NBER WORKING PAPER SERIES

RESTAURANT EMPLOYMENT, MINIMUM WAGES, AND BORDER DISCONTINUITIES

Arindrajit Dube
Michael Reich
Akash Bhatt
Denis Sosinskiy

Working Paper 32902
http://www.nber.org/papers/w32902

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2024

We are grateful to the Institute for Research on Labor and Employment at UC Berkeley for research support, to Ben Zipperer for his discussion of the JNR paper at ASSA 2024, and to Attila Lindner for his valuable comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Restaurant Employment, Minimum Wages, and Border Discontinuities
Arindrajit Dube, Michael Reich, Akash Bhatt, and Denis Sosinskiy
NBER Working Paper No. 32902
September 2024
JEL No. J20, J39, J88

ABSTRACT

Dube, Lester and Reich (2010, DLR), using state minimum wage discontinuities across bordering counties and Quarterly Census of Employment and Wages data, did not detect negative minimum wage effects on restaurant employment. Jha, Neumark and Rodriguez-Lopez (2024, JNR) claim that looking within multi-state commuting zones and using County Business Patterns data provides a superior approach to DLR and does find disemployment effects. We show that JNR's results are confounded by parallel trends violations in the 1990s, when minimum wage events were rare and small in magnitude; JNR's outmoded two-way-fixed-effects model amplifies the biases introduced by these violations. Our estimates using their specifications and data on only post-2000 data fail to detect disemployment effects. The same results hold using QCEW and ACS datasets. Our preferred event study difference-in-differences approach, which analyzes only data that fall clearly within an event's window, also does not detect negative employment effects. This result holds whether we compare across all states, look within commuting zones or within border county pairs, and regardless of the data set or time period.

Arindrajit Dube
Department of Economics
University of Massachusetts
Crotty Hall
412 N. Pleasant Street
Amherst, MA 01002
and NBER
adube@econs.umass.edu

Michael Reich
2521 Channing Way
Institute for Research on Labor and Employment
University of California, Berkeley
Berkeley, CA 94720
mreich@econ.berkeley.edu

Akash Bhatt
UMass Amherst
akashbhatt@umass.edu

Denis Sosinskiy
University of California, Davis
dsosinskiy@ucdavis.edu

1 Introduction

Dube et al. (2010) (DLR) advanced the minimum wage literature by providing a bridge between the local-area case study approach of Card and Krueger (1994, 2000)—who compared minimum wage effects in the border areas of New Jersey and Pennsylvania—and the two-way fixed effects (TWFE) panel regression proposed by Neumark and Wascher (1992). DLR argued that nearby areas are more likely to share similar underlying shocks and therefore offer a better approach to forming counterfactuals, compared to the TWFE assumption that parallel trends hold for all state-by-state comparisons. Using BLS' Quarterly Census of Employment and Wage (QCEW) data, DLR found that the TWFE model with log minimum wage and common time effects suggested a negative employment effect. However, estimates of the same TWFE, but using only variation within border county pairs (BCPs), suggested that restaurant employment was largely unaffected by minimum wages, even when restaurant earnings rose strongly. Furthermore, based on leading effects in distributed lag version of the regression models, the parallel trends assumption appeared much more likely to be satisfied under the BCP strategy than under the simple TWFE model.

In response, Neumark et al. (2014) argued that nearby areas without a minimum wage increase do not provide better counterfactuals. Allegretto et al. (2017) then offered a host of additional results suggesting that local areas were indeed better controls. For example, neighboring counties share more similar employment levels and changes than counties that are randomly chosen from across the U.S.

Subsequently, econometricians have discovered possible biases arising from TWFE-style estimators; these biases include applying negative weights to some events and amplifying violations of parallel trends (Callaway and Sant'Anna (2021); Cengiz et al. (2019); Goodman-Bacon (2021); Sun and Abraham (2021)). It is now widely recognized that the TWFE regression does not provide a reliable approach to aggregating difference-in-differences estimates. As a result, a large set of minimum wage papers have abandoned the TWFE approach and instead applied a properly aggregated difference-in-difference event study approach to studying minimum wage effects. The event study approach examines outcome dynamics for clearly delineated minimum wage events and event windows. Cengiz et al. (2019) (CDLZ), for example, proposed a stacked event study approach, explicitly classifying admissible control groups as those that did not experience minimum wage increases during the event window.[1] The large and growing set of recent papers that use an event-based approach includes Godoey and Reich (2021), Wiltshire et al. (2024), Hampton and Totty (2023), Vergara (2023), Clemens and Strain (2021), and Rao and Risch (2024).[2]

Unfortunately, rather than engaging with this more recent literature, Jha et al. (2024) (JNR) have chosen to re-examine DLR while continuing to rely on the outmoded TWFE approach. In this round, JNR acknowledge the potential advantages of looking at nearby areas, but assert that simply using a different geographical unit (cross-state commuting zones, CZs) and a different dataset (annual County Business Patterns, CBP) is superior to using BCPs and QCEW data. With just these two changes—and using TWFE specifications—JNR claim to find strongly negative effects of minimum wages on restaurant employment.

In this paper, we show that the evidence does not support JNR's contentions. We find that JNR's results are highly susceptible to well-known issues with 1990s data, specifically related to their TWFE-style panel regressions. We start by presenting a simple comparison of trends between two groups of states: those that raised their minimum wages at some point between 1980 and 2019, and those that did not. Restaurant employment significantly decreased in the 1980s and 1990s in "minimum wage states" compared to "non-minimum wage states," despite the limited prevalence and size of state minimum wages during that time. This pattern indicates a violation of the parallel trends assumption during the 1980s and 1990s, before stabilizing around 2000. After 2000, the employment gap between these two groups of states remained fairly stable, even as minimum wage differences became more common and pronounced, particularly after 2010. This pattern suggests that minimum wage increases did not substantially affect restaurant employment. Importantly, the

[1] Importantly, Appendix G of CDLZ also showed that TWFE models were highly influenced by shocks during the 1980s and 1990s that especially affected states that later raised minimum wages. These violations appear to subside for overall employment in the early 1990s. As we show in this paper, the violations for restaurant employment extended to the late 1990s.

[2] Other studies, such as Wiltshire et al. (2024), have applied the synthetic control to conduct event studies, also avoiding the TWFE issues.

same patterns persist when we repeat the comparison, but narrow the focus to within-CZ differences between these two groups of states.

Consistent with the parallel trends violations in the 1990s, we show that the significant disemployment effects implied by JNR's preferred regression specification are not robust when we exclude data from the 1990s. Unlike a difference-in-differences design with clearly defined pre- and post-treatment periods, the TWFE model makes a stronger assumption: that parallel trends hold throughout the entire sample. As a result, TWFE outcomes are likely to be sensitive to violations of the parallel trends assumption that occur at any point in the sample, not just around the time of the actual minimum wage events.

Notably, when we analyze the data since 2000—when most of the minimum wage increases occurred—all the TWFE estimates indicate that minimum wages have minimal effects on restaurant employment. This finding holds whether we use the CBP data preferred by JNR or the QCEW data used by DLR, or whether we examine variations within CZs, within BCPs, or across all states. We also reach the same conclusion using a completely different data source, the American Community Survey (ACS).

Having established the fragility of JNR's TWFE estimates, we then provide what we regard as the best and cleanest evidence. Drawing on Cengiz et al. (2019), we use a proper difference-in-differences event study design that: 1) uses clean control units that did not receive minimum wage changes within the event window, 2) pools over as many as 45 events over the 1990-2019 period (combining a total of 135 minimum wage increases), 3) considers different approaches to constructing the counterfactual potential outcome (BCP, CZ, or all states), and 4) considers outcomes up to 3 years before a policy change and 6 years after.

When we pool across these 45 events, the transparent difference-in-differences design tells the same story as DLR: minimum wage increases have had minimal effects on restaurant employment. These results are highly robust, regardless of how we form the counterfactual. This conclusion remains the same whether we use all untreated states, multi-state commuting zones, or border county pairs. For the own-wage elasticity (OWE) of employment—which scales the employment effect by the wage effect—estimates using QCEW data range between -0.39 and 0.49, depending on the specification, sample, and the use of population weights. None are statistically distinguishable from zero, the confidence intervals rule out an OWE of -1.5 from JNR's preferred specification, and the point estimates fall in the "small or positive" range described in Dube and Lindner (2024). Event study estimates using the CBP and ACS also do not show any statistically significant or sizable reductions in restaurant employment resulting from minimum wage increases.

Let us be clear: looking within commuting zones (CZs) can be a useful way of forming counterfactuals—as we did in a 2009 working paper (Allegretto et al. (2009)). Similarly, looking within border county pairs can also provide valuable insights. Both approaches can be effective tools for studying minimum wages or other policies, but their ability to offer a reliable counterfactual depends on the context.[3]

However, as we have mentioned, the particular tool used by JNR—a variant of the TWFE model that regresses log minimum wages on employment together with a set of fixed effects—can lead to biased and opaque estimates. JNR's set-up—a TWFE panel regression, a sample that stretches back to a period with little minimum wage variation, and using within-commuting zone comparisons—does not sufficiently guard against biases arising from violations of parallel trends from a distant past. In contrast, DLR's BCP estimator performs better, and is less sensitive to the inclusion of data from the 1990s.

At the same time, it is best to dispense with TWFE-style regressions, however one chooses to model the time effects. With hindsight, if the difference-in-differences literature demonstrating the pitfalls of the TWFE estimator had existed in 2010, DLR would not have used a TWFE-style regression to pool cross-border comparisons (two of us were original coauthors of DLR). Instead, we would have produced pooled difference-in-difference estimates that compare across state borders, as we report here. As it turns out, doing so yields the same conclusion as in DLR. Whether we use border county pairs or multi-state commuting zones, or simply compare across all states, minimum

[3] Both approaches can also have shortcomings, such as limited external validity and potential cross-border spillovers. JNR's control ring approach (similar to the method used in Boone et al. (2021)) is a valuable contribution that can be fruitfully applied to assess such spillovers when using border discontinuity designs.

wages so far have had small effects on restaurant employment.

2 A simple comparison of high- and low-minimum wage states

We illustrate the main challenges facing JNR's identification strategy, as well as the conclusions that should follow from any transparent event-study design, by making a simple set of comparisons over time between two groups of states—those that have ever raised their minimum wages versus those that have not.[4] Since the mid-1980s, 35 states have instituted a state-level minimum wage at one time or another, while 15 states never have. The minimum wage differences between these "ever treated" states ($D_s = 1$) and the "never treated" states ($D_s = 0$) capture most of the state-level minimum wage variation in the U.S.[5]

Figure 1 presents descriptive trends of differences in outcomes between these two groups of states, where the outcomes are the log of the minimum wage (greater of the federal or state floor), and the log of per-capita restaurant employment. The data come from the state-level QCEW for the restaurant industry, using a bridge between the post-1990 NAICS classifications and the pre-1990 SIC classifications.[6] Formally, we can express the state-level differences at date t as:

$$\Delta E(Y_{st}|t) = E(Y_{st}|t, D_s = 1) - E(Y_{st}|t, D_s = 0). \tag{1}$$

As shown in Panel A of Figure 1, state-level minimum wage events have occurred in three main waves, each following a prolonged period of federal inaction on minimum wages. The first wave, beginning in the late 1980s, created modest policy differentials among a small number of states. (This variation helped spawn the "new minimum wage" literature of the early 1990s.) A second, larger wave began in the early 2000s. The third wave, starting around 2014, has led to the largest

[4] These patterns are also discussed in detail in the forthcoming *Handbook of Labor Economics* chapter by Dube and Lindner (2024).

[5] Here we focus on the 48 contiguous states following JNR, additionally excluding Rhode Island and Delaware due to difficulties bridging the industry coding, as explained in Appendix C. This results in a comparison of 31 "ever treated" and 15 "never treated" states.

[6] The appendix in Dube and Lindner (2024) provides details on the data construction for the NAICS/SIC bridge. In this paper the alignment is conducted at the county level for the within-CZ analysis.

and most durable minimum wage differentials ever experienced within the U.S, and which persists today.

Panel A of Figure 1 also reveals that the employment rate gap between these two groups of states declined continuously and substantially between 1980 and 2000, despite the minimal changes in the minimum wage gap during that time. The minimum wages gap was close to zero in both 1980 and 1996, even as the restaurant employment gap was around 0.12 log points lower in the latter year. This pattern strongly indicates a violation of parallel trends. In contrast, after 2000, when the minimum wage gap increased substantially, the restaurant employment gap remained largely stable. Taken together, these patterns suggest that parallel trends violations occurred before 2000 and that there has been little correlation between minimum wages and employment since 2000, even as minimum wage levels and their variation have increased substantially.

To illuminate the implications of these patterns for JNR's findings, Panel B of Figure 1 shows the same outcomes, but for within-CZ minimum wage and employment differences. We use the same multi-state commuting zone (MSCZ) pairs as JNR, but use our QCEW data. The within-CZ differences in outcomes by D_s status at any point in time can be expressed as:

$$\Delta_{CZ} E(Y_{szt}|t) = E(Y_{szt}|z, t, D_s = 1) - E(Y_{szt}|z, t, D_s = 0) \quad (2)$$

Panel B reveals that the patterns within CZs are remarkably similar to those between the two groups of states in panel A. The minimum wage gaps in the late 1980s and 1990s were quite modest, even more so when comparing within MSCZs. In contrast, in more recent decades, the minimum wage gaps within CZs have become much larger. In the 1980s, when the minimum wage gap did not change, restaurant employment trends diverged sharply, again indicating a violation of parallel trends. This pattern continued in the 1990s, when very small gaps in minimum wages were accompanied by a continued decline in employment. However, since the late 1990s, even as minimum wages increased much more in the "ever treated" states, the within-CZ restaurant employment gap between the two groups of states remained quite stable. For instance, between

2013 and 2019, the log minimum wage gap between the two sides of cross-state commuting zones grew from 0.04 to 0.18. Yet the gap in restaurant log employment rate changed little, from -0.04 in 2013 to -0.05 percent in 2019. These comparisons yield a minimum wage elasticity (MWE) of $\frac{-0.05-(-0.04)}{0.18-0.04} = -0.07$, much smaller than in JNR, whose preferred estimate imply a MWE of -0.24.

If we instead compare the 2019 gaps with those in 1990—when the log minimum wage and the log employment gaps were 0.01 and 0.07, respectively—the estimated MWE is $\frac{-0.05-(0.07)}{0.18-0.01} = -0.71$, which is ten times larger in magnitude. The estimate for the early baseline yields a larger negative value even within commuting zones. Yet the early period exhibits minimal minimum wage variation and strong indications of parallel trends violations.The larger negative estimate comes entirely from using a 1990 baseline instead of a 2013 baseline for evaluating minimum wage changes occurring after 2013.

The TWFE regression produces an estimate that is a mixture of all such pairwise comparisons among different time periods (though with possibly negative weights). As a result, a TWFE estimate is influenced by data from the earlier period.

To summarize, Figure 1 and our simple comparisons yield six important insights:

1. Relative restaurant employment in "ever treated" states fell substantially during the 1980s and 1990s, even as relative minimum wages changed very little during this period, indicating violations of parallel trends.

2. In contrast, post-2000, employment gaps remained relatively stable despite much larger differences in minimum wages.

3. Both of the above conclusions hold in our within-CZ comparisons.

4. A TWFE regression considers all pairwise comparisons across time periods (see Goodman-Bacon (2021)). As a result, we expect it to be particularly sensitive to the inclusion of pre-2000 data, even though this period contributes little usable minimum wage variation. Including the early years mainly affects the baseline employment level used to compare

employment during periods of high minimum wages after 2000. Moreover, we expect these biases to be present in within-CZ comparisons as well.

5. We expect that a well-defined difference-in-differences event study—where treated and control units are compared within clear event windows rather than across decades—would be less sensitive to violations of parallel trends before 2000. The event window limits the use of very old data for the pre-treatment baseline, and most of the state minimum wage increases occurred after 2000.

6. To be convincing, any estimate suggesting significant disemployment effects in the restaurant sector due to minimum wage increases would need to explain why such results conflict with the findings from these straightforward comparisons.

3 The fragility of JNR's TWFE results

The TWFE regression model has a prominent history, including in the study of minimum wage policies. The classic TWFE study is Neumark and Wascher (1992). Building on this work, DLR adapted the TWFE model to a border discontinuity design by allowing the time effects—typically modeled as being common across units—to vary by local area (border county pairs). This approach enabled DLR to compare contiguous counties across state borders with different minimum wage levels, thereby controlling for unobserved regional heterogeneity that might otherwise confound the results. This method aligns with the spirit of local area case studies, such as those conduced by Card and Krueger (1994).

In their paper, JNR similarly estimate a TWFE-style regression, allowing the time effects to vary by geography (multi-state commuting zones). The key independent variable in their preferred regression specification is the log of the minimum wage. The general estimating equation can be written as follows:

$$\ln y_{it} = \beta \ln(MW_{it}) + X_{it}\Lambda + \gamma_i + \delta_{jt} + \nu_{it} \tag{3}$$

Here i is the local area, which may be a state, a county, or a commuting-zone-by-state group, and γ_i is the unit fixed effect. Additionally, δ_{jt} is a j-specific time fixed effect, where j may be (1) the US as a whole, (2) a commuting zone pair, or (3) a contiguous county pair. When common time effects are used, we refer to the specification as TWFE-logMW; when CZ or BCP pairs are used, we refer to it as TWFE-logMW-CZ or TWFE-logMW-BCP. The outcome $\ln y$ is either log restaurant employment or log of average restaurant earnings. When estimating the TWFE models, we also follow JNR in using a vector of controls, X_{it}, which includes the log of working-age population and the log of non-restaurant employment or earnings.[7] Finally, β is the minimum wage elasticity.

While the TWFE regression has been used extensively in empirical work, it suffers from two well-known problems. First, except in the case of a single treatment event or when all minimum wage increases occur simultaneously, the TWFE estimate does not correspond to one derived from a difference-in-differences design. Although the TWFE estimate represents a weighted average of treatment effects, it may assign potentially negative weights, leading to serious biases. Specifically, with staggered adoption, effects from early-adopting units might enter with negative weights (Goodman-Bacon, 2021).

Second, the results can be particularly sensitive to violations of parallel trends, including those occurring long before the time of adoption (Cengiz et al., 2019). Additionally, the dynamic effects implied by the distributed lag version of the TWFE estimator can be highly misleading and do not provide a reliable assessment of the parallel trends assumption (Sun and Abraham, 2021). In other words, using variation from outside the event window can produce biased and opaque inferences. Generally speaking, while a difference-in-differences event study assumes parallel trends within the studied event window, a TWFE regression imposes a much stronger assumption: parallel trends over the entire sample period. This approach thus can amplify biases resulting from parallel trends violations.[8]

[7]Some technical challenges arise when including time-varying covariates in a difference-in-differences estimation, as discussed in Dube and Lindner (2024). To avoid those issues, in our preferred event study estimation, we use the log of per-capita employment as the outcome and we do not include any covariates. However, this modification makes little substantive difference in practice.

[8]For a discussion of the stronger versus weaker forms of parallel trends assumptions in alternative difference-in-difference estimators, see Roth et al. (2023).

It is important to keep these biases in mind as we examine the empirical estimates from the TWFE model and assess their sensitivity to various sample selections.

3.1 Estimates using the TWFE approach

The first three columns of Table 1 display the core findings in JNR, using their County Business Patterns replication data.[9] The last three columns show analogous estimates, but focusing on the 2000-2016 sample, when most of the minimum wage variation in the data occurred. We present results using both the minimum wage elasticity and the own-wage elasticity. MWE is the metric reported in JNR. Recall that the OWE scales the employment effect by the wage effect, providing a more economically meaningful measure to understand the welfare effects of the policy (see Dube and Lindner (2024)).

JNR's preferred specification (TWFE-logMW-CZ, without weights) yields a MWE of -0.242 (s.e. 0.126) for the 1990-2016 sample (column 2), statistically significant at the 10 percent level. However, when we narrow the focus to the 2000-2016 period, the estimate drops by more than half, to -0.114 (s.e. 0.123), and is no longer statistically distinguishable from zero (column 4).[10] Similarly, the estimated minimum wage employment elasticity from the TWFE-logMW specification (with common time effects) changes from -0.338 (s.e. 0.089) to -0.077 (s.e. 0.063) when we focus on the later period (columns 1 and 4, unweighted).[11] In contrast, the TWFE-logMW-BCP estimates (columns 3 and 6, unweighted) tend to be closer to zero in both samples, shifting from -0.081 (s.e. 0.065) to 0.044 (s.e. 0.059).

Figure 2 plots the MWE estimates using the TWFE-logMW-CZ specification but using rolling

[9] One minor difference between these estimates and those reported in JNR concerns how we handle standard errors. JNR use double-clustering, following the approach in DLR, clustering on both state and border segment. However, since DLR was written, it has become clear that clustering just at the state level is more appropriate, as it is the unit where treatment occurs and accounts for any duplication of areas in the pair design. Therefore, in this paper, we cluster only at the state level. (Double-clustering tends to slightly reduce JNR's standard errors for their preferred specification.)

[10] With population weights, the estimate declines in magnitude from -0.141 (s.e. 0.077) for the 1990-2016 sample to -0.057 (s.e. 0.047) and not statistically significant for the sample starting in 2000.

[11] With population weights, the TWFE-logMW estimate is small and not statistically significant regardless of the sample period (columns 1 and 4, weighted by population).

start years between 1990 and 2005. The figure also shows estimates using the QCEW, separately for samples that end in 2016 (as in JNR) and for samples that extend through 2019. Columns 1-9 of Table 2 display a complete set of QCEW estimates with different specifications and samples.[12]

Figure 2 presents four key results. We begin with the first three, which are based on CBP and QCEW data: 1) Estimates from both datasets are more negative when the data from the 1990s are included. In contrast, results from 2000 and later suggest a minimum wage employment elasticity that is much smaller in magnitude and not statistically distinguishable from zero. 2) Simply extending JNR's sample through 2019 (using QCEW data) leads to a minimum wage employment elasticity estimate that is substantially smaller and statistically indistinguishable from zero. 3) The results from the CBP are somewhat more negative, as are estimates that are not weighted. But again, this result arises mostly from the inclusion of earlier years. Estimates starting in 2002 or later are smaller than -0.1 in magnitude and statistically indistinguishable from zero when we use weights—and whether we use the QCEW or the CBP.[13]

Since the CBP cannot be extended past 2016, it is of limited use for studying the large minimum wage increases during the more recent period.[14] In response to JNR's concerns about the coverage of QCEW data, we provide additional evidence from the more recent period using American Community Survey (ACS) data, which has been used to study commuting-zone level outcomes (e.g., Autor and Dorn, 2013).[15] Since the ACS data allows us to construct CZs back to 2005, our ACS sample comprises 2005-2019. We construct MSCZ and BCP pairs just as we did using QCEW and CBP data. The outcomes here are log of hourly earnings and log employment for restaurants

[12] We note that for JNR's preferred sample period (1990-2016) and specification (within-CZ, unweighted), the QCEW data also produces a statistically significant, and sizable, MWE estimate of -0.151, (s.e. 0.078); so none of the points that follow are driven by an inability to find a substantial negative effect using QCEW data along with JNR's preferred specification and sample period.

[13] The estimated CBP wage effects are weaker in the later period, when the minimum wage increases were stronger. This pattern, which does not arise with the QCEW data, raises questions about the quality of the CBP data. In addition, CBP data are not available after 2016, making it less useful for researchers studying minimum wage policies.

[14] JNR state that changes in Census confidentiality practice after 2016 make it impossible to implement the employment-imputation algorithm they use to create their version of the CBP data.

[15] Although the QCEW restaurant data cover fewer areas due to confidentiality-related suppression—with balanced panels of 430 county pairs and 95 MSCZ pairs, compared to 1,165 and 151 pairs in the CBP, respectively—the overwhelming majority of the population resides in areas with unsuppressed QCEW data. Our QCEW sample includes 83 percent of the BCP sample and 88 percent of the MSCZ pairs sample from the CBP data (see Appendix Table A3) Additionally, none of the MSCZ pairs are dropped in the ACS sample.

(NAICS code 722).

The last 3 columns of Table 2 show the results using ACS data. The ACS data and TWFE models estimate substantial wage effects, ranging between 0.124 and 0.292; all but one are statistically significant at the 5 percent level (the unweighted MSCZ-pair estimate is significant at the 10 percent level). In contrast, the minimum wage employment elasticities range between -0.080 and 0.053, with JNR's preferred specification (unweighted, MSCZ-pair, in column 11) producing a MWE of -0.028 (s.e. 0.095); the 95 percent confidence interval rules out the JNR central estimate of -0.242. The OWEs range between -0.329 and 0.520 across the six specifications (none are statistically significant), suggesting limited or no job loss. In summary, using a completely different data source from a more recent period, while applying JNR's preferred specification in the same multi-state commuting zones, leads to a markedly different conclusion.

These patterns are consistent with expectations from the straightforward two-group comparisons in the previous section. The TWFE regression is particularly sensitive to the inclusion of data from the 1990s—a period with few minimum wage changes but significant violations of the parallel trends assumption. Unfortunately, using MSCZs does not sufficiently guard against these parallel trends violations from the 1990s.

3.2 TWFE-binary is not a solution

JNR also report results using a variant of the TWFE model in which they code events using binary changes instead of continuous ones (TWFE-binary). These results are similar to those from their TWFE-logMW models, leading them to infer that their TWFE results are comparable to what one would obtain from a clean difference-in-differences event study design.

However, using discrete instead of continuous treatments does not resolve the underlying problems. First, JNR's regression estimates do not address the widely-recognized negative weighting problem of the TWFE regression.[16] Second, it fails to address the event window problem. Compar-

[16] Freyaldenhoven et al. (2021), whom JNR cite to motivate their specification, explicitly state: "Under staggered adoption, a two-way fixed effects estimator of β_m in (1) can be represented as a weighted average of policy effects for different event times and cohorts (Sun and Abraham 2021). However, this estimator can put nonzero weight on effects

ing outcomes outside a clearly defined event window—such as making comparisons with periods far in the past—can be problematic, as we have seen. Specifically, the presence of "binned up" lag and lead terms ($I_{i,t-M}$ and $I_{i,t+M}$) in their equation (8) means that the estimates are based on comparisons from outside the event windows. This type of distributed lag TWFE regression (which was once common in empirical work) does not yield unbiased estimates of outcome dynamics around the event date, contrary to what JNR claim.[17]

What is needed instead is a standard difference-in-differences design, with clearly defined events and admissible controls, and with comparisons restricted to within the event window. Not doing so brings in comparisons among all data points, including periods far in the past and previously treated units; and a distributed lag regression does not properly capture the dynamics.

Fortunately, there is a widely-accepted solution to these problems: pooling across proper difference-in-differences event studies.[18] We turn next to this approach.

4 The clean evidence: difference-in-differences event study

We begin by discussing our difference-in-differences specifications. Next, we present our event study results using QCEW data, with counterfactuals formed either by comparing with all untreated states or by restricting comparisons to within CZ or border county pairs. We then repeat these exercises using CBP and ACS data.

4.1 Specifications

A pooled event study can be understood as an extension of individual case studies, like the one by Card and Krueger (1994). The stacked event study method introduced by Cengiz et al. (2019) illustrates this point. Cengiz et al. (2019) examine all prominent minimum wage changes from 1979 to 2016. For each state-level minimum wage change, they create an event-specific dataset, using

at event times other than m, and can put negative weight on the effect at event time m for some cohorts."

[17] Dube and Lindner (2024) provide further illustration of how the TWFE-binary approach encounters the same problems when the analysis sample includes pre-2000 data.

[18] For example, Callaway and Sant'Anna (2021); Cengiz et al. (2019).

data from the treated state for the years surrounding the policy change (e.g., from three years prior to five years after the policy change). The event-specific dataset also incorporates observations from all "clean control" states that did not enact any major state-level minimum wage changes within the event window. Thus, similar to an individual case study, each event-specific dataset contains one treated state, with the counterfactual outcome derived by averaging data across all untreated (control) states. After assembling all the event-specific data in a stacked dataset, we can apply various analytical approaches.

As it turns out, we can obtain the exact estimate from this stacked difference-in-differences approach without stacking the data (i.e., using the original panel data).[19] Additionally, since many minimum wage changes are phased in over time, we combine multiple changes into a single event when they occur within three years of each other. (See Appendix D for details on our event construction.) Using this definition, we have 45 overall state-level minimum wage events during the 1990-2019 period, reflecting a total of 135 underlying minimum wage increases.[20] Appendix Figure A1 plots these events along with all minimum wage changes over time.

The regression specification is as follows:

$$\overline{\Delta \ln y_{it}} = \beta \times D_{it} + \delta_{jt} + \nu_{it} \tag{4}$$

In this equation, $\overline{\Delta \ln y_{it}} = \frac{1}{6}\sum_{k=0}^{5}\left(\ln y_{i,t+k} - \ln y_{i,t-1}\right)$ represents the average difference in the log outcome between the post- and pre-treatment periods over the post-treatment window. The variable $D_{it} = 1$ is an indicator that takes the value of one if unit i (a state, county, or a CZ-state group) is newly treated at date t. The regression is estimated using a sample that includes only observations when the state is either newly treated at time t ($D_{it} = 1$) or belongs to the "clean control" group, as defined below. The time effects (δ_{jt}) here are allowed to be either common across the U.S. or to vary by CZ or BCP pair, just as before. Critically, the estimate of β here

[19] See Dube et al. (2023) for the general case and Dube and Lindner (2024) for the minimum wage context.

[20] There were a total of 269 state-level increases during the 1990-2019 period. This total includes increases of less than 5% or $0.25, many of which resulted from automatic inflation indexing. Apart from the 135 increases we evaluate, there are 38 events which are prominent state-level increases of more than 5% or $0.25; but 36 of them take place in years with a federal increase, and are hence difficult to evaluate.

represents a weighted average of the individual difference-in-differences estimates, where all the weights are guaranteed to be positive.[21]

We can also estimate separate regressions for each k using equation (4) with long-differenced outcomes to recover the same dynamic responses (a regression analog to Callaway and Sant'Anna (2021)):

$$\ln y_{i,t+k} - \ln y_{i,t-1} = \beta_k \times D_{it} + \delta_{jt} + \nu_{it} \tag{5}$$

where the estimation sample consists of observations that are either newly treated or clean control units.

We consider states to be clean controls only if they did not experience a state minimum wage in the three-year pre-treatment period or in the (up to six) year post-treatment period. We also exclude from the event windows any time periods with federal minimum wage changes occurring either in the year prior to the increase or during the post-treatment period. On average, the 45 events have a 4.2 year post-treatment period. We can study almost all these events using state-level QCEW data, which do not have any data suppression issues.[22]

We can also use counties with balanced panels to study the effects of these events within multi-state CZ pairs and within BCPs. When we use only variation within multi-state CZs, we can use 37 events. When we limit the variation to be within BCPs, we can study 43 events.[23]

These state-level events produced on average a 24.6% increase in the minimum wage (taking the maximum increase over the post-treatment period). The average within-BCP and within-MSCZ increases are almost identical, at 24.3 and 24.6 percent, respectively.

[21] As explained in Dube et al. (2023), the observations can be re-weighted to produce an equally weighted estimate of the Average Treatment Effect on the Treated (ATT). The equally weighted version of the LP-DiD is numerically equivalent to the DiD estimate from Callaway and Sant'Anna (2021).

[22] The exceptions are 1999 events in Rhode Island and Delaware, which are missing some state level NAICS restaurant employment data for the 1990s, as explained in Appendix C. As a result, when using state-level QCEW data, we can estimate 43 event studies out of the 45.

[23] Cengiz et al. (2019) as well as Wursten and Reich (2023) use individual minimum wage events, separately studying some increases that are parts of multi-phased increases. This approach requires controlling for other events within the window. However, pooling events as we do here involves fewer assumptions, makes it unnecessary to control for other events and allows studying larger cumulative changes.

4.2 Difference-in-differences event study results

Table 3 and Figure 3 present estimates from the event study analyses using QCEW data. - we first present event study figures for log earnings and employment using state-level data. Since JNR's preferred specification does not use weights, we show event study figures without weights for comparability; the tables report both weighted and unweighted estimates. The top panel of Figure 3 displays the state-level evidence. These results pool all the minimum wage events, use only clean control units, and aggregate the individual event estimates. The aggregation avoids the negative weighting pitfalls of the TWFE model as well as any parallel trends violations that occurred outside the event window.

Panel A of Figure 3 shows clear increases in restaurant earnings, which stabilize by the third year after a minimum wage increase. The employment estimates (Panel B) do not reveal any parallel trends violations; they do indicate a small reduction that is not statistically distinguishable from zero. As column 1 of Table 3 shows, the associated (unweighted) average post-treatment earnings effect is 0.023 (s.e. 0.005), while the employment effect is -0.005 (s.e. 0.005). The associated minimum wage employment elasticity is -0.021 (s.e. 0.019). The associated own-wage elasticity (OWE), which divides the effect on log employment by the effect on log wages is small: -0.213 (s.e. 0.191), similar to OWEs in the minimum wage research literature (Dube and Lindner (2024)).[24] The population-weighted estimates are similar: the minimum wage employment elasticity is 0.011, and the OWE is 0.091; neither is statistically significant. Averaging the weighted and unweighted OWE estimates in column 1 yields an OWE of -0.014.[25]

Our estimates using alternative counterfactuals that look within commuting zones (Panels C, D of Figure 3 and column 2 of Table 3) and border county pairs (Panels E, F of Figure 3 and column 3

[24] The OWE is the labor demand elasticity in the competitive model.

[25] Dube and Lindner (2024) provide analogous event-based estimates using state-level data, but for 61 events in the 1980-2019 period; in this paper we use the post-1990 period because it is the focus of JNR. Dube and Lindner (2024) also show the event-by-event estimates that underlie the averaged estimate, highlighting the transparency of the pooled event study design. They additionally provide an estimate of the ATT that equally weights each event, instead of using variance weighting. When we do so for our 43 events, variance weighting make little difference. Without population weights, the equally-weighted OWE is -0.20 while the variance-weighted OWE is -0.17. With population weights, the equally weighted and variance weighted OWEs are 0.13 and 0.07, respectively.

of Table 3) yield very similar estimates. In all cases, the wage effects are substantial—the estimates range between 0.020 and 0.031—and are statistically significant. However, Panels D and F show that employment is (1) flat prior to the event, indicating no pre-existing trends and (2) remains flat for up to 6 years after the event. As columns 2 and 3 of Table 3 report, the minimum wage employment elasticities range between -0.030 and -0.000, while the OWEs range between -0.385 and -0.002, and none are statistically significant. Importantly, these conclusions still hold six years after the events. As columns 4-6 show, these conclusions all remain the same when we consider events from 2000 or later, in contrast to the sensitivity of the TWFE estimates to sample period (except for a positive and marginally significant within-CZ unweighted OWE estimate of 0.493.)

These results differ substantially from our TWFE-logMW-MSCZ estimates (unweighted, QCEW 1990-2016 sample), where the minimum wage employment elasticities is -0.151 and the OWE is -0.712. In other words, when we use standard difference-in-differences techniques and take a proper average of estimates from all usable state minimum wage increases between 1990 and 2019, we find no indication of job losses, whether we look within CZs, within border county pairs, or across all states.

To emphasize, when we update DLR using a properly pooled difference-in-differences design that restricts comparisons to within border county pairs, we obtain the unweighted column (3) MWE estimate of -0.028 (s.e. 0.039). Alternatively, when we follow JNR's suggestion to look within CZs but apply a proper difference-in-differences design, we obtain a MWE estimate of -0.000 (s.e. 0.047). Both estimates strongly contradict JNR's core conclusion.

These findings do not change when we conduct the event studies using JNR's preferred CBP data or ACS data—and using their preferred geographic comparisons (common or CZ-specific time effects). Appendix Table A1 reports the state-level (column 3) and CZ-level (column 4) estimates using County Business Pattern data. In all cases, the MWEs are small (between -0.059 and 0.003), with standard errors that rule out JNR's central estimates.[26] Columns 1 and 2 present estimates using ACS data (2005-2019); MWEs range between -0.074 and 0.366. Three of the four MWE

[26]The wage effects are less precise and smaller in the CBP data (statistically significant in the all-state sample, but not in the MSCZ sample), similar to our findings above, and raising questions about the quality of the CBP data.

estimates are statistically indistinguishable from zero, while one (within-CZ, unweighted) estimate (column 2) is significant and positive.

The event study estimates are thus much less sensitive to (1) choice of period, and (2) choice of counterfactual (all states, within-CZ, within-BCP). These consistent results are reassuring. They suggest that the difference-in-differences design obtains credible and robust estimates, locally within the event window and for the average event. The employment estimates for these events, estimated using a clean difference-in-differences design, contrast sharply with the TWFE estimates. They are also consistent with the two-group comparisons displayed in Section 2.

5 Conclusion

Minimum wage researchers have made considerable progress in better understanding employment effects, especially in the often-studied restaurant sector. The most transparent evidence—using a difference-in-differences design with clear control groups and an event window—suggests very modest (and sometimes positive) employment effects. This conclusion holds across various approaches to forming counterfactuals, whether by comparing to all states that did not raise minimum wages or by limiting comparisons to within multi-state commuting zones or border county pairs. Estimates from our event study using the QCEW suggest an OWE between -0.39 and 0.49, indicating a small negative or positive effect that is statistically indistinguishable from zero.

JNR use TWFE-style estimators and data starting in 1990 to compare employment outcomes within commuting zones. Their approach produces pronounced negative employment estimates, especially when using their preferred specification and sample years. However, their results are confounded: their TWFE model amplifies parallel trends violations from the pre-2000 years, which experienced little minimum wage variation. Even using their preferred specification, JNR's conclusions do not hold for the period since 2000, whether the data set is the CBP, the QCEW or the ACS. Had JNR used a proper difference-in-differences design instead of a TWFE estimator, they would have reached a very different conclusion, regardless of the data set, control group, or time period used. Further exploration of the sensitivity of different TWFE-style estimators to data from the 1980s and 1990s is unlikely to advance our understanding of minimum wage policies.

References

Allegretto, Sylvia, Arindrajit Dube, and Michael Reich (2009) "Spatial Heterogeneity and Minimum Wages: Employment Estimates for Teens Using Cross-State Commuting Zones,"Technical report, Institute of Industrial Relations, UC Berkeley.

Allegretto, Sylvia, Arindrajit Dube, Michael Reich, and Ben Zipperer (2017) "Credible research designs for minimum wage studies: A response to Neumark, Salas, and Wascher," *ILR Review*, 70 (3), 559–592.

Autor, David and David Dorn (2013) "The growth of low-skill service jobs and the polarization of the US labor market," *The American Economic Review*, 103 (5), 1553–1597.

Boone, Christopher, Arindrajit Dube, Lucas Goodman, and Ethan Kaplan (2021) "Unemployment insurance generosity and aggregate employment," *American Economic Journal: Economic Policy*, 13 (2), 58–99.

Brummund, Peter and Michael R Strain (2020) "Does Employment Respond Differently to Minimum Wage Increases in the Presence of Inflation Indexing?" *Journal of Human Resources*, 55 (3).

Callaway, Brantly and Pedro HC Sant'Anna (2021) "Difference-in-differences with multiple time periods," *Journal of econometrics*, 225 (2), 200–230.

Card, David and Alan B. Krueger (1994) "Minimum Wages and Employment: A Case Study of the New Jersey and Pennsylvania Fast Food Industries," *American Economic Review*, 84 (4), 772–793.

Card, David and Alan B Krueger (2000) "Minimum wages and employment: a case study of the fast-food industry in New Jersey and Pennsylvania: reply," *American Economic Review*, 90 (5), 1397–1420.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019) “The effect of minimum wages on low-wage jobs,” *The Quarterly Journal of Economics*, 134 (3), 1405–1454.

Clemens, Jeffrey and Michael R Strain (2021) “The heterogeneous effects of large and small minimum wage changes: evidence over the short and medium run using a pre-analysis plan,”Technical report, National Bureau of Economic Research.

Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M Taylor (2023) “A Local Projections Approach to Difference-in-Differences,” Working Paper 31184, National Bureau of Economic Research.

Dube, Arindrajit, T. William Lester, and Michael Reich (2010) “Minimum wage effects across state borders: estimates using contiguous counties,” *The Review of Economics and Statistics*, 92 (4), 945–964.

Dube, Arindrajit and Attila Lindner (2024) “Minimum Wages in the 21st Century,” in *Handbook of Labor Economics, volume 5-6*.

Freyaldenhoven, Simon, Christian Hansen, Jorge Pérez Pérez, and Jesse M Shapiro (2021) “Visualization, identification, and estimation in the linear panel event-study design,”Technical report, National Bureau of Economic Research.

Godoey, Anna and Michael Reich (2021) “Are Minimum Wage Effects Greater in Low-Wage Areas?” *Industrial Relations: A Journal of Economy and Society*, 60 (1), 36–83.

Goodman-Bacon, Andrew (2021) “Difference-in-differences with variation in treatment timing,” *Journal of econometrics*, 225 (2), 254–277.

Hampton, Matt and Evan Totty (2023) “Minimum wages, retirement timing, and labor supply,” *Journal of Public Economics*, 224, 104924.

Jha, Priyaranjan, David Neumark, and Antonio Rodriguez-Lopez (2024) “What’s across the border?

Re-evaluating the cross-border evidence on minimum wage effects," *Journal of Political Economy Microeconomics*, Forthcoming.

Neumark, David, JM Ian Salas, and William Wascher (2014) "Revisiting the Minimum Wage-Employment Debate: Throwing Out the Baby with the Bathwater?" *ILR Review*, 67 (3_suppl), 608–648.

Neumark, David and William Wascher (1992) "Employment effects of minimum and subminimum wages: panel data on state minimum wage laws," *ILR Review*, 46 (1), 55–81.

Rao, Nirupama and Max Risch (2024) "Who's Afraid of the Minimum Wage? Measuring the Impacts on Independent Businesses Using Matched US Tax Returns."

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe (2023) "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 235 (2), 2218–2244.

Sun, Liyang and Sarah Abraham (2021) "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 225 (2), 175–199.

Vaghul, Kavya and Ben Zipperer (2016) "Historical state and sub-state minimum wage data," *Washington Center for Equitable Growth Working Paper*, http://cdn.equitablegrowth.org/wp-content/uploads/2016/09/02153029/090716-WP-Historical-min-wage-data.pdf.

Vergara, Damian (2023) "Minimum wages and optimal redistribution: The role of firm profits."

Wiltshire, Justin C, Carl McPherson, Michael Reich, and Denis Sosinskiy (2024) "Minimum wage effects and monopsony explanations," *UC Berkeley, IRLE Working Paper*.

Wursten, Jesse and Michael Reich (2023) "Racial inequality in frictional labor markets: Evidence from minimum wages," *Labour Economics*, 82, 102344.

Tables and Figures

Figure 1: Difference in minimum wages and restaurant employment between "ever-treated" and "never-treated" states, and between "ever-treated" and "never-treated" areas within multi-state commuting zones

Notes: Panel A plots the difference in restaurant log per-capita employment (left axis) and the difference in log minimum wages (right axis) between 31 ever-treated and 15 never-treated states. Ever-treated states had at least one state minimum wage increase, over and above any federal increases between 1980 and 2019. Never-treated states did not have any such increases. Panel B repeats this exercise, but calculates the difference within multi-state commuting zone pairs only.

Figure 2: Minimum wage employment elasticity in restaurants from a TWFE-logMW-CZ specification: different start and end years using CBP, QCEW, and ACS data

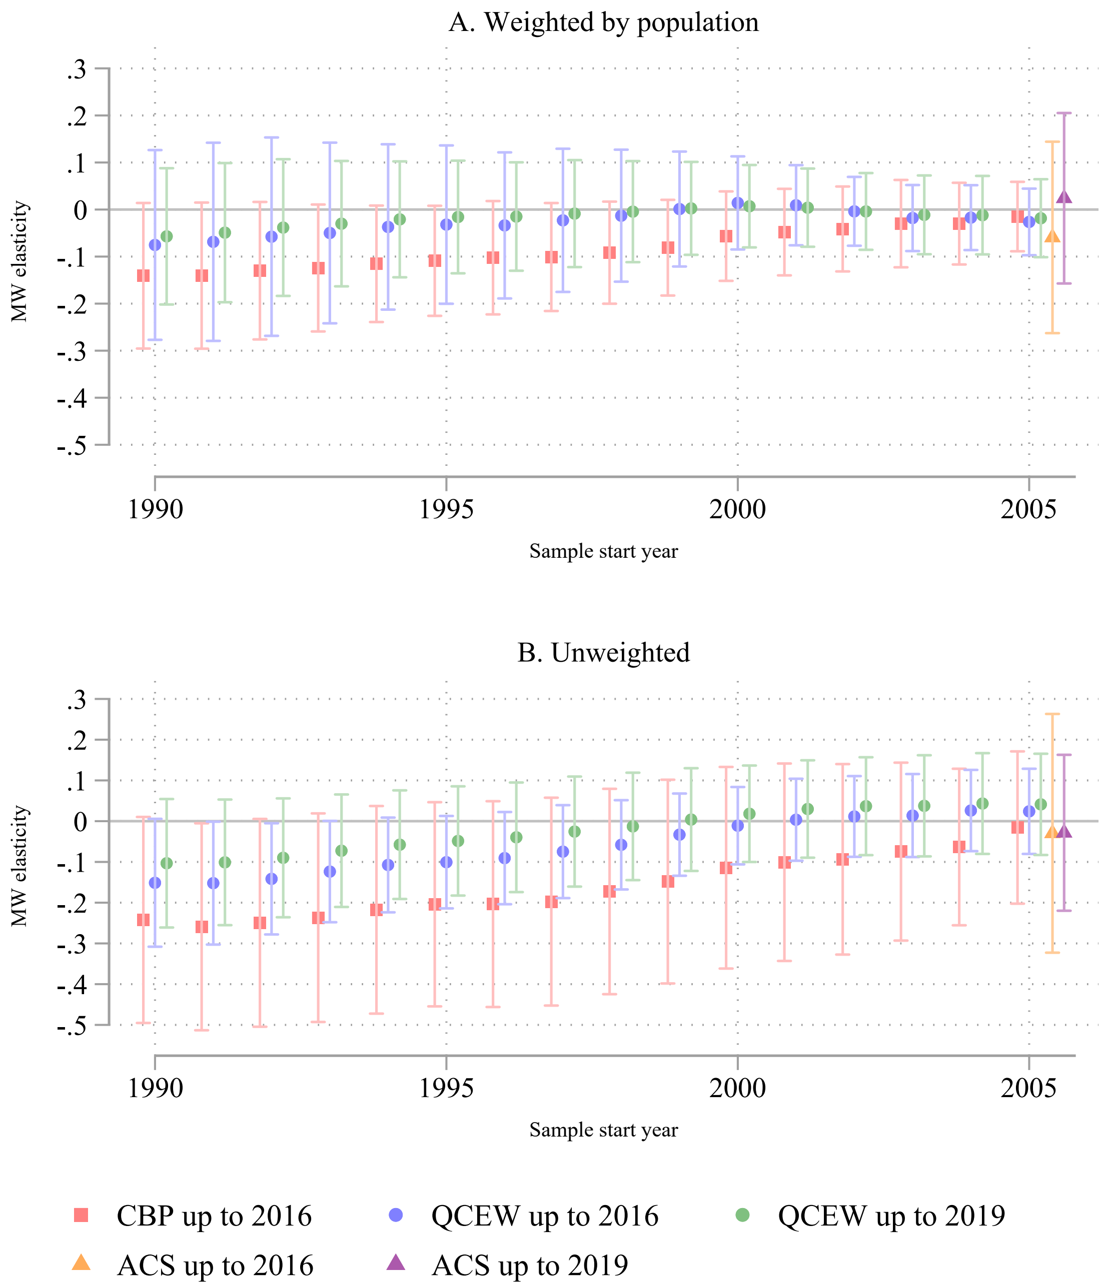

Notes: Estimated using the TWFE-logMW-CZ specification described in equation (3). Each point corresponds to an estimate from a regression using a particular dataset, starting in year shown by y-axis and ending either in 2016 or 2019. The horizontal axis reports the first year in the sample. The vertical axis depicts the effect of log minimum wage on log employment. Each color-shape combination refers to a particular dataset used and the final year in the sample, as described by the legend. Lines show 95% confidence intervals. All estimates use multi-state commuting-zone-by-state pairs sample and include fixed effects for CZ-state (repeated if a CZ-state is part of multiple pairs) and pair-year fixed effects. Estimates in Panel A are weighted using the working-age population, while Panel B is unweighted. All specifications include log of earnings and log of employment outside of the restaurant sector, and log of working-age population as control variables. Standard errors are clustered at the state level.

Figure 3: Event study estimates of MW effects on log employment and log average earnings in restaurants, 1990-2019

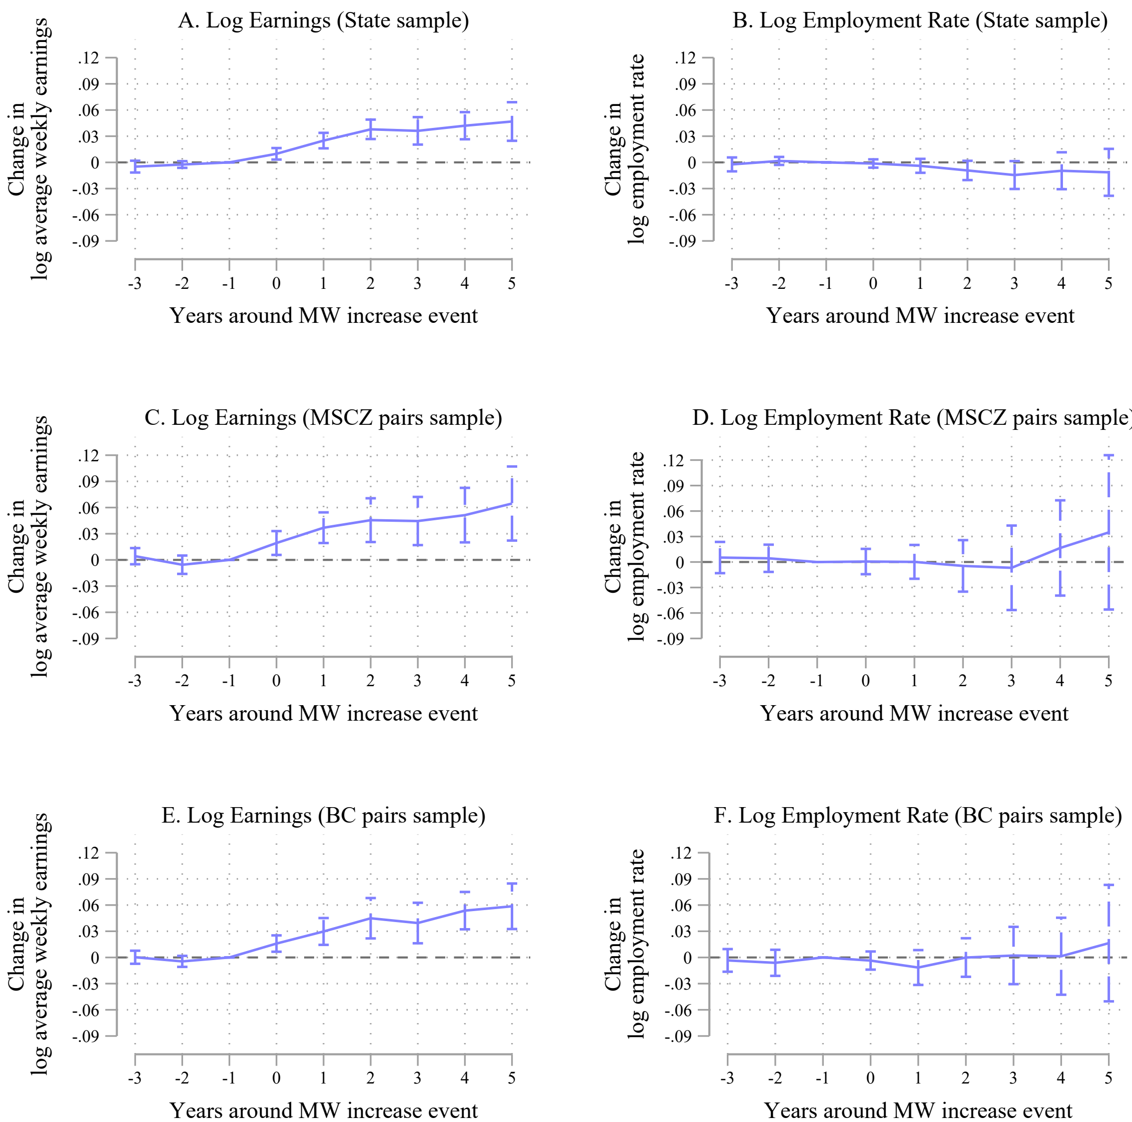

Notes: The plotted points here represent coefficients from regressions described in equation (5) using data from Quarterly Census of Employment and Wages data for 1990-2019. The top row uses the state-level sample and includes year fixed effects. The middle row uses multi-state commuting-zone-by-state pairs and includes pair-year fixed effects. The bottom row uses a sample of border county pairs and also includes pair-year fixed effects. The left column shows results for a log of weekly earnings; right column shows results for a log of employment rate. All estimates are unweighted. Standard errors are clustered at the state level. Lines indicate 95% confidence intervals.

Table 1: MW effects on log average earnings and log employment in restaurants: TWFE estimates for different samples and specifications using CBP data

	1990-2016			2000-2016		
	(1)	(2)	(3)	(4)	(5)	(6)
A. Weighted by population						
Log wages	0.192***	0.141***	0.174***	0.120***	0.098*	0.131***
	(0.021)	(0.048)	(0.019)	(0.020)	(0.058)	(0.039)
Log employment	-0.054	-0.141*	0.029	0.061	-0.057	0.079
	(0.072)	(0.077)	(0.090)	(0.056)	(0.047)	(0.060)
OWE	-0.340	-1.013*	0.170	0.518	-0.580	0.609
	(0.319)	(0.536)	(0.473)	(0.507)	(0.461)	(0.381)
B. Unweighted						
Log wages	0.215***	0.163***	0.156***	0.100***	0.086	0.109*
	(0.037)	(0.050)	(0.044)	(0.029)	(0.071)	(0.064)
Log employment	-0.338***	-0.242*	-0.081	-0.077	-0.114	0.044
	(0.089)	(0.126)	(0.065)	(0.063)	(0.123)	(0.059)
OWE	-1.531***	-1.517	-0.492	-0.723	–	0.385
	(0.371)	(0.945)	(0.430)	(0.611)	–	(0.588)
N	23,361	8,134	62,228	14,703	5,116	39,146
Fixed effects						
CZ-year		Y			Y	
BCP-year			Y			Y
Sample						
All CZ-state	Y			Y		
MSCZ pairs		Y			Y	
Border county pairs			Y			Y

Notes: Estimated using specification described in equation 3 and County Business Patterns data. Columns (1)-(3) use the years 1990-2016, and (4)-(6) use 2000-2016. Columns (1) and (4) use the sample of all commuting-zones-by-state, and include CZ-state and year fixed effects. Columns (2) and (5) use multi-state commuting-zone-by-state pairs sample and include fixed effects for CZ-state (repeated if a CZ-state is part of multiple pairs) and pair-year fixed effects. Columns (3) and (6) similarly use the sample of border county pairs and include fixed effects for border counties (each time a county is in a pair) and pair-year fixed effects. Outcomes include log weekly earnings, log employment, and own wage elasticity (OWE). OWE is the IV estimate from regressing log employment on log earnings instrumented by log minimum wage. OWE estimates are only reported when there is a strong first-stage effect. Estimates in Panel A are weighted using the working-age population, while Panel B is unweighted. All specifications include log of earnings and log of employment outside of the restaurant sector, and working-age population as control variables. Standard errors, reported in parentheses, are clustered at the state level. Stars indicate statistical significance as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: MW effects on log average earnings and log employment in restaurants:
TWFE estimates for different samples and specifications using QCEW and ACS data

	QCEW									ACS		
	1990-2019			1990-2016			2000-2016			2005-2019		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A. Weighted by population												
Log wages	0.214***	0.229***	0.201***	0.195***	0.215***	0.171***	0.146***	0.162***	0.179***	0.124***	0.203***	0.255***
	(0.024)	(0.038)	(0.030)	(0.025)	(0.043)	(0.024)	(0.022)	(0.049)	(0.029)	(0.030)	(0.062)	(0.032)
Log employment	0.025	-0.057	0.080	-0.035	-0.075	0.091	0.066	0.014	0.188**	0.053	0.024	0.028
	(0.078)	(0.072)	(0.078)	(0.068)	(0.100)	(0.082)	(0.060)	(0.049)	(0.071)	(0.048)	(0.090)	(0.067)
OWE	0.079	-0.260	0.364	-0.220	-0.347	0.486	0.456	0.068	1.050**	0.520	0.119	0.115
	(0.366)	(0.292)	(0.352)	(0.311)	(0.441)	(0.430)	(0.446)	(0.260)	(0.473)	(0.425)	(0.440)	(0.264)
B. Unweighted												
Log wages	0.254***	0.229***	0.243***	0.223***	0.198***	0.218***	0.211***	0.220***	0.225***	0.176***	0.207*	0.292***
	(0.028)	(0.056)	(0.032)	(0.029)	(0.061)	(0.039)	(0.027)	(0.029)	(0.021)	(0.047)	(0.108)	(0.072)
Log employment	-0.264***	-0.103	0.031	-0.308***	-0.151*	0.025	-0.078	-0.011	0.106**	-0.080	-0.028	0.043
	(0.088)	(0.078)	(0.057)	(0.084)	(0.078)	(0.056)	(0.053)	(0.047)	(0.049)	(0.076)	(0.095)	(0.075)
OWE	-1.047***	-0.424	0.128	-1.360***	-0.712	0.109	-0.365	-0.048	0.476**	-0.329	-0.093	0.157
	(0.302)	(0.360)	(0.230)	(0.319)	(0.448)	(0.248)	(0.237)	(0.201)	(0.228)	(0.392)	(0.457)	(0.263)
N	21,570	5,700	25,800	19,413	5,130	23,220	12,223	3,230	14,620	13,088	4,530	34,950
Fixed effects												
CZP-year		Y			Y			Y			Y	
BCP-year			Y			Y			Y			Y
Sample												
All CZ-state	Y			Y			Y			Y		
MSCZ pairs		Y			Y			Y			Y	
Border county pairs			Y			Y			Y			Y

Notes: Estimated using specification described in equation 3. Columns (1)-(9) use Quarterly Census of Employment and Wages data, (10)-(12) use American Community Survey. Columns (1)-(3) use the years 1990-2019, (4)-(6) use 1990-2016, (7)-(9) use 2000-2016, and (10)-(12) use 2005-2019. Columns (1), (4), (7), and (10) use the sample of all commuting-zones-by-state, and include CZ-state and year fixed effects. Columns (2), (5), (8), and (11) use multi-state commuting-zone-by-state pairs sample and include fixed effects for CZ-state (repeated if a CZ-state is part of multiple pairs) and pair-year fixed effects. Columns (3), (6), (9), and (12) similarly use the sample of border county pairs and include fixed effects for border counties (each time a county is in a pair) and pair-year fixed effects. Outcomes include log weekly earnings, log employment, and own wage elasticity (OWE). OWE is the IV estimate from regressing log employment on log earnings instrumented by log minimum wage. OWE estimates are only reported when there is a strong first-stage effect. Estimates in Panel A are weighted using the working-age population, while Panel B is unweighted. All specifications include log earnings and log employment outside of the restaurant sector, and working-age population as control variables. Standard errors, reported in parentheses, are clustered at the state level. Stars indicate statistical significance as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Effects of increased MW - event study estimates for different samples and specifications

	1990-2019			2000-2019		
	(1)	(2)	(3)	(4)	(5)	(6)
A. Weighted by population						
Log wages	0.033***	0.020*	0.020*	0.034***	0.024**	0.026***
	(0.007)	(0.011)	(0.010)	(0.006)	(0.010)	(0.010)
Log employment	0.003	-0.008	-0.000	0.008	-0.001	0.003
	(0.006)	(0.006)	(0.005)	(0.006)	(0.007)	(0.006)
MWE	0.011	-0.030	-0.000	0.027	-0.005	0.009
	(0.020)	(0.024)	(0.018)	(0.019)	(0.024)	(0.023)
OWE	0.091	-0.385	-0.003	0.232	-0.055	0.101
	(0.170)	(0.300)	(0.240)	(0.157)	(0.267)	(0.243)
B. Unweighted						
Log wages	0.023***	0.031***	0.028***	0.025***	0.032***	0.032***
	(0.005)	(0.008)	(0.006)	(0.005)	(0.008)	(0.006)
Log employment	-0.005	-0.000	-0.006	0.002	0.016*	-0.005
	(0.005)	(0.010)	(0.009)	(0.004)	(0.009)	(0.011)
MWE	-0.021	-0.000	-0.028	0.009	0.068*	-0.019
	(0.019)	(0.047)	(0.039)	(0.017)	(0.040)	(0.044)
OWE	-0.213	-0.002	-0.236	0.085	0.493*	-0.153
	(0.191)	(0.336)	(0.331)	(0.174)	(0.257)	(0.348)
N	714	2,132	10,414	470	1,220	6,170
Fixed effects						
CZP-year		Y			Y	
BCP-year			Y			Y
Sample						
States	Y			Y		
MSCZ pairs		Y			Y	
Border county pairs			Y			Y

Estimated using specification described in equation 4 using Quarterly Census of Employment and Wages data. Columns (1)-(3) use the years 1990-2019, and (4)-(6) use 2000-2019. Columns (1) and (4) use a states sample, and include year fixed effects. Columns (2) and (5) use multi-state commuting-zone-by-state pairs sample and include pair-year fixed effects. Columns (3) and (6) similarly use the sample of border county pairs and also include pair-year fixed effects. Outcomes include log weekly earnings, log employment rate, and own wage elasticity (OWE). OWE is the IV estimate from regressing log employment rate on log earnings instrumented by log minimum wage. Estimates in Panel A are weighted using the working-age population, while Panel B is unweighted. Standard errors, reported in parentheses, are clustered at the state level. Stars indicate statistical significance as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendix A Additional Figures and Tables

Figure A1: All minimum wage increases from 1990-2019

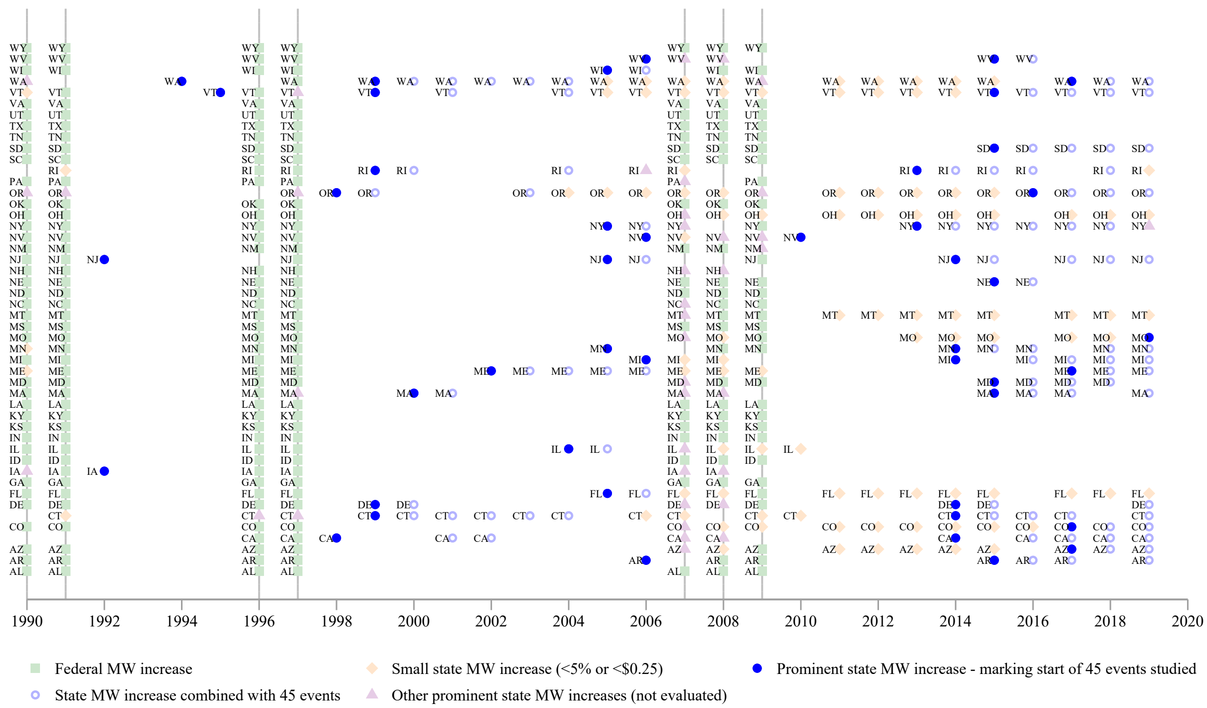

Notes: This figure plots all instances of the minimum wage increasing in any state. The blue solid circles represent the start of 45 events that we study in our main analysis. The blue hollow circles are state MW increases that are evaluated as part of the post-periods of these 45 events. There can also be state MW increases that are neither the start of an event, nor are evaluated as part of another event. These are either very small increases (some of them due to indexation) - represented by orange diamonds; or they are prominent increases, but mostly in years when the federal MW also increased (making it difficult to ascertain their impact) - represented by purple triangles (the only two exceptions are RI 2006 and NY 2019, both close enough to a prominent increase to not be counted as the start of a separate event, but more than 6 periods away from the last event-start). The green squares show instances of increases due to a federal minimum wage increase. Any state that does not have a green square in a federal MW increase year either had its own increase in that year (small or prominent) which made its binding MW higher than the federal, or had a MW higher than the federal in the previous year as well, which made the federal increase redundant. More details are in Appendix Section D

Figure A2: Cumulative response of log employment and log average earnings in restaurants from a log-point increase in MW; TWFE model with distributed lags (1990-2019; unweighted)

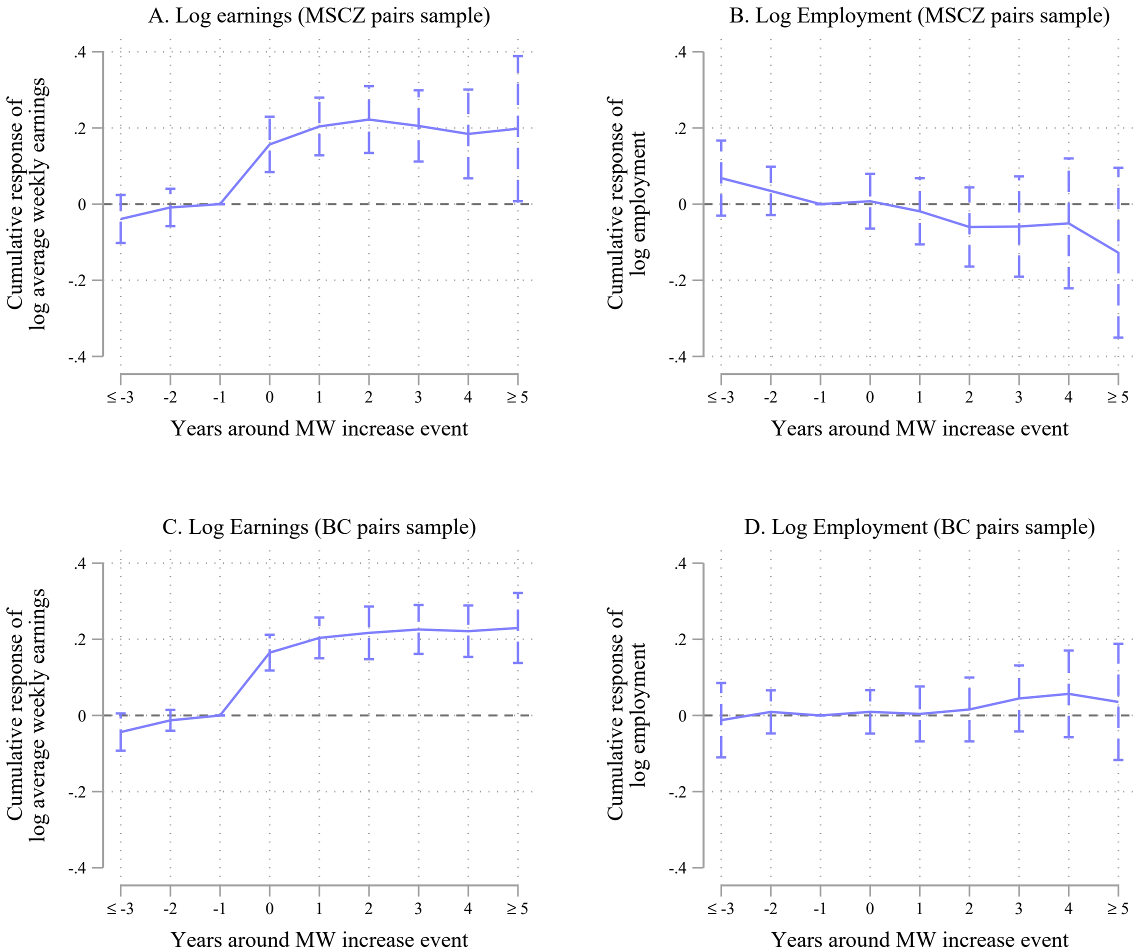

Notes: This figure plots the cumulative response of log employment and log average earnings for restaurant workers to a an increase in the minimum wage using a distributed lag model with 2 leads and 5 lags of log minimum wage. The cumulative responses are normalized relative to event date (-1). 95% confidence intervals are based on standard errors clustered by state, and all regressions use state population weights. Panels A and B use multi-state commuting-zone-by-state pairs data and include fixed effects for CZ-state (repeated if a CZ-state is part of multiple pairs) and pair-year fixed effects. Panels C and D similarly use sample of border county pairs and include fixed effects for border counties (each time a county is in a pair) and pair-year fixed effects. All estimates are unweighted. All specifications include log earnings and log employment outside of the restaurant sector, and working-age population as control variables.

Table A1: Effects of increased MW on log average earnings and log employment in restaurants: Event study estimates for different samples and specifications using ACS and CBP data

	ACS 2005-2019		CBP 1990-2016	
	(1)	(2)	(3)	(4)
A. Weighted by population				
Log wages	0.029**	0.081***	0.026***	0.009
	(0.012)	(0.026)	(0.007)	(0.011)
Log employment	0.010	0.033	0.001	-0.006
	(0.014)	(0.060)	(0.005)	(0.006)
MWE	0.030	0.121	0.003	-0.031
	(0.043)	(0.211)	(0.020)	(0.031)
OWE	0.329	0.406	0.024	–
	(0.494)	(0.774)	(0.174)	–
B. Unweighted				
Log wages	0.031*	0.066	0.020***	0.013
	(0.017)	(0.054)	(0.004)	(0.021)
Log employment	-0.020	0.090**	-0.007	-0.011
	(0.013)	(0.043)	(0.006)	(0.014)
MWE	-0.074	0.366**	-0.034	-0.059
	(0.049)	(0.175)	(0.028)	(0.076)
OWE	-0.645	–	-0.346	–
	(0.614)	–	(0.283)	–
N	245	1,064	643	3,488
Fixed effects				
CZP-year		Y		Y
Sample				
State	Y		Y	
MSCZ-state		Y		Y

Notes: Estimated using specification described in equation 4. Columns (1)-(3) use ACS data for the years 1990-2019, columns (4)-(6) use CBP data for 2000-2019. Columns (1) and (3) use the state-level sample and include year fixed effects. Columns (2) and (4) use multistate commuting-zone-by-state pairs and include CZ-state-pair-year fixed effects. Outcomes include log hourly wage for columns (1)-(3), log weekly earnings for columns (4)-(6), log employment rate, minimum wage elasticity (MWE), and own wage elasticity (OWE). The post-treatment period includes up to six years following the event. The pre-treatment period is the year prior to the event. Outcomes include log weekly earnings, log employment rate, minimum wage elasticity (MWE), and own wage elasticity (OWE). MWE is the IV estimate from regressing the change in log employment rate on the change in log minimum wage instrumented by an indicator of minimum wage treatment. OWE is the IV estimate from regressing the change in log employment on the change in log earnings instrumented by an indicator of minimum wage treatment. OWE estimates are only reported when there is a strong first-stage effect. Estimates in Panel A are weighted using the working-age population, while Panel B is unweighted. Standard errors, reported in parentheses, are clustered at the state level. Stars indicate statistical significance as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Number of events evaluated in different event study specifications

	QCEW 1990-2019			CBP 1990-2016		ACS 2005-2019	
	State	MSCZ pairs	BC pairs	State	MSCZ pairs	State	MSCZ pairs
Combined MW increases	43	37	43	39	33	22	18
All underlying MW increases	130	106	125	118	97	83	69
Pairs in regression	.	91	420	.	147	.	119

This table shows the number of events evaluated and the number of border county and MSCZ pairs in different samples' event study regressions. The numbers reported are for the full period regressions for each sample. The first row reports the total number of combined events, while the second row gives the number of all MW increases evaluated as part of the combined events included. Finally, the third row reports the number of MSCZ and border county pairs included in the regressions for the respective samples.

Table A3: Coverage and population shares of QCEW samples

	Total units (counties/pairs)	Units in sample (counties/pairs)	Percentage of units in sample	Percentage of population covered
Counties	3109	1704	54.81	91.49
MSCZ pairs	151	95	62.91	87.85
Border county pairs	1165	430	36.91	83.36

This table shows the coverage of QCEW data both in terms of number of units (pairs/counties) covered, and in terms of population covered. The QCEW analysis only keeps a balanced sample - this drops a number of counties (row 1). We also only keep complete pairs for the border county and MSCZ pairs regressions (rows 2 and 3). The first column reports the total number of units/pairs in JNR's sample. Samples are exclusive of Alaska, Hawaii, and DC.

Table A4: Details for the 45 combined minimum wage events, by year of the initial increase

	Number of events	Post period length(years)	No. of clean control states	Mean increase in log MW
1992	2	4	39	.131
1994	1	2	41	.142
1995	1	1	44	.057
1998	2	6	39	.249
1999	5	6	38	.269
2000	1	6	33	.251
2002	1	5	29	.271
2004	1	3	30	.233
2005	5	2	31	.245
2006	4	1	32	.2
2010	1	6	17	.089
2013	2	6	22	.368
2014	6	6	24	.302
2015	7	5	24	.26
2016	1	4	24	.196
2017	4	3	24	.305
2019	1	1	24	.091
Overall	45	4.2	29.4	.246

This table provides information by cohort on the 45 events used in the event-study analysis. The first column reports the number of events in that cohort/year. The second column gives the length of the post-period for that cohort in years. Post-periods are a maximum of 6 years unless interrupted by a federal MW increase year or end-of-sample. The third column reports number of states that serve as clean controls for the cohort - a clean control needs to have no state MW increase in the three years before the cohort year, and also no state MW increase in the relevant post-period. Finally, the last column reports the average log MW increase (average of the difference between log MW in the last year of the post-period and the log MW in year *t-1*, where *t* is the event year) across all events in a cohort. The last row gives the total number of events and means *by event* (not by cohort) of post-period length, clean control states, and log MW increase.

Appendix B Construction of ACS data

This Appendix describes our construction of the American Community Survey (ACS) data.

B.1 Obtaining data at the level of geographies of interest

We downloaded ACS data from IPUMS for the years 2005-2019. Detailed geographic data are available for this period. However, the most detailed geographic variable for all individuals in the ACS is for PUMAs, which differ from counties or commuting zones. Moreover, the definition of PUMAs changed in 2012[27]. Hence, we undertook a series of crosswalks to obtain data at our geographic level of interest. Since none of the crosswalks are one-to-many or many-to-one, we use population-based weights at each stage.

We start by assigning a weight to each individual in the years 2005-2011 based on their probability of being in PUMA10, given their observed PUMA00. We use the IPUMS crosswalk file, which is available at https://usa.ipums.org/usa/volii/pumas10.shtml. To obtain the relevant weights, we use variable *pPUMA00_Pop10*, the estimated percent of the 2000 PUMAs 2010 population that lies in the area of intersection. We further divide the variable by 100.

We then perform a similar procedure to obtain a Commuting Zone-by-State level dataset from the PUMA10 level, using the crosswalk described in Autor and Dorn (2013). The crosswalk provides population-based weights. The crosswalk is available from https://www.ddorn.net/data.htm as "E5" or by request to the authors. However, unlike the procedure described in the Autor and Dorn (2013), we collapse data to the CZ-by-state level, using ACS state identifiers.

Finally, we use an equivalent procedure to obtain county-level data from PUMA10, using the 2012 PUMA to 2010 County crosswalk provided by Geocorr (2014). Weights are based on the 2010 Census population. The crosswalk is available at https://mcdc.missouri.edu/applications/geocorr2014.html.

Once we obtain a weight of interest for each county or CZ, we combine it with a personal weight

[27] The definition used prior to 2012 is called PUMA00, while the one adopted in 2012 is called PUMA10

provided in the ACS. This procedure collapses the data to the geography of interest. Additionally, we construct a state-level dataset using the ACS state identifiers. No additional datasets are required.

B.2 Constructing samples and variables of interest

We use the years 2005-2019. For each year and geography of interest (county, CZ-state, state), we restrict the sample to employed individuals, using variable *empstat*. We obtain a total count of workers as a measure of employment in all industries and a count of workers in the restaurant industry as employment in the industry of interest. We classify a worker as working in the ”restaurant industry” if a reported variable *indnaics* is equal to ”722Z”, Restaurants and other food services. The code is universal for all years of interest.

We then construct a wage measure, using *incwage*, wage and salary income; *wkswork2*, binned weeks worked last year uhrswork; and usual hours worked per week. We drop individuals who are missing any of these variables. We also drop individuals with wage and salary income or usual hours worked below the first percentile or above the 99th percentile. Additionally, we use the midpoint to transform the weeks worked last year variable from a reported interval to a number.[28] We then construct an hourly wage variable as wage and salary income divided by weeks worked and divided by usual hours worked in a week. We compute this variable for all employed workers and for all workers employed in the restaurant industry. Additionally, we obtain an estimate of the working-age population, defined as all individuals with reported ages between 15 and 64, using variable *age*.

Lastly, we obtain border county pairs (BCP) and multi-state-comutting zones-pairs (MSCZP) samples to use for the analyses. For that, we use geographic definitions, including borders, specified in JNR paper.

[28] The ACS’ best measure of weeks worked is reported in six bins: 1-13, 14-26, 27-39, 40-47, 48-49, and 50-52.

Appendix C Construction of QCEW data

In this appendix, we describe the process of cleaning and using data from Quarterly Census of Employment and Wages.

We downloaded annual county-level data from QCEW, and merged in working age population from National Cancer Institute's SEER program, state level minimum wage data from Vaghul and Zipperer (2016), and commuting zones from JNR's replication files. We keep only fully balanced counties - that is, counties for which we have restaurant employment and earnings data for all years from 1990-2019. Our CZ-state dataset then collapses these balanced counties to the CZ-state level. For our pairs samples, we start with JNR's replication files (i.e. their pair definitions), and merge in our data to only keep the complete pairs with balanced panels. In all our analysis, we exclude Alaska, Hawaii, and Washington, DC.

In this paper, we mostly use data from 1990-2019, but we do use data from 1980-1990 as well in constructing Figure 1.

The QCEW data from the 1980s is based on SIC classifications, and NAICS data is available only since 1990. To address the comparability issue, we create a harmonized series for earnings and employment based on the method described in Dube and Lindner (2024)'s Appendix C.

We also use state level annual data for some exhibits. The harmonization process between NAICS/SIC for this data follows the same method.

The state level data has one additional issue. The reconstructed NAICS restaurant earnings and employment data (for NAICS code 722) was missing or incomplete for two states in the 1990s: Delaware and Rhode Island. The lack of an overlapping period with SIC and NAICS data prevented us from being able to impute NAICS-based values for the 1980s and 1990s for restaurant earnings and employment. Therefore, for the comparison between ever-treated and never-treated states (Figure 1), we remove all observations from these two states (including post-2000 observations).

Appendix D Construction of the minimum wage events

D.1 Identifying events

Our event study analyses are based on 45 combined state minimum wage increases (the exact set of usable events varies by dataset.) We depict these in the filled, dark blue circles in Figure A1. This appendix details the steps we followed to identify each of the state-level events.

1. We identified all instances of a *state* minimum wage increase of at least $0.25 and 5%.

2. Of these , we exclude those that occurred in in years with a federal minimum wage increase. We do so because an event that begins in a federal increase year cannot have a "clean" counterfactual—there are no states that did not experience minimum wage increases in that year or in the "event period." Thus, we remove events in 1990, 1991, 1996, 1997, 2007, 2008 and 2009. We label the remaining events as our set of "admissible events."

3. We then classify admissible events as "provisional combined events" if they do not include any admissible events in the preceding three years. This step allows us to pool multi-year minimum wage increases into a single combined event.

4. In some cases, this procedure does not properly capture an event's start date. Specifically, if a multi-phased increase begins with too small an increment, we would incorrectly code the start year later than when the event actually began. To remedy this possibility, we use the following algorithm:

    - For each of the provisional combined events that we identified in the previous step, we check for the presence of any minimum wage increases (large or small) in the year prior to the provisional start year, while also checking that a federal increase did not occur in the prior year.

    - If we find such an increase, and the state did not have indexation at that time, we consider that a legislated (albeit small) minimum wage increase. For such events, we assign the start year as one year prior. (We obtained minimum wage indexation years from Brummund and Strain (2020).)

- If the state did have indexation, we check if the small increase was legislated. (Some small increases were legislated , rather than due to indexation, even in states that had indexation at the time). If it was a small legislated increase, we designate the event start date as the year prior to the provisional start year.
- We then repeat this process until none of our start years have a legislated, non-federal-year minimum wage increase in the preceding year.

5. The steps above yield 49 events and a maximum post-period of 6 years. In four states, however, an event fell within the post-period of an earlier event in the state: Delaware 2014 and 2019; New Jersey 2014 and 2019; Oregon 1998 and 2003; and Rhode Island 1999 and 2004. In these cases, we removed the later event and subsumed it with the earlier one, provided that the later event's post-period minimum wage increase was smaller than for the earlier event. In practice, this condition held in all four cases. As a result, our final list includes 45 events. We display these in Figure A1

D.2 Post-periods and clean controls

To implement Equation 4, we need to take the difference in outcomes averaged over (up to) 6 years following the event (event time 0 through 5) and the year prior to the event (event time -1). However, we do not have 6 "clean" post-periods for every event. Events in some years either have a federal minimum wage increase within 6 years, or the 6 year period runs into the end of the sample. We therefor implement a *maximum* 6 year post-period. For example, events occurring in 2013 have a full 6 year post-period (2013-2018). In contrast, events in 1994 have only a 2 year post-period (because of the federal increase in 1996), and events in 2017 have a 3 year post-period, since our sample ends in 2019. Table A4 summarizes the distribution of post-periods by year of the initial event increase.

Post-periods are also important in constructing clean controls. A state qualifies as a clean control for an event if it has *no* state minimum wage increases (large or small) in the three years preceding the event's start year, and *no* state minimum wage increases in the full post-period of that event. Using our example above, a clean control for an event in 2013 needs to have no state minimum wage increases from 2010 to 2018, while a clean control for an event in 1994 needs to have no increases between 1991 and 1995.