David Neumark
UC-Irvine

J.M. Ian Salas
UC-Irvine

January 2013

**Employment
Policies
INSTITUTE**

MINIMUM
WAGES

Evaluating New Evidence
on Employment Effects

The Employment Policies Institute (EPI) is a nonprofit research organization dedicated to studying public policy issues surrounding employment growth. Among other issues, EPI research has quantified the impact of new labor costs on job creation, explored the connection between entry-level employment and welfare reform, and analyzed the demographic distribution of mandated benefits. EPI sponsors nonpartisan research that is conducted by independent economists at major universities around the country.

David Neumark is Chancellor's Professor of Economics and Director of the Center for Economics & Public Policy at the University of California, Irvine. He is also a Research Associate of the National Bureau of Economic Research, and a Research Fellow of IZA. He has held prior positions at the Public Policy Institute of California, Michigan State University, the University of Pennsylvania, and the Federal Reserve Board. He is a labor economist with broad public policy interests including: age, sex, and race discrimination; the economics of aging; affirmative action; minimum wages, living wages, and other anti-poverty policies; the economics of education; youth labor markets; and local economic development. Neumark is the editor of the *IZA Journal of Labor Policy*, and a co-editor of the *Journal of Urban Economics*.

J.M. Ian Salas is a doctoral candidate in Economics at the University of California, Irvine. His fields of specialization are in development economics and in labor and demographic economics. He has published papers on many topics, including detecting collusive behavior in setting gasoline pump prices, validating a central bank's adherence to its avowed goal of managing inflation, tracing disparities in firm productivity and regional development to differences in investment environment, and measuring the aggregate burden of providing for a relatively large young dependent population. His current research is primarily focused on identifying the causal explanations for high fertility in many developing countries.

MINIMUM WAGES

Evaluating New Evidence on Employment Effects

David Neumark
UC-Irvine

J.M. Ian Salas
UC-Irvine

Employment
Policies

INSTITUTE

1090 Vermont Avenue, NW
Suite 800
Washington, DC 20005

MINIMUM WAGES

Evaluating New Evidence on Employment Effects

Table of Contents

MINIMUM WAGES

Evaluating New Evidence on Employment Effects

Executive Summary

The fierce political debate over raising the minimum wage, which is repeated yearly in legislatures across the country, has at times been matched by a strong academic debate on the subject. Specifically, economists have argued over whether a higher minimum wage reduces the employment of less-skilled jobseekers.

The published research on the subject points overwhelmingly in one direction: A summary of the last two decades of literature on the minimum wage, co-authored by the lead economist on this study, concluded that most of the evidence points to job loss following wage hikes. Economists have detected this job loss using state variation in minimum wages, with states that do not raise their minimum wage acting as a "control group" for states that do.

But today, a small group of economists (listed at right) has mounted an aggressive challenge to the existing academic consensus on minimum wages. In a series of studies first published through the organized labor-aligned Institute for Research on Labor and Employment (and later in the journals *Review of Economics and Statistics*

Academics behind the "new" minimum wage debate

Arindrajit Dube
University of Massachusetts-Amherst

Sylvia Allegretto
University of California-Berkeley

Michael Reich
University of California-Berkeley

T. William Lester
University of North Carolina

and Industrial Relations), they've argued that prior studies on the minimum wage were incorrect in blaming the policy for a drop in employment opportunities among less-skilled employees (like teens) or in service-intense industries. Rather, they claim, these employment declines are due to unrelated changes in states' economies—in particular, unexplained downturns in employment of

unskilled workers that just happen to coincide with dozens of state minimum wage increases.

To get around this purported problem, they toss out most of the labor market data available to detect the effects of wage increases, and restrict their analysis to either neighboring state border counties or states in the same Census division. Using this highly-restrictive model, they claim, provides better control groups and shows that a higher minimum wage has no negative effect on employment.

With the encouragement of the economists themselves, who have a history of working in support of progressive causes, these studies (henceforth *IRLE papers*) have gained prominence among activist groups who leverage them to claim that mandated wage hikes will have no adverse impact on employment. (One of the economists even said explicitly at a 2010 conference in Atlanta that this research should "help to pave the way" for higher mandated wages.)

> ### Neumark and Salas are unequivocal in their conclusion: "[N]either the conclusions of these studies nor the methods they use are supported by the data."

Or so they hope. But in this new study, University of California-Irvine labor economist David Neumark worked with UC-Irvine Ph.D. student J.M. Ian Salas to determine whether the IRLE papers have merit. Specifically, do the studies make good on the claim their authors' have put forth, of overturning the decades of research preceding them?

Neumark and Salas report that the evidence presented in the IRLE papers only runs contrary to earlier studies because the authors' empirical models rely on inappropriate control groups, and toss out the economic data necessary to detect the impact of a minimum wage increase. They are unequivocal in their conclusion: "[N]either the conclusions of these studies nor the methods they use are supported by the data."

The authors of the IRLE papers provide no direct evidence to justify their highly-restrictive study design, instead speculating that nearby states or counties constitute ideal control groups against which to measure the effects of the minimum wage. But Neumark and Salas demonstrate that the premise of the IRLE papers is wholly incorrect: If you examine the characteristics of the control counties and states used in the IRLE papers (which the authors of those papers failed to do), you find that they're generally **very poor** control groups.

- For instance, the authors' preferred control for Leon County, FL—home of the Florida state capital in Tallahassee, with a population of roughly 275,000 people—is Grady County, GA, which has barely 25,000 people and no major cities.

- Similarly, one control state for Connecticut—a vibrant northeastern state with 3.5 million people and $237 billion in annual economic output—is Vermont, a state with roughly 1/6th of Connecticut's population and one-tenth its economic output.

Instead of speculating about which states represent ideal controls, Neumark and Salas closely examine the economic characteristics of all states in each Census Division, and find that it's mostly states *outside* the Census Division that serve as better control groups. (A similar pattern holds for nearby counties, which they also examine individually.) Yet all of this identifying data is discarded by the authors of the IRLE papers.

> **"[T]he robust set of controls the authors use actually aren't robust at all—indeed, they're less suited to the task at hand than studies that have come before."**

In other words, the robust set of control groups the authors use actually aren't robust at all—indeed, they're less suited to the task at hand than studies that have come before. And in the small number of cases where nearby states or counties *are* appropriate controls, the data in these cases show that employment did fall after a minimum wage increase.

Given their numerous methodological problems, it's not surprising that the evidence in favor of the empirical approach advocated in the IRLE papers is "weak or non-existent." When the analysis is not restricted to these inappropriate control groups, the data clearly show that wage hikes *do* cause job loss. Indeed, in some cases Neumark and Salas find that the IRLE authors omitted evidence that exposed the weaknesses in their approach.

Neumark and Salas end with a strong admonishment to the authors of the IRLE papers: "[P]rior 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 convincing than the established research literature that finds disemployment effects, and understand why the studies in that literature would have generated misleading evidence."

It's a warning that the economists themselves should heed, as should legislators eager for studies (no matter their accuracy) that validate their ideological preferences.

—Employment Policies Institute

Introduction

A large number of time-series studies conducted in the 1960s and 1970s 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 et al., 1982). As a result, 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* (Neumark and Wascher, 2008), which 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 report, in part, extends this

evaluation and summary to the present by describing research that evaluates some recent studies that have questioned the empirical methods and conclusions in much of the recent literature (Allegretto et al., 2011; Dube et al., 2010).

The key issue raised by these recent studies is how researchers can reliably 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 spuri-

ous 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 argue 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).

The principal goal of the research described in this report 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 variation is associated with unobserved economic shocks to areas that can confound the estimation of minimum wage effects. Consequently, much of the analysis focuses on the validity of this criticism, and on the approaches these studies take to address this potential problem. In particular, the overriding concern is that the research designs used in these new studies, 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."

The findings, in a nutshell, are that neither the conclusions of these studies nor the methods they use are supported by the data. The report also briefly discusses other recent research on the employment effects of minimum wages—research that also departs from the state panel-data approach that has been used in much of the new minimum wage research. This recent research also support the conclusion that minimum wages reduce employment of low-skilled workers

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 particular, ADR focus on standard 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 Sur-

vey 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 analyses that are similar to ADR, estimating 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, which includes county pair × period interactions, identifies the effect of minimum wages from differences in employment changes in paired counties on either side of a state border, with the county pair × period interactions intended to control for shocks common to both counties. In this specification, the narrowing of identification to this within-county-pair comparison causes the employment effects to go from negative and significant to small and insignificant.

Closely related findings are reported in Addison et al. (forthcoming). They also use QCEW data, focusing on nearly the same period (1990-2005) and on nearly the same sector. With regard to DLR's first set of analyses, about the only difference is that the authors include county-specific linear trends. However, the qualitative conclusion is the same: they find negative employment effects when they include only county and quarter fixed effects, but no evidence of employment effects when they include county-specific trends. Similarly, in a 2009 paper these same co-authors estimate these models for various parts of the retail sector with county-specific time trends, and also doing the same border-county analysis as in DLR, with similar findings.[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 was 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 example, it has repeatedly been emphasized that the evidence from the earlier research is strongest for individuals most di-

---

[1]In a new paper (Addison et al., 2011), the authors use similar methods to try to estimate the effects of minimum wages in the United States during the Great Recession. To some extent this paper appears to be a response to Even and Macpherson (2010), who estimate standard state panel-data models (with fixed state and year effects) for the 2005-2009 period and find disemployment effects for teenagers. Addison et al. conclude that it is hard to detect disemployment effects when similar kinds of "spatial heterogeneity" extensions to the panel data methods are made. However, perhaps because of the short time period studied and the large movements in aggregate labor market outcomes during that period, the standard errors on the point estimates are so large as to be uninformative—for example, in some cases (like for teens) the authors find negative employment effects that are in the earlier "consensus" range but that are not statistically significant. Estimating the effects of minimum wages during this exceptional period is likely to be extraordinarily difficult and that applying demanding empirical methods to the data from this period will probably be futile. There are similar concerns about the results reported in Hirsch et al. (2011), who study the effects of the federal minimum wage increases on employment at "quick-service" restaurants in Georgia and Alabama. They report estimated employment elasticities, with respect to the wage cost increases induced by the higher minimum wage, for each year of the increase that are sometimes positive and sometimes negative, but that also have large standard errors (about 0.4-0.5). They also estimated the effects of the cumulative increase—which are not reported in the paper but were supplied in a personal communication from Barry Hirsch—and found large and negative effects (ranging from -0.8 to -1.7) that were statistically insignificant, suggesting that the main problem may be uninformative data.

rectly affected by the minimum wage, and that many workers within an industry sector earn well more than the minimum wage. In addition, the minimum wage can lead employers to substitute higher-skilled 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 because 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 et al., 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 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 wages. The report turns to a more thorough explanation of their analysis next, as well as an evaluation of it; it then does the same for the DLR study.

Evaluation of the Evidence
Allegretto et al. (2011)

As noted above, ADR find that the negative effects estimated from standard 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 x period interactions (or both). This leads them to conclude that models with only state and year 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).

---

[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. For example, Fairris and Bujanda (2008) find evidence of labor-labor substitution by city contractors in response to the Los Angeles living wage ordinance – a different kind of mandated wage floor. In addition, in a study of personnel records from over 600 establishments of a single large retail firm, Giuliano (forthcoming) finds that large minimum-wage induced increases in the wage rates of teenagers relative to adults were associated with increases in the teen share of employment. These increases were fueled by teenagers from higher socioeconomic status zip codes, and were accompanied by decreases in the employment of young adults (ages 20-22). Moreover, data on the individual workers who were laid off and on store performance indicated that at some stores the teens that were hired were of higher quality than teens already employed at the stores, and of higher quality than the young adults at the stores. Similar evidence consistent with this substitution from adults to teens (in food-service occupations) is reported in Lang and Kahn (1998). Finally, Neumark and Wascher (1996) find evidence of this substitution among teens, with a higher minimum wage drawing those enrolled in school and working part-time into full-time work, while pushing those working full-time and not enrolled in school out of jobs into "idleness" (neither working nor employed).

[3] The focus on teenagers is, to some extent, a vestige of the old time-series literature. Because labor economists had to 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.

These findings are re-examined with the same CPS data they use, using a specification with the same aggregate variables they include. The sample is extended to take account of newer data (which does not change the basic conclusions), and, whereas ADR use individual-level data with clustering at the state level, here the data are aggregated 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, the results closely mirror what ADR found (their Table 3). Using the model with the standard labor market controls they use and only state and time fixed effects,[5] the estimated employment elasticity with respect to the minimum wage is -0.165, significant at the 1% level (ADR estimate an elasticity of -0.12, significant at the 5% level).[6] When either state-specific linear trends are added, region × quarter interactions are added, or both are added simultaneously, the estimated elasticities become considerably smaller—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 are statistically insignificant and the estimated elasticities range from -0.036 to 0.047.[7]

This evidence indicates 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. What ADR fail to establish is whether the latter two sources of variation provide better estimates of the effects of minimum wages than the first type of variation. In particular, it is important to explore the implications of including state-specific trends or region × time interactions and ask 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 has become 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. Indeed, Neumark and Wascher's more recent studies of minimum wage effects on employment (2004, 2011) included specifications with these trends, and in both studies the evidence of negative employment effects was robust to including state-specific trends.[8] However, if the estimates turn out to be sensitive to this specification check, it is important to try to assess both why this might be the case, and to ask whether a single way of doing this—by choosing a specification in which these longer-run influences are restricted to be linear—uncovers the right answer.

ADR's analysis is flawed on both accounts. First, the ar-

---

[4]Aggregated data are used because this form is more convenient for some of the analyses that follow, which focus on the time-series patterns in the data by state. 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.

[5]Because ADR use micro-data rather than state-level data, they also include controls for individual demographic characteristics.

[6]The specification is in logs so the estimated coefficient is the elasticity; in contrast, they estimate linear probability models for employment with the level 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.

[7]The specifications were also estimated adding 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 to those reported in Table 1.

[8]The 2011 paper (which ADR cite as an earlier unpublished paper, from 2007) focuses on the United States, and the 2004 paper studies many countries and hence the analogue to state-specific trends is country-specific trends.

guments about the need to account for such long-run influences to correctly identify minimum wage effects are unconvincing. To argue that there may be differential trends in teen employment by state that are not captured by the standard set of control variables, ADR cite research by Smith (2010) 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). Smith does present some results suggesting that the increase in the share of adults in "youth occupations" rose more in states that had a larger baseline share of labor in routine tasks that were more likely to have been affected by computing, consistent with the possibility that state-specific trends are capturing the effects of technological change. But a higher minimum wage might also lead to substitution of low-skilled adults for teens, suggesting that state-specific trends may be capturing some of the variation that was induced by the minimum wage. Indeed, Smith says this: "If there is a binding minimum wage, and, for the same wage, employers prefer less-educated adult labor to teen labor, then the effect of polarization on youth employment relative to adult employment would further be enhanced" (p. 14). Such uncertainty about the reasons for differences in state-specific trends suggests that it would be better to incorporate data on the factors for which the state-specific trends are intended to control, rather than simply including them and interpreting the results as necessarily reflecting these factors.

ADR also cite research by Aaronson et al. (2007) and the CBO (2004) as providing evidence for omitted trends in teen employment unrelated to minimum wages—related, for example, to factors such as changes in financial aid for college students, the attractiveness of college, or technological shifts that have lowered market wages for teens.

They show, in their Figure 1 and Table 1, that "employment rates by teens vary by Census division and differentially so over time" (p. 206), and appeal to these studies—the CBO and Aaronson et al. studies, especially—to argue that these "differences are not captured simply by controls for business cycles, school enrollment rates, relative wages of teens, unskilled immigration, or by the timing of federal minimum wage increases" (p. 206). But the appeal to the evidence in these studies is misleading. For the most part, the two studies cited only discuss aggregate time-series evidence, and time-series changes in teen employment relative to adult labor market indicators for the United States as a whole would be captured in the time-period dummies.[9] As a result, these studies do not make a strong case that there are important unexplained trends by state (or changes by region) that would not be captured by state and time dummies.

The second problem with ADR's analysis is that they ignore the possibility that the sensitivity of the estimates to including state-specific trends may be attributable to the failure to incorrectly capture longer-run influences on teen employment. In particular, when the dependent variable (teen employment) is strongly affected by the business cycle, the use of restrictive linear trends can sometimes give highly misleading estimates. We illustrate this problem in a number of steps, and in so doing show that ADR's conclusion is wrong—and that controlling for longer-run influences on teen employment does not alter the conclusion that minimum wages reduce teen employment.

Neumark and Wascher's (2011) finding that the estimated effects of minimum wages on employment are negative and significant in specifications that include state-specific trends was based on data from 1994-2007. This raises the

---

[9]The only exception is that Aaronson et al. (2007) suggest that increases in some states in merit aid for higher education reduced teen employment (explaining perhaps one-tenth of the aggregate decline in teen employment during the economic expansion of the 2000s).

question as to whether there is something different about the sample period ADR studied that makes it problematic to estimate 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; the effects of these recessions on labor market outcomes for the U.S. as a whole are illustrated in Figures 1 and 2.[10] 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. 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 it can be difficult to sort out the longer-term trends in teen employment from the effects of the business cycle.

As a first indication that this may be a problem, Table 2 estimates the same models as in Table 1 for different sample periods. Columns (1) and (2) repeat the estimates from Table 1. Columns (3) and (4) leave out the first three years of the sample period, hence excluding the early peak in the unemployment rate; columns (5) and (6) similarly leave out the last four years, and columns (7) and (8) leave out both. When either recessionary period is included, the results are fairly similar, with the inclusion of state-specific trends substantially reducing the estimated disemployment effect. But when the sample is restricted to exclude both of these recessionary periods, the evidence of disemployment effects is not diminished by the inclusion of state-specific trends—rather, it gets stronger.[11]

The sensitivity of the estimates to including state-specific trends when the sample period includes either of these recessions may be attributable to large residuals in the recessionary periods exerting a strong influence on the state-specific trends. That, in turn, can substantially alter the classification of periods in which teen employment was high or low relative to the predicted values *net of* the minimum wage, and hence can influence the estimated minimum wage effects. In particular, if the state-specific trends are strongly influenced by the recessionary periods, then they are *not* reflecting the longer-term trends in teen employment that ADR are attempting to capture, but rather are picking up cyclical fluctuations that can, in other periods, generate spurious correlations that obscure the true effects of the minimum wage on teen employment.

To assess whether this concern applies to ADR's results, the model was estimated for the whole period and then the residuals were computed for three subperiods—the 1994-2007 period used earlier, and separate periods that include the recessions that bracket this period. For each state, the share of residuals that was positive in each recessionary period was computed, as were the mean residuals for those periods. Because the residuals should be randomly distributed about zero in a correctly specified model, cases where the share of residuals was sharply different from 0.5 (or the mean residual is sharply different from zero) in the recessionary periods are indicative of misspecification. For many states there is noticeable evidence of misspecification. As but one example, in Connecticut, for the estimation using the whole sample period, 75% of the residuals were positive in the 1990-1993 period, with a mean of 0.072, and in the 2007:Q3-2011:Q2 period 81% or the residuals were positive, with a mean of 0.065. In addition, for the intervening period (1994-2007:Q2) the share of positive residuals is considerably below one (31%), which suggests that the influence of the recessionary peri-

---

[10] The grey bars indicate recessionary quarters based on NBER business cycle dates. Note that Figure 2 points to a sharp decline in the teen employment rate relative to adult rate in the last decade. However, this kind of aggregate trend will be captured in the time dummies included in the regression and has no implications for the question of whether state-specific trends should be included.

[11] These results also were very similar including 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.

ods on the state-specific trends may lead to a worse fit in non-recessionary periods.

To more clearly illustrate the problem of estimating state-specific linear trends when the sample period includes these recessionary periods, Figures 3-6 display more information on four states that had these kinds of patterns in the residuals. First, the model is estimated, again including the state-specific trends, for the period excluding the recessions (1994-2007:Q2). The upper-left panel of each figure shows the actual residuals for this period, and then the prediction errors for the two recessionary periods. These figures illustrate how the sizable deviations of teen employment from predicted values during these two recessionary periods could have a strong influence on the estimated state-specific trends in samples that included those periods. For example, the upper-left panel of Figure 3 shows that, in California, teen employment was considerably higher than would have been predicted by the regression model in the period of the early-90s recession, whereas the opposite was true during the Great Recession. This difference presumably reflects the different industries affected by the two recessions. In contrast, for Indiana (Figure 4), teen employment fell by relatively more in both the early-90s recession and the Great Recession, whereas in Tennessee (Figure 5) there are large negative residuals during the Great Recession. Finally, in Pennsylvania (Figure 6) teen employment remained relatively strong during both recessions. The implication of these data plots in the upper-left panels of the figures is twofold: first, they show that recessions can cause teen employment to deviate significantly from the predictions of the control variables of the model; and second, this pattern can differ substantially across states (and recessions).

These results, in turn, imply that estimating state-specific trends for sample periods that include recessionary periods can result in trend estimates that are influenced by the effect of a recession in a state. In that case, the estimated

trends may not adequately capture the underlying trends associated with longer-term factors that are not captured by the other controls included in the regression model. This point is illustrated in the remaining three panels of Figures 3-6. In these panels, the estimation period is extended to include one or both recessionary periods. The residuals from the fitted regression model are then plotted, showing that the deviations of teen employment in the recessionary periods influence the regression model. For example, for California, in Figure 3, note that in the remaining three panels (lower-left, and both panels on the right), the deviations of the actual values of teen employment from the fitted values (i.e., the residuals) are much smaller for the recessionary periods than in the upper-left panel, and these residuals are more centered on zero. The same is generally true for the states shown in the remaining figures: see, as good examples, Indiana (Figure 4, especially for the Great Recession), Tennessee (Figure 5), and Pennsylvania (Figure 6, especially for the recession in the early 1990s).

Correspondingly, the estimation of the state-specific trends over the non-recessionary period can be importantly influenced by the inclusion of the recessionary periods, which can, in turn, influence the other estimates once the trends are included. In the upper-left panel of each figure, where the model is estimated for 1994-2007:Q2, a regression of the residuals for this period on a time trend will, by construction, yield a zero coefficient. (The model includes a separate time trend for each state, and OLS residuals are orthogonal to the regressors.) However, in the other three panels, the regression of the residuals on a time trend for this subperiod need not have a zero slope; indeed, if the residuals from the recessionary period influence the estimated trends then this slope will differ from zero. This is illustrated by the fitted lines in the middle segment of each of the other three panels, which are the fitted regression lines of the residuals on the time trend for the 1994-2007:Q2 subperiod. The trend for this peri-

od—which still comprises most of the sample period—is usually quite different from zero.

In sum, including the strong recessionary periods of the early 1990s and the Great Recession in the estimation is evidently biasing at least some of the state-specific trend estimates away from their longer-run values and thus potentially confounding the estimates of the employment effects of minimum wages. Two solutions to this problem are proposed.

One is simply to exclude the recessionary periods. As shown in Table 2, columns (7) and (8), including the state-specific time trends for the 1994-2007:Q2 sample period causes the estimated employment effects to become more negative, in sharp contrast to the results in the other columns that include the recessionary periods. These estimates provide one indication that when one adds state-specific linear trends to the model, for a period when these trends are more likely to accurately capture the long-run trends in teen employment, the conclusion is no longer—to quote ADR (2011) again—that "Lack of controls for spatial heterogeneity in employment trends generates biases toward negative employment elasticities in national minimum wage studies" (p. 206).

As another way of illustrating the sensitivity of the specifications with state-specific trends to the recessionary periods at the beginning and end of the sample period, the models are estimated allowing the state-specific trends to be of a higher order than linear. As the preceding figures show, linear trends are ill-suited to capturing the variation induced by the recessions, whereas higher-order trends should be better suited to this—especially third-order and higher polynomials that allow for multiple inflection points. The estimates for the full sample period are reported in columns (1)-(4) of Table 3. (The comparable estimates with no state-specific trends and linear trends are in columns (1) and (2) of Table 2.) The table shows that as long as third-order polynomials or higher are used, there are negative and significant effects of the minimum wage on teen employment, with point estimates very similar to the estimates for the standard model with just state and year fixed effects in Table 2—and similar to the estimates for the subsample excluding the beginning and ending recessionary periods. As shown in columns (5)-(8), the same answer is obtained using the slightly shorter sample period that ADR use.[12] Thus, ADR's claim that underlying trends that vary by state generate spurious evidence of negative minimum wage effects on teen employment is clearly not true. Rather, only with a very specific form of controlling for this spatial heterogeneity—and one that is sensitive to recessionary periods at the ends of the sample—does the evidence of negative minimum wage effects fail to appear.[13]

The analysis thus far indicates that simply including state-specific linear trends in a panel data analysis like the kind used to estimate the effects of minimum wages (and many other policies) can do a bad job of capturing other sources

---

[12] The estimates without state-specific trends or with linear trends are reported in columns (1) and (2) of Table 2. In addition, the point estimates are 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 using ADR's sample).

[13] The possibility was also explored that using more flexible functions of the regression controls would serve to reduce the sensitivity of the state-specific trends to the recessionary periods. Specifically, second- and third-order terms of the controls and cross-products of the controls were added to the regressions. The addition of these variables did not qualitatively alter the estimates reported in Table 2. The estimates without the state-specific linear trends were similar to those in Table 2, as were the (smaller, in absolute value) estimates with the trends added. Given what was shown in Figures 3-6 and discussed more generally, this is not surprising. The movement of teen employment rates in the recessionary periods, conditional on the other controls, is not always consistent, so adding these non-linear terms in the regression controls in a manner that still imposes the same structure on all states is not a substitute for more flexible ways of capturing state-specific trends. If we continue to let the strong movements in teen employment in the recessionary periods drive the state-specific trends, then recession-induced movements are still apparently confounded with minimum wage effects.

of long-term trends in the variables of interest. In particular, in the present context, we have seen that sharp business cycle movements can be confounded with these long-term trends in ways that generate misleading evidence. It is more constructive, perhaps, to think about alternative approaches to capture these long-term trends in a way that avoids these problems, aside from simply dropping from the data years with sharp cyclical movements, which can throw out important identifying information.

One possible approach is to modify the empirical specification to estimate the state-specific linear trends using only the data from the subsample that excludes the recessionary periods and then include those trend estimates in a regression over the full sample period. The appeal of this approach is that it allows the state-specific trend estimates to more accurately reflect unmeasured determinants of the longer-run trends in the relative employment rates of teenagers (such as technological change). Indeed ADR say as much: "A state-specific linear trend variable provides a second means of controlling for heterogeneity in the underlying (long-term) growth prospects of low-wage employment and other trends in teen employment" (p. 216). Clearly, letting the strong recessions at the beginning and end of the sample period influence these trends violates the goal of including the trends.

To implement this approach, note that the standard regression model with state-specific linear trends is:

$$(1) \quad Y_{it} = \alpha + \pi MW_{it} + X_{it}\beta + S_i\gamma + T_t\delta + S_i \cdot t \cdot \theta + \varepsilon_{it} \ .$$

We can simply estimate this model with OLS. If we knew θ, however, we could instead form the dependent variable $Y_{it} - S_i \cdot t \cdot \theta$, and estimate the model

$$(2) \quad Y_{it} - S_i \cdot t \cdot \theta = \alpha + \pi MW_{it} + X_{it}\beta + S_i\gamma + T_t\delta + \varepsilon_{it} \ .$$

This specification will give us the same estimates, although the standard errors would be understated because the sampling variation in estimating θ would be ignored.

Following this intuition, we can define a subperiod of the sample, t_1 to t_2, and estimate equation (1) for this subsample. The estimated state-specific trends for this subsample can be retained, denoted by the vector $\hat{\theta}_{12}$. The dependent variable for the full sample period can then be detrended using $\hat{\theta}_{12}$, and equation (2) estimated over the full sample for the detrended data. This removes the state-specific linear trends estimated for the subperiods excluding the steep recessions, rather than the whole sample period, while still using the whole sample to estimate minimum wage effects. The only complication is that the standard errors have to be corrected to account for the sampling variation in $\hat{\theta}_{12}$. This is done by using block-bootstrapping by state, to account for the non-independence of observations within states.

A slightly more complicated version of this estimator is also estimated, which allows different state-specific trends with a break beginning in 2000:Q4,[14] restricting the trend lines to be joined in that quarter (so that the trend can shift, but not the state effects). Denoting this period $t_{1'}$, and defining a dummy variable $D_t^{1'2}$ for the period $t_{1'}$ to t_2, the model estimated is

$$(3) \quad Y_{it} = \alpha + \pi MW_{it} + X_{it}\beta + S_i\gamma + T_t\delta + S_i \cdot t \cdot \hat{\theta}_{12} + S_i \cdot t \cdot D_t^{1'2}\theta_{1'2} + S_i \cdot D_t^{1'2}\gamma_{1'2} + \varepsilon_{it}$$

for the period t_1 to t_2. The analysis then proceeds in the same way, detrending Y_{it} by constructing

$$(4) \quad Y_{it} - \{S_i \cdot t \cdot \hat{\theta}_{12} + S_i \cdot t \cdot D_t^{1'2}\hat{\theta}_{12} + S_i \cdot D_{t1'2}\hat{\gamma}^{'2}\} \ .$$

---

[14] This is near the peak of the economy before the 2001 recession. It is also around the period in which a trend break in teen labor force participation is identified in Aaronson et al. (2007), although here the trend break can occur differently in each state.

As reported in columns (1) and (2) of Table 4, when the panel data model with state-specific trends is estimated in this way—however the trends are treated—the estimated effects of minimum wages are much more strongly negative and nearly always statistically significant. The estimates in columns (1) and (2) can be compared to those for the full period (1990—2011:Q2) in Table 2. The important result is that detrending in this way, which is not sensitive to the business cycle, the estimated minimum wage effect remains negative and significant, ranging from −0.178 to −0.207 (the latter significant at the ten-percent level), vs. −0.165 in Table 2 without the state-specific linear trends, and −0.074 (and insignificant) with them. Thus, removing state-specific trends in a way that is not sensitive to the recessions at the beginning and end of the sample if anything leads to stronger evidence of disemployment effects.

Another possible approach is to estimate the model over the full sample period after pre-filtering the data to remove state-specific trends. This pre-filtering is done in two ways: (1) calculating the trend in each variable as a linear spline between consecutive business cycle peaks as defined by the NBER (and extrapolating the trends over the relevant range of the sample before the first business cycle peak and after the last business cycle peak);[15] and (2) passing each data series (by state) through a Hodrick-Prescott filter using the standard smoothing parameter

for quarterly data and extracting the non-trend component of the series.[16] Estimates of the model using the data detrended in each of these ways are shown in columns (3) and (4) of Table 4. Similar to the results in columns (1) and (2), the coefficients on the minimum wage variable are more strongly negative than when linear trends are included in the model estimated over the full sample period (as in the second column of Table 2).

The general conclusion from these results is quite apparent. The estimated effects of minimum wages on teen employment are negative and significant with or without the inclusion of controls for long-term trends in teen employment when those long-term trends are estimated ways that are not highly sensitive to the business cycle. The only way to generate the results in ADR (2011)—that inclusion of state-specific time trends eliminates the negative effects of minimum wages—is to include in the sample period the recessionary period of the early 1990s or the recent Great Recession, and to let these periods have a strong influence on the estimated trends by use of a highly restrictive specification for those trends. The evidence suggests that the state-specific trends for these sample periods are influenced by the recessions in ways that do not reflect long-run trends and that apparently contaminate estimates of minimum wage effects on teen employment.[17]

---

[15]The rationale for this is perhaps best illustrated by the problem demonstrated earlier in Figures 3-6, namely that recessions can produce quite different patterns in the teen employment variable. Doing this calculation from peak-to-peak is a simple way to avoid this problem.

[16]The Hodrick-Prescott (or HP) filter is a common technique used in the empirical macroeconomics literature to remove stochastic trends from aggregate time-series data prior to estimation by statistical procedures that assume stationarity (see, for example, Hodrick and Prescott, 1997). Effectively, the HP filter is a two-sided symmetric moving average filter that extracts the trend by smoothing the time series, with the degree of smoothing specified in advance by the researcher. The smoothing parameter (λ) penalizes variations in the growth rate of the trend component and can range from zero (in which case there is no smoothing) to ∞ (in which case the smoothed trend is linear). For quarterly data, λ is typically set to 1600, based on research by Ravn and Uhlig (2002) on how the optimal smoothing parameters changes for data of varying frequencies. One issue that can arise in using the HP filter relates to the time-series properties of the original data. In particular, Cogley and Nason (1995) find that the HP filter performs well for time-series that are stationary or trend stationary, but can generate spurious business cycle periodicity for time series that are difference stationary. Accordingly, each series was tested at the state level for the presence of a unit root using the standard Dickey-Fuller test, and in most cases were able to characterize the time series as stationary or trend stationary. That said, the use of the HP filter to detrend the data is better viewed as providing supporting evidence on the robustness of the results to alternative methods of detrending rather than as necessarily the best technique for accounting for state-specific trends in this particular dataset.

Division × Period Interactions. The preceding analysis suggests that there is little reason to be concerned about long-run trends in teen employment contaminating the estimated effects of minimum wages on teenagers. Table 1 also shows, though, that the inclusion of Census division × period interactions renders the estimated minimum wage effect smaller and statistically insignificant. As a prelude to delving into what to make of these estimates, it is useful to consider how ADR see this specification as accounting for the spatial heterogeneity that they think needs to be controlled for in order to obtain unbiased estimates of minimum wage 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 minimum wages increases are endogenous in a manner that generates a bias towards finding negative employment effects. As evidence, they cite Reich (2009), who ADR claim show 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, italics added). This is the type of mechanism that could be controlled for by adding the Census division x period interactions.

One problem with this argument is that 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). So Reich provides no evidence of spurious negative correlations between minimum wage increases and employment shocks. Rather, he shows that there is generally growth during periods of approval and implementation of minimum wages. ADR's argument is also problematic because the regression model already controls 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.

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. 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 can be attributed to controlling for spatial heterogeneity, as ADR argue.

From one perspective, ADR's results are not especially surprising. The inclusion of Census division × period interactions adds over 1,900 dummy variables (there are 20 years of monthly data for 9 divisions) to the specification. An obvious concern is that this extensive set of controls

---

[17]Addison et al. (2009) discuss this issue in somewhat more detail. They note that Sabia (2008) estimates models for retail employment (using CPS data from 1979-2004) and finds negative employment effects without state trends, but not when they are included (they become positive). They write that Sabia argues against including these trends, quoting him as saying that the trends may reduce "potentially important identifying information" (p. 88, cited in footnote 18). However, they then argue that because the standard errors fall when the trends are included, this is not a concern. This argument is disputable. First, if the state-specific trends are mis-estimated because of the recessionary periods, they may still lead to more precise estimates – even though these may not be reliable estimates of minimum wage effects. Second, they omit the continuation of Sabia's argument, that "[S]tate trends may, in fact, be capturing retail employment variation that the model seeks to explain" (p. 88). That is a legitimate concern about the inclusion of state-specific trends, raised earlier with regard to Smith (2010). It is also acknowledged as a problem by DLR: "The sensitivity of the estimates from the traditional specification ... to the inclusion of a linear time trend does not necessarily imply that it is biased. Inclusion of parametric trends may "overcontrol" if minimum wages themselves reduce the employment trends of minimum wage workers ..." (p. 954).

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. This is an issue that is never explored by ADR.

More specifically, 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. Figure 7 provides a picture of the identifying information that is available within Census divisions. The top panel plots the average quarterly minimum wage by division from 1990-2011, and the nine lower graphs plot the minimum wage by state, within division; all values are shown relative to the federal minimum. The graph shows that within some Census divisions there is virtually no variation relative to the federal minimum (e.g., East South Central), in some divisions there is more variation, but little within-division identifying information (e.g., Mountain), and in other divisions there is both variation over time and across states (e.g., New England).[18] Conversely, as the top panel shows, the across-division variation is rather extensive.

What this discussion highlights is that when ADR estimate a much more saturated model, the identifying information is substantially reduced, with identification coming from a much more restricted set of comparisons. Moreover, Figure 7 shows that much of the within-division variation arises during the period coinciding with the Great Recession—exactly the period when the state-specific linear trends appear to misattribute variation in the data to the unexplained trends. Thus, 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 simply estimate the mod-

els separately by Census division, equivalent to estimating the model with a full set of Census division interactions? After all, if divisions have different patterns of employment change that are associated with, but not caused by, minimum wage changes—as implied by ADR's concern with spatial heterogeneity at the division level—then it makes sense to look within divisions, and not across divisions, to identify the effects of state variation in minimum wages. Moreover, 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, equivalent to estimating the model Census division by Census division?

The results, which are reported in Table 5, 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 two other divisions—East North Central and East South Central—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 South Atlantic and Pacific divisions the estimates are near zero.

Looking at the standard errors, and having in mind as a plausible range of elasticities from prior evidence some-

---

[18]Much of the variation in the South Atlantic division is driven by Delaware and Washington, DC.

thing like -0.1 to -0.2, 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. And, in the three divisions that yield precise estimates, the estimated elasticities are negative and significant (and roughly in this range).[19] 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.

The least restrictive approach of estimating the models separately by Census division often leads to very imprecise estimates, as does adding all of the division × interactions. This naturally raises the question of whether, in estimating models that control for spatial heterogeneity, it is really necessary to throw out information on other potential comparison states on the *presumption* that states in other divisions cannot serve as effective controls because of spatial heterogeneity.

The assumption that ADR make, but do not test in any direct way, 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, ADR 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). One might have expected them to provide convincing evidence that the counterfactual employment growth that comes from states in other Census divisions does not provide a good control, yet they fail to do so.[20] 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.

The question of which states are good controls for the states with minimum wage increases is addressed in two ways. The first approach uses the initial step of the Abadie et al. (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 in others. The latter—which are the potential control units, and are referred to as "donor" units—are selected based on a matching estimator, with the choice of variables on which to match 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.[21]

To draw a comparison, 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

---

[19]This would not necessarily have been predicted from looking at the variation in Figure 7. On the one hand, New England and West North Central exhibit a fair amount of within-division variation in minimum wages. But West South Central has only one change, for Arkansas. In contrast, other divisions that display substantial variation – such as South Atlantic – do not yield precise estimates.

[20]Even if one accepted the notion that geographically-proximate states provide better controls, it raises the following question: Are states within a Census division better controls than closer states in other Census divisions? ADR, like DLR, do present some other indirect evidence that their specification better captures spatial heterogeneity. This evidence is considered later, after discussing both papers, and shows that this indirect evidence is much less persuasive than the authors claim.

state raised its minimum wage, the standard 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 even a 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. Nothing in ADR's study, however, establishes that states in the same division are better controls (just as nothing in the Card and Krueger study establishes that Pennsylvania is a good control for New Jersey). The synthetic control approach provides us with empirical evidence on which states are the best control states.[22]

This approach was applied to the CPS data used in the preceding analyses. In particular, a set of treatment observations in the data set are defined as state-quarter observations with minimum wage increases in that quarter and no minimum wage increase in the previous four quarters; this yields a set of 129 potential minimum wage treatments to analyze. For each one of these, a set of potential donor units (state-quarter observations) is defined as states 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. In these analyses, a set of variables is chosen on which to match over the four quarters prior to the treatment, 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. A finding that most of the estimation weight is on donor states in the same Census division as the treatment state would rationalize ADR's approach; states in other divisions would not match as well because those other regions have different prior trends. Conversely, if only a little weight is put on state-quarter pairs in the same division, this would tell us that there is no good rationale to restrict (or even focus) attention on the states in the same Census division, either because spatial heterogeneity is not important, or if it is, because it is not specific to the Census divisions used by ADR.[23]

As noted above, this approach also requires a choice of variables on which to match. Four different alternatives were used. Three of these involve matching on forms of the dependent variable; the log of the teen employment-to-population ratio was used, 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.[24] Finally, matching was done on the residuals from the standard panel data estimator for teen employment (Table 1, column (1)), again for the prior four quarters. This is not a standard type of variable to match on, primarily because there typically is not a regression model underlying the application of the synthetic control approach; rather, the

---

[21]After ascertaining the match quality of each potential control unit, the method is then used to estimate the treatment effect by weighting control units based on quality of the match. In the present case, however, the focus is only on the identification of which states are better matches as controls, because in the minimum wage setting there is continuous variation in the treatment, and there are multiple increases, so it is unclear how to implement the second step.

[22]Sabia et al. (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 estimates are very similar, because the approach puts much of the weight on the geographically proximate states. However, they do not match on lagged values of the dependent variable.

[23]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 clearly 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, whether there is substantial weight on donor states in the same division is considered only when there are such donor states, to avoid overstating of the extent to which donor states come from other divisions.

[24]As explained in the notes to Table 6 when the matching on the one- and four-quarter changes, treatment observations are lost at the beginning of the sample period.

synthetic control estimator is typically used *instead* of a regression model.[25] However, matching on the residuals is informative about the spatial heterogeneity arguments that ADR advance, 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 will tell us whether the residuals for states in the same region share these commonalities and hence whether these states are good controls—or, in ADR's approach, should be controlled for with the period × division dummy variables.

The results are summarized in Table 6. As noted, there are 129 unique treatments as they are defined for this analysis. Of these, 50 have potential donors in the same Census division, covering six divisions. The key results are reported in columns (1)-(4); these are the weights from the matching process on states in the same division. 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.[26] Columns (5)-(7) report on the average number of divisions and states in the donor pool, and the average number of states in the same division. One thing we see from these columns is 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, East North Central has the second highest average number of potential donor states from the same division but weights close to zero, while West North Central has one of the lowest number of potential donor states from the same division but the highest weights.

These results provide striking evidence against ADR's choice to restrict the control states to be states 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, it may well be that the conventional panel data estimator provides more reliable estimates of minimum wage effects.

A second—and perhaps more transparent—method is also used to address the question of which states are good controls for the states with minimum wage increases, labeled the "ranked prediction error" approach. The synthetic control approach finds a weighted average of the potential donor counties to best match the treatment unit. Comparing the weight that this method assigns to the control units used by ADR—in particular, the weight on states in the same Census division as the treatment state—is informative about whether the same-division states in fact provide a good control. The second method, instead, matches up the treatment unit to each potential donor one-by-one. For each of these potential controls, the root mean squared prediction error (RMSPE) is 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), for the exact same matching variables used for the synthetic control approach. 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 and other division states.

Some notation helps to clarify the method and the difference between the two approaches. Denote a specific treatment unit by T,[27] the potential donors in the same

---

[25]However, Abadie et al. (2011) show that their estimator is similar to a regression in that a regression-based approach can also be expressed as a weighting estimator, where the weights are "unrestricted and may take on negative values or values greater than one" (p. 8).

[26]Code in R was used to do these calculations. The software is available at http://www.mit.edu/~jhainm/synthpage.html (viewed July 30, 2012).

division $D^s_1, ..., D^s_J$, and the potential donors in other divisions $D^o_1, ..., D^o_K$. The synthetic control approach finds a weight for each donor, $w^s_1, ..., w^s_J$ and $w^o_1, ..., w^o_K$, 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.

What is done now is to calculate the RSMPE for each potential donor, denoted r^s_j, j = 1,..., J, or r^s_k, k = 1,..., K, for the same-division and other-division donor states respectively. These r's are then ordered based on how well they match the treatment unit based on lower RMSPE. 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.[28]

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

tograms for these percentiles are constructed to see if they are in fact clustered higher than would be expected if the same-division states were on average no better or worse controls than other states (or, equivalently, if the distributions of the percentiles appear approximately uniform).

Figure 8 displays an example of the first step of this process. This figure simply focuses on 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., Hawaii in the lower-right panel), or they can provide quite bad matches, with relatively high RMSPEs (e.g., Hawaii in the upper-left panel).

Figure 9 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, and in fact 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.

Finally, Figure 10 plots, for each of these analyses, the medians of these percentiles for each of the Census divi-

---

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

[28]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 $(100 \times \text{rank})/(N+1)$.

sions. The figure shows that 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 6, notice that the synthetic control approach 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 that identifying information be confined to states in the same division. There is one notable exception—and that is for the West North Central division.

Finally, a comparison of the estimates in Table 5 with the results from the synthetic control or ranked prediction error approach is instructive. Table 6 and Figure 10 both 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 5 shows that when minimum wage effects are estimated for this division in isolation, there is statistically significant evidence of a negative employment effect of minimum wages, with an estimated elasticity -0.19—very much in line with much of the existing evidence.

Furthermore, Table 6 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 5 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 6 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. These are two cases that, in Table 5, do not provide any evidence of disemployment effects. Thus, while this analysis does not pin down one "best" estimator, it does indicate that (a) there is generally 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 is generally quite similar to the prior evidence suggesting that teen minimum wage employment elasticities fall into or near the -0.1 to -0.2 range.

Dube et al. (2010)

As noted earlier, Dube et al. (DLR) focus on restaurant employment with county-level QCEW data from 1990-2006. They show that the standard 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 this paper is on a research design based on cross-border county pairs. When DLR 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 state-by-year specifications in ADR relies on the within-Census division variation. Given that this identification strategy is the key contribution of this paper, the focus is on their cross-border analysis of the effects of minimum wages on restaurant employment rather than on their analyses that more closely parallel ADR.

The key estimates from this approach are reported in Table 7, replicating the results in DLR (Table 2, specifications 5 and 6); the estimates are nearly identical to theirs. The first two columns use the subset of counties in the contiguous border county analysis, but include only

county and period 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 near zero, positive, and statistically insignificant.

Prior to getting to the main line of inquiry regarding this paper, it is worth noting that DLR substantially overstate the number of cross-border county pairs that are used to identify the effects of minimum wages in this approach. Their Figure 2 claims to show all the state borders—and counties along them—that are used in their analysis. This figure is replicated in Panel A of Figure 11, and includes 81 distinct state border pairs. 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. Panel B shows the corrected figure, which clearly includes far fewer state borders. For example, while the top panel suggests that there is identifying information along the Michigan, Indiana, and Ohio borders, the bottom panel shows that there is in fact no minimum wage variation along these borders. All told, there are only 48 distinct state border pairs with identifying information.[29]

However, the main question concerns the underlying assumption in DLR's identification strategy—that the cross-border contiguous county in the state that did not raise its minimum wage is the best control group for the county in the state that did raise its minimum. As they point out, this has close parallels to the type of analysis in Card and Krueger (1994), who studied the 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. How strong is the evidence for their assumption that the cross-border contiguous county is the appropriate control group? Like in ADR, the authors present no direct evidence validating this research design. Instead, they only speculate:

> "Contiguous border counties represent good control groups for estimating minimum wage effects if ... a county is more similar to its cross-state counterpart than to a randomly chosen county. In contrast, panel and period fixed-effects models used in the national-level estimates implicitly assume that one county in the United States is as good a control as any other" (Dube at al., 2010, pp. 949-50).

To address this question of the quality of cross-border contiguous counties as controls, the synthetic control matching estimator is again used—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 pre-

---

[29]Note that Panel B has the counties on the North and South Dakota border shaded, whereas Panel A does not. There was a higher minimum wage in North Dakota (by five cents) in the first three months of 1990. There is also another slight discrepancy. DLR's data includes a higher minimum wage in Maryland in the first six months of 2006 ($6.15, instead of the $5.15 federal minimum wage that actually prevailed). This correction would eliminate two additional state border pairs (Maryland's border with Virginia and West Virginia). However, the research described in this report uses their original data and making this correction to the data did not materially change the results. Finally, note that, as DLR do, a balanced sample of counties with data in all periods is used. That restriction drops some counties along the borders shown in the figures.

vious analysis.[30] 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 other analysis, the set of donor counties is restricted to simplify the synthetic control calculation. In particular, the first step is to calculate, for each treatment and the potential donor counties, the RMSPE of each donor county for the four quarters following the minimum wage increase. Then the 50 donor counties with the lowest RMSPE are used as potential donors, adding DLR's contiguous cross-border counties if they are not 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 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. 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 8. Panel A reports the

analysis with the smaller donor pools. For seven of the nine Census divisions there are counties with potential contiguous cross-border donors. However, the weight put on the contiguous cross-border counties is almost always very small. For example, for the regression residuals (which are most pertinent to the critiques in ADR and DLR), the weights average less than 0.05 for five of these divisions, and 0.033 overall. (In one case—column (3) for the Mountain division—the weight averages 0.216.) The results using the full donor pools are reported in Panel B. The conclusions are similar. In particular, the weight put on the contiguous cross-border counties is again almost always very small—and almost always quite a bit smaller than in Panel A.

This analysis reveals that the weight put on the cross-border contiguous counties as controls—the only controls DLR use—are very small. For example, in Panel A, where there are typically just over 50 possible controls for each treatment, the weight put on the cross-border contiguous counties is 0.033 in column (1), and of a similar if not smaller magnitude in the other columns. Given that there are on average 1.7 cross-border contiguous counties in the donor pool, and the donor pool on average includes 51.3 counties,[31] if we put equal weights on each county in the donor pool that weight would be 0.033 (1.7/51.3). In other words, within these donor pools the control counties that DLR use appear no better than a random draw. The same is true in Panel B. In column (1) the weight on DLR's control counties is 0.007, only slightly above the share of these latter counties in the potential donor pool (0.005).

Next, the ranked prediction error approach is again used to assess whether contiguous cross-border counties are

---

[30]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).

[31]Note that the latter number need not be 51.7 (i.e., 50 + 1.7), because the contiguous cross-border counties are sometimes among the 50 counties with the lowest RMSPE.

better controls than other counties. The method is the same as before, but now simply looking at the percentiles for the matched contiguous cross-border counties (analogous to the previous examination of same-division states). Figure 12 gives an example of the distribution of RMSPEs for one particular treatment unit (San Bernardino County, California in 2001:Q1), with the RMSPEs of the contiguous cross-border counties highlighted.[32] Again, the figure reveals that some cross-border counties can provide good matches, while others may not.

Figure 13 then 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 then, there is little in the data to support the assumption made by DLR that the contiguous cross-border counties are the appropriate controls.[33]

One might argue that it is not surprising that the contiguous cross-border counties get so little weight or fail to stand out as the best control areas. After all, as shown in column (5) of Table 8, there is 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?

Is there other evidence that justifies the approach in ADR and DLR?

ADR and DLR present two 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. The first is an analysis of employment changes prior to the implementation of minimum wages, and the second is pitched as a falsification test showing that county employment appears to respond to cross-border minimum wage changes.

In the first analysis, DLR estimate dynamic models for employment and earnings with long leads (and sometimes lags), and then present figures showing cumulative effects at these leads and lags based on the coefficients from these models. For example, in their Figure 2, ADR plot the cumulative effects from leads of two years to lags of four years for the standard panel data specification with fixed state and period effects, as well as for the specification that adds both the state-specific linear time trends and the Census division × interactions. The replication of their graphs for employment effects appear in the upper-left and upper-right panels of Figure 14. According to ADR, the evidence of leading effects in the upper-left panel, in comparison with the evidence in the upper-right panel, provides "strong evidence against the model without controls for heterogeneity across states ..." (p. 220).

---

[32]In some cases, like this one, DLR match a treatment county to multiple adjacent counties across a state border. The percentiles for each of these are collected and displayed in the figure that follows.

[33]To their credit, Addison et al. (2009) at least address this issue. They note an example of a cross-border county match that is quite bad, with a 3.5% unemployment rate for one observation and a 7.7% unemployment rate for a matched county, and suggest that "such examples of poor matches across state borders could be rather common" (p. 406). Nonetheless, they still use this method to estimate minimum wage effects. Another possible reason that cross-border contiguous counties may not be good controls is that they often are affected by the minimum wage on the other side of the border. For example, if disemployment effects in the treated county lead workers to cross the border for jobs, then the disemployment effects would be overstated. Conversely, if the minimum wage increase induces workers from the cross-border control county to come to the treated county to search for jobs, then employment in the control county can fall.

However, this evidence appears to be both overstated and misleading. First, note that the figure also suggests that when the state-specific linear time trends and the division × period interactions are included, there appear to be fairly substantial positive effects of the minimum wage on teen employment—with elasticities of about 0.2—three to four years after the minimum wage increase. Although ADR do not remark on this, it seems unlikely that those estimates represent the real effects of minimum wages and hence they could be viewed as providing evidence against the model that includes those controls.

Second, even though the regression estimates for their main analyses are based on quarterly data, the graphs they show in their Figure 2 (replicated here in the top row of Figure 14) are generated from a model specified with leads and lags on an annual basis. The bottom two panels of the figure show the graphs from a less constrained model that specifies the leads and lags at a quarterly frequency. Not surprisingly, the plots are noisier. The quarterly graphs appear to show three things: (1) they do not give a clear indication of any kind of pre-trend in the standard panel data model; (2) they point to negative employment effects in the two years after the minimum wage increase that look quite distinct from anything occurring in the data prior to the minimum wage increase; and (3) they show significant and positive employment effects (with elasticities in the 0.4 to 0.5 range) more than three years after the minimum wage increases in the model with state-specific linear trends and division × period interactions, which seems to be strong evidence against that model.

DLR also present evidence suggesting that the inclusion of state-specific linear trends and Census division × pe-

riod interactions in panel data eliminates spurious negative estimates of the effects of minimum wage for periods without a minimum wage increase. In this case, their analysis is captured succinctly in their Table 3, which—for the model with just county and period fixed effects—reports the "effect" three years prior to the minimum wage increase, one year prior, and the "pre-trend" based on the difference between the cumulative effects at these two points.[34] The replication of their results is reported in column (1) of Table 9; again, the results are nearly identical to theirs. The evidence points to a fairly large cumulative negative "effect" at a four-quarter lead, and also a negative trend 12 quarters to four quarters prior to the minimum wage increase; both estimates are significant at the ten-percent level. These same results, in somewhat more detail, are reported in their Figure 4, which shows a growing cumulative negative effect up to the minimum wage increase, although this cumulative leading "effect" is never statistically significant. This is replicated in the upper-left panel of Figure 15.

Note, however, their figure also shows a positive cumulative leading effect for the model with state-specific linear trends and division × period interactions, and in this case the cumulative effect is sometimes statistically significant; see the upper-middle panel of Figure 15, which also replicates their Figure 4. Not only do DLR fail to remark on this result; they essentially deny it, claiming (on p. 956) that this specification shows "relatively stable coefficients for the leads centered around 0."[35]

DLR also omit the estimates of this specification from their Table 3—estimates on the basis of which they criticize the "canonical" panel data model as showing a pre-

---

[34]These models include a contemporaneous effect of the minimum wage, two leads at three years (12 quarters) and one year (4 quarters), but no lagged minimum wage effects. The results discussed below for this specification and variants thereof are very similar if the intervening semi-annual leads (which they include in the specification on which the figure discussed below is based) are included as well.

[35]To be precise, this quote refers to a different specification (specification 6), but they then say that the results in question are similar: "Intermediate specifications … with coarser controls for heterogeneity in employment show similar results to the local specification (6)" (p. 956).

existing negative trend. The corresponding estimates for the model including the state-specific linear trends and the division × period interactions (their specification 3) were computed, and are reported in column (2) of Table 9. These estimates show a significant positive "pre-trend" of the same magnitude as the negative pre-trend obtained from the standard model. And since negative anticipatory effects of a minimum wage are at least in principle plausible (as acknowledged by ADR, p. 220), one might argue that the positive pre-trend raises particular doubts about the specification with state-specific trends and division × period interactions.

Moreover, there is also no obvious rationale for focusing only on the pre-trend calculated between the 4th and the 12th quarters. If the results were robust to which interval was used, it would be of little consequence. However, Panel A of Table 10, which retains the specification with 12-quarter leads, but computes the pre-trend between the 2nd, 6th, 8th, and 10th quarters as well, shows that this is not the case. As columns (1) and (2)—which otherwise correspond to the specifications in Table 9—show, it is only for the 4th-to-12th quarter interval that the "canonical" model (column (1)) generates a statistically significant pre-trend, and for many of the other intervals it is quite small. In contrast, for the model with county and period fixed effects, state-specific linear trends, and division × period interactions, the pre-trend is significant and positive in every case. Column (3) shows the similar estimates for the contiguous county-pair sample. These estimates are all statistically insignificant; however, the standard errors are large, and in many cases the coefficient estimates are of roughly the same absolute magnitude as the estimates in column (1).

As Panel B of Table 10 shows, the evidence for their claim is even weaker if regression models that correspond to their Figure 4—with leads only up to 8 quarters—are estimated. In this case, the canonical model shows no evidence of pre-trends; the estimates are small and statistically insignificant. However, for the model with county and period fixed effects, state-specific linear trends, and division × period interactions, there is again robust evidence of a positive pre-trend. And for the contiguous county-pair sample, the point estimates in two of three cases are large and positive, although insignificant given the very large standard errors.

The table highlights in bold and italics the two estimates reported in DLR's Table 3 for this specification.[36] It is quite clear that the estimates DLR emphasized are the ones that most strongly make their case, and 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.

Finally, Figure 4 in DLR—replicated in the top three panels of Figure 15—parallels the figure in ADR by using data at a smoother semi-annual frequency than the quarterly frequency used in all their model estimates. As shown in Panel B of Figure 15, when the figures are replicated using the data on a quarterly basis, the figure for specification 1 is much less suggestive of any kind of pre-treatment trend.[37] In fact, the only pronounced negative estimate is one quarter before the treatment. In contrast, the figure for specification 3 still shows a pronounced upward trend prior to the treatment. And, to us at least, it is not obvious that the pre-trend apparent for specification 6—DLR's preferred contiguous county-pair analysis—is less problematic. It is true that the estimates are much less precise, as indicated by the wider confidence intervals. But there is a noticeable increase in employment in

---

[36]There are alternative specifications with a control for private-sector employment added, but the results are similar.

[37]How one views these graphs is partly subjective. But this interpretation appears to be more accurate than DLR's claim that "Using leads and lags for every quarter, as opposed to every other quarter, produces virtually identical results" (p. 956, footnote 24).

the two quarters preceding the minimum wage increase, and the decline subsequent to that is roughly the same order of magnitude as the decline in the figure for specification 1.

In a second analysis, DLR present results from what they refer to as a falsification test, which in their view proves that spatial heterogeneity is responsible for the apparent evidence of minimum wage effects in conventional panel data estimates (p. 958). In particular, they define a narrow sample of all border counties where the minimum wage was never above the federal minimum wage in the sample period, and then estimate the standard 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.[38] These estimates are replicated in column (1) of Table 11.

However, one possible reason why DLR find an unexpected negative employment effect in their placebo sample is that they do not actually have 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.[39] 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, there is not evidence of a placebo effect when the sample used for DLR's falsification test is modified to avoid having a contaminated placebo sample. Specifically, the sample is restricted to observations after the federal minimum wage increase in 1997, so that there is no federal minimum wage variation in the placebo counties that is captured by the counties matched to them.[40] As shown in column (2) of Table 11, 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.[41] 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

---

[38]See their Appendix B for further information on the implementation of their falsification test.

[39]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.

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

[41]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.

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; it is this variation that is informative about the falsification test. When this is done, in column (3), 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 yet additional evidence refuting DLR's claim that spatial heterogeneity generates spurious evidence of disemployment effects of minimum wages.

Discussion

Echoing long-standing concerns in the minimum wage literature, the studies by ADR and DLR 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. There is some intuitive appeal to what they do, as their analyses have parallels to the "case-study" method of estimating minimum wage effects best typified by Card and Krueger (1994).[42] The problem is that they do nothing to assess the validity of the different controls they effectively obtain by using these methods. Indeed, the analysis suggests that the reason they obtain evidence at odds with disemployment effects is that they throw out lots of valid identifying information and likely end up with less valid controls.

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 puz-

zling 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 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 and thus what types of potentially identifying information are valid, it is important to ask the following question of all of their results: Out of their concern about avoiding minimum wage variation that is potentially confounded with other sources of employment change, have they thrown out so much useful and potentially valid identifying information that their estimates become uninformative or invalid? That is, have they thrown out the "baby" along with—or worse yet, instead of—the contaminated "bathwater"? The analysis described in this report suggests they have. Moreover, despite the claims made by ADR and DLR, the evidence that their approaches provide more compelling identifying information than the standard panel data estimates that they criticize is weak or non-existent. Indeed, some evidence suggests that the standard panel data model provides better identification than the methods they use.

---

[42]Indeed, DLR describe the estimation using contiguous cross-border county pairs, with county pair × period interactions, as "producing a pooled estimate from individual case studies" (2010, p. 957).

Other Recent Work on Minimum Wages

There are a number of other significant studies of the employment effects of minimum wages that have been written in recent years, some of which also depart from the standard state panel data analysis that was often used in the new minimum wage research. While it is beyond the scope of this report to provide a detailed analysis and critique of these papers, it is noteworthy that the evidence from these studies tends to point to *stronger* evidence of disemployment effects of minimum wages. Moreover, these studies raise substantively important issues and, in some respects, provide more compelling reasons or evidence to pursue their extensions of the standard state panel data approach than do DLR and ADR.

For example, Thompson (2009) notes that past studies estimating minimum wage effects at the state level ignore considerable within-state heterogeneity in wage levels and local labor market conditions. This heterogeneity implies that minimum wages may be binding for teenagers—and have negative employment effects—in many regions within a state, but that the aggregation to the state level can mask these negative effects.[43]

Thompson argues that counties better represent labor markets for teens than do entire states, given constraints that keep teens close to where they live. He uses data from the Census Bureau's Quarterly Workforce Indicators, which is based on the Longitudinal Employer-Household Dynamics (LEHD) program. One cannot calculate a teen employment rate from these data, so Thompson instead uses as the dependent variable the teen share of total employment.[44] Numerous controls are constructed from these data, and average quarterly earnings of teenagers in

the quarter before a minimum wage increase are used to capture how binding minimum wages are likely to be in a county and period. The heterogeneity in this average across counties within states is quite striking, providing a strong rationale for looking at differences in minimum wage effects in low- versus high-earnings counties.

The key estimates are based on a standard panel data, difference-in-differences estimator, which expands on the usual state-level analysis by allowing for differences in effects between high- and low-earnings counties. In particular, the analysis focuses on the 1996 and 1997 increases in the federal minimum wage, using data from the first quarter before and after each of the increases only on the states for which the increase in the federal minimum wage raised the effective minimum wage in the state. The specification for the teen employment share includes, like in other analyses, adult labor market controls, the teen population share, a dummy variable for the period after the federal increase, a dummy variable for low-earnings counties, and—to identify the differential effect of the federal minimum wage increase on the low-earnings (or high-impact) counties—an interaction between the dummy variable for low-earnings counties and the dummy variable for the period after the federal minimum wage increase.[45] In addition, some specifications include state fixed effects, in which case the effect of the minimum wage in low- versus high-earnings counties is identified only from the within-state differences in the impact of the minimum wage increase on these counties.

For both of the federal minimum wage increases, the estimates uniformly show an adverse employment effect on teenagers in the counties where teenagers have low

---

[43]This is one of very few papers that focus on variation in minimum wage effects depending on how binding minimum wages are, although of course the emphasis on teens or other low-skilled workers in much of the literature is a nod in this direction. An earlier attempt to account for how binding minimum wages are at the state level is Neumark and Wascher (2002). Thompson's approach is less structural, and focuses (importantly, it would appear) on within-state variation in how binding minimum wages are.

[44]This dependent variable has been used in other research on minimum wages as well.

[45]There are also continuous versions of this specification.

average earnings. Note, also, that Thompson's analysis considers the types of factors that DLR and ADR suggest can lead to spurious evidence of disemployment effects of minimum wages. First, by identifying the effects from differences between counties with low versus high teen earnings, they control 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 across-county variation within states, any endogeneity of state minimum wages is unlikely to be a confounding influence. Second, in a "placebo test," he uses the same methodology to estimate minimum wage effects a couple of years later when the federal minimum wage did not change. In this case Thompson finds no effect, suggesting that differential trends for low- and high-earnings counties do not drive the results. An implication of these analyses is that we should be skeptical of the claim made by DLR and ADR, that spatial heterogeneity leads to spurious evidence of disemployment effects. The problem, instead, may be that the poorly-motivated estimators DLR and ADR use do not correct for spatial heterogeneity, but instead yield spurious evidence that minimum wages do not reduce employment of low-skilled workers.

Of course, as Thompson points out, these results do not imply that state level analyses give misleading estimates of the state-level impact of minimum wages (p. 343). However, the results do show that the workers for whom the minimum wage would be expected to have the greatest impact—and for whom, therefore, minimum wages potentially offer the largest wage gains—are often adversely affected by the minimum wage. As Thompson suggests, this is a "geographic" perspective on the same point made by Neumark and Wascher (2008) that the disemployment effects of minimum wages are more apparent when the focus is on those individuals most likely to be affected.

A second example of the latest research is Baskaya and Rubinstein (2011). This study follows the more conventional state-level panel data framework in estimating the effects of minimum wages on teenagers, but explicitly focuses on the endogeneity of state minimum wage changes—in particular the variation that is not driven by federal minimum wages. The core of the authors' approach is to recognize that in states that have tended to let the federal minimum wage be binding (by having a lower state minimum wage or no state minimum wage), the variation in the effective minimum wage is primarily federal. In these states, the minimum wage variation used to identify minimum wage effects is less likely to be influenced by state economic conditions, and is therefore more likely to reveal the causal effects of minimum wages. Conversely, in states where the effective minimum wage tends to be determined by the state minimum wage, endogeneity is more likely to be a problem.

They take a couple of empirical approaches, but their main estimates seek to eliminate the endogeneity of minimum wages by instrumenting for the effective state minimum wage (the usual variable used in state-level panel data analyses) with interactions between the federal minimum wage and the state's propensity to let the federal minimum wage be binding.[46] The variables that underlie this propensity include the political ideology of the state in the 1960s, per capita income in the 1960s, and lagged data on whether the federal minimum wage was binding in the state (sometimes time-invariant, by considering this propensity only in a fixed period before their analysis sample). The idea behind these instruments is that the federal minimum wage is exogenous to state economic conditions, and will have a greater impact on the effective minimum wage in states that have historically let the federal minimum wage be binding; this latter variation is also likely to be exogenous because it is inferred from

---

[46]Note that this instrument varies by state and year, so that they can estimate this model including fixed state effects.

historical rather than the current decision about which minimum wage is binding, which could be endogenous.

In the data set the authors use, which is constructed from May CPS files extending from 1977-2007, the conventional panel data estimator yields negative and significant effects of minimum wages on teen employment. But these weaken if Census region × year interactions are added (paralleling some of ADR's results), and become insignificant when richer year effects are added that have some parallels to state-specific linear trends (see their Table 3). But when they implement their instrumental variables procedure, they find robust evidence of a strong negative effect of minimum wages on teen employment, with an elasticity of around -1. The authors also present a number of other analyses that bolster the validity of their identification strategy.

In a sense, by worrying about the endogeneity of state minimum wages, the Baskaya and Rubinstein study is concerned with the same issues as ADR and DLR. However, ADR and DLR simply posit—with neither supporting evidence nor a compelling argument—alternative identification strategies. In contrast, Baskaya and Rubinstein both present a compelling argument for their identification strategy, and present auxiliary evidence that backs it up. First, they present evidence from regressions of effective state minimum wages on lagged labor market conditions and find that state minimum wages tend to be set lower when the labor market is weak. Second, this effect is attenuated (essentially completely) in states that let the federal minimum wage be binding.[47] These results support their contention that state minimum wage increases are procyclical, so that failure to account for endogeneity masks the negative effect of minimum wages, biasing the estimates toward zero. And they provide some

evidence that bolsters their strategy of exploiting federal minimum wage variation to identify the effects of the minimum wage on teen employment, showing that it is the minimum wage variation induced by the interaction between the federal minimum wage and its propensity to be binding that generates their negative employment effects, rather than additional state-level variation in the minimum wage that is correlated with the instrument.

The identification strategy that Baskaya and Rubinstein use merits further consideration. Nonetheless, the paper provides an example of a study that considers what ADR and DLR would term "spatial heterogeneity" in minimum wage effects and finds strong evidence of disemployment effects of minimum wages. This is directly at odds with the conclusions from the problematic methods that ADR and DLR use to address this issue.

Conclusions

In two recent studies, Dube et al., 2010 and Allegretto et al., 2011 present evidence and a forceful critique of much of the prior research on the employment effects of minimum wages. They argue that the evidence of negative employment effects for low-skilled workers is spurious, and generated by other differences across geographic areas that were not adequately controlled for by researchers. And they put forth a self-proclaimed "fourth generation" of minimum wage studies that control for this spatial heterogeneity and conclude that once one does this, there are "no detectable employment losses from the kind of minimum wage increases we have seen in the United States" (Dube et al., 2010, p. 962).

However, the analysis described in this report suggests that their methods are flawed and give misleading answers. In particular, neither study makes a compelling ar-

---

[47]Of course this does not establish the endogeneity of teen employment rates conditional on adult unemployment rates, or the direction of the bias that this endogeneity would create.

gument that its methods isolate more reliable identifying information (i.e., a better counterfactual). In one case— the issue of state-specific trends—the research described in this report explicitly demonstrates the problem with their methods and shows how more appropriate ways of controlling for unobserved trends that affect teen employment lead to evidence of disemployment effects that is similar to 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—the same research shows that the exclusion of other regions or counties as potential controls is not supported by the data. Moreover, when it is supported by the data, the evidence is again consistent with past findings of disemployment effects.

In addition to the flaws in the studies by ADR and DLR that the empirical analysis in this evaluation identifies, other recent studies more convincingly address potential biases in state-level panel data that the ADR and DLR studies were intended to address, and find negative employment effects of minimum wages. Together, the evidence invalidates 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 convincing than the established research literature that finds disemployment effects, and understand why the studies in that literature would have generated misleading evidence.

The research described in this report indicates that the studies by Allegretto et al. (2011) and Dube et al. (2010) claiming to overturn evidence that minimum wages reduce employment—like earlier prominent studies making this claim (e.g., Card and Krueger, 1995, relying heavily on their 1994 study)[48]—fail to meet these standards. Instead, the research record 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.

---

[48]See the critique of this study in Neumark and Wascher (2000).

References

Aaronson, Daniel, Park, Kyung-Hong, 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. 2011. "Comparative Politics and the Synthetic Control Method." Working Paper No. 2011-25, MIT Political Science Department.

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, June, pp. 493-505.

Addison, John T., McKinley L. Blackburn, and Chad D. Cotti. "The Effect of Minimum Wages on Labour Market Outcomes: County-Level Estimates from the Restaurant-and-Bar Sector." Forthcoming in *British Journal of Industrial Relations*.

Addison, John T., McKinley L. Blackburn, and Chad D. Cotti. 2011. "Minimum Wage Increases in a Soft U.S. Economy." Economic Series No. 273, Institute for Advanced Studies, Vienna.

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.

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, April, pp. 205-240.

Ashenfelter, Orley, and Robert S. Smith. 1979. "Compliance with the Minimum Wage Law." *Journal of Political Economy*, Vol. 82, No. 2, April, pp. 333-50.

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, June, 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, October, 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, October, 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, September, pp. 772-93.

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

Cogley, Timothy, and James M. Nason. 1995. "Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research." *Journal of Economic Dynamics and Control*, Vol. 19, January-February, pp. 253-278.

Congressional Budget Office. 2004. "What Is Happening to Youth Employment Rates?" Available at http://cbo.gov/sites/default/files/cbofiles/ftpdocs/60xx/doc6017/11-18-youthemployment.pdf (viewed March 12, 2011).

Douty, H. M. 1960. "Some Effects of the $1.00 Minimum Wage in the United States." *Economica*, Vol. 27, No. 196, May, pp. 137-47.

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, November, pp. 945-64.

Even, William E., and David A. Macpherson. 2010. "The Teen Employment Crisis: The Effects of the 2007-2009 Federal Minimum Wage Increases on Teen Employment." Employment Policies Institute, Washington, DC.

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, October, pp. 473-96.

Giuliano, Laura. "Minimum Wage Effects on Employment, Substitution, and the Teenage Labor Supply: Evidence from Personnel Data." Forthcoming in *Journal of Labor Economics*.

Hirsch, Barry T., Bruce E. Kaufman, and Tetyana Zelenska. 2011. "Minimum Wage Channels of Adjustment." Andrew Young School of Policy Studies Research Paper Series No. 11-34.

Hodrick, Robert J., and Edward C. Prescott, 1997. "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit and Banking*, Vol. 29, No. 1, February, pp. 1-16.

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, October, 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, July, pp. 67-82.

Neumark, David, J. M. Ian Salas, and William Wascher. 2012. "Revisiting the Minimum Wage-Employment Debate: Throwing Out the Baby with the Bathwater?" Unpublished paper.

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

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, July, pp. 712-46.

Neumark, David, and William Wascher. 2000. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Comment." *American Economic Review*, Vol. 90, No. 5, December, pp. 1362-96.

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

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, January, pp. 223-46.

Neumark, David, and William Wascher. 2002. "State-Level Estimates of Minimum Wage Effects: New Evidence and Interpretations from Disequilibrium Methods." *Journal of Human Resources*, Vol. 37, No. 1, Winter, pp. 35-62.

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

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, October, pp. 55-81.

Ravn, Morten O., and Harald Uhlig. 2002. "On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations." *The Review of Economics and Statistics*, Vol. 84, No. 2, May, pp. 371-380.

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

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, March, 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, April, pp. 350-76.

Smith, Christopher L. 2010. "The Polarization of the U.S. Adult Labor Market and its Effects on the Demand for Teenage Labor." Unpublished paper, Federal Reserve Board of Governors.
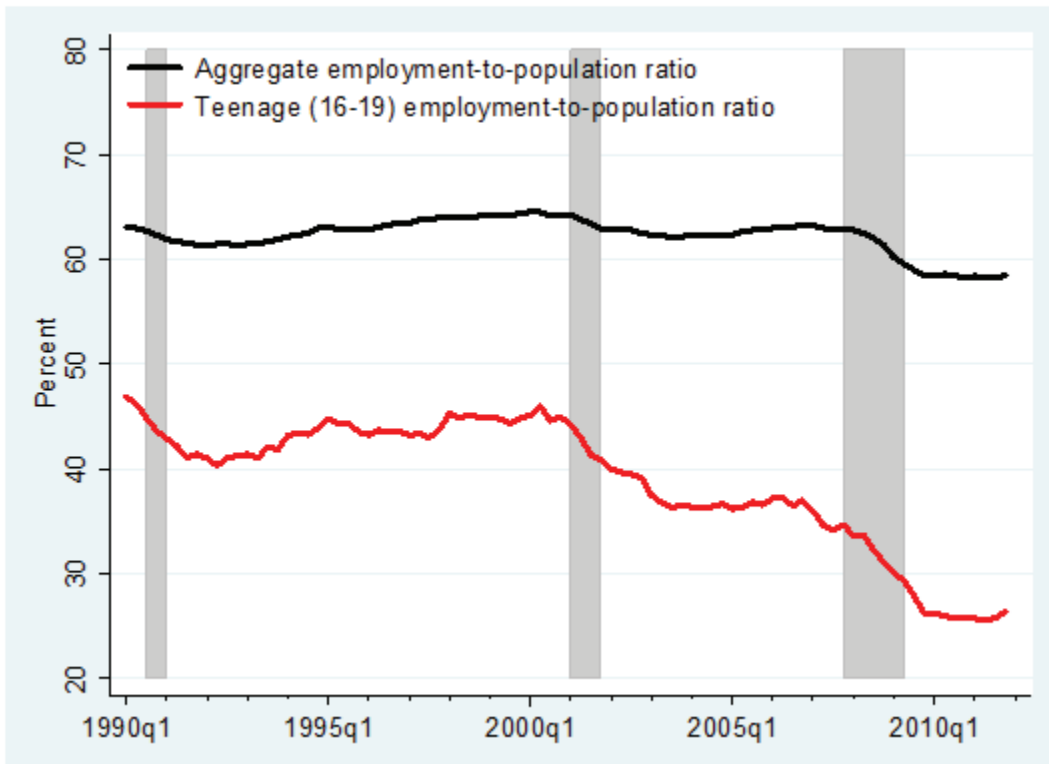
Thompson, Jeffrey P. 2009. "Using Local Labor Market Data to Re-Examine the Employment Effects of the Minimum Wage." *Industrial and Labor Relations Review*, Vol. 62, No. 3, April, pp. 343-66.

Figure 1: Aggregate and Teenage Unemployment Rates
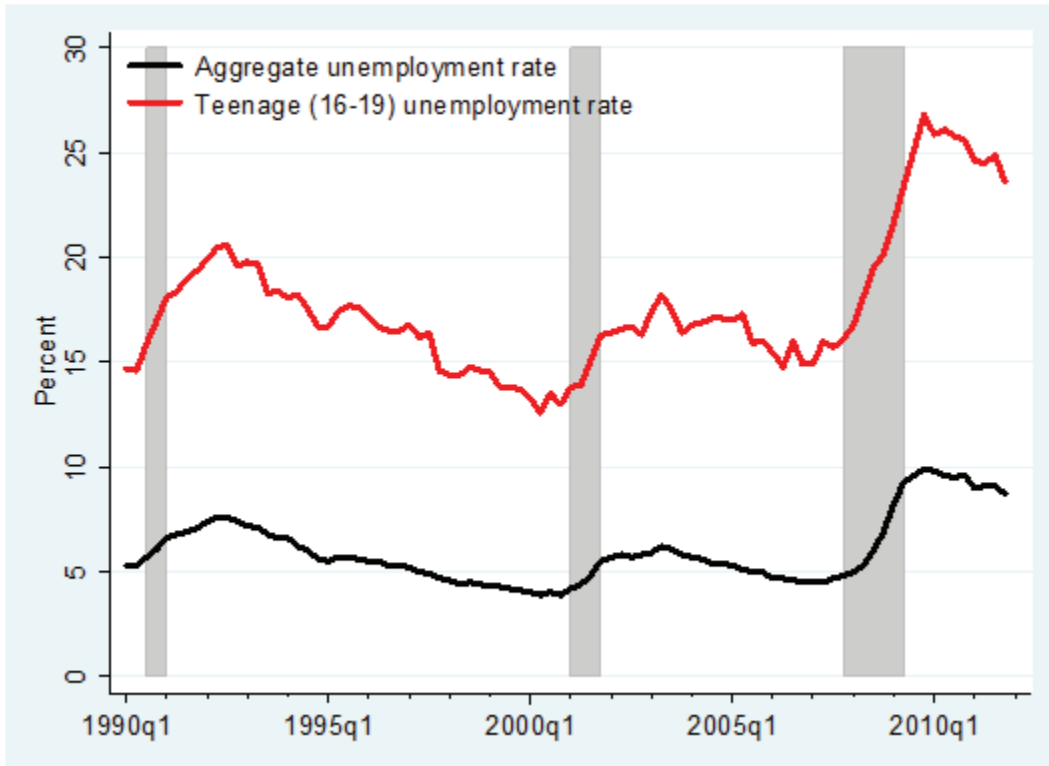
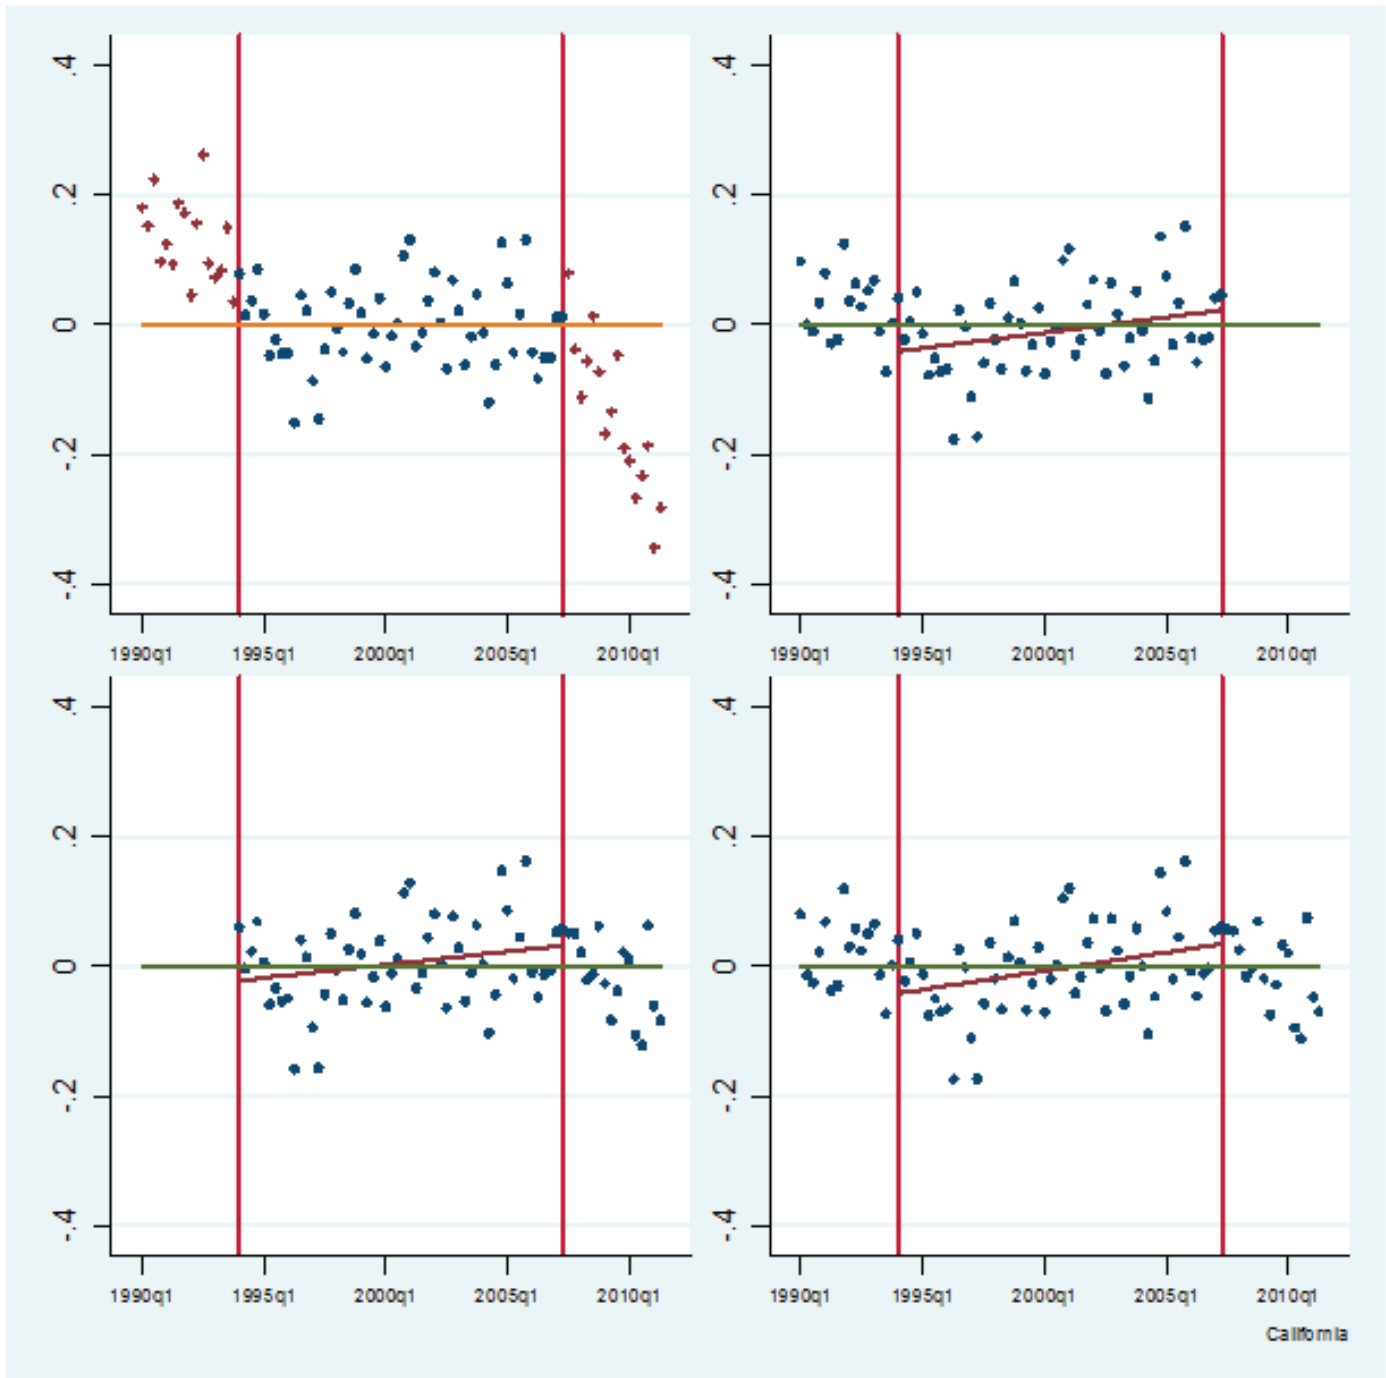The grey bars indicate recessionary quarters based on NBER business cycle dates.

## Figure 2: Aggregate and Teenage Employment Rates

See notes to Figure 1.

Figure 3: Residual Plots for California

The upper-left panel shows the residuals from estimating the model with state-specific linear trends for 1994-2007:Q2; for the quarters outside this period prediction errors are shown. The lower-left, upper-right, and lower-right panels, respectively, show the residuals for the following estimation periods: 1990-2007:Q2; 1994-2011:Q2; and 1990-2011:Q2. These three panels also display the fitted regression lines of the residuals on the time trend for the 1994-2007:Q2 subperiod.
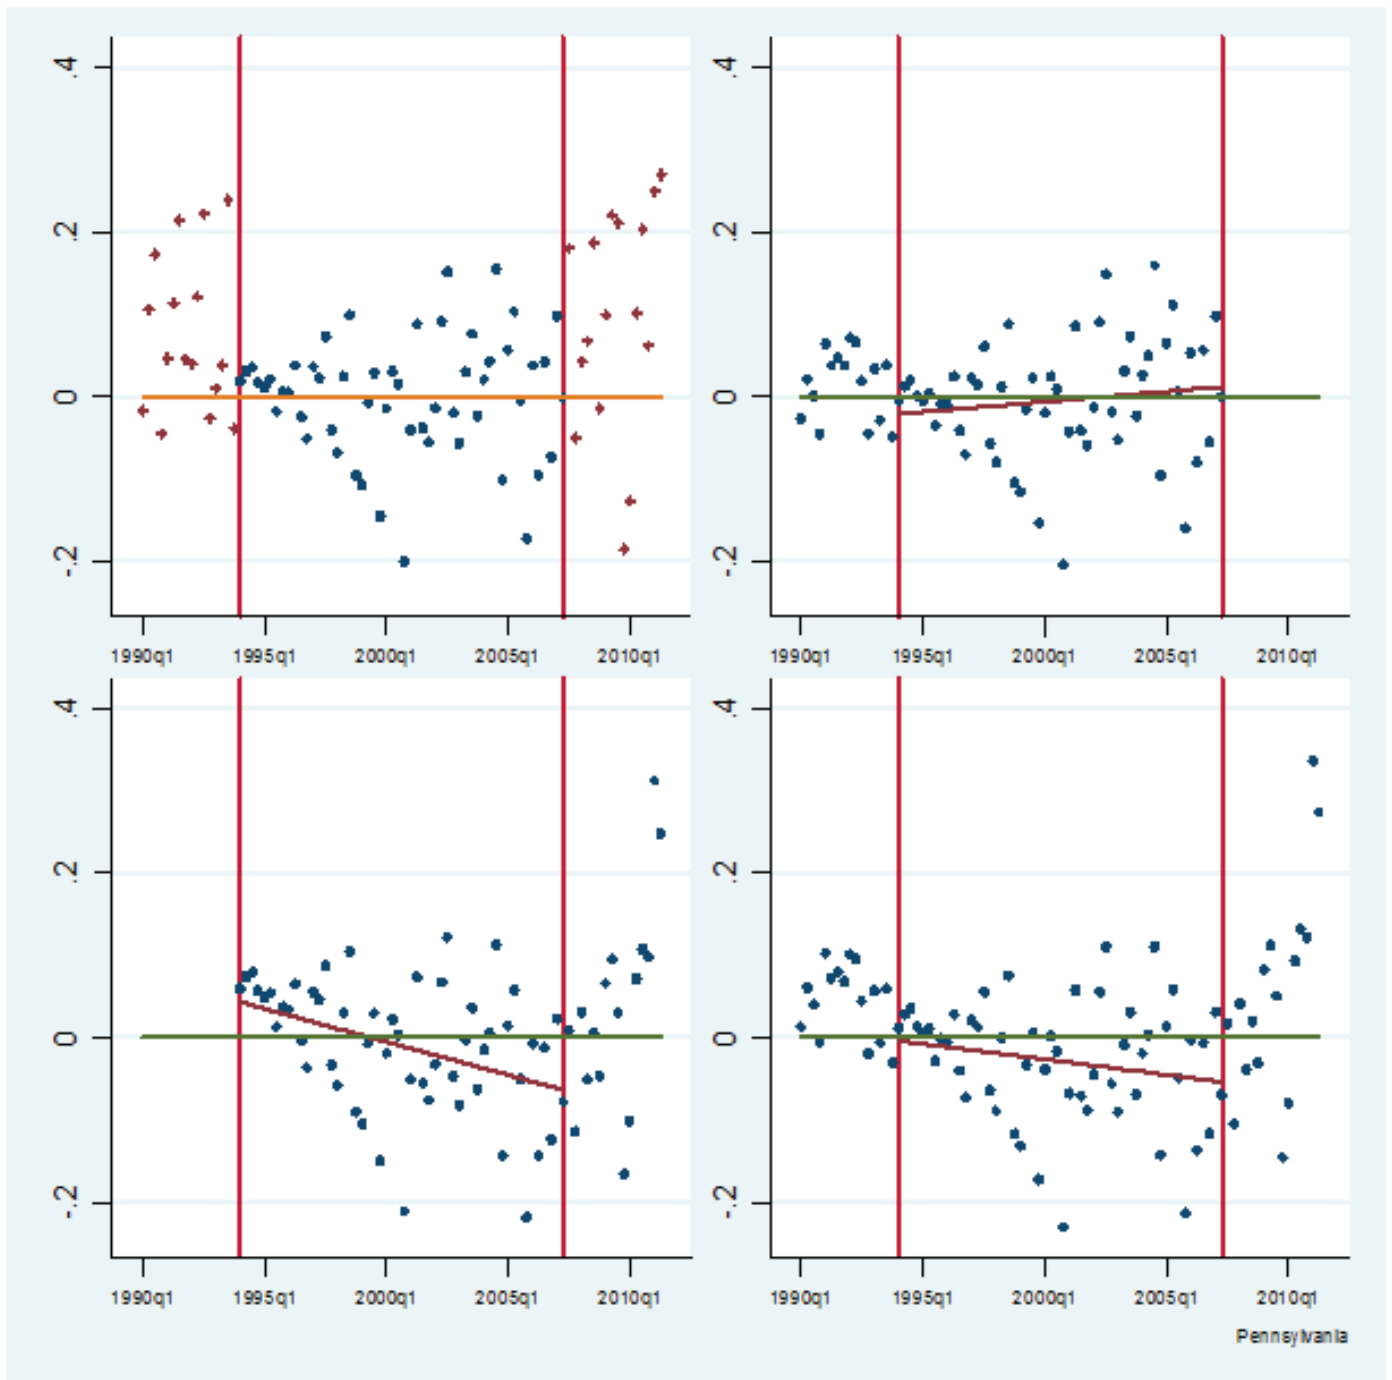
Figure 4: Residual Plots for Indiana
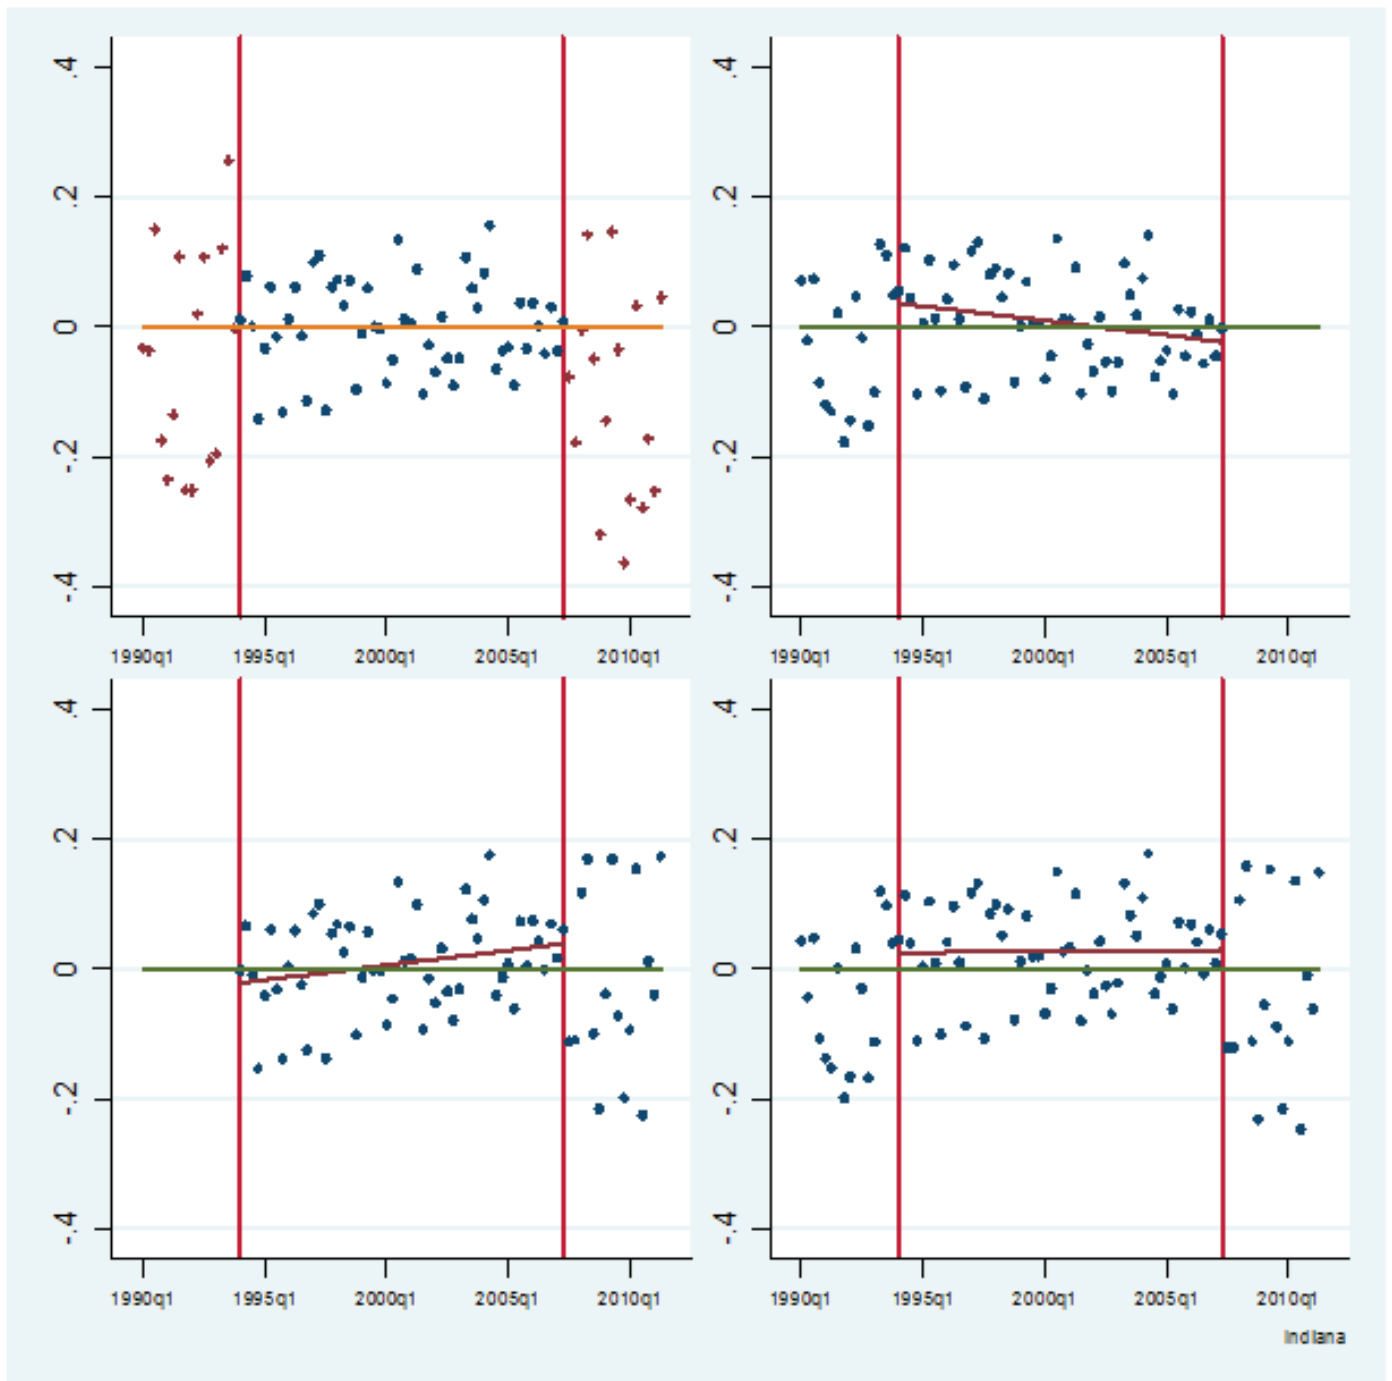
See notes to Figure 3.

Figure 5: Residual Plots for Tennessee
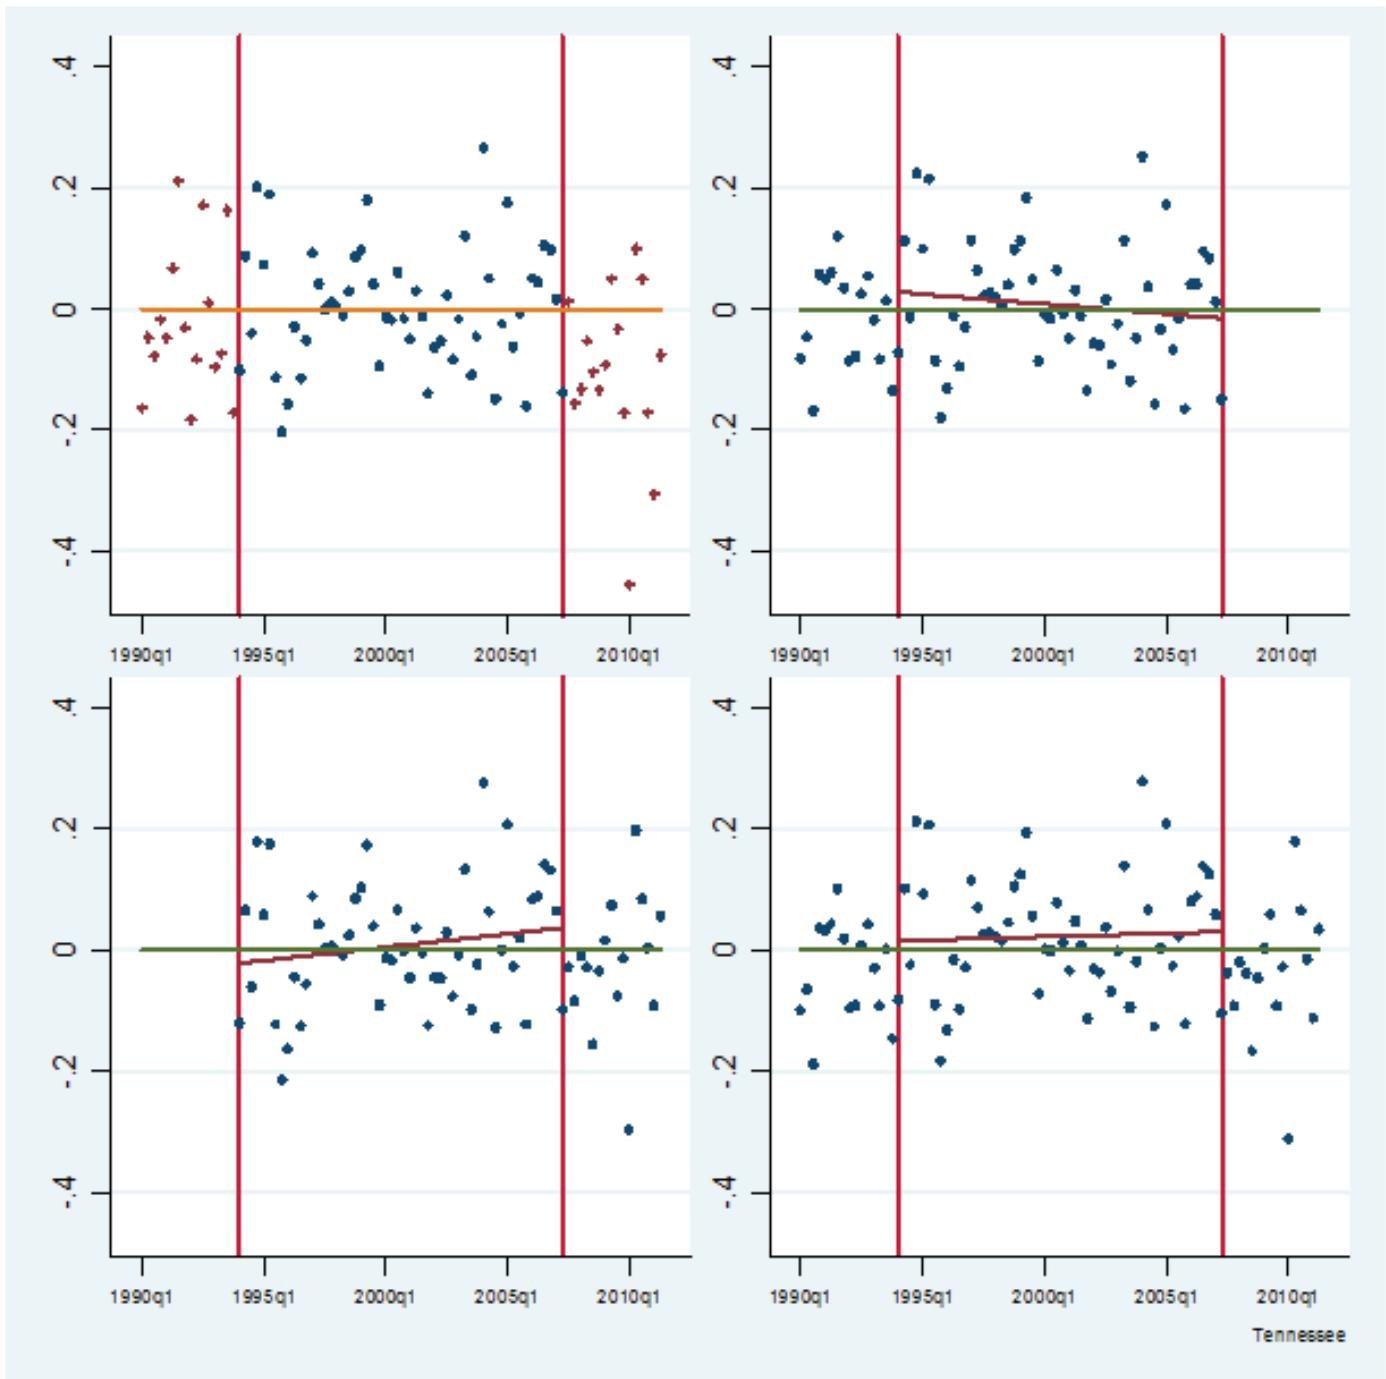
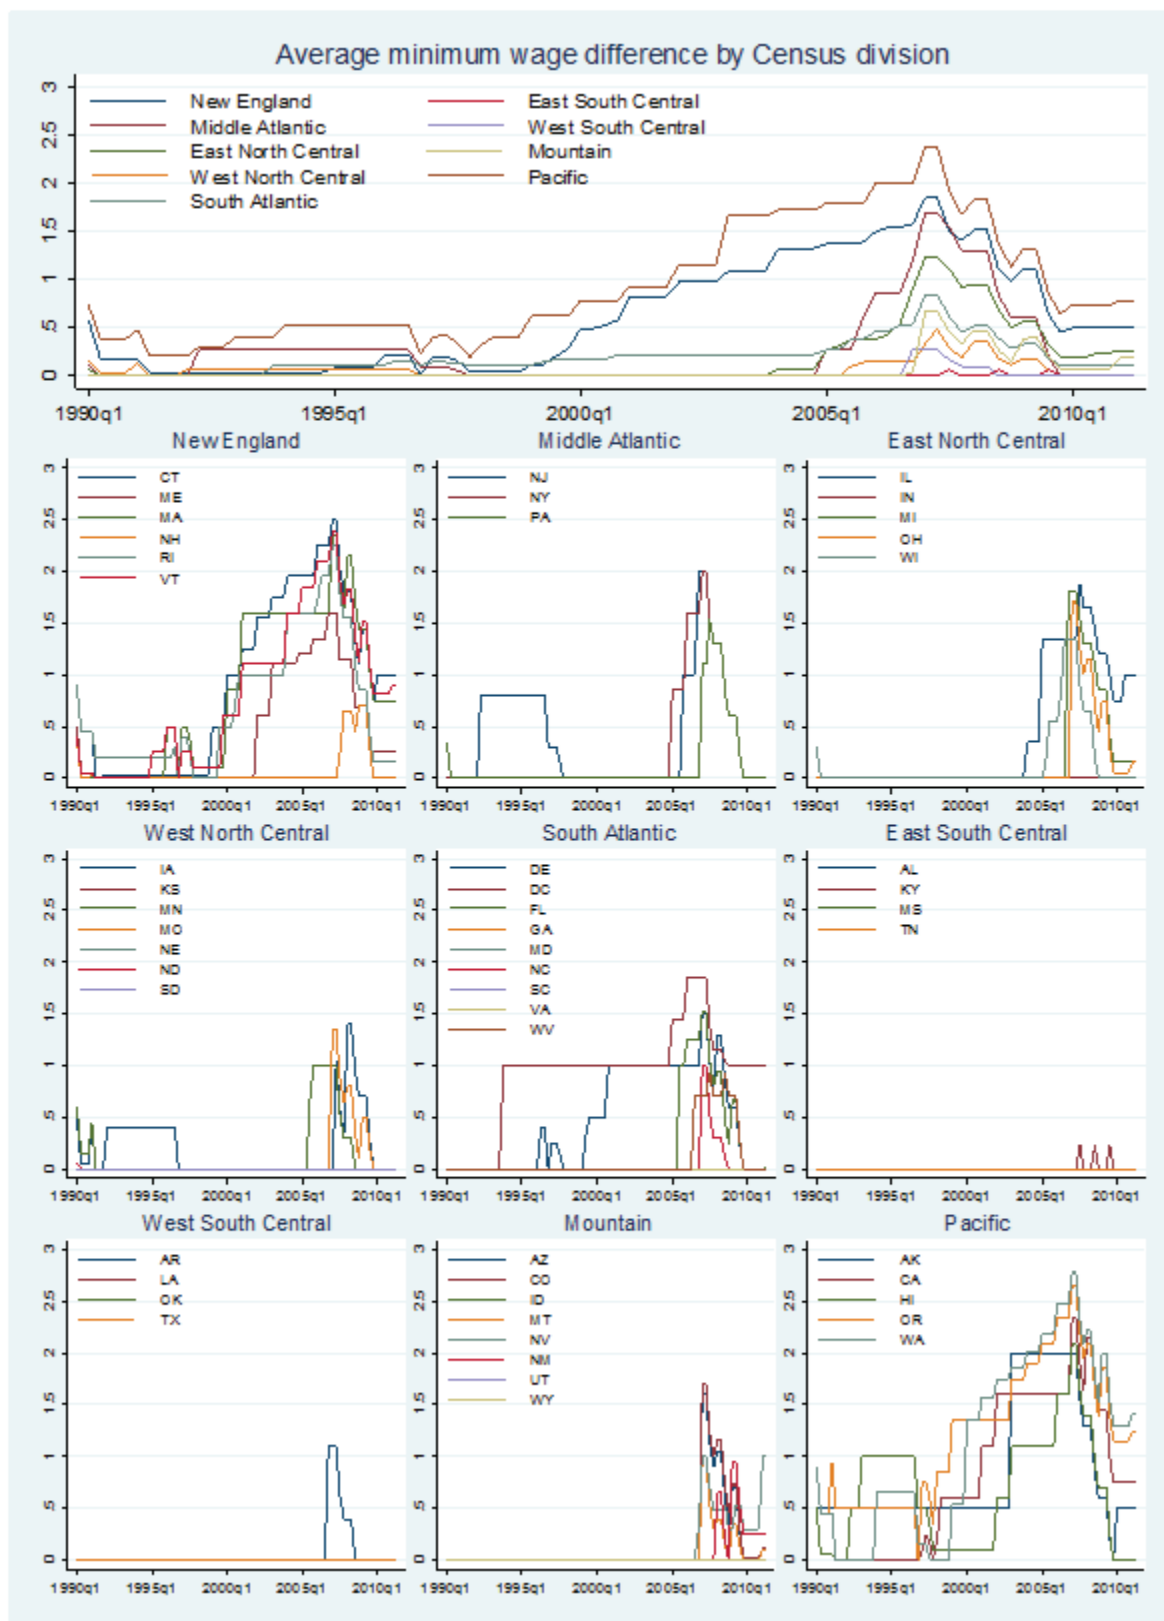

See notes to Figure 3.

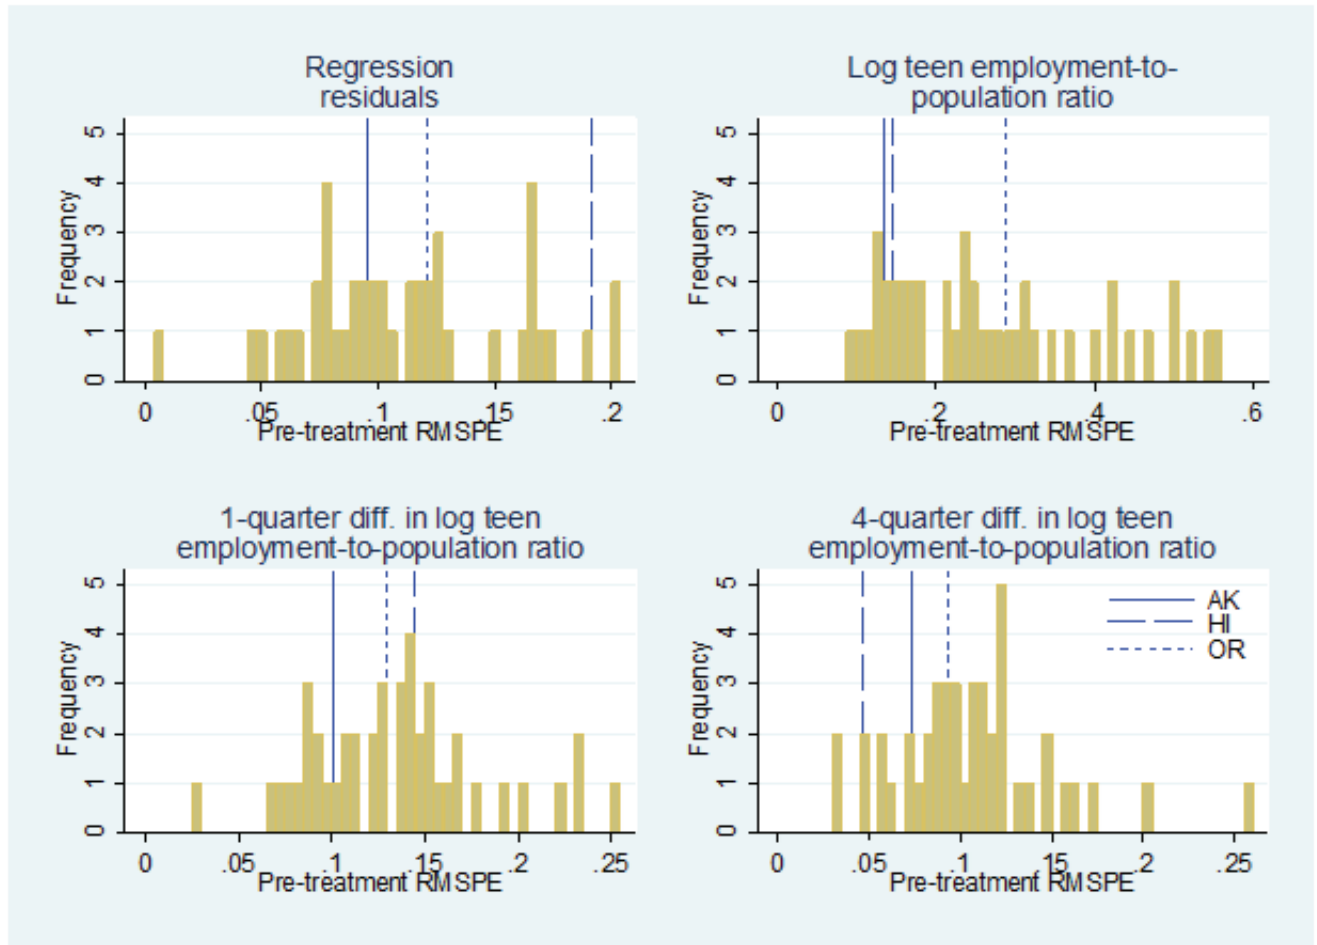## Figure 6: Residual Plots for Pennsylvania

See notes to Figure 3.

Figure 7: Between- and Within-Census Division Variation in State Minimum Wages (Difference Relative to Federal Minimum Wage, in Dollars per Hour)

Figure 8: Example of RMSPE Calculation at State Level

Treatment state: CA, 2001:Q1. Same-division control states are AK, HI, and OR. (WA is excluded because it had a minimum wage increase in same quarter.) The thick bars are the histogram. The thin vertical lines extending to the top of the graph show the placement of the RMSPEs for each control state in the same division as the treatment state.

Figure 9: Distributions of Percentiles of Same-Division States' RMSPEs
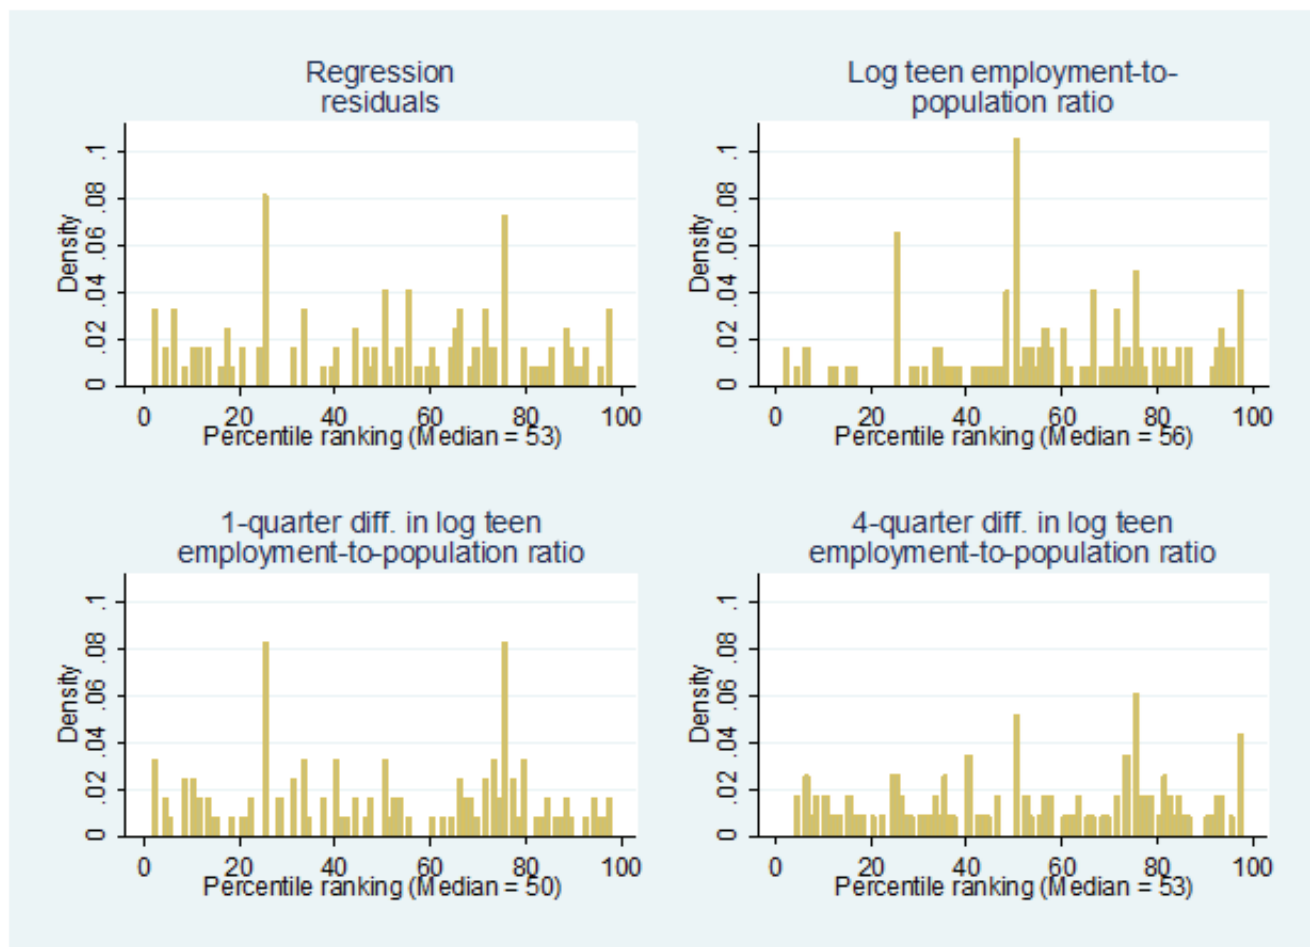
Ranks are converted to percentile rankings using the Weibull rule described in the text.

Figure 10: Medians of Percentiles of Same-Division States' RMSPEs, by Census Division
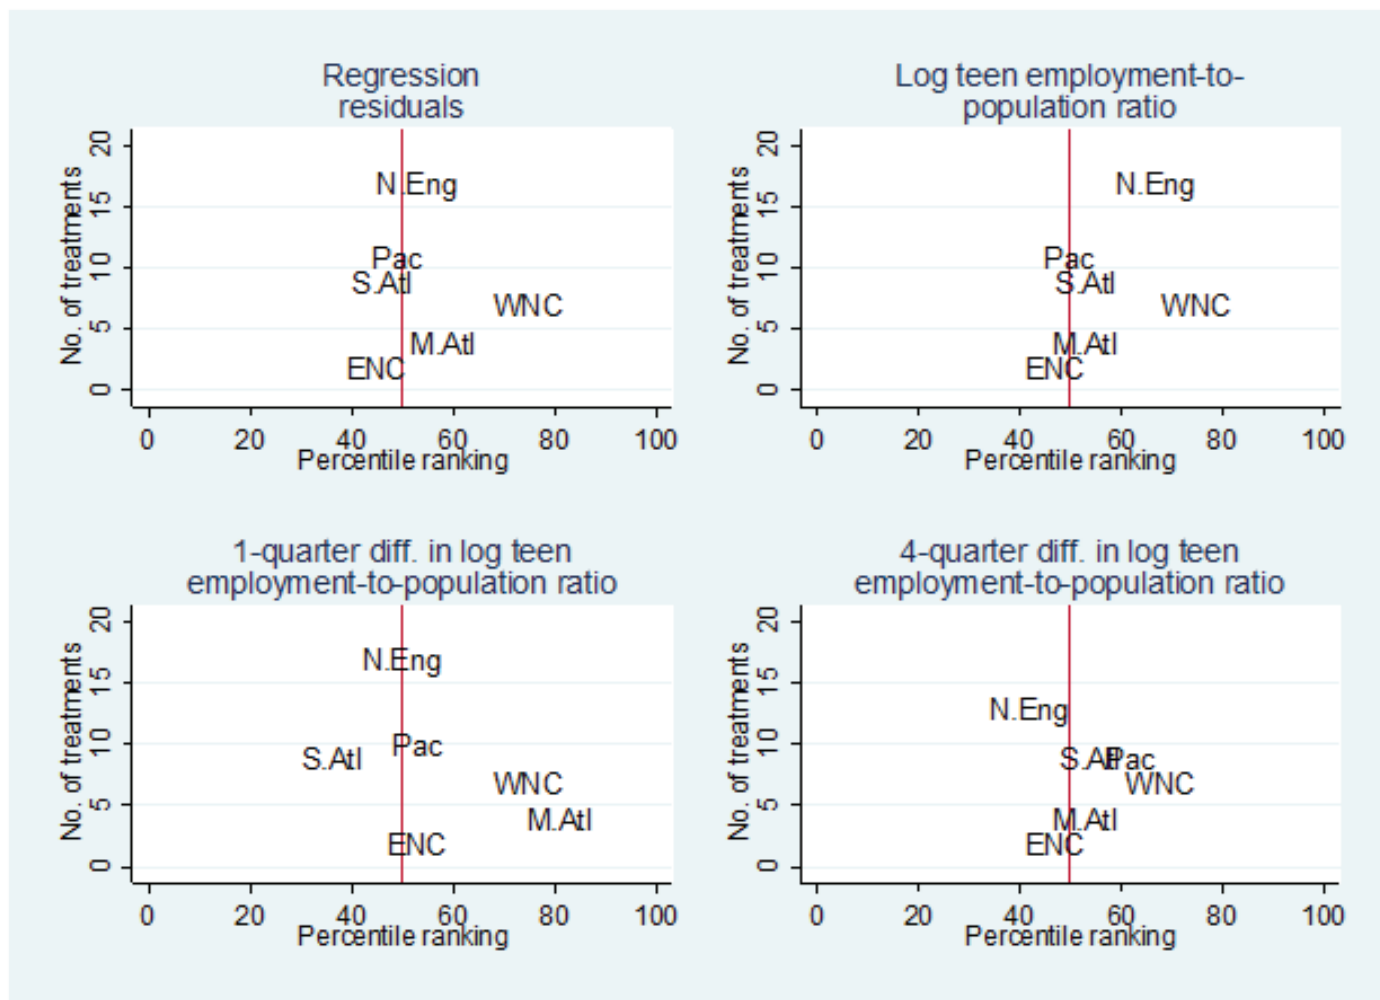
The Census division "plotting symbol" shows the median of the RMSPE percentile ranking for the states in the indicated division.

Figure 11: Matched County Pairs Along State Borders for DLR Analysis

A. DLR's Figure 2, eyeballed
(81 state borders with MW differential over 1990:Q1-2006:Q2
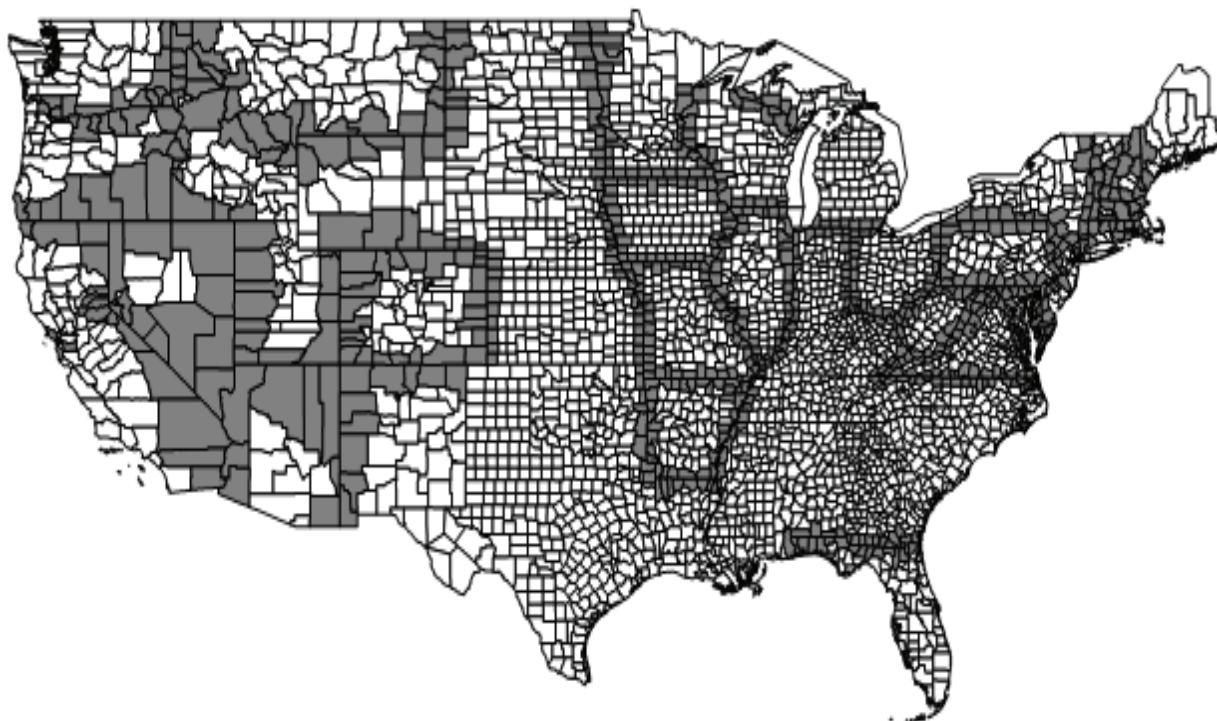
B. Corrected version of DLR's Figure 2, using their minimum wage data
(48 state borders with MW differential over 1990:Q1-2006:Q2)
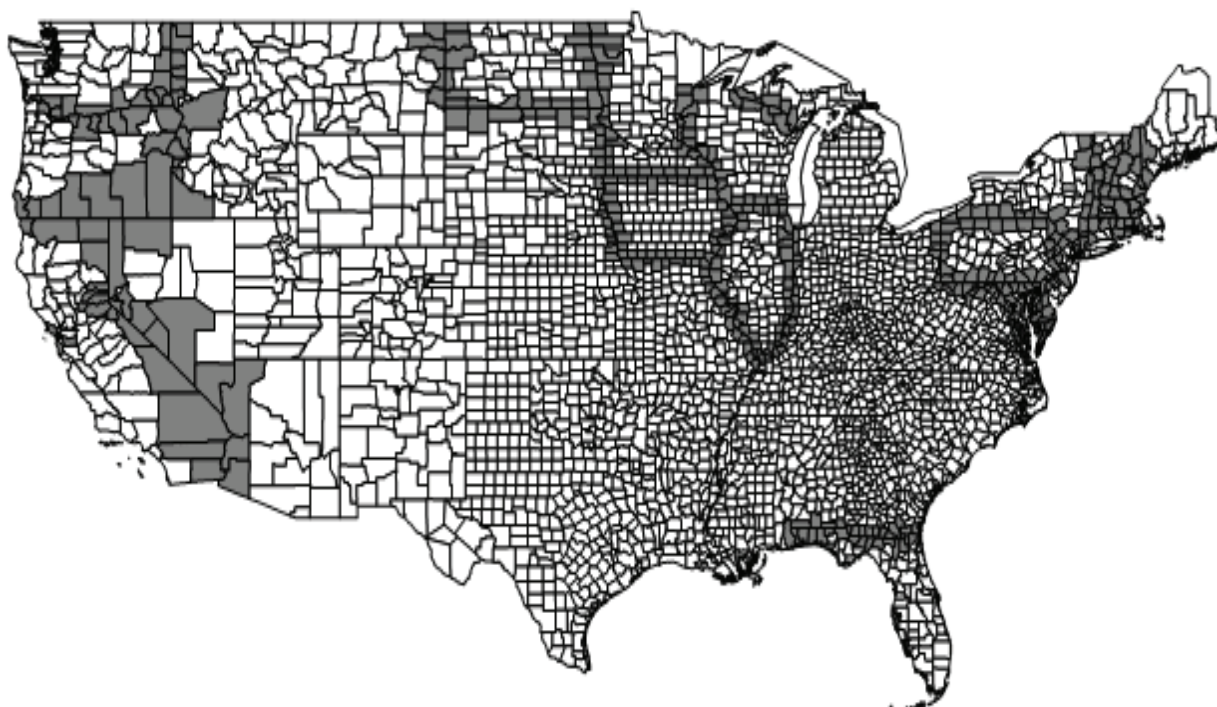
Figure 12: Example of RMSPE Calculation at County Level
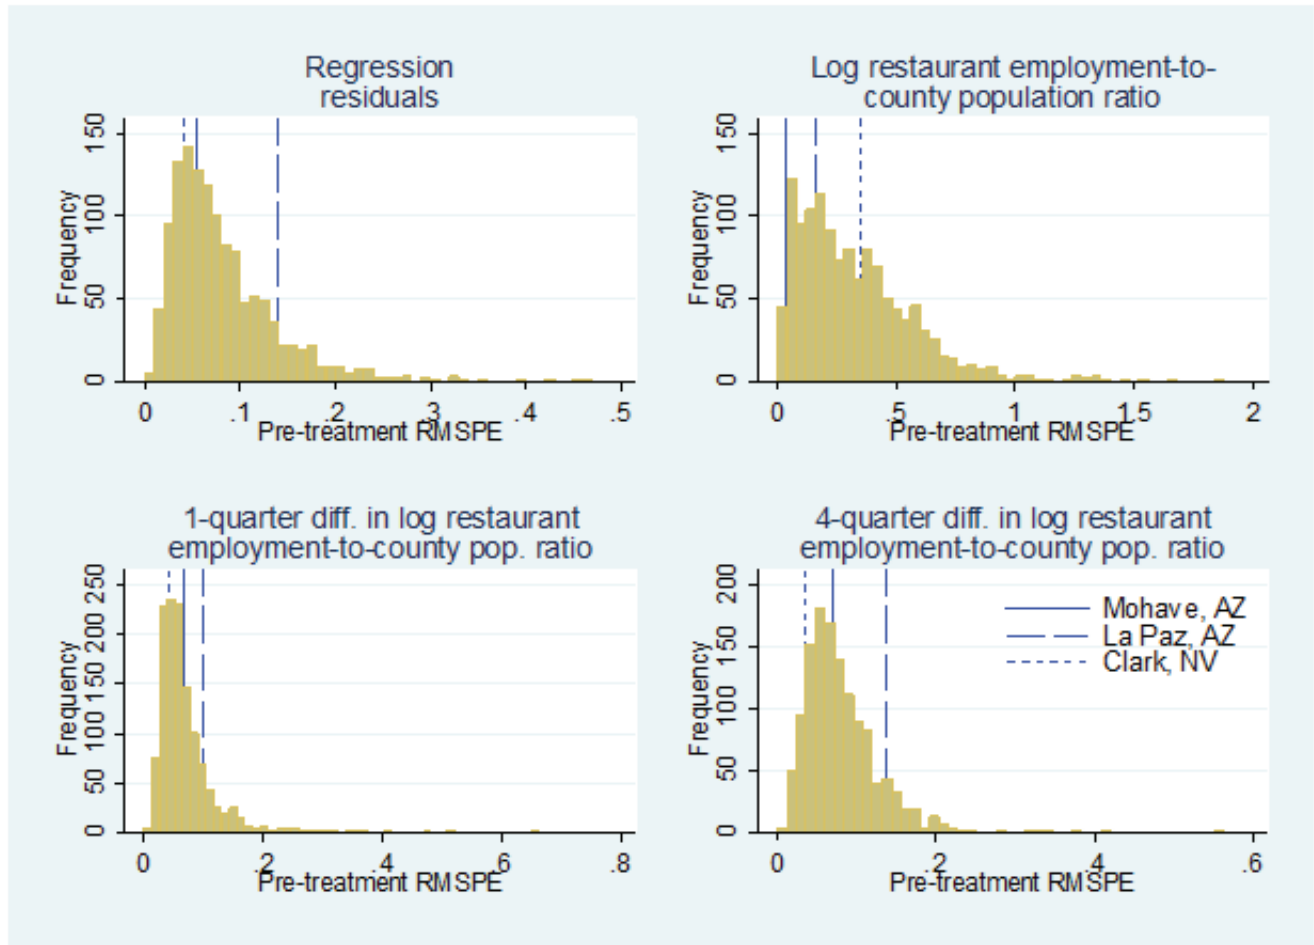
Treatment county: San Bernardino, CA, 2001:Q1. Contiguous cross-border control counties are La Paz and Mohave, AZ, and Clark, NV. The thick bars are the histogram. The thin vertical lines extending to the top of the graph show the placement of the RMSPEs for the control counties.

Figure 13: Distributions of Percentiles of Contiguous Cross-Border Counties' RMSPEs

Ranks are converted to percentile rankings using the Weibull rule described in the text.

Figure 14: Time Path of Cumulative Employment Elasticity in Response to Minimum Wage Change, ADR
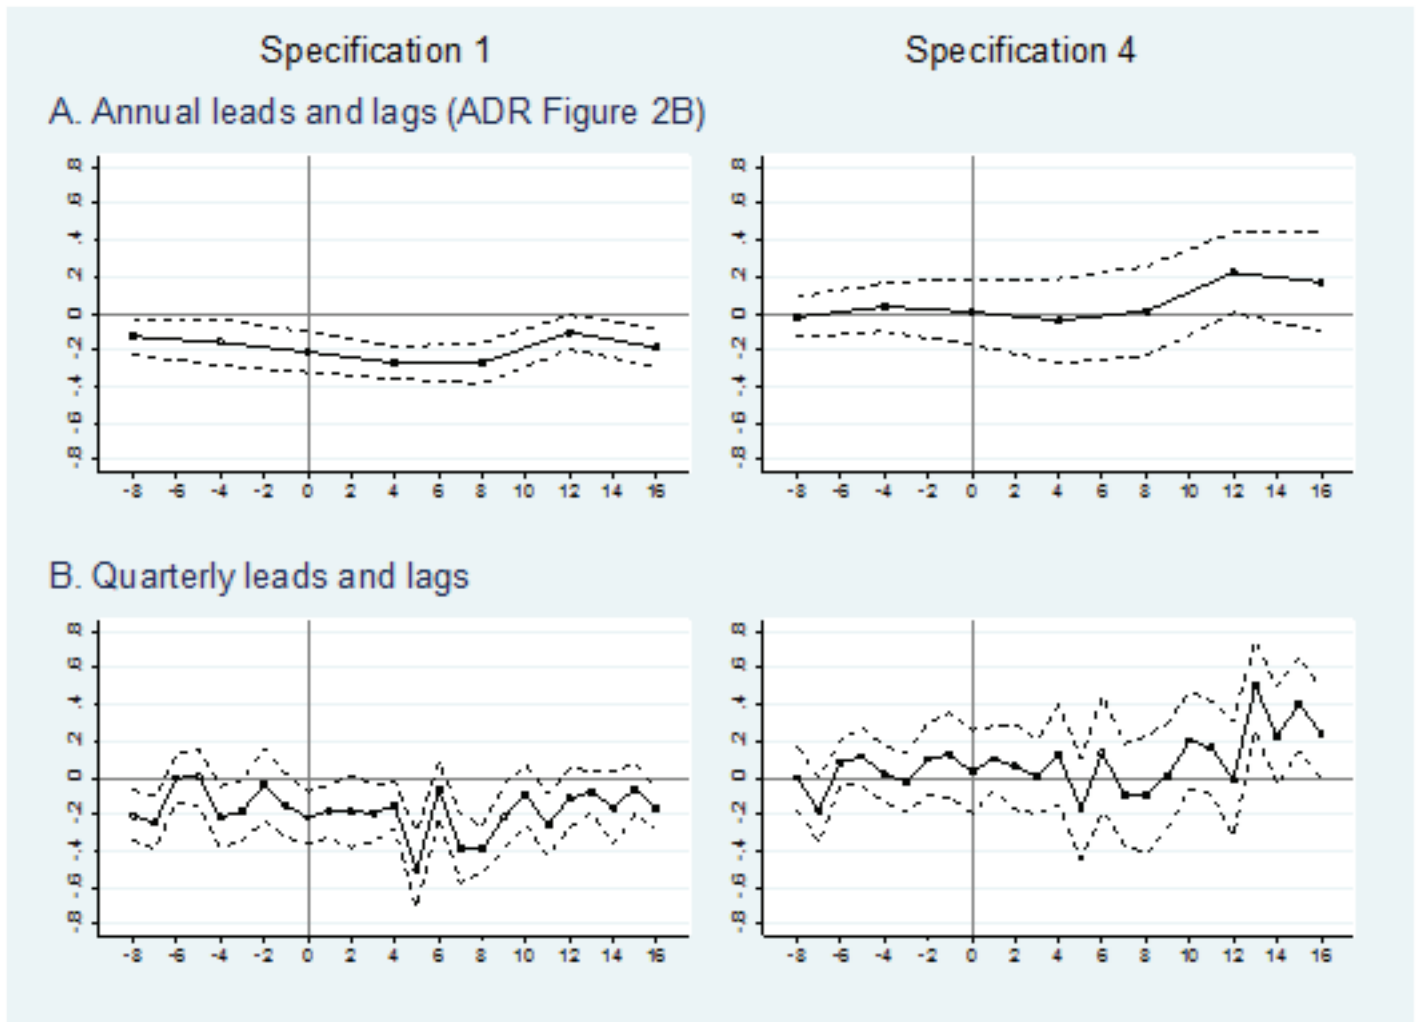
"Specification 1" is the specification with state and period fixed effects. "Specification 4" also includes the state-specific linear trends and the division × period interactions. The top graphs use the data from ADR (2011) and replicate the employment results in Figure 2 from that paper, with annual leads and lags of the minimum wage variable. The bottom graphs use the same data but include leads and lags at a quarterly frequency, which corresponds to the frequency of the minimum wage variable used in all of the regression analyses in ADR. The dashed lines show 90-percent confidence intervals.

Figure 15: Time Path of Cumulative Employment Elasticity in Response to Minimum Wage Change, DLR
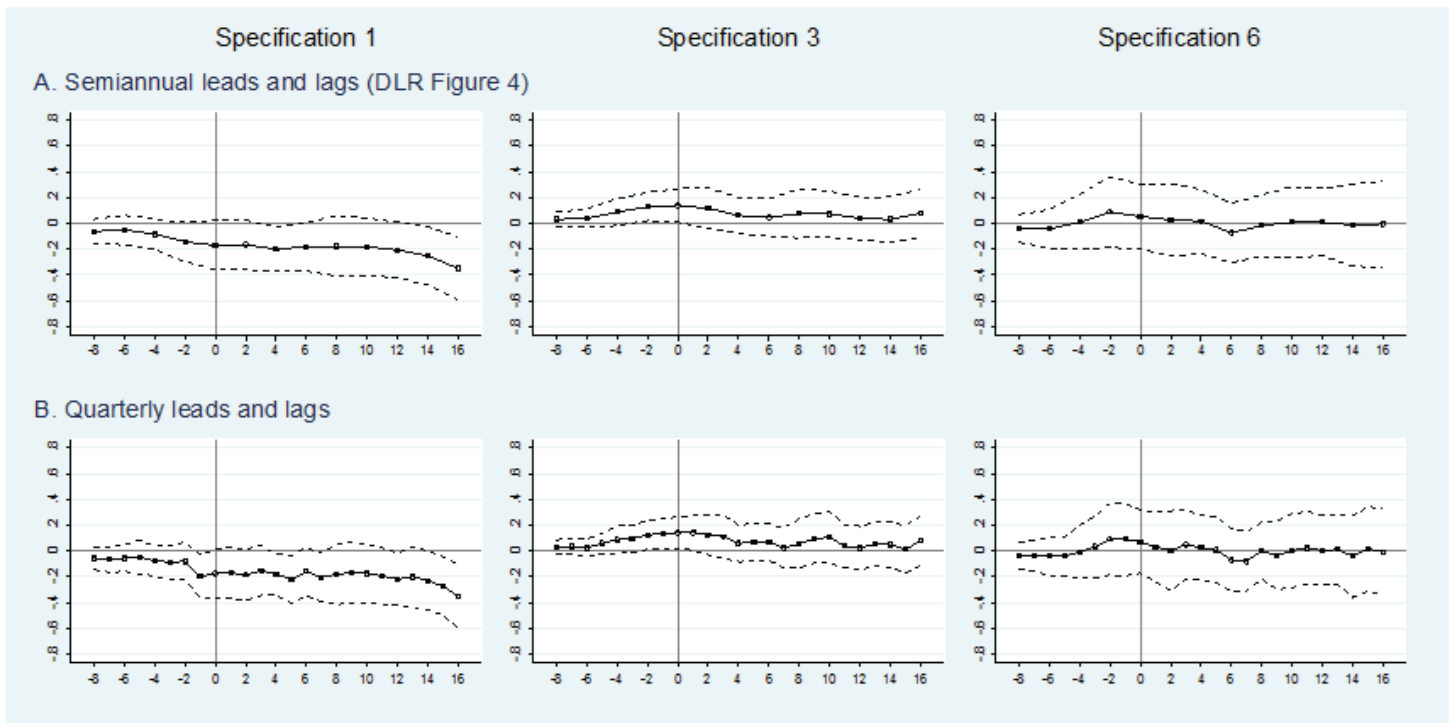
"Specification 1" is the specification with county and period fixed effects. "Specification 3" also includes the state-specific linear trends and the Census division × period interactions. "Specification 6" includes county-pair × period interactions. Specifications 1 and 3 use the all-county sample, while Specification 6 uses the contiguous border county-pair sample. The top graphs use the data from DLR (2010) and replicate the employment results in Figure 4 from that paper, with semi-annual leads and lags of the minimum wage variable. The bottom graphs use the same data but include leads and lags at a quarterly frequency, which corresponds to the frequency of the minimum wage variable used in all of the regression analyses in DLR. The dashed lines show 90-percent confidence intervals.

Tables

Table 1: The Effects of the Minimum Wage on Teen (16-19) Employment 1990-2011: Q2				
	(1)	(2)	(3)	(4)
	Dependent Variable: Log (E/P)			
Log (MW)	-.165*** (.041)	-.074 (.102)	-.098 (.097)	.009 (.058)
Unemployment Rate	-4.20*** (.427)	-3.83*** (.387)	-3.86*** (.403)	-3.12*** (.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
Region-Specific Time Effects	No	No	Yes	Yes
R^2	.877	.893	.911	.921
N	4386	4386	4386	4386

Estimates are weighted by teen population. Standard errors are clustered at the state level. ***, **, and * indicate estimates that are statistically different from zero at the one-, five-, and ten-percent levels, respectively.

Table 2: Sensitivity of Minimum Wage Effects to Sample Period								
	1990-2011:Q2		1994-2011:Q2		1990-2007:Q2		1994-2007:Q2	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Dependent Variable: Log (E/P)							
Log (MW)	-.165*** (.041)	-.074 (.102)	-.110** (.047)	-.085 (.081)	-.195*** (.049)	-.109 (.099)	-.148** (.060)	-.229** (.095)
Unemployment Rate	-4.20*** (.427)	-3.83*** (.387)	-4.35*** (.441)	-3.83*** (.441)	-3.48*** (.436)	-3.69*** (.427)	-3.60*** (.473)	-1.40*** (.469)
Relative Size of Youth Population	.100 (.316)	.218 (.336)	.169 (.380)	.428 (.359)	.526 (.344)	.514 (.322)	.732 (.409)	.267 (.377)
State Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State Trends	No	Yes	No	Yes	No	Yes	No	Yes
R^2	.877	.893	.883	.898	.845	.857	.851	.863
N	4386	4386	3570	3570	3570	3570	2754	2754

Estimates are weighted by teen population. Standard errors are clustered at the state level. ***, **, and * indicate estimated minimum wage effects that are statistically different from zero at the one-, five-, and ten-percent levels, respectively.

	1990-2011:Q2				1990-2009 (ADR sample)			
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
	Dependent Variable: Log (E/P)							
Order of Polynomial for State-specific "Trends"	2nd	3rd	4th	5th	2nd	3rd	4th	5th
Log (MW)	-.051 (.85)	-.230*** (.073)	-.180** (.069)	-.185** (.073)	-.071 (.086)	-.219*** (.066)	-.169** (.078)	-.201** (.094)
Unemployment Rate	-3.591*** (.494)	-2.571** (.454)	-2.376** (.461)	-2.378*** (.492)	-3.303*** (.477)	-2.425*** (.429)	-2.495*** (.446)	-2.412*** (.455)
Relative Size of Youth Population	.490 (.296)	.402 (.280)	.412 (.291)	.354 (.308)	.471* (.281)	.409 (.280)	.333 (.267)	.271 (.289)
R²	.899	.903	.906	.908	.881	.886	.889	.892
N	4386	4386	4386	4386	4080	4080	4080	4080

Table 3: Sensitivity of Minimum Wage Effects to Polynomial Order of State-Specific "Trends"

Estimates are weighted by teen population. Standard errors are clustered at state level. Models include state dummy variables interacted with a polynomial in time, with order of polynomial as indicated. ***, **, and * indicate estimated effects that are statistically different from zero at the one-, five-, and ten-percent levels, respectively.

Table 4: The Effects of the Minimum Wage on Teen (16-19) Employment, Controlling for State-Specific Trends Estimated from Subperiod Excluding Severe Recessions

	(1)	**(2)**	**(3)**	**(4)**
	Dependent Variable: Log (E/P)			
Method	Single trends, estimated from subperiod excluding severe recessions	Split trends, estimated from subperiod excluding severe recessions	Peak-to-peak trends	HP filter (= 1600)
Log (MW)	-.178** (.090)	-.207* (.106)	-.319** (.126)	-.184*** (.068)
Unemployment Rate	-3.06*** (.307)	-3.20*** (.568)	-3.52*** (.760)	-2.38*** (.468)
Relative Size of Youth Population	.112 (.285)	.521* (.308)	.650 (.486)	.304 (.280)
N	4386	4386	4386	4386

Estimates are weighted by teen population. In columns (1) and (2), state-specific trends are estimated from 1994:Q1-2007:Q2 and then extrapolated to 1990:Q1-2011:Q2. For split trend specification, trend is a spline with a knot in 2000:Q4. (These estimates are based on equations (1)-(4) in the text.) In column (3), the peaks are based on NBER business cycle dates. In column (4), as noted in the text,   is the smoothing parameter for the HP filter and is set to the value commonly used for quarterly data. For each column, standard errors are block bootstrapped by state using 200 replications. ***, **, and * indicate estimated effects that are statistically different from zero at the one-, five-, and ten-percent levels, respectively, based on the normal approximation.

	(1)	(2)
Table 5: The Effects of the Minimum Wage on Teen (16-19) Employment, By Division		
	Dependent Variable: Log (E/P)	
	1990:Q1-2011:Q2	**1990:Q1-2009:Q4 (ADR Sample)**
New England	-.390*** (.052)	-.384*** (.058)
Mid-Atlantic	.166 (.143)	.105 (.162)
East North Central	-.208 (.284)	-.166 (.272)
West North Central	-.191** (.082)	-.194*** (.067)
South Atlantic	-.150 (.242)	-.152 (.281)
East South Central	-2.24 (1.41)	-2.02 (1.51)
West South Central	-.217*** (.062)	-.147** (.053)
Mountain	-.598*** (.139)	-.638*** (.187)
Pacific	-.002 (.133)	.016 (.143)

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. ***, **, and * indicate estimated effects that are statistically different from zero at the one-, five-, and ten-percent levels, respectively.

	Proportion of weight on states in same division						Average # states in donor pool in same division
	Matching on:						
	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	Average # divisions in donor pool	Average # states in donor pool	
Division	(1)	(2)	(3)	(4)	(5)	(6)	(7)
New England	0.168	0.209	0.163	0.185	6.9	30.4	1.9
Middle Atlantic	0.073	0.134	0.455	0.168	5.5	20.0	1.0
East North Central	0.0	0.0	0.016	0.015	9.0	39.5	3.5
West North Central	0.547	0.823	0.698	0.464	3.7	7.7	1.7
South Atlantic	0.123	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	0.233	0.323	0.264	0.251	6.1	24.0	2.5

Table 6: Weights on States in Same Census Division from Synthetic Control Estimator, State-Level CPS Data

Results are reported for the 50 unique minimum wage treatments (out of a total of 129) 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 treatments and donors 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 7: The Effects of the Minimum Wage on Restaurant Employment, 1990-2006:Q2				
	(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

Standard errors are two-way clustered at the (non-nested) state and the border segment levels; the border segment is the set of all counties on both sides of a border between two states. ***, **, and * indicate estimates that are statistically different from zero at the one-, five-, and ten-percent levels, respectively.

| | Proportion of weigh on contiguous cross-border counties | | | |
| | Matching on: | | | |
Distribution	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
	(1)	(2)	(3)	(4)
A. Donor pools restricted to 50 counties with lowest RMSPE for four quarters following minimum wage increase				
Minimum	.000	.000	.000	.000
10th Percentile	.000	.000	.000	.001
25th Percentile	.002	.002	.002	.004
Median	.008	.010	.009	.010
75th Percentile	.029	.027	.027	.019
90th Percentile	.075	.1001	.065	.035
Maximum	.723	.659	.496	.336
Mean	.033	.038	.035	.017
B. Full donor pools				
Minimum	.000	.000	.000	.000
10th Percentile	.000	.000	.000	.000
25th Percentile	.001	.000	.001	.001
Median	.001	.001	.001	.001
75th Percentile	.002	.002	.002	.002
90th Percentile	.006	.007	.012	.006
Maximum	.228	.308	.474	.393
Mean	.006	.007	.015	.007

Table 8: Weights on Contiguous Cross-Border Counties from Synthetic Control Estimator, County-Level QCEW Data

County results are reported for the 121 unique minimum wage treatments at the state-by-quarter level 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 959.6 in Panel B. There are somewhat fewer treatments and donors 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.

	(1)	(2)
Table 9: The Effects of the Minimum Wage on Restaurant Employment, with Leading Effects, 1990-2006:Q2		
	Dependent Variable: Log (restaurant employment)	
	All counties (DLR, Table 3, specification 1)	All counties, with state-specific linear trends and Census division × period interactions (DLR specification 3)
Cumulative Effect of Log (MW)		
12-Quarter Lead	-.069 (.058)	.070 (.047)
4-Quarter Lead	-.192* (.113)	.188* (.094)
Pre-Trend: 4-Quarter lead–12-Quarter Lead	-.122* (.070)	.117** (.054)
County Effects	Yes	Yes
Period Effects	Yes	No
N	82,800	82,800

Estimates correspond to Table 3 of DLR. Standard errors are clustered at the state level. ***, **, and * indicate estimates that are statistically different from zero at the one-, five-, and ten-percent levels, respectively.

	(1)	(2)	(3)
Table 10: The Effects of the Minimum Wage on Restaurant Employment, with Leading Effects, 1990-2006:Q2			
	Dependent Variable: Log (restaurant employment)		
	All counties (DLR, Table 3, specification 1)	All counties, with state-specific linear trends, and Census division × period interactions (DLR specification 3)	DLR contiguous border county-pair sample, with county-pair × period interactions (DLR specification 6)
A. 12-Quarter Leads			
Pre-trend:			
10-Quarter Lead – 12-Quarter Lead	-.004 (.042)	.083*** (.031)	.049 (.072)
8-Quarter Lead – 12-Quarter Lead	-.047 (.064)	.127** (.051)	.057 (.095)
6-Quarter Lead – 12-Quarter Lead	-.067 (.060)	.124** (.049)	.056 (.121)
4-Quarter Lead – 12-Quarter Lead	*-.122** (.070)*	.117** (.054)	*.040* *(.134)*
2-Quarter Lead – 12-Quarter Lead	-.105 (.071)	.088* (.052)	.058 (.143)
B. 8-Quarter Leads			
6-Quarter Lead – 8-Quarter Lead	-.015 (.038)	.058** (.022)	.062 (.092)
4-Quarter Lead – 8-Quarter Lead	-.055 (.065)	.098** (.039)	.119 (.125)
2-Quarter Lead – 8-Quarter Lead	-.040 (.066)	.106*** (.036)	.137 (.125)

Specifications in Panel A, columns (1) and (2) correspond to Table 11, but with the pre-trend estimated over different periods. Specifications in column (3) do the same for DLR's specification 6. Specifications in Panel B are the same, but only include leads up to 8 quarters. The two highlighted estimates are the ones reported in DLR's Table 3 for restaurant employment. These specifications exclude the private-sector employment control, although they include this control in the specifications on which their Figure 4 is based. Standard errors are clustered at the state level. ***, **, and * indicate estimates that are statistically different from zero at the one-, five-, and ten-percent levels, respectively.

	(1)	(2)	(3)
Table 11: The Effects of the Minimum Wage on Restaurant Employment, "Falsification Tests"			
	Dependent Variable: Log (restaurant employment)		
	DLR Sample: 1990:Q1-2006:Q2 (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 in Sample Period	17.8	22.3	100.0
County Effects	Yes	Yes	Yes
Period Effects	Yes	Yes	Yes

These specifications include controls for population and private-sector employment. Following DLR's code, sample is restricted to counties that have an area less than 2,000 square miles. In columns (2) and (3), a balanced panel of counties is used, as in DLR's other analyses; some counties that are not included in column (1) can be included in the samples in these columns. In column (3), the subset of county pairs in column (2) that had one or more minimum wage difference in the period always had at least two minimum wage differences. Standard errors are clustered at the state level. ***, **, and * indicate estimates that are statistically different from zero at the one-, five-, and ten-percent levels, respectively.

Notes

Employment
Policies

INSTITUTE