Schweitzer, and Wascher 2004). Nonetheless, a negative effect of minimum wages on employment of the lowest skilled ought to imply negative effects for at least some groups of teenagers, and for the sample period they study, ADR do find that a panel data model with *only* state and year fixed effects (plus standard controls) produces evidence of disemployment effects in the range of past estimates. Thus, their finding that this conclusion is sensitive to whether state-specific linear trends or region × period interactions are included in the specification poses a challenge to the conventional view of minimum wage employment effects. We next turn to a more thorough explanation of their analysis, as well as an evaluation of it; we then do the same for the DLR study.

---

[2] DLR acknowledge this possibility, noting that "… [O]ur data do not permit us to test whether restaurants respond to minimum wage increases by hiring more skilled workers and fewer less skilled ones" (p. 962). However, the existing literature suggests that such labor-labor substitution is important. See, for example, Fairris and Bujanda (2008),Giuliano (2013), Lang and Kahn (1998), and Neumark and Wascher (1996).

[3] The focus on teenagers is, to some extent, a vestige of the old time-series literature. Because labor economists used aggregate employment data by age group, it made sense to look mainly at teenagers because minimum wage workers comprised such a small share of older age groups.

Evaluation of the Evidence

Allegretto et al. (2011)

As noted above, ADR find that the negative effects estimated from state-level panel data specifications of the effects of minimum wages on the employment of teenagers are sensitive to including either state-specific linear trends or Census division × period interactions (or both). This leads them to conclude that models with only state and year fixed effects "fail to account for heterogeneous employment patterns that are correlated with selectivity among states with minimum wages. As a result, the estimates are often biased and not robust to the sources of the identifying information" (p. 205). More specifically, they argue that "Lack of controls for spatial heterogeneity in employment trends generates biases toward negative employment elasticities in national minimum wage studies" (p. 206).

We re-examine these findings with the same CPS data, using a specification with the same aggregate variables they include but not the individual-level controls. The sample is extended to take account of newer data, and, whereas ADR use individual-level data with clustering at the state level, we aggregate the data to the state level by quarter, also clustering at the state level.[4] The minimum wage, here and throughout, is defined as the higher of the state and federal minimum.

As can be seen in Table 1, Panel A, the results closely mirror what ADR found (their Table 3). In the model that includes the standard labor market controls along with state and time fixed effects, the estimated employment elasticity with respect to the minimum wage is −0.165,

---

[4] Aggregated data are used because this form is more convenient for some of the analyses that follow. Moreover, because the identifying information is the state-level minimum wage variation, the use of state-level data for the other variables should be inconsequential. Nonetheless, it is possible that the individual-level controls ADR use (sex, race, age, education, and marital status) could lead to some differences in the results.

significant at the 1% level (ADR estimate an elasticity of −0.12, significant at the 5% level).[5]

When either state-specific linear trends are added, region × month interactions are added, or both are added simultaneously, the estimated elasticities become smaller, sometimes considerably (ranging from −0.098 to 0.009) and are statistically insignificant. The same is true in the ADR results (their Table 3), where the estimates for these specifications are statistically insignificant and the estimated elasticities range from −0.036 to 0.047.[6]

{{Place Table 1 about here}}

This evidence, taken at face value, suggests that conclusions about the effects of minimum wages on teenagers may not always be robust to the type of identifying variation used to estimate these effects: differences in within-state variation associated with minimum wage changes relative to other states in the same year; differences in within-state variation relative to other states in the same year that is also net of state-specific linear trends; or (essentially) differences in within-state variation relative to states in the same Census division.

Interestingly, the only time ADR question the validity of their approach is with regard to their evidence of statistically significant negative effects on hours of Hispanic teens. In response to these findings, they write "the puzzling and somewhat fragile evidence for Hispanic teens may be driven by the concentration of Hispanic teens in a small number of Census divisions, on the one hand, and the small number of Hispanic teens in most states at the beginning of the sample period. These patterns reduce the ability to estimate effects for this group robustly within our methodology" (2011, p. 234). Similarly, they argue that "[I]ncluding spatial controls renders the

---

[5] The specification is in logs so the estimated coefficient is the elasticity; in contrast, ADR estimate linear probability models for employment with the log of the minimum wage on the right-hand side. They report the magnitude and statistical significance of the estimated linear probability coefficients, and then the implied elasticity.
[6] We also experimented with a specification that added controls for the adult wage and adult employment-to-population ratio, and defining these variables (and the adult unemployment rate) for skilled adults aged 25-64 with more than a high school education. The estimated minimum wage effects were very similar.

estimates for Latinos particularly imprecise and fragile" (p. 208). But in their Table 7, on which this discussion is based, the estimates are actually more precise for Hispanics than for blacks, yet they conclude that "controlling for spatial heterogeneity by using within-Census division variation is particularly important when looking at African-American employment effects" (p. 234).

Rather than judgmentally deciding where and when to include area-specific time trends or region × period dummies based on unclear criteria like these, researchers should examine what sources of variation provide better estimates of the effects of minimum wages. In the context of this paper, this preferred approach entails exploring the implications of including state-specific trends or region × time interactions and whether doing so results in more or less reliable estimates of minimum wage effects.

State-specific trends.

We first focus on the evidence regarding state-specific trends, which are intended to control for longer-run influences not captured in the other control variables. It is standard practice to assess the robustness of panel data estimates of state policy effects to the inclusion of state-specific trends, including in the minimum wage literature (e.g., Neumark and Wascher 2004, 2011). If these kinds of analyses deliver results that are insensitive to the inclusion of these trends, then they can clearly bolster the evidence. If, however, they point to different conclusions, then the researcher has to seriously explore which analysis is most convincing, rather than simply relying on a priori hunches.[7]

---

[7] To support their inclusion of linear trends, ADR point to various longer-term influences on teen employment that are not included in the specification. Specifically, they cite research by Smith (2011) suggesting that technological change may have led to an increased supply of adult workers for low-skill jobs that had been commonly held by youth (their footnote 2). And they cite research by Aaronson, Park, and Sullivan (2007) and the Congressional Budget Office (2004) suggesting that teen employment may have been influenced by changes in financial aid for college students, the attractiveness of college, or technological shifts that have lowered market wages for teens. Of

The first thing to note is that Neumark and Wascher (2011), using data from 1994-2007, found that the estimated effects of minimum wages on teen employment are negative and significant in specifications that include state-specific trends.[8] This result raises the question as to whether there is something different about the sample period ADR studied that makes it problematic to include linear state-specific trends. An obvious candidate is the severe recession at the end of their sample period, as is the recession at the beginning of their sample period (in 1990 and 1991). In models that include state-specific trends, the recessions at the beginning and end of ADR's sample period could have a large influence on the estimated state-specific trends – so-called "endpoint bias." If the recessions have purely an 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 misclassification 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.[9]

One approach to this problem is to allow the state-specific trends to be of a higher order than linear. Higher-order trends should be better than linear trends at capturing the variation

course, it would clearly be preferable to incorporate data on the potentially omitted factors rather than simply including trends and interpreting the results as reflecting these factors

[8] Although not reported in a table, we verified that with the data used here, this result still holds. Estimating the model from Table 1, Panel A, with data from 1994-2007:Q2, the minimum wage effect (standard error) is –0.148 (0.060) without state linear trends, and –0.229 (0.095) with them.

[9] In a longer version of this paper (Neumark, Salas, and Wascher 2013), we show how including the recessionary periods at the beginning and end of the sample period biases the state-specific trend estimates away from the values associated with the longer-term factors not captured by the other controls in the model. Note that the recession in 2001 is less problematic in estimating the trends because the sample period includes lengthy periods of expansion both before and after that recession.

induced by the recessions – especially third-order and higher polynomials that allow for multiple inflection points.[10] Estimates for the full sample period that include these higher-order trends are reported in Panel B of Table 1. (The comparable estimates with no state-specific trends and with linear trends are in columns (1) and (2) of Panel A.) The table shows that as long as third-order polynomials or higher are used, the estimated effects of the minimum wage on teen employment are negative and significant, with point estimates very similar to the estimates for the model that excludes state-specific trends – and similar to the estimates for the subsample excluding the beginning and ending recessionary periods (as well as the period after these recessions formally ended when unemployment continued to rise).[11] In each column, we tested the statistical significance of the higher-order terms added relative to the previous column (in column (1) we tested the significance of the squared time interactions). These were significant for the 2nd-, 3rd-, 4th- and 5th-order terms (p-values < 0.001). Thus, the evidence indicates that linear state-specific trends are too restrictive and invalidates ADR's (2011) conclusion that "Lack of controls for spatial heterogeneity in employment trends generates biases toward negative employment elasticities in national minimum wage studies" (p. 206).[12]

Division × period interactions.

The preceding analysis suggests that ADR's claim that underlying trends that vary by state generate spurious evidence of negative minimum wage effects on teen employment is

---

[10] Macroeconomists frequently use higher order polynomials to account for trends in time-series studies. See, for example, Aruoba, Diebold, and Scotti (2009).

[11] The results are very similar using the slightly shorter sample period that ADR use. The point estimates are also similar, although a bit less precise, with 6th- or 7th-order polynomials; for these specifications, the estimated minimum wage elasticity ranges from −0.14 to −0.17, and is significant at the ten-percent level in three out of four cases (including both specifications and either ADR's sample period or our longer period).

[12] We also found negative and statistically significant minimum wage effects when we estimated the state-specific linear trends using only the data from the subsample that excludes recessionary periods and when we pre-filtered the data using a Hodrick-Prescott filter, as is suggested by Ionides, Wang, and Granados (2012). For details, see Neumark et al. (2013).

clearly not true. However, Table 1 also shows that the inclusion of Census division × period interactions renders the estimated minimum wage effects smaller and statistically insignificant. As a prelude to delving into what to make of this result, we start by considering the arguments that ADR use to support their view that including these interactions is necessary to account for the spatial heterogeneity that they think biases estimates of minimum wage effects in the panel data specification with only state and period fixed effects. In this regard, they appeal to Figure 1 and Table 1 (from their paper) to argue that "employment rates for teens vary by Census division and differentially so over time" (p. 206).

One particular hypothesis they suggest is that the endogeneity of minimum wages increases generates a bias towards finding negative employment effects. As evidence, they cite Reich (2009), who ADR claim shows that minimum wages "are often enacted when the economy is expanding and unemployment is low. But, by the time of implementation, the economy may be contracting and unemployment increasing, possibly leading to a spurious time series correlation between minimum wages and employment" (p. 212).[13] However, ADR's argument does not speak to biases in the estimated effect of minimum wages on teen employment in the kinds of models estimated in the literature (which they critique), because these models already control for aggregate state-level labor market conditions via the unemployment rate, and includes time dummies that will capture aggregate changes not picked up in the state-level controls. Instead, what is relevant is whether shocks to teen employment conditional on aggregate labor market changes affect minimum wages.

With regard to this point, the evidence points to *positive* endogeneity bias. A recent

---

[13] In fact, Reich actually argues the opposite, based on his evidence: "Minimum-wage increases are voted almost without exception and are mostly implemented in times of growing employment. This pattern holds for both federal and state increases" (p. 366, italics added). Thus, if anything, Reich's evidence points to possible spurious positive correlations that would bias estimated minimum wage effects toward zero assuming the true effects are negative.

study by Baskaya and Rubinstein (2011) that looks specifically at the endogeneity problem using an instrumental variables approach – and which conditions on the adult male unemployment rate – supports the hypothesis that state minimum wage increases respond positively to increases in teen employment, so that failure to account for endogeneity biases the estimated employment effect towards zero, masking the negative effect of minimum wages.[14] Additional evidence in Thompson (2009) supports the same conclusion. His analysis is at the county level, because there is considerable within-state heterogeneity in wage levels and local labor market conditions, and because, he argues, counties better represent labor markets for teens than do entire states, given constraints that keep teens close to where they live. This analysis considers the types of factors that DLR and ADR suggest can lead to spurious evidence of disemployment effects of minimum wages, yet finds substantial negative effects of federal minimum wage increases in the low-earnings (or high-impact) counties.[15]

That said, there could be other omitted factors that drive patterns of teen employment differentially by Census division, and these could be correlated with minimum wage changes in any direction. And the sensitivity of the estimates to the inclusion of the division × period interactions is something that is important to understand further – in particular whether the disappearance of minimum wage effects when these interactions are included is evidence of the need to control for spatial heterogeneity, as ADR argue.

---

[14] Earlier work by Watson (1996) looking at state minimum wages reaches the same conclusion: "[T]een employment rates positively affect legislator's decisions on reforming minimum wage legislation" (p. 29). However, Watson only conditions on the adult wage and year fixed effects, not an adult employment or unemployment rate in the state, so it is possible that her results reflect responses to overall state economic conditions, rather than the teen labor market per se. Regardless, the results contradict ADR's claims (but not Reich's actual conclusion).

[15] By identifying the effects from differences between counties with low versus high teen earnings, Thompson controls for state-specific changes or trends that could be correlated with minimum wage changes. Moreover, because the minimum wage variation comes from federal legislation, and the identification comes from cross-county variation within states, any endogeneity of state minimum wages is unlikely to be a confounding influence.

An important concern with their approach, though, relates to the implications of augmenting the specification to include more than 1,900 Census division × period interactions (there are 20 years of monthly data for 9 divisions). When the Census division × period interactions are included, all of the identifying information about minimum wage effects comes from within-division variation in minimum wages. An obvious concern is that this extensive set of controls captures a lot more than just the unobserved regional variation, and in particular may remove a good deal of valid identifying information on the effects of minimum wages; moreover, the identifying information that remains is not obviously better than the identifying information that has been removed, in the sense of providing a more valid counterfactual.

Before concluding that this more restricted identification provides more convincing evidence on the effects of minimum wages, a number of questions should be asked. To begin, what do we find if we estimate the models separately by Census division? In ADR's saturated specification the effect of the minimum wage (and all other controls) is constrained to be the same in each division. However, if we think that the patterns of unobserved shocks to divisions differ (or that the effects of the same observed shocks – like technological change – differ), why not also allow the effects of the observed variables to differ by division?

The results, which are reported in Table 2, reveal that the estimated effects of the minimum wage differ substantially across Census divisions, and – just as importantly – that our ability to obtain a precise estimate of the minimum wage effect varies substantially across divisions. In particular, for the New England, West North Central, West South Central, and Mountain divisions there are significant disemployment effects, with elasticities ranging from −0.15 to −0.64, a rather large range. For three other divisions – East North Central, East South Central, and South Atlantic – the estimated effects are also negative, although these are not

statistically significant, and the estimates for East South Central are implausibly large. Finally, for the Mid-Atlantic division the estimated elasticities are positive but insignificant, and for the Pacific division the estimates are near zero.

{{Place Table 2 about here}}

Looking at the standard errors, and given a plausible range of elasticities from prior evidence of −0.1 to −0.3 or perhaps somewhat larger, it is clear that only three divisions – New England, West North Central, and West South Central – yield sufficiently precise estimates to detect a statistically significant elasticity in this range. Thus, looking division-by-division, which in the spirit of ADR's study seems like the best way to control for spatial heterogeneity across Census divisions, yields one of two things – either precise estimates that point to disemployment effects, or estimates too imprecise to be informative. In and of itself, these results lead to a very different conclusion than the one reached by ADR.

Our finding that estimating the models separately by Census division often leads to very imprecise estimates naturally raises the question of whether, in controlling for spatial heterogeneity, it is really necessary to throw out information on other potential comparison states. The assumption that ADR make is that the states within a Census division are better controls for states where minimum wages increase than are states in other Census divisions. In particular, they argue that because minimum wage changes are correlated with economic shocks at the regional level, the models should include "…Census division-specific time effects, which sweeps out the variation across the nine divisions and thereby control for spatial heterogeneity in regional economic shocks …" (p. 206). However, they do not test this hypothesis directly and, indeed, provide no convincing evidence that the counterfactual employment growth that comes from states in other Census divisions does not provide a good control. Moreover, there is

considerable heterogeneity among states within Census divisions (e.g., Maryland vs. South

Carolina, West Virginia vs. Florida, or Connecticut vs. Maine), and some divisions have many

states and cover huge areas (e.g., the Mountain division), so the a priori argument for why the

within-division states provide better controls is unclear.[16]

To address this shortcoming in ADR's analysis, we use two related empirical approaches

to determine which states are good controls for the states with minimum wage increases in any

particular period. Our first approach uses the initial step of the Abadie, Diamond, and

Hainmueller (2010) synthetic control approach to estimating treatment effects. This method can

be applied to simple settings when a discrete treatment (like implementing a program) is applied

to one unit (such as a geographic area) and not to others. The latter – which are the potential

control units (referred to as "donor" units) – are then sorted according to a matching estimator

intended to measure the quality of each unit as a control. The choice of variables on which to

match is subject to the choice of the researcher; most typical, perhaps, would be to match on

prior values or percent changes (where there are level differences) in the outcome of interest.

To draw a comparison with the literature on minimum wages, suppose we want to

estimate the effect of a specific state minimum wage increase in this setting. If we had a time

period when only one state raised its minimum wage, the panel data approach would use the

other 49 states as controls. In contrast, the "case study" approach, typified by Card and Krueger

(1994), would choose another state (or subset of that state) based on geographic contiguity. In

this context, ADR's approach essentially restricts the set of control states to those in the same

Census division. However, nothing in ADR's study establishes that states in the same division

are better controls (just as nothing in the Card and Krueger study establishes that Pennsylvania is

---

[16] For example, even if one accepted the notion that geographically-proximate states provide better controls, there is a question of whether states within a Census division are better controls than closer states in other Census divisions.

a good – let alone the best – control for New Jersey). The synthetic control approach provides us with empirical evidence on which states are the best control states.[17]

However, the application of the synthetic control approach is not straightforward in the minimum wage setting. In the Abadie et al. (2010) paper, a single treatment (a tobacco control program) was implemented in one state (California), and there were long histories on potential control states without such a policy. In contrast, state minimum wages are enacted repeatedly and at high frequency, so we can only match on a short period prior to any minimum wage increase and must drop many potential control states that increased their minimum wage around the same time. In addition, the approach does not allow for policy treatments that vary across regions (like a minimum wage). In this section of the paper, we therefore only examine the first step of the synthetic control approach – the estimation of weights on potential control states. Later in the paper, we propose a procedure to implement the second step in the more challenging context of minimum wages, and provide estimates of minimum wage effects based on that procedure.

To apply the first step of the synthetic control approach to the CPS data used in the preceding analyses, we first define the set of treatment observations as state-quarter observations with minimum wage increases of at least 5 cents in that quarter and no minimum wage increase in the previous four quarters, yielding a set of 129 potential minimum wage treatments to analyze. For each of these treatments, we define a set of potential donor units (state-quarter observations) as states with no minimum wage increase in the same quarter and the succeeding

---

[17] Sabia, Burkhauser, and Hansen (2012) use the synthetic control approach to estimate the effects of the increase in the minimum wage in New York in 2004. In particular, they compare estimates using geographically-proximate states to those that instead use control states picked by the synthetic control method. In their case the approach puts much of the weight on the geographically-proximate states, although some states outside of the Census division get substantial weight in the employment regression (Maryland and Ohio). Jones (2012) uses a related approach to estimating the effects of city-specific living wages, using propensity score matching on characteristics of both individuals and cities.

three quarters, and no minimum wage increase in the previous four quarters. In these analyses, a variable is chosen on which to match over the four quarters prior to the treatment,[18] and then the weights that the matching estimator assigns to each of these donor units are computed.

The validity of ADR's approach can be assessed by looking at how much of this weight is assigned to states in the same Census division. If, as ADR believe, there are common shocks within Census divisions, then the synthetic control method should put most of the weight on same-division states, since they are the states that will match the prior variation in the treatment states. In contrast, if the analysis does not put much weight on same-division states, this would imply that those states are not good controls, so using the within-division variation in minimum wages to identify minimum wage effects is less likely to identify the true effect of minimum wages.

The synthetic control approach requires a choice of variables on which to match, and we used four different alternatives to implement the procedure. Three of these involve matching on forms of the dependent variable: the log of the teen employment-to-population ratio, as well as the one-quarter change and the four-quarter change in that variable, each of these defined over the four pre-treatment quarters.[19] We also matched on the residuals for teen employment from the panel data estimator with only period and state fixed effects (Table 1, Panel A, column (1)), again for the prior four quarters. This is not a standard type of variable on which to match,

---

[18] In the matching process for each treatment case, the relative importance of each value of the matching variable over the four previous quarters is included in the optimization routine, but it can also be specified in advance. To obtain standardized results across all treatment cases, we treated lags one through four of the matching variable as equally important, although relaxing this restriction does not affect the results.

[19] When matching on one-quarter or four-quarter changes over the four pre-treatment quarters, there are sometimes minimum wage increases in the treatment or potential donor states 5 to 8 quarters prior to the treatment quarter. Because excluding these cases would significantly reduce the number of eligible minimum wage treatments and potential donors, we retained them for this analysis. However, restricting the analysis to the treatments that did not include minimum wage increases for the treatment or potential donor states for the preceding 8 quarters produces very similar results.

primarily because there typically is not a regression model underlying the application of the synthetic control approach; rather, the synthetic control estimator is typically used *instead* of a regression model. However, matching on the residuals is informative about the spatial heterogeneity arguments that ADR put forward, as their contention is that the residuals for states in the same Census division share common features that are correlated with minimum wage changes. Consequently, matching on the residuals provides information on whether the residuals for states in the same region share these commonalities and hence whether these states are good controls – and as in ADR's approach, should be isolated as the control states by including division × period dummy variables.[20]

The results are summarized in Table 3. As noted, there are 129 unique treatments that are defined for this analysis. Of these, 50 have potential donors in the same Census division, covering six divisions.[21] The key results are reported in columns (1)-(4); these are the weights from the matching process on states in the same division.[22] With the exception of West North Central, these weights are generally well below one. In 14 out of the 24 cases they are below 0.25, and in some cases they are quite close to zero, implying that most of the weight chosen by the synthetic control method is on states outside the division.

{{Place Table 3 about here}}

---

[20] Given that the estimate of the minimum wage coefficient is in question, we also constructed the residuals from a specification that omits the minimum wage as a control variable (restricting its effect to be zero). The results shown in column (1) of Table 3 were robust to this change in specification. The same is true for the similar county-level analysis described below.

[21] In some cases, there were no potential donor units in a division because all other states in the division had a minimum wage increase in the current quarter, the next three quarters, or the previous four; these cases were thrown out since no weight can be assigned to state-quarter pairs in the Census division if there are no donors in the division for that particular treatment. As a result, we look to see whether there is substantial weight on donor states in the same division only when there *are* such donor states, to avoid overstating the extent to which donor states come from other divisions. Also, as explained in the notes to Table 3, when matching on the one- and four-quarter changes, treatment observations are lost at the beginning of the sample period.

[22] Code in Stata was used for the state analysis, and code in R for the similar county analysis described later (because of a far greater number of potential donors). The software is available at http://www.mit.edu/~jhainm/synthpage.html (viewed July 30, 2012).

Columns (5)-(7) report the average number of divisions and states in the donor pool, and the average number of states in the same division, and help illustrate that the low weight on states in the same division is not attributable to a small number of potential donor states from the same division. For example, Pacific has a low number of potential donor states from the same division relative to all potential donors, but relatively high weight, and South Atlantic has a high number of potential donor states from the same division relative to all potential donors, but relatively low weight.

Finally, calculations based on this analysis demonstrate that there is generally little reason to prefer the same-division states as controls relatively to randomly-chosen states. For the analysis in column (1), the average weight per same-division donor state is higher than the random threshold of 1/(number of potential donors) in only 18 of 50 cases. For columns (2)-(4), the corresponding numbers are 24, 17, and 19 (these last two are relative to 49 and 44 unique minimum wage treatments, owing to some loss of observations from the lagged variables).

These results provide striking evidence against ADR's choice to restrict the control states to those in the same Census division. For most Census divisions, states *outside* the Census division tend to be better controls for treatment observations, whether matched on regression residuals or on levels or growth rates of teen employment. In cases where most of the weight is on states outside the division, the conventional panel data estimator should provide more reliable estimates of minimum wage effects than the specification that includes division × period controls.

A second method, which we term the "ranked prediction error" approach, is also used to address the question of which states are good controls for the states with minimum wage increases. The synthetic control approach finds a weighted average of the potential donor states

to best match the treatment unit. The second method, instead, matches up the treatment unit to each potential donor unit one-by-one. For each of these potential controls, the root mean squared prediction errors (RMSPE) of the exact same matching variables used for the synthetic control approach are calculated for the donor unit relative to the treatment unit in the pre-treatment period (the four quarters prior to the MW change in the treatment unit).[23] The analysis then asks whether the donors in the same division are better controls than the donors outside the division by comparing the RMSPEs for the same-division states with the RMSPEs for other states.

Some notation helps to clarify the method and the difference between the two approaches. Denote a specific treatment unit by T,[24] the potential donors in the same division $D^s_1, \ldots, D^s_J$, and the potential donors in other divisions $D^o_1, \ldots, D^o_K$. The synthetic control approach finds a weight for each donor, $w^s_1, \ldots, w^s_J$ and $w^o_1, \ldots, w^o_K$, to best match the treatment unit during the pre-treatment period, using an RMSPE criterion. What was done before, then, was to sum up the weights for the donor states in the same division, $W^s = \Sigma_j w^s_j$, and ask whether this weight was large.

In the ranked prediction error approach, we calculate the RMSPE separately for each potential donor, for the same-division and other-division donor states respectively (denoted r^s_j, j = 1, …, J, or r^o_k, k = 1, …, K). These r's are then pooled, and ordered based on how well they match the treatment unit. Finally, a percentile in this ranking is assigned to each donor unit, denoted P_m, m = 1, …, (J + K), where the highest rank is given to the donor with lowest r.[25]

The percentile assigned to a donor state is defined as the percentage of donor states with a

---

[23] The RMSPE here refers to the root-mean-squared prediction error when using the donor observations to predict the treatment observations for the four quarters prior to the minimum wage increase.

[24] The treatment unit is a particular state in a particular quarter; the time subscript is omitted.

[25] For example, if there are 50 donor units, then the unit with the lowest RMSPE gets a rank of 50. The Weibull rule is used to convert ranks to percentiles. With N units, the percentile is (100 × rank)/(N+1).

higher RMSPE – i.e., a worse match. Thus, a percentile of 100 (or near 100 with a smaller number of states) would imply that a particular donor unit provides the best match. A percentile near zero would imply that it provides the worst match. And a percentile near 50 would suggest that it provides about as good a match as a randomly-chosen control unit.

If ADR are correct that same-division states provide better controls than states in other divisions, then the percentile ranking should be higher, on average, for states in the same division as a treatment unit than for states in other divisions. To test this, the percentiles for same-division states are collected after doing this analysis for every possible treatment unit and the associated matching variables (exactly as in the synthetic control analysis), and histograms for these percentiles are constructed to see if the same-division states are clustered at higher percentiles than would be expected if these states were, on average, no better or worse controls than other states – or, equivalently, if the distributions of the percentiles appear approximately uniform.

To help explain, Figure 1 displays an example of the first step of this process for one treatment unit – California in 2001:Q1. The potential same-division donor states are Alaska, Hawaii, and Oregon; Washington is also in the Pacific division but had a minimum wage increase in the same quarter. For each of the four matching variables, the corresponding figure is the histogram of RMSPEs for all potential donor states, with the three same-division states highlighted with the thin vertical lines that extend to the top of the box. As the figure reveals, states within the same division can provide quite good matches, with low RMSPEs relative to other states (e.g., Alaska in the upper-right panel), or they can provide quite bad matches, with relatively high RMSPEs (e.g., Hawaii in the upper-left panel).[26]

---

[26] The figure also shows that, for a particular minimum wage treatment, whether a particular state provides a good

{{Place Figure 1 about here}}

Figure 2 then presents the analysis aggregating across all of the treatment units, plotting

the histogram for the percentiles in the RMSPE distribution for each same-division state that ever

appears as a potential donor in this analysis. The figure indicates no tendency of these

percentiles to be clustered towards the upper end of the distribution for any of the matching

variables. Instead they look much closer to uniform, and the medians are around 50. The

implication is that the same-division control states are, on average, no better than the control

states from other divisions, contrary to ADR's identification strategy.

{{Place Figure 2 about here}}

We also examined the medians of these percentiles for each of the Census divisions. The

only division where other states in that division consistently stand out as generally providing the

best controls – in the sense that the percentile rankings are above the median – is the West North

Central region. Looking back at Table 3, notice that the synthetic control approach also

indicated a high weight on same-division states *only* for this division. Thus, the two analyses

lead to a qualitatively similar conclusion: Both raise doubts about the validity of ADR's

restriction confining identifying information to states in the same division, with one notable

exception – the West North Central division.

Finally, it is instructive to compare the division-specific minimum wage estimates in

Table 2 with the results from the synthetic control or ranked prediction error approach. Both

approaches show that the only Census division for which there is a strong indication that most of

the control states should come from within the division is West North Central. Table 2 shows

that when minimum wage effects are estimated for this division in isolation, there is statistically

---

match can depend on the matching variable, not surprisingly. This is much less of an issue for the aggregated
analysis below.

significant evidence of a negative employment effect of minimum wages, with an estimated

elasticity of −0.19, very much in line with much of the existing evidence on minimum wages.

Furthermore, Table 3 also indicates that New England and the Pacific regions assign non-

negligible weight to states in the same region. Of the two, the estimates in Table 2 for New

England are precise, as noted above, and these estimates also point to negative employment

effects (with a larger disemployment effect). In contrast, Table 3 indicates that especially for the

matching on residuals – which seems most pertinent to ADR's argument – states in the same

division get essentially no weight for the Middle Atlantic and East North Central divisions.

These are two cases that, in Table 2, do not provide any evidence of disemployment effects.

Thus, while this analysis does not pin down one "best" estimator, it does indicate that (a) in most

cases, there is little rationale for ADR's choice to focus only on the within-division variation to

identify minimum wage effects; and (b) when there *is* a good rationale for doing this, the

evidence shows negative and statistically significant effects of minimum wages on teen

employment, with elasticities that are in or near the −0.1 to −0.2 range.

Dube et al. (2010)

As noted earlier, DLR focus on restaurant employment with county-level QCEW data

from 1990-2006. They show that the panel data model with county and period fixed effects

yields negative employment effects, with elasticities in the conventional range, whereas these

effects become small and insignificant when either state-specific linear time trends or Census

division × quarter interactions (or both) are added. As noted above, the inclusion of these

additional controls is problematic.

However, the main focus of DLR is on a research design based on cross-border county

pairs. When they include unique dummy variables for cross-border contiguous county pairs

interacted with period, they identify the effect of minimum wages from differences in

employment changes in these paired counties on either side of a state border – using the within-county-pair variation in the same way that including division × period dummy variables in the specifications in ADR relies on the within-Census division variation. Given that this identification strategy is the key contribution of this paper, we focus on their cross-border analysis of the effects of minimum wages on restaurant employment.

The key estimates from this approach are reported in Table 4, replicating the results in DLR (Table 2, specifications 5 and 6); the estimates are nearly identical to theirs. The first two columns use the balanced panel of the subset of counties in the contiguous border county-pair analysis, but include only county and period (quarter) fixed effects. As in DLR, two specifications are reported – with and without a control for total private-sector employment. The estimated minimum wage effects on restaurant employment are negative and in the old "consensus range," with the first significant at the ten-percent level. In columns (3) and (4), county-pair × period interactions are added to replicate DLR's method of controlling for spatial heterogeneity. As the table shows, the estimated minimum wage effects are slightly positive and statistically insignificant.[27]

{{Place Table 4 about here}}

The main question concerns the underlying assumption in DLR's identification strategy – that the cross-border contiguous county in the bordering state that did not raise its minimum wage is the best control for the county in the state that did raise its minimum.[28] As they point out, this has close parallels to the type of analysis in Card and Krueger (1994), who studied the

---

[27] As we show in Neumark et al. (2013), DLR substantially overstate the number of cross-border county pairs that are used to identify the effects of minimum wages in their approach. Their Figure 2 claims to show all the state borders – and counties along them – that are used in their analysis. However, many of the borders highlighted in this figure are for pairs of states that did not have a minimum wage higher than the federal minimum during their sample period, and their sample actually includes only 48 distinct state border pairs with identifying information, rather than the 81 pairs depicted in their figure.

[28] For some treatment counties, there are multiple cross-border contiguous counties.

effects of the 1992 minimum wage increase in New Jersey by comparing employment changes in the fast food industry in that state to areas in Pennsylvania – where the minimum wage stayed fixed – on or near the border with New Jersey.

There is some intuitive appeal to the idea that cross-border counties are good controls because of their geographic proximity. And we might expect that, on a priori grounds, the case for using contiguous counties as controls is stronger than for using states in the same Census division (as in ADR).[29] Nonetheless, it is not obvious, without evidence, that cross-border counties are appropriate controls. For example, spillover effects can certainly contaminate the control observations. If workers displaced by the minimum wage find jobs on the other side of the border, employment will expand in the control areas; this, however, implies that we should find larger disemployment effects when restricting attention to cross-border controls. The same would apply to price effects, if higher prices in the county with the minimum wage led to higher product demand in the cross-border county. But there are other possibilities. Workers could cross into the state that raises its minimum wage to search for jobs (à la Mincer (1976)), reducing employment in the cross-border state as well as in the state where the minimum wage increased, muting the effects. And, in principle, price-setting could also be affected: when the minimum wage goes up across the border, firms may raise prices because the elasticity of product demand has effectively fallen. Moreover, geographic proximity does not necessarily imply that cross-

---

[29] One might also reasonably ask why ADR use same-division states as controls rather than cross-border states, paralleling the DLR paper that was published earlier. Interestingly, as reported in Neumark et al. (2013), if we estimate the model using CPS data adding interactions between dummy variables for pairs of bordering states and quarters, rather than interactions between dummy variables for divisions and quarters, the estimated minimum wage effect is unchanged relative to the panel data estimator with fixed state and period effects. In particular, adding the bordering state-pair × quarter interactions yields an estimate of −0.162 (standard error of 0.113) – very close to the panel data estimate of −0.168 (0.046), albeit less precise. (Note that the panel data estimate differs slightly from that shown in Table 1, Panel A, column (1) because states can appear multiple times when they border multiple states, just like in DLR's county-level analysis.) On a priori grounds the same-border analysis seems more compelling than the same-division analysis, if there is, as ADR/DLR believe, important spatial heterogeneity. After all, states in the same division often do not share borders and can be quite far apart.

border counties experience the same shocks,[30] especially given that the relevant shocks in DLR are those that affect restaurant employment conditional on aggregate economic activity.

The bottom line is that the assumptions underlying the vast narrowing of potential controls in DLR's analysis should be tested. Is there evidence in support of DLR's assumption that cross-border contiguous counties provide appropriate control groups? As before, we address the question of the quality of cross-border contiguous counties as controls using the synthetic control matching estimator – this time to calculate the weight that the matching puts on the contiguous cross-border counties relative to the weight it puts on other potential control counties.

The analysis exactly parallels the state-level analysis. Potential donors to the control group are all counties in the states that were identified as potential donors in the previous analysis.[31] The estimator is then implemented, and the weights put on the cross-border control counties that DLR actually use are computed. The criteria for defining treatments and controls are the same as before, but now done at the county level. In particular, the set of treatment counties are border counties with a minimum wage increase and where there was no minimum wage increase in the previous four quarters. Potential donor units are county-quarter observations with no minimum wage increase in the same quarter and the succeeding three quarters, and similarly no minimum wage increase in the previous four quarters. Two different analyses are done. The first includes all potential donor counties. In the second analysis, we restrict the set of donor counties on which the synthetic control calculation has to match. In particular, we first calculate, for each treatment and the potential donor counties, the RMSPE of

---

[30] Addison et al. (2009), who do a similar analysis, acknowledge this issue. Although they still use this method to estimate minimum wage effects, they give an example of a cross-border county match that is quite bad, with a 3.5% unemployment rate for one treatment county and a 7.7% unemployment rate for the contiguous cross-border county, and they suggest that "such examples of poor matches across state borders could be rather common" (p. 406).
[31] The set of potential donors is restricted to the counties in the balanced sample of counties (with non-missing employment data) that DLR use.

each donor county for the four quarters prior to the minimum wage increase. We then use the 50

donor counties with the lowest RMSPE as potential donors, adding in DLR's contiguous cross-

border counties if they are not already in this set of 50.

The match is done on the same types of variables as before defined over the four previous

quarters: regression residuals (from a regression of the log of the ratio of restaurant employment

to county population on the log of the minimum wage, the ratio of total private-sector

employment to county population, and state and period (quarter) dummy variables); the log of

county restaurant employment relative to county population; and the one-quarter and four-

quarter differences in logs of restaurant employment relative to county population.[32] As before,

in some cases there were no potential contiguous cross-border donor units for a county, and these

cases are thrown out.

The results are reported in Table 5, and the analysis reveals that the weight assigned to

the cross-border contiguous counties as controls – the *only* controls DLR use – is very small. In

Panel A, where there are typically just over 50 possible controls for each treatment, the weight

on the cross-border contiguous counties is, on average, 0.021 in column (1), and of a similar or

smaller magnitude in the other columns. (As for ADR, we view the results for the regression

residuals as most pertinent to the critique in DLR.) The table also shows the percentiles of the

distribution of weights. The distribution is highly skewed. For example, in column (1) of Panel

A, the median is only 0.008, compared with the mean of 0.021, and the 75[th] percentile is below

the mean (0.019). In Panel B (also column (1)), the 90[th] percentile is below 0.01.

{{Place Table 5 about here}}

---

[32] Note that this dependent variable differs from that used in Table 4, by dividing by county population. We do this because we do not want to match on levels of county restaurant employment, which can vary tremendously across counties.

Going through a similar calculation to what we did for states, for the analysis in column (1), in only 16 of 121 cases are the average weights per contiguous cross-border county higher than the random threshold of 1/(number of potential donors). For columns (2)-(4), the corresponding numbers are 13, 19, and 13 (these last two are relative to 114 and 109 unique minimum wage treatments, owing to some loss of observations from the lagged variables). We reach the same conclusion if we do these computations for the restricted set of donors in Panel A; indeed in this case the evidence is even stronger that the contiguous cross-border counties do not provide better controls. Thus, Table 5 and these computations indicate quite clearly that for almost all treatment counties, the cross-border county is a poor match – no better than a county chosen at random from the list of all potential comparison counties. Given this evidence, it seems difficult to argue that throwing out the information on other potential comparison counties, as DLR do, is necessarily preferable to using the panel data estimator.

Because there is some *a priori* appeal to the idea that cross-border counties are useful controls (even if the evidence to this point does not bear this out), we also looked more closely at cross-border counties for MSAs that straddle two states, to see whether within these more-integrated labor markets the cross-border counties were better controls. Out of our 121 unique minimum wage treatments considered in Table 5, 31 have cross-border controls that are in other states but in the same MSA.[33] We compared the weights we calculated for the cross-border counties in this subsample to the weights implied by a randomly-chosen county. Specifically, if there was only one cross-border county in the MSA we compared the weight on that county to 1/(number of potential donors), and if there was more than one cross-border county we averaged the weight across these counties and did the same comparison. In only five cases out of 120

---

[33] Examples are Nez Perce County, ID, in the Lewiston, ID-WA MSA, and Westchester County, NY, in the New York-Newark-Edison, NY-NJ-PA MSA.

possible matches was there a much higher weight for the cross-border county; three of these were in the New York-Newark-Edison, NY-NJ-PA MSA, and two were in the Lewiston, ID-WA MSA. In seven other cases the average weight for the cross-border county was higher than for a randomly-chosen county, but by less than 0.001 (relative to within-MSA average weights ranging from 0.001 to 0.019). Thus, even within MSAs there is no evidence that cross-border counties provide better controls.

We also used the ranked prediction error approach to assess whether contiguous cross-border counties are better controls than other counties. The method is the same as before, but now examines the percentiles for the matched contiguous cross-border counties (analogous to the previous examination of same-division states). Figure 3 shows the percentiles for these matched counties. There is perhaps a slight tendency for these percentiles to be clustered above 50, but, in general, the histograms seem fairly close to uniform. Again, there is little in the data to support the assumption made by DLR that the contiguous cross-border counties are the appropriate controls.

{{Place Figure 3 about here}}

One might argue that it is not surprising that in our synthetic control analysis the contiguous cross-border counties get so little weight or fail to stand out as the best control areas. After all, there are typically a very large number of potential donor counties for any one county's minimum wage increase. But that is precisely the point: With the large set of potential donor counties, why throw away so much potential identifying information without assessing which counties are in fact the best controls?

Nonetheless, one might be concerned that when there is a very large number of potential donors, noise in the data will cause the weight assigned by the synthetic control method to

"valid" controls (the contiguous cross-border counties, according to DLR) to be smaller than the weight assigned to the "invalid" controls. That is, with a large number of potential donor counties, a weighted average of "invalid" controls picked out by the synthetic control method could provide a better fit to the pre-treatment observations on the treated county – even if these invalid controls would not be good predictors of the post-treatment "counterfactual." However, we have already seen two types of evidence indicating that this is not what is driving our results. First, we found similar results from the synthetic control analysis for states as for counties, even though there are far fewer potential donors in the state-level analysis. Similarly, the analysis of the 50 counties with the best fit led to the same conclusion as the analysis of all counties. Second, the ranked prediction error approach does one-to-one comparisons between each control that ADR and DLR use and other potential controls, so nothing in this approach leads to putting inordinate weight on what – according to ADR and DLR – would be invalid controls.

Is There Other Evidence that Justifies the Approach in ADR and DLR?

ADR and DLR present additional analyses that are intended to show that their identification strategy is valid and that the more conventional panel data approach, which uses a much broader set of controls (states or counties), leads to spurious evidence of negative minimum wage effects because of spatial heterogeneity. One key analysis in this regard is pitched as a falsification test showing that county employment appears to respond to cross-border minimum wage changes.[34]

In particular, they define a narrow sample of all border counties where the minimum

---

[34] ADR and DLR also present some evidence suggesting that the panel data specification with only period and state or county fixed effects produces leading effects of the minimum wage on employment, which they argue provides "strong evidence against the model without controls for heterogeneity across states …" (p. 220). However, as we show in Neumark et al. (2013), this evidence is both overstated and misleading. In particular, the estimates that ADR and DLR emphasized in their paper are the ones that most strongly make their case, and we highlight that there are many more equally plausible analyses of the issue of pre-trends that produce much weaker evidence, and indeed sometimes support the opposite conclusion.

wage was *never* above the federal minimum wage in the sample period, and then estimate the panel data specification with county and period fixed effects for this sample, substituting the cross-border counties' minimum wages. When they do this, their estimated "placebo" minimum wage effect is negative, albeit smaller than its standard error, and it is about 60 percent of the estimated minimum wage effect for the counties bordering the placebo sample but using their actual minimum wages.[35] These estimates are replicated in column (1) of Table 6.

{{Place Table 6 about here}}

However, this is not a valid falsification test. For most county pair-quarter observations in the sample they use (96 percent), both the cross-border minimum wage and the own-county minimum wage are the same – equal to the federal minimum wage. Thus, in most cases the placebo minimum wage assigned to the county is equal to the actual minimum wage prevailing in the county, which of course can affect employment.[36] In other words, DLR assume that the null hypothesis of no spatial heterogeneity implies that the effect of the placebo minimum wage is zero, and then reject this null because their estimated placebo minimum wage effect is negative. But because the "placebo" minimum wage they use is often the same as the actual minimum wage, we would expect a negative minimum wage effect in their placebo analysis even if there is no spatial heterogeneity – invalidating their falsification test.

Confirming this problem, we do not find a placebo effect when the sample used for DLR's falsification test is modified to avoid having a contaminated placebo sample. Specifically, we restrict the sample to observations after the federal minimum wage increase in

---

[35] See DLR's Appendix B for further information on how they implemented their falsification test.
[36] Note that the effect of federal minimum wage variation is still identified in their placebo sample when county and period fixed effects are included as long as the federal minimum wage is not binding in some states. The minimum wage change induced by federal variation will vary across placebo counties depending on the level of the state minimum wage in the cross-border county.

1997, so that there is no federal minimum wage variation in the placebo counties that is captured by the counties matched to them.[37] As shown in column (2) of Table 6, in this case the estimated minimum wage effect in what DLR term the "actual minimum wage" sample is large, negative, and statistically significant, while the estimate for the placebo sample is much smaller and statistically insignificant.[38] In this placebo sample, there are still many counties paired with cross-border counties that have the same federal minimum wage (although now it does not vary); only 7 percent of the county pair-quarters have a minimum wage difference. It is possible to further restrict attention to an even more informative placebo sample by focusing on county pairs where there is at least one minimum wage difference in this sample period between the true minimum wage in the placebo county and the cross-border minimum wage that is used in the falsification test; after all, it is this variation that is informative about their falsification test.[39] As shown in column (3), in this case the estimated minimum wage effect in the placebo sample falls to zero, while the estimated minimum wage effect in the actual minimum wage sample is little changed. These estimates provide additional evidence refuting DLR's claim that spatial heterogeneity generates spurious evidence of disemployment effects of minimum wages.

What is the Right Estimate of the Minimum Wage Effect?

The state-level panel data estimate of the employment effect of the minimum wage for teenagers, using state and period fixed effects, is negative and significant, with an estimated elasticity of −0.165. In contrast, the alternative estimates that ADR report are insignificant, and

---

[37] Their sample ends before the most recent round of federal increases beginning in 2007. In this exercise, we use data beginning in 1998:Q3, one year after the last federal minimum wage increase, to avoid lagged effects of federal minimum wages. But the results were very similar if the sample starts in 1997:Q4, the first quarter after the last federal increase.

[38] The standard errors in both samples are a good deal smaller, likely because there is much more state minimum wage variation in the latter part of the sample.

[39] This restricts the comparisons between cross-border county pairs with and without a minimum wage difference along their shared state border to a homogeneous set of state borders that are sometimes observed with a cross-border minimum wage difference, and sometimes without a difference.

often close to zero. The same is approximately true for the estimates for restaurant workers from the panel data model with fixed county and period effects and those reported by DLR. We have shown that: (i) the assumptions underlying the approaches used by ADR and DLR are not supported by the data; and (ii) in the isolated cases when the assumptions ADR use are supported by the data, the estimates of the employment effects of minimum wages from the panel data specification with fixed state and period effects are again negative and statistically significant. In addition, we have shown that the evidence that ADR and DLR report to argue that this panel model is misspecified and leads to negatively-biased estimates is overstated and misleading. In and of itself, this evidence indicates that there is no reason to discard existing panel data estimates that lead to minimum wage employment elasticities very much in line with the earlier consensus.

But having raised the issue of what the best controls *are* for the states that raise their minimum wage, and having suggested how we can use the data to answer this question, it is of interest to ask whether we can extend these methods to arrive at an estimated elasticity that uses the data to construct the controls. In this section, we discuss such an approach and the resulting estimates. We do not claim to have completely resolved the econometric issues involved, but we do think our approach has some intuitive, heuristic appeal.

Before plunging into the details, note that we have already presented some casual evidence, with regard to the analysis in ADR, that using the information on the weights from the synthetic control method leads to evidence of negative employment effects of minimum wages. Specifically, Tables 2 and 3 showed that for the Census division for which the synthetic control approach strongly indicated that states in the same Census division provide valid controls (West North Central), the within-division estimate of the employment elasticity is −0.19, and

statistically significant. Here, we suggest an approach that makes use of the entire data set incorporating information on appropriate controls from the synthetic control approach.

Our approach to some extent follows Autor, Donohue, and Schwab (2006). In studying the effects of the adoption of wrongful-discharge laws on employment, they restrict attention to a subset of observations across states and over time, including some periods before and after the adoption of these laws in some states, and the same periods for other states that did not adopt such laws.[40] The idea is to identify the effects of laws from relatively short-term changes in treatment states relative to control states. We adopt the same approach here. Using the same window from the earlier synthetic control analysis, we focus on four-quarter treatment periods defined by the state and quarter in which the minimum wage is increased. Given that we want to incorporate information from the synthetic control analysis discussed earlier, we use the period four quarters prior to a minimum wage increase as the pre-treatment period.

Specifically, we first define the set of potential treatment observations as states and quarters when the minimum wage increased, and the four preceding and three succeeding quarters for such states. (We discuss our methods in the context of the state $\times$ quarter analysis in ADR; it carries over fully to the county $\times$ quarter analysis in DLR, to which we return below.) Each observation in this sample is characterized by the employment rate, the minimum wage, and the control variables, $O_{it} = (E_{it}, MW_{it}, X_{it})$.[41] We want to construct a control or counterfactual observation for each state and quarter in this subsample, defined as $C_{it} = (E^c_{st}, MW^c_{st}, X^c_{st})$, where the c superscript indicates that the data represent an average, possibly weighted, of other states that serve as controls; C_{it} could also be a single state. The observations

---

[40] They actually study three different kinds of laws, but this does not alter the basic approach.
[41] Note that this implies that when there are minimum wage increases in a state fewer than four quarters apart, observations on treatment states can be repeated, unlike in the panel data model.

on O_{it} and C_{it} corresponding to each experiment span eight quarters, as just mentioned; denote the vectors corresponding to these eight quarters for any particular minimum wage increase (experiment "e") as O_i^e and C_i^e. We will return in a moment to how C_i^e might be constructed. The data set is now doubled in size, and includes all of the original observations for an identified set of treatments, plus a counterfactual observation for each original observation.

Stacking the data to combine these observations, we can then run the minimum wage employment regression on all of these observations. In these regressions, we include quarter dummy variables as in the panel data model. In addition, we include dummy variables for each set of eight observations on a treatment state in each unique experiment e (O_i^e), and each set of 16 observations on either the treatment states or the counterfactual observations in each unique experiment (O_i^e and C_i^e). The latter sets of dummy variables allow the counterfactual observations to have their own intercepts in each experiment, whether or not the composition of the counterfactual states varies for different minimum wage increases for the same treatment state. In addition, the experiment-specific treatment state dummy variables allow the level of the dependent variable to differ between the treatment and counterfactual states.

The key question, of course, is how to construct the counterfactual observations. The usual panel data estimator effectively treats each state as an equally valid control. Thus, with the only difference that we now restrict attention to the period of a minimum wage increase and the four preceding and three succeeding observations for states with an increase, the natural analogue is to construct the counterfactual as the equally-weighted average of other states.[42] Given that the sample includes the 50 states and

---

[42] This does not exactly parallel the panel data estimator when the data are weighted. However, as explained below, we are ultimately going to use the observations weighted by the synthetic control weights to construct the counterfactual observation, in which case one does not want to use population weights but rather let the data dictate which states – large or small – best match the pre-treatment observations. Having constructed the observations, though, in the regression estimation we use teen population to weight the treatment observations; we apply the same weights to the corresponding counterfactual observations, because these are supposed to estimate what would have happened in the treatment observations absent the treatment.

Washington, DC, these weights would equal 1/50 if we do not weight the state observations.

In contrast, the analysis in ADR effectively restricts the identifying information to states in the same Census division (by including Census division × period interactions). We can accommodate this restriction in our approach by constructing the counterfactual observations only from the same-division states, with weights equal to one divided by the number of states in the division.

Finally, we can also use the weights obtained from the synthetic control analysis to construct the counterfactual observations. To do this, we first consider the 129 "clean" minimum wage treatments and their associated controls that we used in the analysis in Table 3, consisting of state-by-quarter observations with average quarterly minimum wage increases of at least 5 cents, with no increase in the previous four quarters, and controls with no minimum wage increases in the previous four quarters or the next three.[43] It would be ideal if, using our matched panel data estimator, these minimum wage increases were similar to those from the usual panel data estimator, since we would then have a baseline estimator for which we could vary the weights on counterfactual observations and observe the implications for the estimated employment effect. As it turns out, however, this subset of minimum wage increases (there are a total of 544 minimum wage increases in our sample period) appears to be unusual, in that it does not generate a significantly negative minimum wage effect using the modified panel data estimator described above.

We therefore instead use the synthetic control approach for *all* of the minimum wage

---

[43] In that table, because we were focusing on the validity of restricting controls to states in the same Census division (as in ADR), we used a subset of 50 of these 129 minimum wage increases where there was at least one potential control in the same division.

increases in the sample period.[44] This means that the counterfactual observations can include

observations with minimum wage increases. However, because these counterfactual

observations contribute to both the estimated employment rate and the estimated minimum wage

variables (as well as the other controls), this is not problematic. When there is less minimum

wage variation between the treatment and counterfactual observations, we would

correspondingly expect less of an employment response (if there is one). Note that this is no

different from what ADR do when they focus on minimum wage variation within the same

Census division, in which case all of the states in the same division serve as potential controls,

contributing no identifying information if the minimum wage change is the same, and less

identifying information when the minimum wage change is similar but not equal to the treatment

state. Because we want to weight the counterfactual observations by how well they match the

treatment observations in the four quarters prior to the minimum wage increase, and the

counterfactual observations can now have minimum wage increases of their own, we use the

synthetic control weights that we get from the matching on residuals; these residuals will be net

of minimum wage effects and the effects of the other controls.

There are three econometric issues with our matched panel data estimator that are

potentially problematic. First, there is a potential circularity here, because we need an estimated

minimum wage effect to compute these residuals, and that effect is in dispute. We therefore

compute residuals in two ways – first using the estimated minimum wage effect from the panel

data estimator with fixed state and period effects, and then restricting the minimum wage effect

---

[44] We actually exclude the minimum wage increases occurring near the beginning of the sample period for which there are not 4 quarters in the pre-treatment period that can be used in this analysis. In addition, we also exclude minimum wage increases occurring near the end of the sample period for which there are not 3 quarters after the treatment quarter; however, including the latter set of increases and constructing the counterfactual observations using the available quarters did not affect the results.

to be zero. The method of computing the residuals turns out to have little effect on the estimates, although we recognize that we have not determined a way to simultaneously estimate the minimum wage effect and the weights on the counterfactual observations. Second, there is no way to cluster the standard errors at the state level anymore, because the counterfactual observations can be made up of different states. We therefore instead cluster the standard errors at the level of either the treatment state or the counterfactual observation in each experiment.[45] Third, our counterfactual observations are constructed based on estimated weights, and we have not accounted for this estimation in the construction of the standard errors, which are therefore downward biased. However, the main issue concerns the point estimates, and as we will see, our matched panel data estimator based on the synthetic control weights gives point estimates very similar to those of the panel data estimator with fixed state and period effects.

The results of this analysis are reported in Table 7. For purposes of comparison, column (1) repeats the panel data estimates with fixed state and period effects. Column (2) shows the estimates using the 129 "clean" minimum wage experiments that were the starting point for the Abadie et al. (2010) synthetic control analysis; the states used to construct the counterfactual are equally weighted. One can view this as applying the Autor et al. (2006) methods to this subset of minimum wage increases. In this case the estimated minimum wage effect falls and is not statistically significant. Thus, for this subset of observations, the matched panel data estimator does not provide a useful baseline for contrasting the effects of different weighting on the estimated employment effect of minimum wages, because for these observations the estimate does not mimic the conventional panel data estimator. In contrast, column (3) reports estimates using the same methods, but incorporating all of the minimum wage increases (except those

---

[45] Note that this implies that the treatment states are clustered at a lower level than in the panel data case where we cluster by state. We do this to be symmetric with the counterfactual observations.

occurring near the beginning or ending of the sample) and all of the control states; this is most like the panel data estimator in that all of the minimum wage variation is used. The estimated elasticity is now larger, at −0.122. The estimates in column (3) effectively serve as our baseline for comparing the estimates that result with different weighting of potential control states to construct the counterfactual observations.

{{Place Table 7 about here}}

In column (4), we mimic what ADR did by including the period × division interactions, which restricts the identifying information to minimum wage variation within Census divisions. In our set-up, we do this by using only the states within the same division (equally weighted) to construct the counterfactual observations. Paralleling their findings, the minimum wage effect is diminished and is no longer statistically significant. Of course we have called into question this restriction, because the synthetic control analysis indicated that same-division states were not the appropriate control states.

Finally, we turn to the estimates that use the weights from the synthetic control method. Column (5) reports estimates from a procedure in which the residuals used for the matching and computation of weights are based on the panel data model including the minimum wage. The estimated minimum wage effect is noticeably stronger – an elasticity of −0.143. Column (6) reports estimates when the minimum wage effect used to construct the residuals is restricted to be zero – more in line with the conclusions of ADR and DLR. As the table shows, this restriction has almost no effect on the estimated elasticity of teen employment with respect to the minimum wage, which is now −0. 145. Thus, the evidence indicates that when we let the data determine the appropriate control states to use for estimating the effects of state minimum wage increases, we find evidence of disemployment effects, with teen employment elasticities near

−0.15.[46]

In Table 8 we report estimates from the same approach applied to the QCEW data that DLR use. We report both unweighted and weighted estimates (by county population). DLR do not weight their county-level estimates, although the estimates in ADR, because they are at the individual level, are effectively weighted by state population. Column (1) reports panel data estimates with fixed county and period effects, which yield elasticities in the −0.1 to −0.2 range. Column (2) restricts the sample to the paired (contiguous cross-border) counties that are the focus of DLR (although not of our analysis in Table 5). The estimated employment effects of minimum wages become smaller, and insignificant for the unweighted case. Column (3) reports our Autor et al.-type estimator, restricting attention to the quarter of each minimum wage increase and the four preceding and three succeeding quarters for all paired border counties, and constructing the counterfactual observation as the equally-weighted average across all counties in other states; we restrict attention to the bordering counties so that we can also estimate the counterfactual observations using only the contiguous cross-border counties. The estimates fall to zero or near-zero. In other words, for the county-level analysis we are not successful at finding an Autor et al.-type panel data estimator that matches the panel data estimator, likely reflecting the combination of the restricted sample and the small window surrounding the minimum wage increase, which eliminates identifying information from long differences. [47]

{{Place Table 8 about here}}

Nonetheless, column (4) mimics the DLR research design, weighting the counterfactual

---

[46] We also did this estimation without using population weights for the treatment and corresponding counterfactual observations. The pattern of how the estimated minimum wages effects change across the columns was similar. But all of the estimated minimum wage effects were larger, so that they were negative and significant in columns (3) and (4). For what we regard as the best estimates – in columns (5) and (6) – the estimated elasticities were −0.19.
[47] Other work suggests that these longer differences may be relevant for identifying adverse effects of minimum wages (Baker, Benjamin, and Stanger 1999).

observations to use only the paired cross-border counties as controls. Both the unweighted and weighted estimates increase slightly. Finally, in columns (5) and (6) we report the estimates using the synthetic control weights (based on the "top 50" donors as in Panel A of Table 5). Here, however, we use all counties in other states as potential controls, rather than just the border counties. For the unweighted results the estimates are small and near zero. The weighted estimates, in contrast, are small, but negative and statistically significant, with elasticities of around −0.05 or −0.06.

What do we conclude? First, we clearly do not obtain as strong evidence of disemployment effects for restaurant employment in the QCEW, when using the synthetic control weights, as for teen employment in the CPS. As noted earlier, this is not surprising, nor is it in any way contradictory with the existing literature. Most of the existing evidence focuses on teenagers or other low-skill groups, rather than on any particular industry; and at the industry level, labor-labor substitution seems more likely to mask gross disemployment effects for the least-skilled. This difference in findings may also stem, in part, from difficulties in applying our method to the QCEW data. Second, and perhaps more significantly, when we weight the estimates we find some evidence of disemployment effects when the synthetic control weights are used. These are, arguably, the most defensible estimates.[48]

Conclusions

Throughout the long-running debate about the employment effects of minimum wages, the empirical evidence has focused on similar questions: How does a minimum wage affect employment? Which workers are affected? And how do we ensure that we are getting a valid comparison that isolates the effect of the minimum wage?

---

[48] We do not find these negative estimates if we exclude the private-sector employment control. But that exclusion does not seem defensible.

Given the ongoing ebb and flow of this debate, it would have been shortsighted to think that the 2008 book that two of us wrote (Neumark and Wascher 2008), despite surveying a massive amount of evidence, would have settled the issue. And indeed it has not. In particular, echoing long-standing concerns in the minimum wage literature, Dube et al. (2010) and Allegretto et al. (2011) attempt to construct better counterfactuals for estimating how minimum wages affect employment. When they narrow the source of identifying variation – looking either at deviations around state-specific linear trends or at within-region or within-county-pair variation – they find no effects of minimum wages on employment, rather than negative effects. Based on this evidence, they argue that the negative employment effects for low-skilled workers found in the literature are spurious, and generated by other differences across geographic areas that were not adequately controlled for by researchers.

The analysis in this paper, however, provides compelling evidence that their methods are flawed and lead to incorrect conclusions. In particular, the methods advocated in these studies do not isolate more reliable identifying information (i.e., a better counterfactual). In one case – the issue of state-specific trends – we explicitly demonstrate the problem with their methods and show how more appropriate ways of controlling for unobserved trends that affect teen employment lead to evidence of disemployment effects similar to that reported in past studies. In the other case – identifying minimum wage effects from the variation within Census divisions or, even more narrowly, within contiguous cross-border county pairs – we show that the exclusion of other regions or counties as potential controls is not supported by the data.

We think the central question to ask is whether, out of their concern for avoiding minimum wage variation that is potentially confounded with other sources of employment change, ADR and DLR have thrown out so much useful and potentially valid identifying

information that their estimates are uninformative or invalid. That is, have they thrown out the "baby" along with – or worse yet, instead of – the contaminated "bathwater"? Our analysis suggests they have. Moreover, despite the claims made by ADR and DLR, the evidence that their approaches provide more compelling identifying information than the panel data estimates that they criticize is weak or non-existent.

In addition, when the identifying variation they use *is* supported by the data, the evidence is consistent with past findings of disemployment effects. Moreover, when we let the data determine the appropriate control states to use for estimating the effects of state minimum wage increases in the CPS data, we find strong evidence of disemployment effects, with teen employment elasticities near −0.15. The findings from similar analyses of restaurant employment in the QCEW data are a bit more mixed, but the weighted estimates again point to negative employment effects (with smaller elasticities of around −0.05). Thus, our analysis substantially undermines the strong conclusions that ADR and DLR draw – that there are "no detectable employment losses from the kind of minimum wage increases we have seen in the United States" (DLR 2010, p. 962), and that "Interpretations of the quality and nature of the evidence in the existing minimum wage literature …, must be revised substantially" (ADR 2011, p. 238).

Can one come up with a dataset and an econometric specification of the effects of minimum wages on teen and low-skilled employment that does not yield disemployment effects? As in the earlier literature, the answer is yes. But prior to concluding that one has overturned a literature based on a vast number of studies, one has to make a much stronger case that the data and methods that yield this answer are more believable than the established research literature, and convincingly demonstrate why the studies in that literature generated misleading evidence.

Our analysis demonstrates that the studies by Allegretto et al. (2011) and Dube et al. (2010) fail to meet these standards. As a result, we continue to view the available empirical evidence as indicating that minimum wages pose a tradeoff of higher wages for some against job losses for others, and that policymakers need to bear this tradeoff in mind when making decisions about increasing the minimum wage.

We also believe that a similar set of issues to those we consider carries over to the analysis of essentially any kind of policy that might be studied with panel data on geographic regions over time. When doing these kinds of panel data studies, researchers often make the same choices as ADR and DLR – such as including state-specific linear time trends, or narrowing the scope of the geographic areas used for controls by either restricting the sample or estimating a more saturated model that reduces the identifying information to variation within a smaller region. Our evidence suggests that these kinds of analyses, even if well motivated, can deliver misleading evidence of either the presence or absence of effects. We *do not* advocate ignoring the potential biases introduced by differences in the regions where policies are enacted. We *do*, however, advocate using the data to explore more fully what specifications provide the most reliable counterfactuals, and we discuss some methods for doing this. After all, in other contexts – such as instrumental variables estimation – we generally ask hard questions about the validity of the identifying assumptions that are used.

In particular, if these kinds of sensitivity analyses deliver robust results that are insensitive to detrending or to the narrowing of identifying information by restricting the set of control areas, then they can clearly bolster the evidence. If, however, they point to different evidence, then the researcher has to seriously explore which analysis is most convincing. In the case of removing long-term trends from panel data over time, we have suggested methods that

increase the likelihood that business cycle movements are not confounded with long-term trends. And in the case of restricting the set of control areas, we have shown how to obtain evidence on which areas are better controls, and have suggested a way to incorporate this evidence in estimating policy effects. We believe these kinds of approaches – and others more appropriate to different types of analyses – should be incorporated into what can otherwise be a somewhat blind approach to sensitivity analysis. And we would suggest that these kinds of approaches are imperative in cases where a particular sensitivity analysis is claimed to overturn a large body of existing evidence.

References

Aaronson, Daniel, Kyung-Hong Park, and Daniel Sullivan. 2007. "Explaining the Decline in Teen Labor Force Participation." *Chicago Fed Letter*, Number 234, January. The Federal Reserve Bank of Chicago.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association*, Vol. 105, No. 490, pp. 493-505.

Addison, John T., McKinley L. Blackburn, and Chad D. Cotti. 2009. "Do Minimum Wages Raise Employment? Evidence from the U.S. Retail-Trade Sector." *Labour Economics*, Vol. 16, No. 4, August, pp. 397-408.

Addison, John T., McKinley L. Blackburn, and Chad D. Cotti. 2012. "The Effect of Minimum Wages on Labour Market Outcomes: County-Level Estimates from the Restaurant-and-Bar Sector." *British Journal of Industrial Relations*, Vol. 50, No. 3, pp. 412-435.

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*, Vol. 50, No. 2, pp. 205-240.

Aruoba, S. Boragan, Francis X. Diebold, and Chiara Scotti. 2009. 'Real-Time Measurement of Business Conditions." *Journal of Business and Economic Statistics*. Vol. 27, No. 4, pp. 417-427.

Autor, David H., John J. Donohue III, and Stewart J. Schwab. 2006. "The Costs of Wrongful-Discharge Laws." *Review of Economics and Statistics*, Vol. 88, No. 2, pp. 211-231.

Baker, Michael, Dwayne Benjamin, and Shuchita Stanger. 1999. "The Highs and Lows of the Minimum Wage Effect: A Time-Series Cross-Section Study of the Canadian Law." *Journal of Labor Economics*, Vol. 17, No. 2, pp. 318-50.

Baskaya, Yusuf Soner, and Yona Rubinstein. 2011. "Using Federal Minimum Wage Effects to Identify the Impact of Minimum Wages on Employment and Earnings Across U.S. States." Unpublished paper, Central Bank of Turkey.

Brown, Charles, Charles Gilroy, and Andrew Kohen. 1982. "The Effect of the Minimum Wage on Employment and Unemployment," *Journal of Economic Literature*, Vol. 20, No. 2, pp. 487-528.

Card, David. 1992a. "Using Regional Variation in Wages to Measure the Effects of the Federal Minimum Wage." *Industrial and Labor Relations Review*, Vol. 46, No. 1, pp. 22-37.

Card, David. 1992b. "Do Minimum Wages Reduce Employment? A Case Study of California, 1987-1989." *Industrial and Labor Relations Review*, Vol. 46, No. 1, pp. 38-54.

Card, David, and Alan B. Krueger. 1994. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania." *American Economic Review*, Vol. 84, No. 4, pp. 772-93.

Congressional Budget Office. 2004. "What Is Happening to Youth Employment Rates?" Accessed at http://cbo.gov/sites/default/files/cbofiles/ftpdocs/60xx/doc6017/11-18-youthemployment.pdf.

Dube, Arindrajit. 2011. "Review of *Minimum Wages* by David Neumark and William L. Wascher." *Journal of Economic Literature*, Vol. 49, No. 3, pp. 762-6.

Dube, Arindrajit, T. William Lester, and Michael Reich. 2010. "Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties." *Review of Economics and Statistics*, Vol. 92, No. 4, pp. 945-64.

Fairris, David, and Leon Fernandez Bujanda. 2008. "The Dissipation of Minimum Wage Gains for Workers through Labor-Labor Substitution: Evidence from the Los Angeles Living Wage Ordinance." *Southern Economic Journal*, Vol. 75, No. 2, pp. 473-96.

Giuliano, Laura. 2013. "Minimum Wage Effects on Employment, Substitution, and the Teenage Labor Supply: Evidence from Personnel Data." *Journal of Labor Economics*, Vol. 31, No. 1, pp. 155-94.

Ionides, Edward L., Zhen Wang, and José A. Tapia Granados. 2012. "Macroeconomic Effects on Mortality Revealed by Panel Analysis with Nonlinear Trends." University of Michigan. Forthcoming in *Annals of Applied Statistics*.

Jones, Margaret. 2012. "The Effects of Living Wages on Wages, Employment, and Hours Worked: New Evidence Using Propensity Score Matching." Unpublished paper, Cornell University.

Katz, Lawrence F., and Alan B. Krueger. 1992. "The Effect of the Minimum Wage on the Fast-Food Industry." *Industrial and Labor Relations Review*, Vol. 46, No. 1, pp. 6-21.

Lang, Kevin, and Shulamit Kahn. 1998. "The Effect of Minimum-Wage Laws on the Distribution of Employment: Theory and Evidence." *Journal of Public Economics*, Vol. 69, No. 1, pp. 67-82.

Lester, Richard A. 1960. "Employment Effects of Minimum Wages: Comment." *Industrial and Labor Relations Review*, Vol. 13, No. 2, pp. 254-64.

Mincer, Jacob. 1976. "Unemployment Effects of Minimum Wages." *Journal of Political Economy*, Vol. 84, No. 3, Part 2, pp. S87-S104.

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 No. 18681. Cambridge, MA: National Bureau of Economic Research.

Neumark, David, Mark Schweitzer, and William Wascher. 2004. "Minimum Wage Effects Throughout the Wage Distribution." *Journal of Human Resources*, Vol. 39, No. 2, pp. 425-50.

Neumark, David, and William Wascher. 1992. "Employment Effects of Minimum and Subminimum Wages: Panel Data on State Minimum Wage Laws." *Industrial and Labor Relations Review*, Vol. 46, No. 1, pp. 55-81.

Neumark, David, and William Wascher. 1994. "Employment Effects of Minimum and Subminimum Wages: Reply to Card, Katz, and Krueger." *Industrial and Labor Relations Review*, Vol. 47, No. 3, pp. 497-512.

Neumark, David, and William Wascher. 1996. "The Effects of Minimum Wages on Teenage Employment and Enrollment: Estimates from Matched CPS Data." *Research in Labor Economics*, Vol. 15, pp. 25-64.

Neumark, David, and William Wascher. 2004. "Minimum Wages, Labor Market Institutions, and Youth Employment: A Cross-National Analysis." *Industrial and Labor Relations Review*, Vol. 57, No. 2, pp. 223-46.

Neumark, David, and William L. Wascher. 2008. *Minimum Wages*. Cambridge, MA: MIT Press.

Neumark, David, and William Wascher. 2011. "Does a Higher Minimum Wage Enhance the Effectiveness of the Earned Income Tax Credit?" *Industrial and Labor Relations Review*, Vol. 64, No. 4, pp. 712-46.

Obenauer, Marie L., and Bertha von der Nienburg. 1915. "Effect of Minimum-Wage Determination in Oregon." *Bureau of Labor Statistics Bulletin*, No. 176. Washington, DC: United States Department of Labor.

Peterson, John M. 1957. "Employment Effects of Minimum Wages: 1938-1950." *Journal of Political Economy*, Vol. 65, No. 5, pp. 412-30.

Peterson, John M. 1959. "Employment Effects of State Minimum Wages for Women: Three Historical Cases Re-Examined." *Industrial and Labor Relations Review*, Vol. 12, No. 3, pp. 406-22.

Peterson, John M. 1960. "Employment Effects of Minimum Wages: Reply." *Industrial and Labor Relations Review*, Vol. 13, No. 2, pp. 264-73.

Reich, Michael. 2009. "Minimum Wages in the United States: Politics, Economics, and Econometrics. In Clair Brown, Barry J. Eichengreen, and Michael Reich (Eds.), *Labor in the Era of Globalization*, pp. 353-74. Cambridge, UK: Cambridge University Press.

Sabia, Joseph J. 2009. "The Effects of Minimum Wage Increases on Retail Employment and Hours: New Evidence from Monthly CPS Data." *Journal of Labor Research*, Vol. 30, No. 1, pp. 75-97.

Sabia, Joseph J., Richard V. Burkhauser, and Benjamin Hansen. 2012. "Are the Effects of Minimum Wage Increases Always Small? New Evidence from a Case Study of New York State." *Industrial and Labor Relations Review*, Vol. 65, No. 2, pp. 350-76.

Smith, Christopher L. 2011. "The Polarization of the U.S. Adult Labor Market and its Effects on the Demand for Teenage Labor." Finance and Economics Discussion Series, Federal Reserve Board, 2011-41. Accessed at http://www.federalreserve.gov/pubs/feds/2011/201141/201141pap.pdf.

Thompson, Jeffrey P. 2009. "Using Local Labor Market Data to Re-Examine the Employment Effects of the Minimum Wage." *Industrial and Labor Relations Review*, Vol. 62, No. 3, pp. 343-66.

Watson, Nadine. 1996. "Positive Minimum Wage Effects on Employment – Alternative Explanations." Unpublished doctoral dissertation. University of California, San Diego.

Figure 1. Example of RMSPE Calculation at State Level[a]
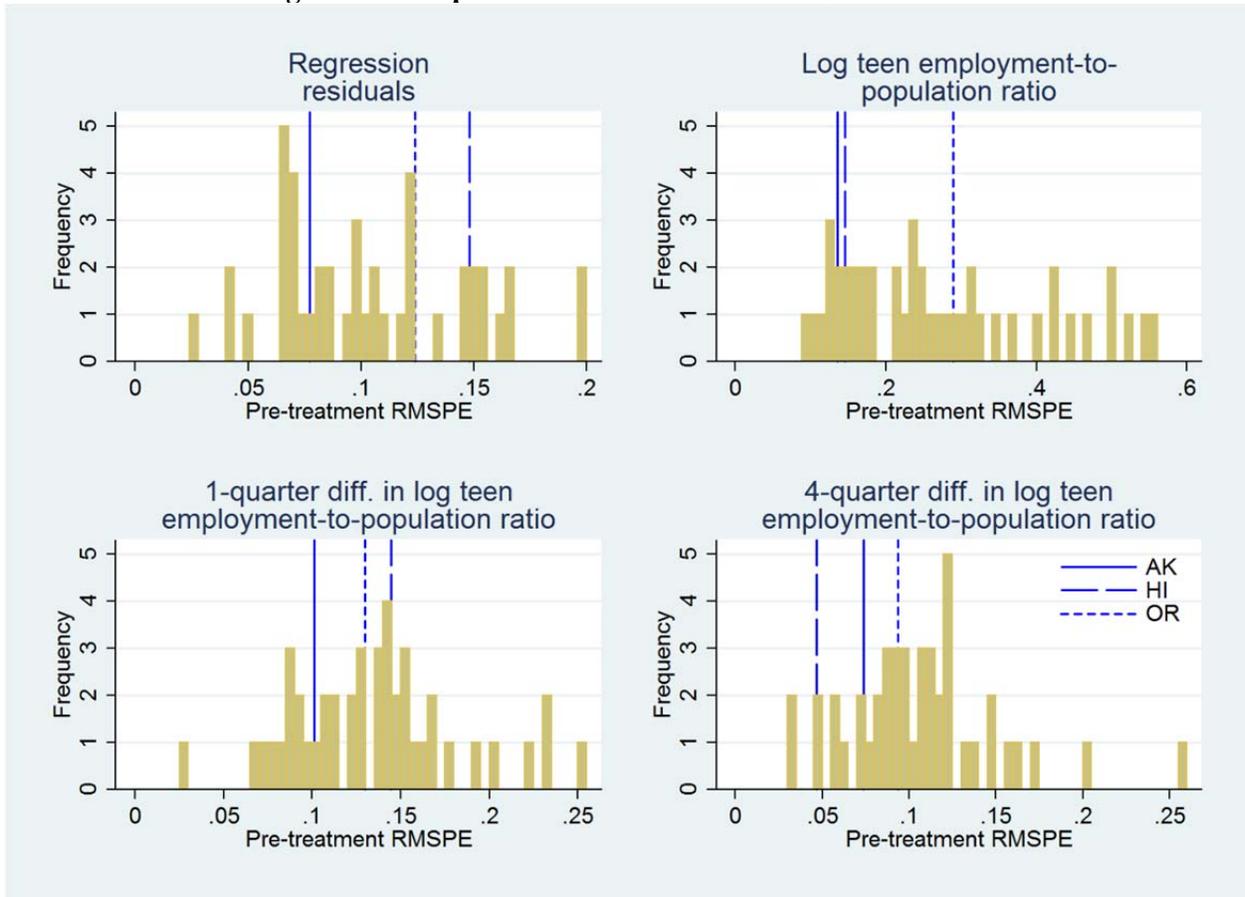
[a]Treatment state: CA, 2001:Q1. Based on CPS data for that quarter and the previous four quarters.
Same-division control states are AK, HI, and OR. (WA is excluded because it had a minimum wage
increase in the same quarter.) The thick bars are the histogram. The thin vertical lines extending to the
top of the graph show the RMSPEs for each control state in the same division as the treatment state.

Figure 2. Distributions of Percentiles of Same-Division States' RMSPEs[a]

Figure 3. Distributions of Percentiles of Contiguous Cross-Border Counties' RMSPEs[a]
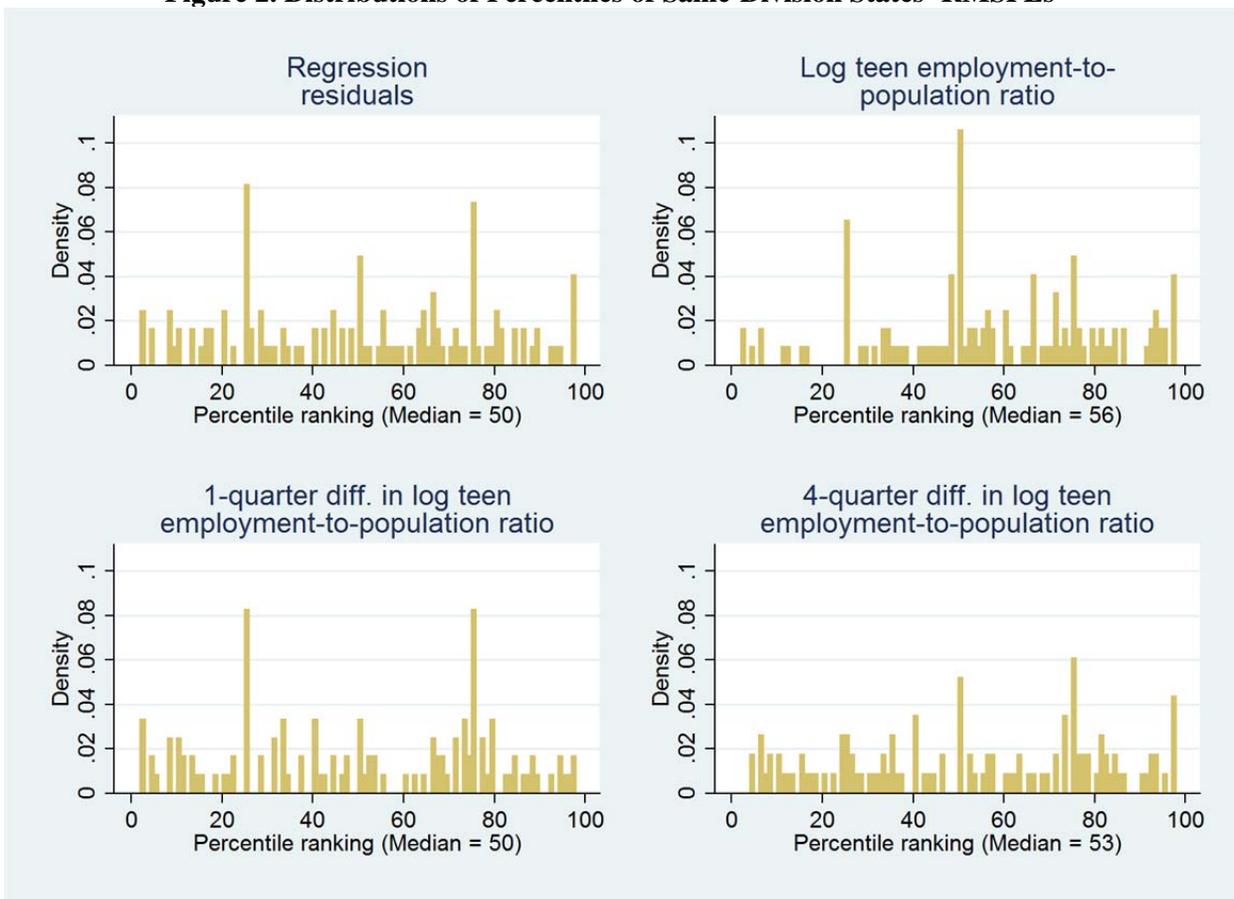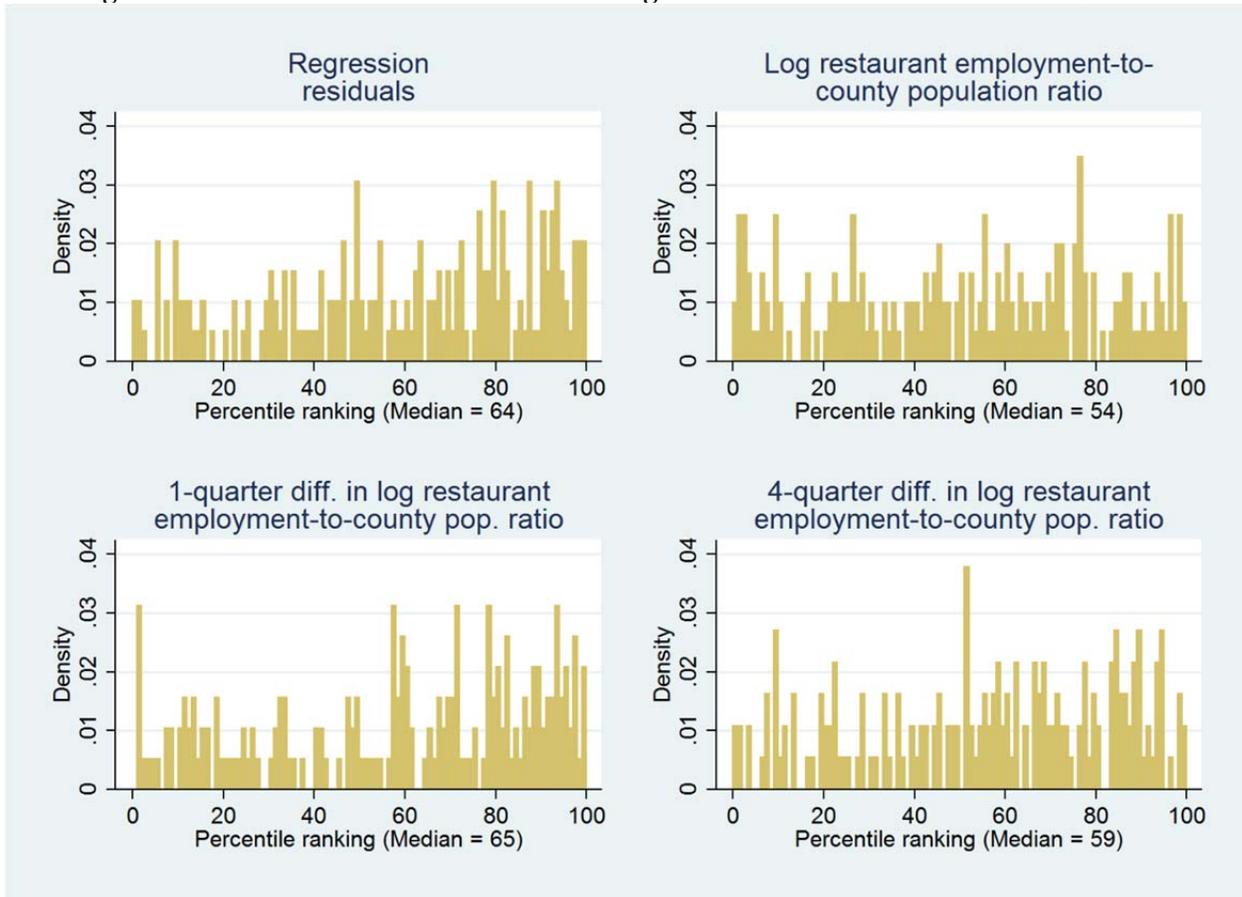
[a]Based on QCEW data for restaurants for 1990-2006:Q2. Ranks are converted to percentile rankings using the Weibull rule described in the text.

Table 1. The Effects of the Minimum Wage on Teen (16-19) Employment, CPS Data at State-by-Quarter Level, 1990 – 2011:Q2[a]

	(1)	(2)	(3)	(4)
	Dependent variable: Log (Employment/Population)			
A. ADR Replication				
Log(MW)	-.165*** (.041)	-.074 (.078)	-.098 (.097)	.009 (.058)
Unemployment rate	-4.195*** (.427)	-3.832*** (.387)	-3.857*** (.403)	-3.118*** (.397)
Relative size of youth population	.100 (.316)	.218 (.336)	.126 (.360)	.161 (.310)
State effects	Yes	Yes	Yes	Yes
Time effects	Yes	Yes	Yes	Yes
State trends	No	Yes	No	Yes
Division-specific time effects	No	No	Yes	Yes
R^2	.877	.893	.911	.921
N	4,386	4,386	4,386	4,386
B. Models with state and time effects, and higher-order polynomials for state-specific trends				
Order of polynomial for state-specific "trends"	2nd	3rd	4th	5th
Log(MW)	-0.051 (0.085)	-0.230*** (0.073)	-0.180** (0.069)	-0.185** (0.073)
Unemployment rate	-3.591*** (0.494)	-2.571*** (0.454)	-2.376** (0.461)	-2.378*** (0.492)
Relative size of youth population	0.490 (0.296)	0.402 (0.280)	0.412 (0.291)	0.354 (0.308)
N	4,386	4,386	4,386	4,386

[a]Estimates are weighted by teen population. In Panel B, models include state dummy variables interacted with a polynomial in time, with order of polynomial as indicated. Standard errors are clustered at the state level. * Statistically significant at the .10 level; ** at the .05 level; *** at the .01 level.

Table 2. The Effects of the Minimum Wage on Teen (16-19) Employment, By Division, CPS Data at State-by-Quarter Level[a]

	(1)	(2)
	Dependent variable: Log (Employment/Population)	
	1990:Q1-2011:Q2	1990:Q1-2009:Q4 (ADR sample)
New England	-.390*** (.052)	-.384*** (.058)
Middle Atlantic	.166 (.143)	.105 (.162)
East North Central	-.208 (.284)	-.166 (.272)
West North Central	-.191* (.082)	-.194** (.067)
South Atlantic	-.150 (.243)	-.152 (.281)
East South Central	-2.235 (1.414)	-2.024 (1.515)
West South Central	-.217** (.062)	-.147* (.053)
Mountain	-.598*** (.139)	-.638** (.187)
Pacific	-.002 (.133)	.016 (.143)

[a]The specification reported in each row, for a division, includes the unemployment rate, the ratio of teen population to total population, and state and time (quarter) fixed effects. Estimates are weighted by teen population. Standard errors are clustered at the state level. * Statistically significant at the .10 level; ** at the .05 level; *** at the .01 level.

Table 3. Weights on States in Same Census Division from Synthetic Control Method, CPS Data at State by Quarter Level, 1990 – 2011:Q2[a]

	Proportion of weight on states in same division						
	Matching on:						
Division	Regression residuals	Log teen employment-to-population ratio	One-quarter difference in log teen employment-to-population ratio	Four-quarter difference in log teen employment-to-population ratio	Avg. # divisions in donor pool	Avg. # states in donor pool	Avg. # states in donor pool in same division
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
New England	0.190	0.209	0.163	0.185	6.9	30.4	1.9
Middle Atlantic	0.088	0.134	0.455	0.168	5.5	20.0	1.0
East North Central	0.000	0.000	0.016	0.015	9.0	39.5	3.5
West North Central	0.575	0.823	0.698	0.464	3.7	7.7	1.7
South Atlantic	0.131	0.290	0.075	0.222	6.9	26.8	4.9
Pacific	0.322	0.339	0.279	0.297	5.3	21.1	2.1
Aggregate	0.246	0.323	0.264	0.251	6.1	24.0	2.5

[a]Results are reported for the 50 unique minimum wage treatments (out of a total of 129 increases based on criteria described in the text) for which there is at least one potential donor state from the same Census division. The numbers in columns (5)-(7) refer to the matching on residuals or the log teen employment-to-population ratio. There are somewhat fewer minimum wage treatments when matching on the one- or four-quarter differences in the employment-to-population ratio because the earliest lags are not available at the beginning of the sample period. The aggregate row reports the means across all treatment units.

**Table 4. The Effects of the Minimum Wage on Restaurant Employment,
County-Level QCEW Data, 1990-2006:Q2[a]**

	(1)	(2)	(3)	(4)
	Dependent variable: log (restaurant employment), DLR contiguous border county-pair sample			
	Without county-pair × period interactions (DLR, Table 2, specification 5)		With county-pair × period interactions (DLR, Table 2, specification 6)	
Log(MW)	-.137[*] (.072)	-.112 (.079)	.057 (.115)	.016 (.099)
Log(population)	.952[***] (.073)	.567[***] (.103)	1.116[***] (.190)	.714[***] (.246)
Log(private-sector employment)	…	.405[***] (.067)	…	.393[***] (.117)
County effects	Yes	Yes	Yes	Yes
Period effects	Yes	Yes	No	No
County-pair × period interactions	No	No	Yes	Yes
N	70,620	70,582	70,620	70,582

[a]Standard errors are two-way clustered at the (non-nested) state and border segment levels; the border segment is the set of all counties on both sides of a border between two states. * Statistically significant at the .10 level; ** at the .05 level; *** at the .01 level.

Table 5. Weights on Contiguous Cross-Border Counties from Synthetic Control Method, County-Level QCEW Data, 1990-2006:Q2[a]

	Proportion of weight on contiguous cross-border counties			
	Matching on:			
	Regression residuals	Log restaurant employment-to-county population ratio	One-quarter difference in log restaurant employment-to-county population ratio	Four-quarter difference in log restaurant employment-to-county population ratio
Distribution	(1)	(2)	(3)	(4)
A. Donor pools restricted to 50 counties with lowest RMSPE for four quarters prior to minimum wage increase				
Minimum	0.000	0.000	0.000	0.000
10th percentile	0.000	0.000	0.000	0.001
25th percentile	0.001	0.002	0.002	0.004
Median	0.008	0.010	0.009	0.010
75th percentile	0.019	0.029	0.027	0.019
90th percentile	0.048	0.101	0.068	0.035
Maximum	0.194	0.659	0.496	0.336
Mean	0.021	0.038	0.035	0.017
B. Full donor pools				
Minimum	0.000	0.000	0.000	0.000
10th percentile	0.000	0.000	0.000	0.000
25th percentile	0.001	0.000	0.001	0.001
Median	0.001	0.001	0.001	0.001
75th percentile	0.002	0.002	0.002	0.002
90th percentile	0.009	0.007	0.012	0.006
Maximum	0.200	0.308	0.474	0.393
Mean	0.009	0.007	0.015	0.007

[a]County results are reported for the 121 unique minimum wage treatments (out of a total of 129 increases based on criteria described in the text) for which there is at least one potential contiguous cross-border donor county. In Panel A, for each treatment the donor pool consists of the 50 counties with the lowest RMSPE for the four quarters following the minimum wage increase. If the contiguous cross-border counties that DLR use as controls are not in this top 50, they are added to the donor pool. Panel B does not impose this restriction. In both panels, the average number of contiguous cross-border counties in the donor pool is 1.7, while the average number of counties in the donor pool is 51.3 in Panel A and 961.7 in Panel B. There are somewhat fewer minimum wage treatments when matching on the one- or four-quarter differences because the earliest lags are not available at the beginning of the sample period.

Table 6. The Effects of the Minimum Wage on Restaurant Employment, "Falsification Tests,"
County-Level QCEW Data[a]

	(1)	(2)	(3)
	Dependent variable: Log (restaurant employment)		
	DLR sample: 1990:Q1-2006:Q2 (DLR Table B1, specification 2)	Sample restricted to 1998:Q3-2006:Q2 (period with no federal MW changes)	Sample restricted to 1998:Q3-2006:Q2, county pairs with minimum wage difference for at least one quarter
Actual MW sample:			
Log(MW)	-0.208 (0.150)	-0.247*** (0.042)	-0.260** (0.097)
N	34,514	21,308	5,180
Placebo MW sample:			
Log(MW)	-0.123 (0.158)	-0.107 (0.068)	0.005 (0.082)
N	33,726	20,768	4,640
% of county pair-quarter observations with minimum wage difference between counties	4.0	7.0	31.2
% of county pairs with minimum wage difference between counties	17.8	22.3	100.0
County effects	Yes	Yes	Yes
Period effects	Yes	Yes	Yes

[a]These specifications include controls for population and private-sector employment. Following DLR's code, the sample is restricted to counties that have an area less than 2,000 square miles, and have data for each quarter in the sample period. In columns (2) and (3), a balanced panel of counties is used, as in DLR's other analyses; we restrict the sample only to counties with data in the shorter sample period studied. In column (3), the subset of county pairs in column (2) that had one or more minimum wage differences in the period always had at least two quarters of minimum wage differences. Standard errors are clustered at the state level. * Statistically significant at the .10 level; ** at the .05 level; *** at the .01 level.

Table 7. The Effects of the Minimum Wage on Teen (16-19) Employment, Matched Panel Data Estimates, CPS Data at State-by-Quarter Level, 1990 – 2011:Q2[a]

	(1)	(2)	(3)	(4)	(5)	(6)
		Dependent variable: Log (Employment/Population)				
		Matched panel data estimator, treatment state and experiment fixed effects				
	Panel data estimator	129 "clean" minimum wage increases, all states in donor pool as controls, equally weighted	All minimum wage increases, all states as controls, equally weighted	All minimum wage increases, same-division states as controls, equally weighted	All minimum wage increases, synthetic control weights based on residuals including minimum wage	All minimum wage increases, synthetic control weights based on residuals excluding minimum wage
Log(MW)	-0.165*** (0.041)	-0.062 (0.106)	-0.122** (0.061)	-0.055 (0.055)	-0.143** (0.061)	-0.145** (0.060)
Unemployment rate	-4.195*** (0.427)	-1.407* (0.823)	-1.656*** (0.419)	-0.921*** (0.342)	-1.554*** (0.348)	-1.487*** (0.344)
Relative size of youth population	0.100 (0.316)	-0.152 (0.673)	0.072 (0.386)	0.070 (0.332)	0.041 (0.322)	-0.021 (0.319)
No. of minimum wage increases	544	129	493	493	493	493
N	4,386	2,064	7,888	7,888	7,888	7,888

[a]The estimates in column (1) are panel data estimates with fixed state and period effects, corresponding to those in column (1) of Table 1, Panel A. The remaining estimates use treatment observations for the quarter of a minimum wage increase and the four preceding and three succeeding quarters, and matched counterfactual observations as described in the text. All specifications beginning with column (2) include dummy variables corresponding to each minimum wage "experiment" (i.e., the matched treatment and counterfactual observations for each minimum wage increase) and each set of treatment observations in each experiment. In columns (3) and higher, minimum wage increases in 1990:Q2, 2011:Q1, and 2011:Q2 are excluded because there are not eight quarters of observations available for these increases. Standard errors are clustered at the level of the treatment or counterfactual observations in each experiment. Observations in columns (2)-(6) are weighted by the teen population in the treatment observations (similarly applied to corresponding counterfactual observations). * Statistically significant at the .10 level; ** at the .05 level; *** at the .01 level.

Table 8. The Effects of the Minimum Wage on Border County Restaurant Employment, Matched Panel Data Estimates, County-Level QCEW Data, 1990 – 2006:Q2[a]

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: Log (Restaurant Employment/County Population)					
			Matched panel data estimator, treatment county and experiment fixed effects			
	Panel data estimator	Panel data estimator, paired border counties only	Paired border counties sample, all minimum wage increases, all border counties in other states as controls, equally weighted	Paired border counties sample, all minimum wage increases, paired border counties as controls, equally weighted	All counties, all minimum wage increases, synthetic control weights based on residuals including minimum wage	All counties, all minimum wage increases, synthetic control weights based on residuals excluding minimum wage
A. Unweighted, with private-sector employment control						
Log(MW)	-0.174[*] (0.100)	-0.080 (0.070)	0.013 (0.022)	0.049[**] (0.023)	0.004 (0.023)	0.008 (0.023)
B. Weighted by county population, with private-sector employment control						
Log(MW)	-0.120[***] (0.042)	-0.104[**] (0.050)	-0.039[*] (0.021)	0.008 (0.024)	-0.052[**] (0.022)	-0.063[***] (0.022)
N	90,948	25,146	29,344	29,344	29,344	29,344

[a]The estimates in column (1) and (2) are panel data estimates with fixed county and period effects. The remaining estimates use treatment observations for the quarter of a minimum wage increase and the four preceding and three succeeding quarters, and matched counterfactual observations as described in the text. All specifications beginning with column (3) include dummy variables corresponding to each minimum wage "experiment" (i.e., the matched treatment and counterfactual observations for each minimum wage increase) and each set of treatment observations in each experiment. In columns (3) and higher, minimum wage increases in 1990:Q2, 2005:Q4, 2006:Q1, and 2006:Q2 are excluded because there are not eight quarters of observations available for these increases. Standard errors are clustered at the level of the treatment or counterfactual observations in each experiment. In Panel B, observations in columns (3)-(6) are weighted by the county population in the treatment observations (similarly applied to corresponding counterfactual observations). In columns (5) and (6) all counties in other states used to construct the counterfactual observations. * Statistically significant at the .10 level; ** at the .05 level; *** at the .01 level.