Modern Difference-in-Differences, Same Old Answer: What Event-Study Estimates Really Tell Us About the Effects of Minimum Wages on Jobs

David Neumark
University of California, Irvine, Hoover Institution, NBER, IZA@LISER, and CESifo

Antonio Rodriguez-Lopez
University of California, Irvine and CESifo

June 2026

Abstract

Two prominent publications in the recent minimum wage literature argue that estimation of the employment effects of minimum wages should use "clean" event-study designs, and that doing so leads to the conclusion that minimum wages have very limited, if any, effects on employment. We explore the use of event-study designs in this context, using the event-study stacked design of Cengiz, Dube, Lindner and Zipperer (2019a) and the related local projections difference-in-differences approach in Dube and Lindner (2024), along with their same data sources and period coverage. We generally find negative and significant employment effects of minimum wages in the United States, both overall and—more strongly—in the restaurant industry. The null results in these two papers are fragile and depend critically on a number of choices regarding variables, events, sample definitions, and weighting; they are not attributable to using an event-study design.

JEL Classification: J23, J38

Keywords: minimum wage, employment, event study

1 Introduction

The effects of minimum wages on employment—mainly of low-skilled workers—have been studied for nearly 100 years. As an instructive early example, Douty (1941) studied minimum wage regulation in the seamless hosiery industry, examining the effects of the Fair Labor Standards Act of 1938. Given such a long history of empirical analysis, coupled with persistent policy and research interest in the employment effects of minimum wages, it is not surprising that developments in the methods economists use to estimate policy impacts generally are closely mirrored in developments in the research literature on minimum wages.

Research until the early 1990s was dominated by time-series methods, as reflected in the predominance of such studies in the Brown et al. (1982) survey in the *Journal of Economic Literature*. Next, in the early 1990s, the emergence of meaningful variation in state minimum wages meshed with the diffusion of panel data econometric methods, in the "new minimum wage research." This new research, whether using panel data estimators across states and time (e.g., Card, 1992a; Neumark and Wascher, 1992), or isolated state minimum wage increases (e.g., Card, 1992b; Card and Krueger, 1994), had the important advantage of relying on explicit control observations to construct counterfactuals for the absence of minimum wage increases, in contrast to the time-series method's reliance on extrapolation. In the 2010s there was pushback on the blind application of what might be called the "experimental paradigm" to these methods, given that minimum wage variation is not random. Most notably, Allegretto et al. (2011) and Dube et al. (2010) argued that minimum wage variation could be associated with underlying trends or changes in employment, which could happen either coincidentally or because of policy endogeneity. They argued that cross-border designs solved this problem by holding these trends or changes fixed along opposite sides of the border.

Finally, and the jumping off point for this paper, Cengiz et al. (2019a) advocate for the use of event-study designs to estimate minimum wage effects on employment (and developed the stacked difference-in-differences approach). In their recent Handbook of Labor Economics chapter, Dube and Lindner (2024) argue that "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" (p. 2). Indeed, this approach has come to be widely adopted in the minimum wage literature (and in the evaluation of other policies).

There are two key claimed virtues of the stacked event-study design. The first is that it avoids comparisons between treated and control areas that may be many years apart—comparisons that could make it more likely that changes in these areas were driven by different trends (Dube and Lindner, 2024, p. 276). Second, it avoids the potential problem—as emphasized in the new difference-in-differences literature spurred by Goodman-Bacon (2021)—that more standard panel data estimators can have contaminated treatment or control areas, perhaps most notably when effects are dynamic but the dynamics are not modeled (as also emphasized by Meer and West, 2016). In contrast, the stacked event-study offers "clearly defined events and admissible controls, and with comparisons restricted to within the event window" (Dube et al., forthcoming, p. 13).

These considerations have led to the following prescription: "[A]n event-study based difference-in-differences design ... is a much more promising approach and should be the default among researchers using a DiD approach. Estimates from a TWFE panel regression with log minimum wage have proved opaque and fragile, and we think they are unlikely to yield convincing evidence going forward" (Dube and Lindner, 2024, p. 288).

To what conclusion does this prescription lead? According to Cengiz et al. (2019a, Appendix D), based on a binned wage approach, the resulting employment elasticity (standard error) is 0.001 (0.022), and the own-wage elasticity is 0.018 (0.056). According to Dube and Lindner (2024, p. 283), for the often-studied restaurant industry the estimated employment effect is indistinguishable from zero, and the own-wage elasticity for the largest number of events is -0.03 (0.15). In contrast, at least when using a longer period and two-way fixed effects, the estimated employment effect for this industry is much larger and negative, with an own-wage elasticity of -0.52 (0.35).

Overall, then, the clear impression given in the research literature is that the use of stacked event-study designs undermines or overturns the conclusion that minimum wages reduce employment. As argued by Dube and Lindner (2024): "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" (p. 18). They also argue that this conclusion holds in the traditional approach of comparing states that did or did not raise minimum wages, or "limiting comparisons to within multi-state commuting zones or border county pairs" (Dube et al., forthcoming, p. 18). Most succinctly, Dube and Lindner (2024) conclude that: "modern difference-in-differences event-study approaches ... yield unbiased estimates and correctly show a null effect" (p. 288).

Our goal in this paper is not to litigate the relative merits of the stacked event-study approach vs. two-way fixed effects. Rather, it is to provide a more thorough assessment of what the stacked event-study approach tells us about the employment effects of minimum wages. We focus, in particular, on what we see as the key contributions that have both advanced this approach and supported the conclusion that it leads to evidence of little or no job loss from higher minimum wages—the evidence in Cengiz et al. (2019a, Appendix D), and in the recent Handbook chapter by Dube and Lindner (2024). In particular, we focus on a number of potential problems with this evidence, including the choice of a flawed dependent variable, a sample definition that ignores the full distribution of wages, and the definition of events. In addition, these core studies reflect other choices that have the same consequence, such as restricting the estimates and specifications reported to those that tend not to find adverse effects of minimum wages.

Our overall conclusion is that the robust and appropriate evidence from estimating employment effects of minimum wages using a stacked event-study research design—and the related local projections difference-in-differences approach—is that higher minimum wages reduce employment, especially in the restaurant sector that has been the focus of so much of the recent research. The opposite conclusion—that there is little or no detectable job loss—is a fragile result that corresponds to a narrow set of choices regarding variables, events, sample, and weighting. We do not take issue with the application of modern difference-in-differences techniques to the question of the employment effects of minimum wages. Nonetheless, one lesson from our analysis is that in the studies we consider, a more thorough analysis tends to lead to the same old answer from these modern methods—that higher minimum wages reduce employment.

2 The Stacked Event-Study Design

Using 138 minimum-wage events in the U.S. during the 1979–2016 period, Cengiz, Dube, Lindner and Zipperer (2019a)—CDLZ hereafter—find no evidence of employment losses: jobs are simply reallocated across wage bins, with the reduction in jobs paying below the new minimum wage almost exactly offset by an increase in jobs paying just above the new minimum wage. In the stacked design, CDLZ uses four \$1 bins below the minimum wage and five \$1 bins above. That is, their analysis focuses on the segment of the wage distribution in the $[-4, 5)$ dollar range around the new minimum wage.[1] Our analysis examines two important—and consequential—specification

[1] Even though the stacked DiD design was introduced in CDLZ (Appendix D) only as a robustness check immune to negative-weighting bias (Goodman-Bacon, 2021; Sun and Abraham, 2021), the reference to this design accounts for a large share of its citations. For example, the word "stacked" appears in 1,740 out of 3,390 citations (51%) to CDLZ in Google Scholar (by May 2026). CDLZ's stacked design confirms their baseline results from a TWFE specification with a binary treatment. Here we focus exclusively on the stacked design, following some of the CDLZ authors who consider TWFE results to be fragile due to negative weights and more prone to pre-trend issues and who strongly advocate for using event-study designs in minimum wage research (see, for example, Dube and Lindner, 2024 and Dube et al., forthcoming).

choices or assumptions: (1) the use of employment-to-population ratios with *varying population* as the dependent variable, and (2) the assumption that minimum wages do not affect employment outside this range (i.e., wages more than \$4 below or more than \$5 above the new minimum wage).

On the first point, we show that their approach is biased toward finding no disemployment effects because minimum wages are negatively related to population: if both employment and population decline after a minimum wage event, their ratio can decrease, increase, or stay the same. CDLZ interpret any such increase—which arises when population falls proportionally more than employment—as a positive (net job creation) effect of the minimum wage, when in fact the number of jobs declined. CDLZ offer no rationale for dividing by population, and doing so is inconsistent with the goal of their paper: to "estimate the effect of minimum wage changes on low-wage jobs..." (Cengiz et al., 2019a, p. 1405). One might argue that dividing by population is useful to apply a consistent scale to employment across large and small labor markets. But then all that is needed is to divide by some base-level population measure that does not change over time, which would avoid introducing bias that can generate the wrong-signed effect. Below, we show that using instead employment-to-population ratios with *constant population*, or log employment, yields significant disemployment effects, while also documenting a negative and significant relationship between minimum wages and population.

On the second point, we show that minimum wages have negative employment effects outside the $[-4, 5)$ dollar range considered by CDLZ. Hence, CDLZ's assumption that minimum wages cannot affect employment outside this range—which underlies their decision to restrict the analysis to the $[-4, 5)$ dollar range—is incorrect. A simple refutation of this assumption comes from the literature documenting that minimum wages increase firm exit (Luca and Luca, 2019; Aaronson, French, Sorkin and To, 2018; Chava, Oettl and Singh, 2019; and Jha and Rodriguez-Lopez, 2021).

2.1 Econometric Specification

The approach is called a "stacked" design because each event and its controls form a panel of 32 quarters—12 quarters for the pre-treatment period and 20 quarters for the post-treatment period—and these 138 panels are then stacked into a single dataset that delivers pooled average treatment effect estimates at yearly horizons. Events vary in the number of control states, with the earliest events having up to 49 control states, whereas the latest events have about 20 control states—the average number of control states per event is 27.6 across the 138 events. CDLZ define control states as those without state minimum wage changes throughout both the 20-quarter post-period window and the 31 quarters before the treatment quarter; federal minimum wage changes do not disqualify a state as a control.

For any bin k, where a bin can be defined as any segment of the wage distribution, the stacked event-study specification is given by

$$y_{hst}^{k} = \sum_{\tau=-3}^{4} \alpha_{\tau}^{k} I_{hst}^{\tau} + \mu_{hs}^{k} + \rho_{ht}^{k} + \Omega_{hst}^{k} + u_{hst}^{k}, \tag{1}$$

where $h \in \{1, \ldots, 138\}$ indexes the event panel, s denotes state, which can be either the treated state s_h^* or one of event h's controls, and t indexes the calendar quarter, with $t \in [q_h - 12, q_h + 19]$ where q_h is event h's treatment quarter. Variable y_{hst}^{k} is the outcome—an employment-to-population ratio or log employment in bin k. The treatment indicator I_{hst}^{τ} is defined as

$$I_{hst}^{\tau} = \mathbf{1}\{s = s_h^* \text{ and } t \in \{q_h + 4\tau, \ldots, q_h + 4\tau + 3\}\}. \tag{2}$$

Here $\tau \in \{-3, -2, \ldots, 4\}$ indexes yearly horizons relative to event h's treatment quarter, so that I_{hst}^{τ} is 1 for the four calendar quarters of treated state s_h^* that fall in year τ of event h's window. The remaining terms μ_{hs}^{k} and ρ_{ht}^{k} are event-state and event-quarter fixed effects, respectively; Ω_{hst}^{k} includes fixed effects absorbing contamination from federal minimum wage changes, overlapping state-level minimum wage changes, and other minimum wage events; and u_{hst}^{k} is the error term. Taking $\tau = -1$ as the omitted reference period ($\alpha_{-1}^{k} = 0$), the seven coefficients of interest are α_{τ}^{k} for $\tau \in \{-3, -2, 0, 1, 2, 3, 4\}$, each measuring the four-quarter average treatment effect in bin k at yearly horizon τ. The average post-treatment effect is the average of the five post-treatment coefficients

$$\alpha_{\text{post}}^{k} = \frac{1}{5} \sum_{\tau=0}^{4} \alpha_{\tau}^{k}.$$

CDLZ define bins as \$1 segments of the wage distribution in the $[-4, 5)$ dollar range around the new minimum wage, but they present aggregate results for estimated percentage changes in total employment for two broader ranges: the $[-4, 0)$ range, which they refer to as "missing" jobs below the new minimum wage, and the $[0, 5)$ range, which they refer to as "excess" jobs above the new minimum wage. When only interested in these two aggregate effects, and not in the individual nine \$1 ranges, it is equivalent to simply estimate (1) for two aggregate bins: the missing bin and the excess bin.[2] CDLZ assume that minimum wage effects are zero outside the $[-4, 5)$ dollar range.

We investigate the robustness of that assumption by looking at employment responses across the entire wage distribution. Thus, with respect to the new minimum wage (MW), we define five more bins in addition to CDLZ's missing and excess bins: "lower" $[-\text{MW}, -4)$, "upper" $[5, \infty)$, "below"

[2]This is a consequence of using the employment-to-population ratio as the outcome variable: the four \$1-range ratios in $[-4, 0)$ sum to the missing ratio, and the five in $[0, 5)$ sum to the excess ratio.

$[-\text{MW}, 0)$, "above" $[0, \infty)$, and "total" $[-\text{MW}, \infty)$. Using the initial letter to identify each bin, we have $k \in \{L, M, E, U, B, A, T\}$. It then follows that, when using the employment-to-population ratio as the outcome variable, for every τ,

$$\alpha_\tau^T = \underbrace{\alpha_\tau^L + \alpha_\tau^M}_{\alpha_\tau^B} + \underbrace{\alpha_\tau^E + \alpha_\tau^U}_{\alpha_\tau^A}, \tag{3}$$

with a similar property for the average post-treatment effect, α_{post}^k.

Following CDLZ, our tables and figures below report estimates not in the raw α-units of (1), but expressed as the implied percentage change in *total* employment attributable to bin k. When the outcome is the bin-k employment-to-population ratio, α_τ^k is the change in that ratio induced by the new minimum wage; dividing by the baseline total employment-to-population ratio, $\overline{\text{EPOP}}_{-1}$, and multiplying by 100 rescales this change so it can be read as bin k's contribution to the percentage change in total employment. That is, for any $k \in \{L, M, E, U, B, A, T\}$,

$$\%\Delta\text{Emp}_\tau^k = \frac{\alpha_\tau^k}{\overline{\text{EPOP}}_{-1}} \times 100, \qquad \%\Delta\text{Emp}_{\text{post}}^k = \frac{\alpha_{\text{post}}^k}{\overline{\text{EPOP}}_{-1}} \times 100, \tag{4}$$

where $\overline{\text{EPOP}}_{-1}$ is the average total employment-to-population ratio across treated states over the four pre-event quarters $t \in \{q_h - 4, \ldots, q_h - 1\}$. Because $\overline{\text{EPOP}}_{-1}$ is a common scalar, the additivity in (3) carries over directly to these percentage changes.

Strictly speaking, the equations in (4) express the change in the bin-k employment-to-population ratio as a percentage of the baseline total employment-to-population ratio, and only approximate bin k's contribution to the percentage change in total employment when population changes are negligible—an implicit assumption that, as we show below, does not hold in CDLZ's data and drives their null results.

In contrast to CDLZ, we also estimate (1) using log employment in bin k as the dependent variable, $\ln \text{Emp}_{hst}^k$. In that case, α_τ^k approximates the proportional change in bin-k employment and the equivalence in (3) no longer holds:

$$\ln(\text{Emp}_{hst}^T) \neq \ln(\text{Emp}_{hst}^L) + \ln(\text{Emp}_{hst}^M) + \ln(\text{Emp}_{hst}^E) + \ln(\text{Emp}_{hst}^U),$$

so the bin-k coefficients cannot simply be summed to recover the total employment effect. Thus, to be able to compare the log employment results to the employment-to-population results, we convert the former's estimates to an approximate percentage-of-total-employment scale. That is, for each bin $k \in \{L, M, E, U, B, A, T\}$, the approximate log employment counterpart of (4) is

$$\%\Delta\text{Emp}_\tau^k = \alpha_\tau^k \left(\frac{\overline{\text{EPOP}}_{-1}^k}{\overline{\text{EPOP}}_{-1}} \right) \times 100, \qquad \%\Delta\text{Emp}_{\text{post}}^k = \alpha_{\text{post}}^k \left(\frac{\overline{\text{EPOP}}_{-1}^k}{\overline{\text{EPOP}}_{-1}} \right) \times 100. \tag{5}$$

Note that for $k = T$, $\overline{\text{EPOP}}_{-1}^{T} = \overline{\text{EPOP}}_{-1}$, so (5) collapses to $\%\Delta\text{Emp}_{\tau}^{T} = \alpha_{\tau}^{T} \times 100$ and $\%\Delta\text{Emp}_{\text{post}}^{T} = \alpha_{\text{post}}^{T} \times 100$.

2.2 Data and Results

CDLZ construct quarterly state-level employment from 1979 to 2016 using NBER's CPS Merged Outgoing Rotation Group (CPS-MORG) and Quarterly Census of Employment and Wages (QCEW) data. They use the CPS-MORG to obtain employment shares of each wage bin, and convert these shares to employment levels by multiplying by QCEW total employment for the state-quarter (they refer to this second step as "benchmarking" to QCEW). We obtain employment levels for each state in each event panel using the code and data in CDLZ's replication package (Cengiz et al., 2019b), but we update the QCEW total employment series.[3] The differences between our QCEW series and CDLZ's are minimal, and only involve QCEW revisions to 2016 data (CDLZ use a 2016 vintage that was later subject to BLS revisions) and a correction in 1979-1980 data for Alaska and the District of Columbia.[4] State-level minimum wage and population data also come from CDLZ's replication package, originally drawn from Vaghul and Zipperer (2016) and the CPS-MORG, respectively.

As mentioned above, in addition to expanding CDLZ's analysis to include the full wage distribution through our seven bins, we examine the sensitivity of CDLZ's findings to their choice of dependent variable. For a bin k, CDLZ use as outcome variable the employment-to-population ratio with a time-varying denominator, $\text{EPOP}_{hst}^{k} = \text{Emp}_{hst}^{k}/\text{Pop}_{st}$, where Pop_{st} is the state-quarter population. Their estimates may therefore be biased by endogenous population responses to minimum wage changes; we discuss the relationship between minimum wages and population in CDLZ's data in Section 3. Here, we first examine how the bin estimates change if instead we use as outcome variables *(1)* an employment-to-population ratio with a constant state-level population in the denominator, $\widetilde{\text{EPOP}}_{hst}^{k} = \text{Emp}_{hst}^{k}/\overline{\text{Pop}}_{s}$, where $\overline{\text{Pop}}_{s}$ is the simple average of Pop_{st} across the full 1979–2016 sample period for state s (fixed within each state and therefore unaffected by any individual minimum wage event), and *(2)* log employment, $\ln \text{Emp}_{hst}^{k}$.

In the full sample with 138 events, which we refer to as the "All-Events sample," the post-treatment period is truncated by the last sample period (2016, fourth quarter) for 40 events oc-

[3]With the exception of the total bin (T), a state that serves as control for two different events will, in general, have different employment counts for the same bin in the same calendar quarter, because bins are defined relative to the new minimum wage of each event's treated state.

[4]BLS did not report QCEW total (all-ownership) employment for these state-quarters, so CDLZ filled the gap with CPS-based imputations, whereas we instead impute the missing total from QCEW private-sector employment (which BLS did report for those state-quarters), scaled by the state-specific total-to-private employment ratio observed in 1981, the first year both series are reported for Alaska and DC.

curring after the first quarter of 2012.[5] Similar to CDLZ, for robustness we also present results for a subsample, the "Pre-2012 sample," which only includes events with full post-treatment windows (98 events).[6] Moreover, while CDLZ only present weighted estimates for their stacked design (they use $\overline{\text{Pop}}_s$ as weights), we also present unweighted estimates, since otherwise the estimates may be driven by a few large states; the weighted-unweighted distinction, for example, is crucial for the results of Dube and Lindner (2024) discussed in Section 4. Hence, in total we estimate 12 variations of specification (1) for each of our seven bins: three dependent variables, two samples, and two weighting schemes ($3 \times 2 \times 2$).

Table 1 reports estimates of the average post-treatment percentage change in total employment attributable to each bin k, $\widehat{\%\Delta\text{Emp}}^{k}_{\text{post}}$. We begin with the shaded cells, which replicate CDLZ's findings in their Table D.1 (columns 2 and 4, top two rows). In both samples, and using weighted specifications, these cells show CDLZ's main finding of pure wage-bin reallocation effects of minimum wages: the decline in jobs in the missing bin—about 1.7% of total employment in the all-events sample and about 1.5% in the pre-2012 sample—is almost exactly offset by the increase in jobs in the excess bin. The unweighted specifications show the same pattern, and although the lower bin shows significant losses across all varying-population specifications, the total effect remains small and insignificant whenever we use CDLZ's employment-to-varying-population dependent variable.

The story changes along several dimensions once we remove the bias caused by a varying-population denominator. With constant population in the denominator, CDLZ's story no longer holds even if we only look at the missing and excess bins. Compared to the varying-population results, in both samples and weighting schemes the decline in the missing bin is larger while the increase in the excess bin is smaller; the offset is no longer almost exact, and a net loss of about 0.2–0.4 percentage points remains within the $[-4, 5)$ dollar range alone.

Expanding to the full wage distribution amplifies the gap: the total employment effect is negative and significant in both samples (see Total row in both panels), with weighted net losses of about -0.6% in the all-events sample and -0.7% in the pre-2012 sample, and unweighted net losses of about -1.3% and -1.6%, respectively. The sharper net losses in the unweighted specifications are the result of significant losses in the upper bin that wipe out the excess-bin gains, rendering the above-bin effect insignificant. Using log employment as the outcome variable confirms these patterns across all bins (recall that in this case the sum of the bin-level estimates no longer matches

[5] As we explain later, while these events are supposed to be "clean," in fact they define a state minimum wage increase as an event even if it follows a recent minimum wage increase in the same state.

[6] The pre-treatment period is also truncated by the first sample period—the first quarter of 1979—for event panels with treatment dates before 1982, but this does not affect the estimation of post-treatment effects.

Table 1: Estimated % Change in Total Employment from Minimum Wage Increases

	Emp/Pop Varying		**Emp/Pop Constant**		**Log Employment**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Panel A: All Events						
Lower	-0.102***	-0.089**	-0.135***	-0.117***	-0.196***	-0.166
	(0.034)	(0.038)	(0.034)	(0.037)	(0.059)	(0.110)
Missing	-1.685***	-1.387***	-1.845***	-1.544***	-2.994***	-2.739***
	(0.246)	(0.168)	(0.256)	(0.154)	(0.379)	(0.218)
Excess	1.661***	1.311***	1.607***	1.185***	1.472***	1.045***
	(0.276)	(0.243)	(0.302)	(0.272)	(0.262)	(0.258)
Upper	-0.008	0.278	-0.239	-0.840**	-0.395	-1.212***
	(0.386)	(0.380)	(0.316)	(0.414)	(0.323)	(0.378)
Below	-1.787***	-1.476***	-1.980***	-1.661***	-3.019***	-2.752***
	(0.240)	(0.181)	(0.251)	(0.164)	(0.365)	(0.213)
Above	1.653***	1.589***	1.368***	0.345	1.263***	0.097
	(0.511)	(0.374)	(0.437)	(0.439)	(0.443)	(0.451)
Total	-0.134	0.113	-0.612**	-1.315***	-0.518*	-1.376***
	(0.387)	(0.313)	(0.275)	(0.376)	(0.274)	(0.350)
Panel B: Pre-2012 Events						
Lower	-0.102**	-0.074*	-0.135***	-0.104**	-0.174**	-0.058
	(0.040)	(0.043)	(0.041)	(0.042)	(0.071)	(0.131)
Missing	-1.491***	-1.207***	-1.640***	-1.379***	-2.508***	-2.398***
	(0.265)	(0.181)	(0.270)	(0.166)	(0.374)	(0.223)
Excess	1.458***	1.212***	1.427***	1.097***	1.323***	0.925***
	(0.283)	(0.267)	(0.320)	(0.306)	(0.283)	(0.280)
Upper	-0.232	0.098	-0.365	-1.187**	-0.492	-1.423***
	(0.418)	(0.418)	(0.361)	(0.464)	(0.362)	(0.412)
Below	-1.594***	-1.281***	-1.776***	-1.483***	-2.575***	-2.452***
	(0.259)	(0.196)	(0.264)	(0.180)	(0.359)	(0.220)
Above	1.225**	1.311***	1.062**	-0.091	0.951**	-0.260
	(0.520)	(0.409)	(0.469)	(0.486)	(0.472)	(0.491)
Total	-0.368	0.030	-0.714**	-1.574***	-0.642**	-1.562***
	(0.398)	(0.350)	(0.307)	(0.421)	(0.300)	(0.380)

Notes: Each cell reports the estimated average post-treatment percentage change in total employment (from minimum wage increases) attributable to the indicated bin; see definition of $\%\Delta\text{Emp}^{k}_{\text{post}}$ in equation (4) for the four employment-to-population columns and in equation (5) for the two log-employment columns. The top header indicates the dependent variable: employment to varying population, employment to constant population, and log employment. The second header indicates whether the estimation of equation (1) is weighted (by constant state population) or unweighted. Panel A uses the All Events sample (138 events, 109,326 observations), and Panel B uses the Pre-2012 sample (98 events with full post-period windows, 93,176 observations). All regressions include event-state and event-quarter fixed effects, and fixed effects accounting for federal minimum wage changes, overlapping minimum wage changes, and other minimum wage events. Standard errors (in parentheses) are clustered at the event-state level. Shaded cells highlight CDLZ's missing/excess results. $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$.

the total effect), with weighted net losses of about −0.5% in the all-events sample and −0.6% in the pre-2012 sample, and unweighted net losses of about −1.4% and −1.6%, respectively.

Table 1 shows average post-period estimates, but specification (1) produces estimates for each yearly horizon in the pre-period and post-period, with $\tau = -1$ as the reference period. Thus, to look at the dynamics of the employment effects of minimum wages, Figure 1 presents the evolution of $\%\widehat{\Delta \text{Emp}}_{\tau}^{k}$—as defined in (4) and (5)—for the below, above, and total bins, using the all-events sample across dependent variables and weighting schemes. The figure is arranged as a three-by-two grid: rows correspond to the three dependent variables (employment-to-varying-population in the top row, employment-to-constant-population in the middle, and log employment in the bottom), and columns to the weighting scheme (weighted on the left, unweighted on the right). Within each subplot, the below-bin path is shown in red diamonds, the above-bin path in blue circles, and the total-bin path in black triangles, with 95% confidence intervals around each estimate.

In the varying-population specifications, panels (a) and (b) of Figure 1 show CDLZ's pure wage-bin reallocation story, with employment gains above the new minimum wage almost perfectly offsetting the employment losses below the new minimum wage, leaving total employment flat over time. The other four panels, however, tell a different story: total employment losses build up over time, reaching about −1% by $\tau = 4$ in the weighted specifications of panels (c) and (e), and about −2% in the unweighted specifications of panels (d) and (f). The declining trend in total employment is mainly driven by dynamics of the above bin, for which initial gains consistently fall over time, even turning into (insignificant) losses in the unweighted specifications; below-bin dynamics are more stable, remaining relatively flat after the initial decline. Therefore, even though in the year of the event most of the jobs below the new minimum wage reappear above the new minimum wage, over time firms adjust employment down, a pattern consistent with slow-moving adjustment margins such as automation (Lordan and Neumark, 2018) and firm exit (Aaronson et al., 2018; Luca and Luca, 2019; Jha and Rodriguez-Lopez, 2021). Appendix Figure A1 shows similar paths in the pre-2012 sample.

2.3 Overlapping Events and Bundling

A potential concern from Figure 1, panels (c)-(f), is that the total-bin path shows a pre-trend: although the estimate is stable from $\tau = -3$ to $\tau = -2$, $\%\widehat{\Delta \text{Emp}}_{\tau}^{k}$ declines from $\tau = -2$ to $\tau = -1$ in all cases, being significant at the 5% level for the unweighted specifications. Although this could be explained by anticipation effects due to minimum wage changes being announced months before their implementation (see Karabarbounis et al., 2022, for evidence on minimum wage anticipation

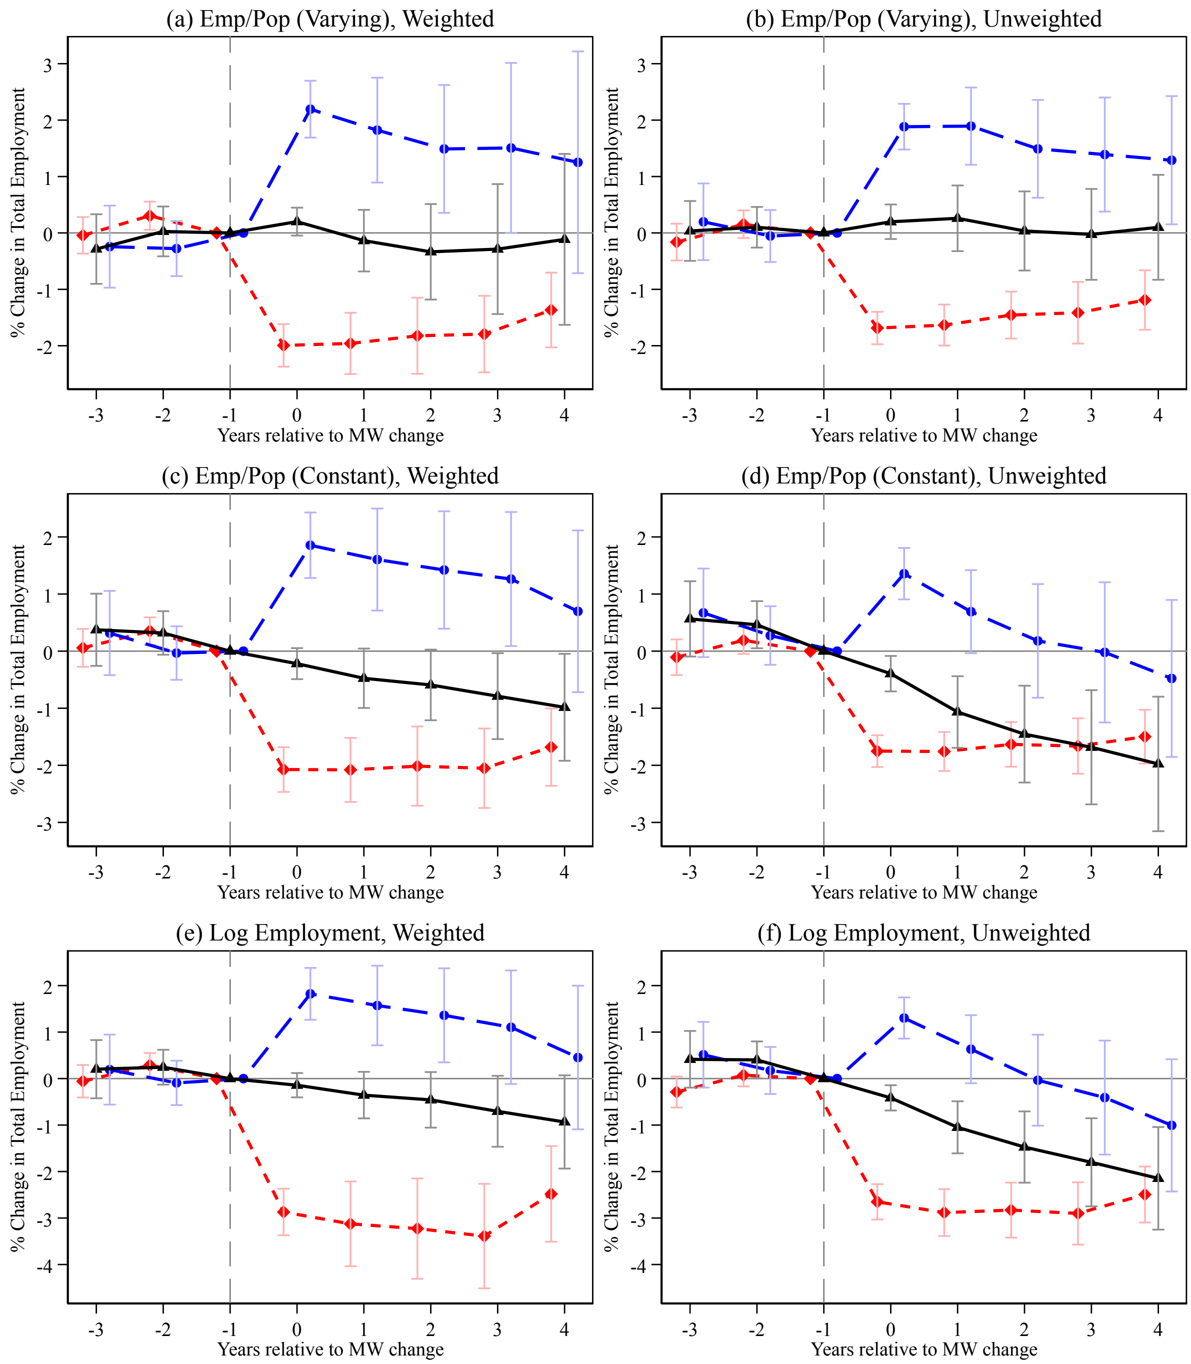

Figure 1: Event-Study Paths for % Change in Total Employment with 95% CIs (All Events): Below (dash, ♦), Above (long dash, •), and Total (solid, ▲)

effects), the explanation is much simpler: some states experience multiple minimum-wage increases within a short window, so the post-treatment horizon of an earlier event overlaps with the pre-treatment horizon of a later event in the same state. We show that this mechanical overlap—despite CDLZ's fixed effects for overlapping events—contaminates pre-trends and understates the magnitude of the long-term total employment losses.

In particular, we follow a bundling-of-events strategy similar to Dube and Lindner (2024) and retain only the first event per state and any subsequent event that occurs at least 20 quarters (5 years) after the previous retained event in that state. This yields 80 events (60 for the pre-2012 sample) down from 138 (98 for the pre-2012 sample); that is, 58 events (38 for the pre-2012 sample) overlap with prior events in the same state. Table 2 presents estimates of the bundled-events version of Table 1. The results are qualitatively similar—both the constant-population and log employment specifications continue to show significant total-employment losses—but with magnitudes that are larger throughout the table: the total-bin losses become more negative in every specification except the unweighted varying-population one, which remains near zero. Notably, the weighted varying-population estimates of the total-bin losses are now much larger in magnitude: approximately −0.7% in the all-bundled-events sample and −1% in the pre-2012 bundled-events sample. In the latter case, the gap between the varying- and constant-population specifications shrinks to essentially zero.

Figure 2 shows the bundled-events version of Figure 1. The evidence of pre-trends disappears. Moreover, by $\tau = 4$ the total employment losses reach about −1.5% in the weighted specifications in panels (c) and (e), and about −2.5% in the unweighted specifications in panels (d) and (f). Even in panel (a)—the varying-population, weighted specification preferred by CDLZ—the total path now drifts to about −1% by $\tau = 4$, though this estimate is insignificant. To sum up, CDLZ's ad hoc approach to control for contaminated (overlapping) events does not work. We note that Dube and Lindner (2024) also avoid this problem by bundling events as we do; however, they introduce a different problem by truncating post-treatment periods.[7]

2.4 Restaurant Employment

Whereas CDLZ focus on all low-wage jobs and we expand the analysis to the entire wage distribution, we can also apply the stacked event-study design to restaurant employment—the most studied sector in the minimum wage literature, and the one where the minimum wage binds most.[8] We

[7]Appendix Figure A2 shows similar results for the pre-2012 bundled sample. We report bundled-events robustness checks for the remainder of this section and for the analyses in Section 3.

[8]CDLZ also look at restaurants in some of their specifications, though not for their stacked design.

Table 2: Estimated % Change in Total Employment from Minimum Wage Increases (Bundled Events)

	Emp/Pop Varying		**Emp/Pop Constant**		**Log Employment**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Panel A: All Bundled Events						
Lower	-0.124***	-0.037	-0.156***	-0.075	-0.211***	0.012
	(0.046)	(0.052)	(0.045)	(0.051)	(0.081)	(0.146)
Missing	-1.434***	-1.123***	-1.622***	-1.333***	-2.542***	-2.430***
	(0.223)	(0.205)	(0.263)	(0.185)	(0.395)	(0.266)
Excess	1.452***	1.263***	1.468***	1.068***	1.336***	0.933***
	(0.356)	(0.294)	(0.406)	(0.335)	(0.336)	(0.338)
Upper	-0.556	-0.004	-0.566	-1.183**	-0.787*	-1.335**
	(0.573)	(0.498)	(0.406)	(0.554)	(0.425)	(0.558)
Below	-1.558***	-1.160***	-1.778***	-1.409***	-2.576***	-2.388***
	(0.227)	(0.218)	(0.269)	(0.199)	(0.386)	(0.260)
Above	0.897	1.259**	0.902*	-0.115	0.704	-0.272
	(0.678)	(0.540)	(0.467)	(0.670)	(0.481)	(0.681)
Total	-0.661	0.098	-0.875**	-1.524**	-0.818**	-1.589***
	(0.579)	(0.478)	(0.351)	(0.598)	(0.346)	(0.539)
Panel B: Pre-2012 Bundled Events						
Lower	-0.141***	-0.026	-0.173***	-0.067	-0.230**	0.081
	(0.051)	(0.057)	(0.052)	(0.058)	(0.097)	(0.173)
Missing	-1.152***	-0.887***	-1.317***	-1.110***	-1.821***	-1.965***
	(0.201)	(0.219)	(0.223)	(0.198)	(0.266)	(0.264)
Excess	1.222***	1.086***	1.263***	0.876**	1.199***	0.738**
	(0.352)	(0.313)	(0.414)	(0.365)	(0.356)	(0.365)
Upper	-0.886	-0.148	-0.749	-1.459**	-0.976**	-1.563**
	(0.615)	(0.546)	(0.468)	(0.621)	(0.482)	(0.608)
Below	-1.293***	-0.913***	-1.490***	-1.177***	-1.938***	-1.989***
	(0.209)	(0.236)	(0.232)	(0.216)	(0.269)	(0.263)
Above	0.336	0.938	0.514	-0.583	0.294	-0.726
	(0.658)	(0.592)	(0.443)	(0.741)	(0.471)	(0.740)
Total	-0.957	0.025	-0.976**	-1.760***	-0.962**	-1.805***
	(0.599)	(0.536)	(0.389)	(0.668)	(0.378)	(0.586)

Notes: Each cell reports the estimated average post-treatment percentage change in total employment (from minimum wage increases) attributable to the indicated bin; see definition of $\%\Delta\text{Emp}^{k}_{\text{post}}$ in equation (4) for the four employment-to-population columns and in equation (5) for the two log-employment columns. The top header indicates the dependent variable: employment to varying population, employment to constant population, and log employment. The second header indicates whether the estimation of equation (1) is weighted (by constant state population) or unweighted. Panel A uses the All Bundled Events sample (80 events), and Panel B uses the Pre-2012 Bundled Events sample (60 events with full post-period windows). All regressions include event-state and event-quarter fixed effects, and fixed effects accounting for federal minimum wage changes, overlapping minimum wage changes, and other minimum wage events. Standard errors (in parentheses) are clustered at the event-state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

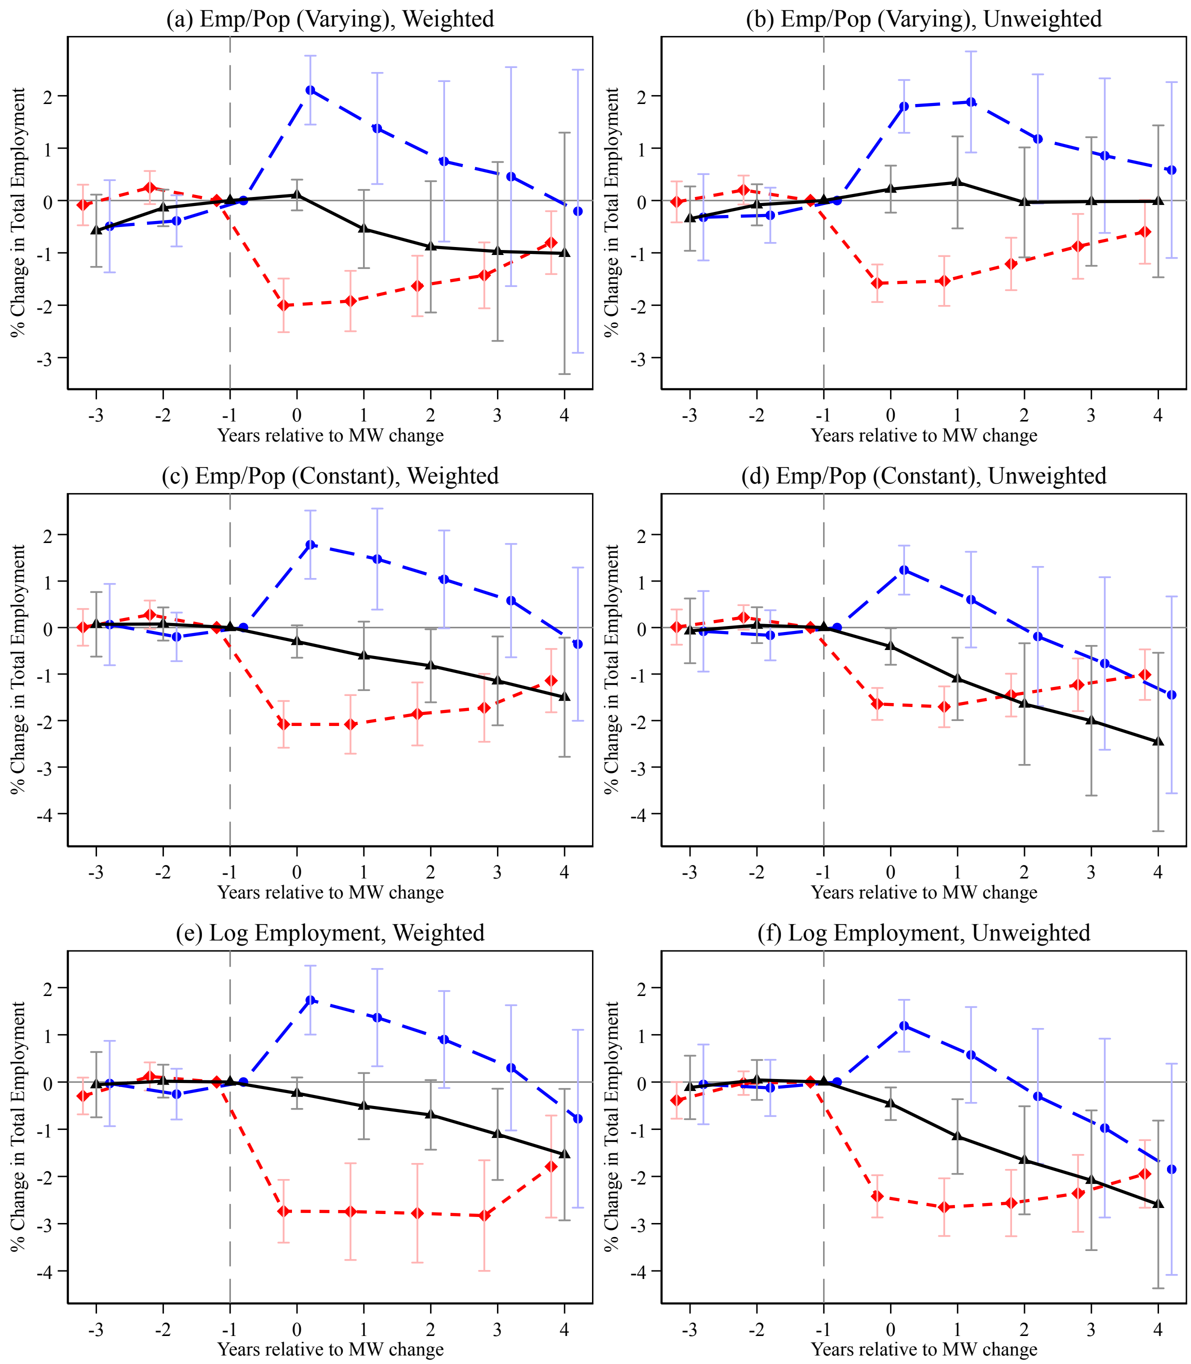

Figure 2: Event-Study Paths for % Change in Total Employment with 95% CIs (Bundled Events): Below (dash, ♦), Above (long dash, •), and Total (solid, ▲)

obtain quarterly restaurant employment counts (NAICS 722) from the state-level QCEW from 1979 to 2019, which lets us extend the sample to the full pre-Covid period. This allows us to construct a pre-2015 sample of events occurring through the first quarter of 2015, for which the full 20-quarter post-period window (running through the fourth quarter of 2019) lies within our data; 119 of the 138 events qualify. We similarly extend the QCEW total employment series through 2019, which lets us trace the event-study path for total employment under this pre-2015 sample.

Rather than scaling estimates of α_τ^R from (1) (where R denotes the restaurant sector) as percentage changes in total employment, we scale them as percentage changes in restaurant employment. For employment-to-population outcomes, this replaces $\overline{\text{EPOP}}_{-1}$ with $\overline{\text{EPOP}}_{-1}^R$ in (4), where $\overline{\text{EPOP}}_{-1}^R$ is the average restaurant employment-to-population ratio across treated states over the four pre-event quarters. For log employment outcomes, $\%\Delta\text{Emp}_\tau^R = \alpha_\tau^R \times 100$ and $\%\Delta\text{Emp}_{\text{post}}^R = \alpha_{\text{post}}^R \times 100$.

Table 3 shows estimates of $\%\Delta\text{Emp}_{\text{post}}^T$ (the total bin, as defined in Section 2.1) and $\%\Delta\text{Emp}_{\text{post}}^R$ for our three dependent variables, two weighting schemes, and three samples. Recall that the all-events sample is truncated at the fourth quarter of 2016, so events after the first quarter of 2012 have incomplete post-period windows; the pre-2012 and pre-2015 samples include 98 and 119 events, respectively, with complete post-period windows. Panel A shows the estimates for total employment, with the all-events and pre-2012 rows exactly matching the "Total" rows in Table 1. The pre-2015 row tells a similar story, so expanding the sample to 119 events with complete post-period windows does not affect our total-employment findings.

Panel B shows the restaurant results. Every estimate is negative and statistically significant—16 at the 1% level and 2 at the 5% level—even those from the varying-population specifications. Across the three samples, the story is the same: minimum wage hikes reduce restaurant employment. The declines are smallest in magnitude under the varying-population denominator, running from about 0.9% to 1.2% over the post-period across samples and weights. With the constant-population denominator they run about 1.4% to 1.5% (weighted) and 1.6% to 1.9% (unweighted); with log employment, about 1.1% to 1.3% (weighted) and 1.8% to 2.0% (unweighted). Hence, the effect on restaurant employment is strong enough that even the varying-population specifications—in which population itself responds to minimum wage changes, as we discuss in Section 3—yield negative and significant declines.

The event-study paths in Figure 3 expand the picture from Table 3 by showing $\widehat{\%\Delta\text{Emp}}_\tau^R$, the estimated percentage change in restaurant employment at horizon τ, across the entire event window ($\tau \in \{-3, -2, \dots, 4\}$) for our three dependent variables: the employment ratio with varying

Table 3: Estimated % Change in Total and Restaurant Employment from Minimum Wage Increases

	Emp/Pop Varying		**Emp/Pop Constant**		**Log Employment**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Panel A: Total Employment						
All Events	-0.134	0.113	-0.612**	-1.315***	-0.518*	-1.376***
	(0.387)	(0.313)	(0.275)	(0.376)	(0.274)	(0.350)
Pre-2012	-0.368	0.030	-0.714**	-1.574***	-0.642**	-1.562***
	(0.398)	(0.350)	(0.307)	(0.421)	(0.300)	(0.380)
Pre-2015	-0.203	0.092	-0.670**	-1.311***	-0.570**	-1.326***
	(0.379)	(0.307)	(0.278)	(0.374)	(0.273)	(0.342)
Panel B: Restaurant Employment						
All Events	-0.929**	-0.930***	-1.405***	-1.624***	-1.102***	-1.824***
	(0.434)	(0.292)	(0.386)	(0.384)	(0.355)	(0.368)
Pre-2012	-1.212***	-1.089***	-1.445***	-1.876***	-1.264***	-2.001***
	(0.468)	(0.330)	(0.427)	(0.434)	(0.381)	(0.394)
Pre-2015	-1.037**	-1.079***	-1.502***	-1.774***	-1.168***	-1.873***
	(0.433)	(0.293)	(0.399)	(0.395)	(0.354)	(0.362)

Notes: Each cell reports the estimated average post-treatment percentage change in total (Panel A) and restaurant (Panel B) employment from a minimum wage increase. The top header indicates the dependent variable: employment to varying population, employment to constant population, and log employment. The second header indicates whether the estimation of equation (1) is weighted (by constant state population) or unweighted. Rows in each panel indicate the sample used: All Events (138 events, 109,326 observations), Pre-2012 (98 events with full post-period windows, 93,176 observations), and Pre-2015 (119 events with full post-period windows, 106,520 observations). The Pre-2015 sample extends through 2019Q4. All regressions include event-state and event-quarter fixed effects, and fixed effects accounting for federal minimum wage changes, overlapping minimum wage changes, and other minimum wage events. Standard errors (in parentheses) are clustered at the event-state level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

population (black triangles), the employment ratio with constant population (red diamonds), and log employment (blue circles). The top row shows the all-events sample, the middle row the pre-2012 sample, and the bottom row the pre-2015 sample; the left column reports the weighted specifications, and the right column the unweighted ones. Across the three samples, restaurant employment declines steadily over time when using the employment ratio with constant population or log employment, reaching about −2% by $\tau = 4$ in the weighted column and approaching −3% in the unweighted column. When using the employment ratio with varying population, by contrast, restaurant employment stabilizes by $\tau = 2$ at about −1% in both columns. Overall, Figure 3 shows strong and persistently negative dynamic effects of minimum wage increases on restaurant employment.

As with total employment, the unweighted paths show a pre-trend from $\tau = -2$ to $\tau = -1$ (the paths are flat from $\tau = -3$ to $\tau = -2$). Again, this is a consequence of overlapping events: the apparent "pre-trend" reflects employment effects of earlier events. Figure 4 presents the paths

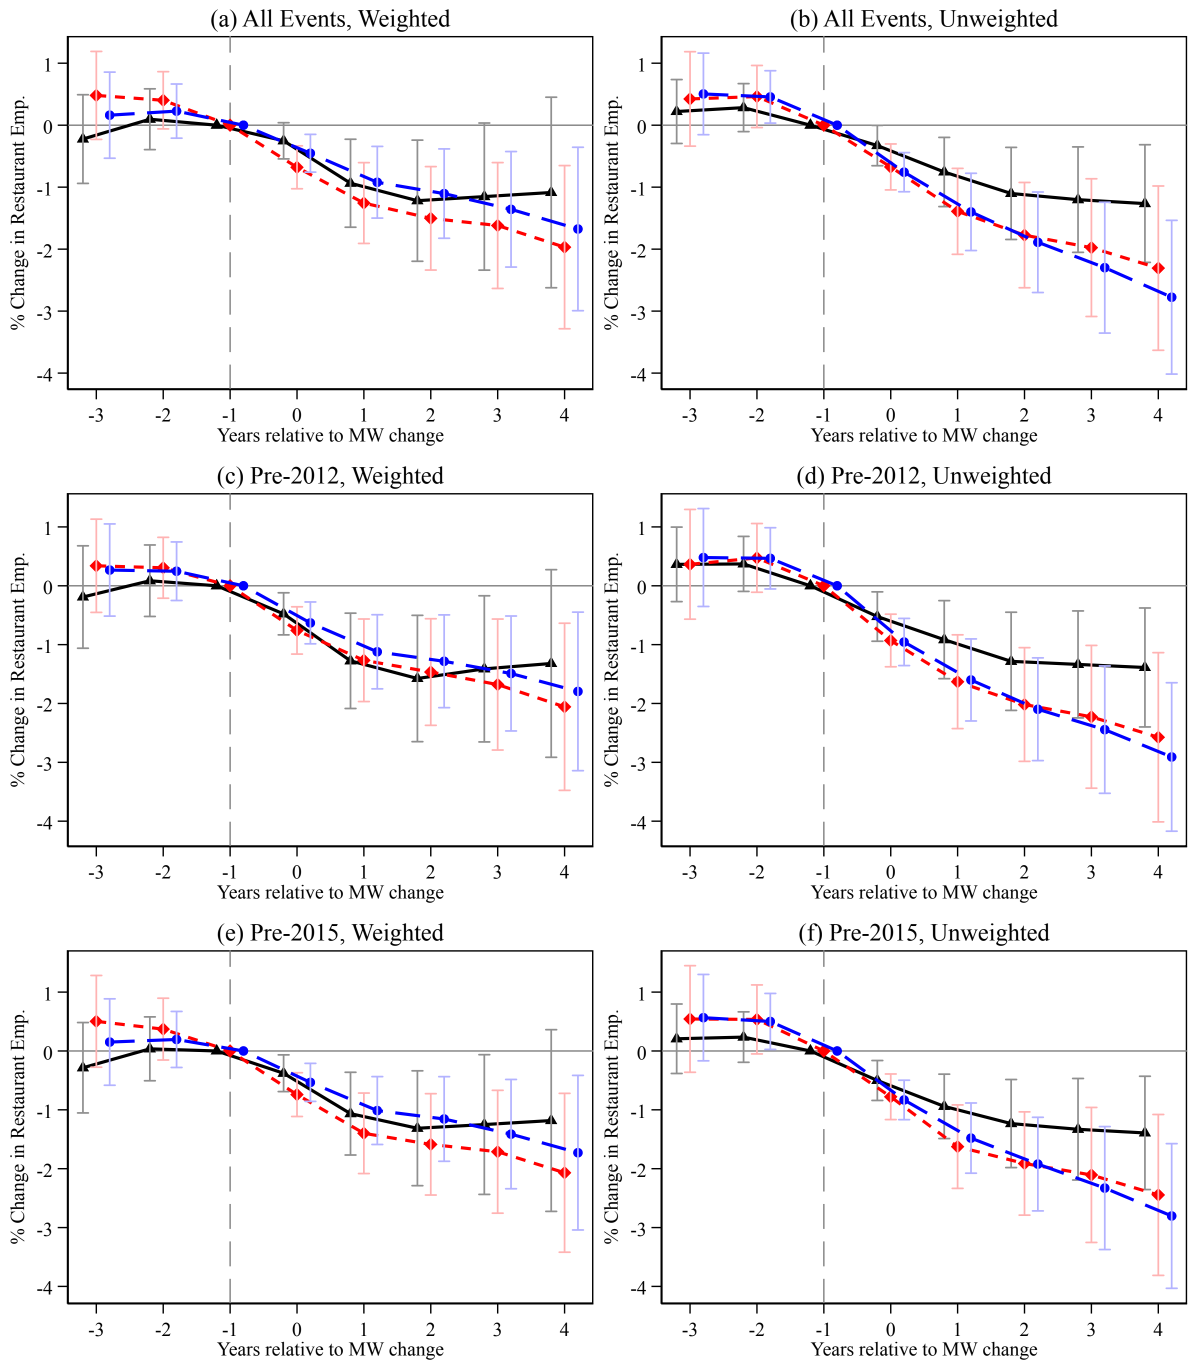

Figure 3: Event-Study Paths for % Change in Restaurant Employment with 95% CIs: Emp/Pop Varying (solid, ▲), Emp/Pop Constant (dash, ◆), Log Employment (long dash, •)

Table 4: Estimated % Change in Total and Restaurant Employment from Minimum Wage Increases (Bundled Events)

	Emp/Pop Varying		**Emp/Pop Constant**		**Log Employment**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Panel A: Total Employment						
All Events	-0.661	0.098	-0.875**	-1.524**	-0.818**	-1.589***
	(0.579)	(0.478)	(0.351)	(0.598)	(0.346)	(0.539)
Pre-2012	-0.957	0.025	-0.976**	-1.760***	-0.962**	-1.805***
	(0.599)	(0.536)	(0.389)	(0.668)	(0.378)	(0.586)
Pre-2015	-0.644	0.125	-0.865**	-1.368**	-0.812**	-1.430***
	(0.567)	(0.439)	(0.341)	(0.561)	(0.337)	(0.505)
Panel B: Restaurant Employment						
All Events	-1.541**	-1.321***	-1.847***	-2.184***	-1.517***	-2.293***
	(0.676)	(0.451)	(0.585)	(0.519)	(0.522)	(0.546)
Pre-2012	-1.909***	-1.447***	-1.899***	-2.401***	-1.739***	-2.483***
	(0.732)	(0.515)	(0.655)	(0.584)	(0.565)	(0.586)
Pre-2015	-1.534**	-1.362***	-1.847***	-2.170***	-1.518***	-2.230***
	(0.659)	(0.425)	(0.579)	(0.508)	(0.511)	(0.520)

Notes: Each cell reports the estimated average post-treatment percentage change in total (Panel A) and restaurant (Panel B) employment from a minimum wage increase. The top header indicates the dependent variable: employment to varying population, employment to constant population, and log employment. The second header indicates whether the estimation of equation (1) is weighted (by constant state population) or unweighted. Rows in each panel indicate the sample used: All Bundled Events (80 events, 65,184 observations), Pre-2012 Bundled Events (60 events with full post-period windows, 56,451 observations), and Pre-2015 Bundled Events (79 events with full post-period windows, 68,515 observations). The Pre-2015 sample extends through 2019Q4. All regressions include event-state and event-quarter fixed effects, and fixed effects accounting for federal minimum wage changes, overlapping minimum wage changes, and other minimum wage events. Standard errors (in parentheses) are clustered at the event-state level. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

under the bundled-events samples—80, 60, and 79 events in the all-events, pre-2012, and pre-2015 samples, respectively—and shows no meaningful evidence of pre-trends. Moreover, by $\tau = 4$, the bundled-events paths show larger declines than those in Figure 3, even reaching -2% under the varying-population specifications. As in Section 2.3 for total employment, overlapping events lead to understatement of the magnitude of the long-term restaurant employment losses. Table 4 confirms this pattern in the average post-treatment estimates: every restaurant coefficient in Panel B is more negative than its counterpart in Table 3, with the bundled estimates being on average 34% larger in magnitude across the 18 cells (ranging from 19% to 66% larger).

2.5 Elasticities

All previous results are presented in terms of percentage employment changes (either for total employment or for the restaurant industry) due to minimum wage events, but typically we are interested in the minimum wage elasticity of employment. The average post-period elasticity of

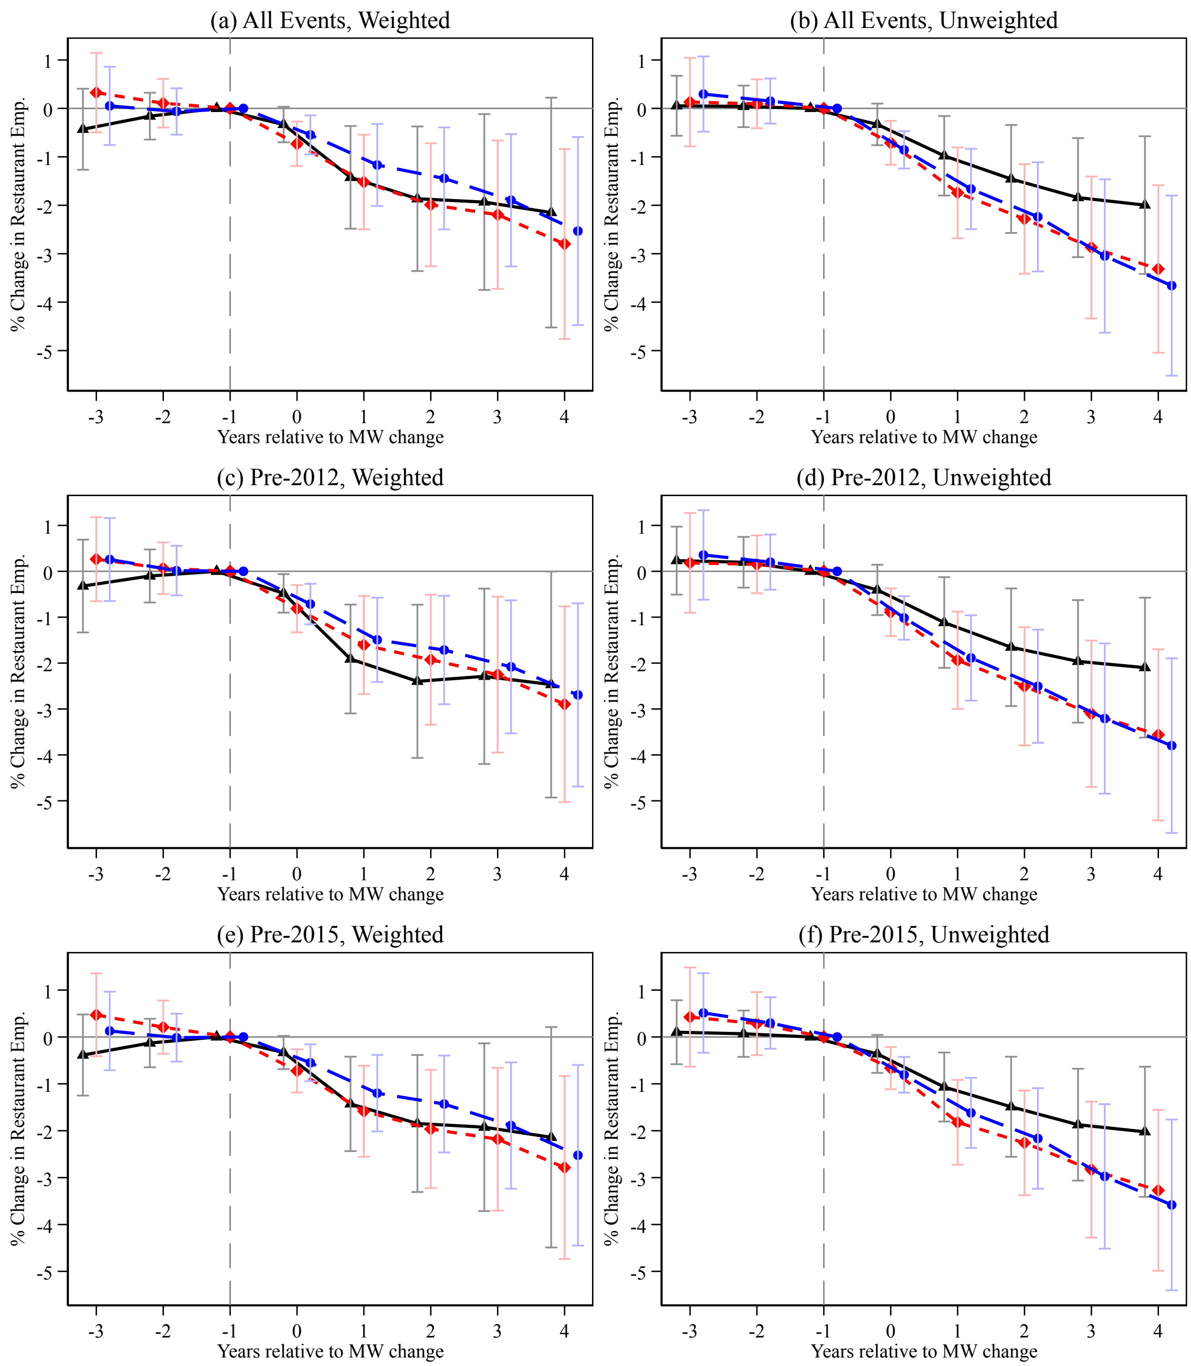

Figure 4: Event-Study Paths for % Change in Restaurant Employment with 95% CIs (Bundled Events): Emp/Pop Varying (solid, ▲), Emp/Pop Constant (dash, ♦), Log Employment (long dash, •)

total employment is given by

$$\varepsilon_{\text{post}}^{T} = \frac{\%\Delta\text{Emp}_{\text{post}}^{T}}{\overline{\%\Delta\text{MW}}},$$

where $\%\Delta\text{Emp}_{\text{post}}^{T}$ is defined as in (4) and (5), and $\overline{\%\Delta\text{MW}}$ is the (weighted or unweighted) average percentage change in the minimum wage across the events in each sample. Analogous definitions hold for the average minimum wage elasticity of restaurant employment, $\varepsilon_{\text{post}}^{R}$, and for the elasticities at any of the individual horizons τ.

Table 5 shows the estimated minimum wage elasticities of employment. Panel A reports the average elasticities over the post-period, and Panel B the long-term elasticities (at $\tau = 4$), with columns matching the structure of Tables 1 and 3. To reproduce CDLZ's approach of assuming no effects of minimum wages in the lower and upper bins, we also present estimated elasticities for CDLZ's "Affected" employment group, which focuses on the missing and excess bins; following CDLZ, elasticities for this group are calculated using the total employment-to-population ratio as the baseline. Each panel's rows are identified by employment group (CDLZ's Affected, Total, and Restaurant) and sample (all events, pre-2012, and pre-2015).[9]

The shaded cells show CDLZ's estimated elasticities from weighted specifications using employment ratios to varying population. Their conclusion is: no matter the sample or the horizon, the minimum wage elasticity of CDLZ's "Affected" group is essentially zero. Using total employment or unweighted specifications does not change the story, as long as population varies in the outcome variable. Once we remove the varying population component, using the employment ratio to constant population or log employment as the outcome variable, the story breaks down. First, all estimated elasticities for the Affected group are negative and larger in magnitude, though statistical significance is limited to 1 out of 8 coefficients in Panel A, and to 3 out of 8 coefficients in Panel B.[10] Second, once we account for all the employment (which adds the lower and upper bins to CDLZ's Affected group), all the estimated average and long-term elasticities across samples and weighting schemes (24 estimated elasticities) are negative and significant with average elasticities in the $[-0.07, -0.05]$ range for the weighted specifications, and in the $[-0.15, -0.13]$ range for the unweighted specifications; the long-term estimated elasticities are in the $[-0.10, -0.09]$ and $[-0.22, -0.19]$ ranges, respectively. These results show not only the upward bias in estimated elasticities when using varying population, but also that focusing only on CDLZ's Affected group

[9] Appendix Table A1 shows the values for $\overline{\text{EPOP}}_{-1}$, $\overline{\text{EPOP}}_{-1}^{R}$, and $\overline{\%\Delta\text{MW}}$ used in the calculation of percentage employment changes and elasticities in Tables 1, 3, and 5. Across cases, $\overline{\text{EPOP}}_{-1}$ ranges between 0.57 and 0.67, $\overline{\text{EPOP}}_{-1}^{R}$ ranges between 0.04 and 0.05, and $\overline{\%\Delta\text{MW}}$ ranges between 10.11% and 10.95%.

[10] We do not obtain estimates for the Affected group with the pre-2015 sample because we do not extend CDLZ's CPS-MORG sample to 2019, and thus we cannot calculate employment counts for the $\{L, M, E, U\}$ bins after 2016.

Table 5: Estimated Minimum Wage Elasticities of Employment

	Emp/Pop Varying		**Emp/Pop Constant**		**Log Employment**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Panel A: Average Post-Period Elasticities (τ=0 through τ=4)						
Affected (M+E), All Events	-0.002	-0.007	-0.023	-0.035	-0.079	-0.165**
	(0.021)	(0.026)	(0.023)	(0.026)	(0.072)	(0.078)
Affected (M+E), Pre-2012	-0.003	0.000	-0.019	-0.026	-0.053	-0.129
	(0.023)	(0.027)	(0.026)	(0.028)	(0.079)	(0.080)
Total, All Events	-0.013	0.011	-0.059**	-0.130***	-0.050*	-0.136***
	(0.037)	(0.031)	(0.026)	(0.037)	(0.026)	(0.035)
Total, Pre-2012	-0.034	0.003	-0.065**	-0.147***	-0.059**	-0.146***
	(0.036)	(0.033)	(0.028)	(0.039)	(0.027)	(0.036)
Total, Pre-2015	-0.019	0.009	-0.062**	-0.125***	-0.052**	-0.126***
	(0.035)	(0.029)	(0.025)	(0.036)	(0.025)	(0.032)
Restaurant, All Events	-0.089**	-0.092***	-0.135***	-0.161***	-0.106***	-0.180***
	(0.042)	(0.029)	(0.037)	(0.038)	(0.034)	(0.036)
Restaurant, Pre-2012	-0.111***	-0.102***	-0.132***	-0.176***	-0.115***	-0.187***
	(0.043)	(0.031)	(0.039)	(0.041)	(0.035)	(0.037)
Restaurant, Pre-2015	-0.095**	-0.102***	-0.138***	-0.168***	-0.107***	-0.178***
	(0.040)	(0.028)	(0.037)	(0.037)	(0.033)	(0.034)
Panel B: Long-Term Elasticities (τ=4)						
Affected (M+E), All Events	-0.003	-0.031	-0.045	-0.079*	-0.124	-0.286**
	(0.032)	(0.043)	(0.034)	(0.043)	(0.096)	(0.119)
Affected (M+E), Pre-2012	-0.003	-0.023	-0.041	-0.071	-0.098	-0.249**
	(0.032)	(0.042)	(0.035)	(0.043)	(0.098)	(0.117)
Total, All Events	-0.011	0.010	-0.094**	-0.195***	-0.089*	-0.212***
	(0.074)	(0.047)	(0.046)	(0.059)	(0.049)	(0.056)
Total, Pre-2012	-0.028	0.002	-0.098**	-0.208***	-0.094**	-0.215***
	(0.070)	(0.047)	(0.046)	(0.060)	(0.048)	(0.054)
Total, Pre-2015	-0.015	0.008	-0.095**	-0.188***	-0.089*	-0.199***
	(0.070)	(0.045)	(0.044)	(0.058)	(0.047)	(0.053)
Restaurant, All Events	-0.104	-0.125***	-0.189***	-0.228***	-0.160**	-0.275***
	(0.075)	(0.048)	(0.064)	(0.067)	(0.064)	(0.063)
Restaurant, Pre-2012	-0.121	-0.130***	-0.188***	-0.241***	-0.164***	-0.272***
	(0.074)	(0.048)	(0.066)	(0.069)	(0.063)	(0.060)
Restaurant, Pre-2015	-0.108	-0.132***	-0.190***	-0.232***	-0.159***	-0.266***
	(0.072)	(0.047)	(0.063)	(0.066)	(0.062)	(0.060)

Notes: Each cell reports the estimated elasticity of employment with respect to the minimum wage. The top header indicates the dependent variable: employment to varying population, employment to constant population, and log employment. The second header indicates whether the estimation of equation (1) is weighted (by constant state population) or unweighted. Rows in each panel indicate the sector (CDLZ's "Affected" bins, Total, and Restaurants) and sample used: All Events (138 events), Pre-2012 (98 events), and Pre-2015 (119 events). The Pre-2015 sample extends through 2019Q4. Standard errors (in parentheses) are clustered at the event-state level. Shaded cells highlight CDLZ's missing/excess results. * p < 0.10, ** p < 0.05, *** p < 0.01.

conceals substantial employment losses in the lower and upper bins (compare the "Affected" and "Total" rows in the third through sixth columns), directly contradicting CDLZ's assumption that minimum wages have no effect outside the $[-4, 5)$ dollar range.

For the restaurant industry, the 36 estimated average and long-term elasticities in Table 5 are negative, with only 3 non-significant (the long-term weighted estimates under varying population). Estimates are stable across the three samples, so cross-sample averages provide a compact summary. Under varying population, weighted and unweighted estimates are close: average post-period elasticities cluster near -0.10 in both cases, while long-term elasticities are -0.11 (weighted, insignificant) and -0.13 (unweighted). Once we remove the varying-population component, magnitudes grow noticeably and a weighted–unweighted gap opens up. Under the employment ratio to constant population, the average post-period elasticity is -0.14 (weighted) and -0.17 (unweighted), with long-term values of -0.19 and -0.23; under log employment, the average post-period elasticity is -0.11 (weighted) and -0.18 (unweighted), with long-term values of -0.16 and -0.27. Therefore, as with total employment, restaurant elasticities are upward-biased under varying population, and unweighted estimates are systematically larger in magnitude than weighted ones once population is held fixed. The message is clear: CDLZ's stacked design, applied to their own events, yields negative and significant elasticity estimates for the restaurant industry that are in line with those from other studies (see Introduction) using alternative approaches.

As we would expect, comparing restaurant to total employment elasticities shows that restaurants absorb a disproportionate fraction of the economy's minimum-wage response. Across the 24 constant-population and log-employment cells, the restaurant-to-total elasticity ratio averages 1.63 (range $[1.16, 2.29]$), with weighted specifications averaging 2.0 and unweighted ones 1.3. Nevertheless, our finding of sizable and significant elasticities for total employment, with the long-term unweighted estimates reaching about -0.21, shows that minimum wage effects on employment extend beyond the restaurant industry to the entire wage distribution.

CDLZ and Dube and Lindner (2024) consider the own-wage elasticity of employment (OWE)—defined as the ratio of the minimum wage elasticity of employment to the minimum wage elasticity of wages or average earnings—more informative than the elasticity of employment, as the OWE captures the employment response to the wage increase driven by the minimum wage, which is akin to the elasticity of labor demand in a competitive model. However, as Dube et al. (forthcoming) put it, "the OWE is only meaningful when there is a clear wage effect": if the minimum wage elasticity of earnings is indistinguishable from zero, the OWE estimates and their standard errors will become arbitrarily large.

Table 6 presents the OWE estimates for total and restaurant employment, constructed from the elasticities in Table 5 and the minimum wage elasticities of average earnings reported in Table 7.[11] As shown in the latter table, there is no wage effect of minimum wages for the total economy: the estimated earnings elasticities are all insignificant and essentially zero across weighting schemes and horizons (the largest cross-sample average, for the unweighted post-period estimates, reaches only 0.02 in magnitude). This is reflected in the OWEs for total employment in both panels, with most estimates showing unrealistic values and extremely large standard errors.[12]

By contrast, the Restaurant OWEs in Table 6 are well-defined, as a consequence of significant restaurant earnings effects of minimum wages; Table 7 shows average cross-sample elasticity estimates of 0.14 (weighted and unweighted) for the average post-period, and 0.12 (weighted) and 0.10 (unweighted) in the long term. Average post-period OWEs are all significant (most of them at the 1% level). When using varying population, they are in the $[-0.95, -0.57]$ range and show broadly similar weighted and unweighted estimates. When using the employment ratio to constant population or log employment, unweighted OWEs are larger in magnitude than weighted OWEs, with the first in the $[-1.44, -1.11]$ range and the second in the $[-1.13, -0.68]$ range; thus, for unweighted estimates, the negative employment effects of minimum wages more than offset their positive effects on average earnings. Comparing Panel B against Panel A, long-term Restaurant OWEs across the 12 constant-population and log-employment cells are on average 1.87 times their average-post-period counterparts (range $[1.63, 2.15]$), reflecting both the deepening employment response over time and the weakening earnings response after $\tau = 2$ (see Appendix Figure A3 for the paths of the earnings elasticities)—in the long term, Restaurant OWEs reach up to -1.90 (weighted) and -2.96 (unweighted).[13]

When restricting the estimation to the bundled-events samples, estimated employment elasticities and OWEs are qualitatively similar to the baseline results (see Appendix Tables A2, A3, A4 and Appendix Figure A4). However, for the restaurant industry, employment effects become stronger and average earnings effects weaker (especially in the long term, see Table 7), which translates into larger estimated Restaurant OWEs both for the average post-period and in the long

[11] Average earnings are computed as the ratio of QCEW's payroll to employment, separately for the entire economy and for the restaurant industry. Also, we do not report OWEs for CDLZ's Affected group, as its bins are defined as wage ranges relative to the new minimum wage, so a \$1 minimum wage increase mechanically moves the group's average wage by close to \$1 by construction. In such a case, the "wage elasticity" reflects the bin definition rather than an economic response, and the OWE for the Affected group would approximate the employment elasticity already reported in Table 5.

[12] Standard errors for OWE estimates are computed using the delta method assuming independence between employment and earnings coefficient estimates.

[13] Note that even when using varying population, unweighted long-term Restaurant OWEs range between -1.41 and -1.22.

Table 6: Own-Wage Elasticities of Employment

	Emp/Pop Varying		**Emp/Pop Constant**		**Log Employment**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Panel A: Average Post-Period OWEs (τ=0 through τ=4)						
Total, All Events	-1.151	0.955	-5.252	-11.088	-4.447	-11.603
	(4.109)	(3.377)	(11.303)	(24.652)	(9.651)	(25.754)
Total, Pre-2012	16.203	0.240	31.387	-12.724	28.218	-12.626
	(200.979)	(2.884)	(388.071)	(29.488)	(348.929)	(29.226)
Total, Pre-2015	-1.610	0.340	-5.327	-4.870	-4.529	-4.923
	(4.382)	(1.184)	(10.747)	(4.846)	(9.201)	(4.862)
Restaurant, All Events	-0.569**	-0.637***	-0.860***	-1.112***	-0.675***	-1.249***
	(0.282)	(0.234)	(0.275)	(0.338)	(0.243)	(0.347)
Restaurant, Pre-2012	-0.950**	-0.782***	-1.133***	-1.347***	-0.991***	-1.438***
	(0.424)	(0.292)	(0.419)	(0.430)	(0.371)	(0.424)
Restaurant, Pre-2015	-0.686**	-0.759***	-0.993***	-1.246***	-0.773***	-1.316***
	(0.312)	(0.254)	(0.318)	(0.369)	(0.272)	(0.361)
Panel B: Long-Term OWEs (τ=4)						
Total, All Events	-5.191	32.758	-44.793	-648.767	-42.455	-704.428
	(80.122)	(4274.077)	(621.519)	(84590.286)	(589.175)	(91847.628)
Total, Pre-2012	3.494	3.646	12.327	-360.210	11.798	-373.599
	(15.941)	(256.869)	(47.173)	(24084.142)	(45.213)	(24979.303)
Total, Pre-2015	-4.919	0.620	-30.223	-14.836	-28.430	-15.697
	(49.613)	(3.975)	(272.161)	(43.908)	(256.106)	(46.394)
Restaurant, All Events	-0.790	-1.222*	-1.432**	-2.228*	-1.218*	-2.681**
	(0.627)	(0.700)	(0.679)	(1.153)	(0.633)	(1.297)
Restaurant, Pre-2012	-1.220	-1.411*	-1.901*	-2.617*	-1.659*	-2.957*
	(0.915)	(0.838)	(1.054)	(1.424)	(0.952)	(1.519)
Restaurant, Pre-2015	-0.924	-1.391*	-1.618**	-2.445*	-1.352*	-2.802**
	(0.698)	(0.784)	(0.787)	(1.280)	(0.710)	(1.381)

Notes: Each cell reports the estimated own-wage elasticity of employment (OWE), defined as the ratio of the estimated minimum wage elasticity of employment to the minimum wage elasticity of average earnings. The top header indicates the dependent variable: employment to varying population, employment to constant population, and log employment. The second header indicates whether the estimation of equation (1) is weighted (by constant state population) or unweighted. Rows in each panel indicate the sector (Total or Restaurants) and sample used: All Events (138 events), Pre-2012 (98 events), and Pre-2015 (119 events). The Pre-2015 sample extends through 2019Q4. Standard errors (in parentheses) are clustered at the event-state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

term; all long-term Restaurant OWEs are larger in magnitude, but only 3 remain significant, as a consequence of the weakened long-term average earnings effect.

To summarize, the results in this section show that one CDLZ choice drives their no-employment-effect conclusion: once we drop the varying-population denominator, that conclusion does not survive their own design, for both total and restaurant employment. Moreover, the second CDLZ choice of ignoring employment effects beyond the $[-4, 5)$ wage bins further obscures the full negative minimum wage effects on total employment.

Table 7: Estimated Minimum Wage Elasticities of Average Earnings: Baseline and Bundled Events

	Baseline Sample		**Bundled-Events Sample**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*
Panel A: Average Post-Period Elasticities (τ=0 through τ=4)				
Total, All Events	0.011	0.012	0.031	0.036
	(0.024)	(0.026)	(0.022)	(0.033)
Total, Pre-2012	-0.002	0.012	0.017	0.035
	(0.026)	(0.027)	(0.023)	(0.034)
Total, Pre-2015	0.012	0.026	0.032	0.046
	(0.023)	(0.024)	(0.021)	(0.030)
Restaurant, All Events	0.156***	0.144***	0.143***	0.128***
	(0.026)	(0.028)	(0.025)	(0.035)
Restaurant, Pre-2012	0.116***	0.130***	0.110***	0.118***
	(0.026)	(0.029)	(0.024)	(0.037)
Restaurant, Pre-2015	0.139***	0.135***	0.136***	0.124***
	(0.025)	(0.026)	(0.025)	(0.033)
Panel B: Long-Term Elasticities (τ=4)				
Total, All Events	0.002	0.000	0.034	0.047
	(0.029)	(0.039)	(0.032)	(0.050)
Total, Pre-2012	-0.008	0.001	0.023	0.045
	(0.030)	(0.039)	(0.032)	(0.050)
Total, Pre-2015	0.003	0.013	0.034	0.055
	(0.028)	(0.037)	(0.032)	(0.048)
Restaurant, All Events	0.132***	0.102**	0.096**	0.080
	(0.043)	(0.044)	(0.043)	(0.056)
Restaurant, Pre-2012	0.099**	0.092**	0.068*	0.072
	(0.042)	(0.043)	(0.041)	(0.055)
Restaurant, Pre-2015	0.117***	0.095**	0.090**	0.075
	(0.042)	(0.042)	(0.043)	(0.055)

Notes: Each cell reports the estimated elasticity of average earnings with respect to the minimum wage, computed from estimating equation (1) with log QCEW average earnings (earnings/employment) as the outcome and dividing by $\overline{\%\Delta \text{MW}}$. Panel A reports the average elasticities over the post-period, and Panel B the long-term elasticities (at τ=4). The top header indicates the events sample: the baseline samples have 138 events (All Events), 98 events (Pre-2012), and 119 events (Pre-2015); the bundled-events samples retain only the first event per state and any subsequent event occurring at least 20 quarters after the previous retained event in that state, yielding 80 events (All Events), 60 events (Pre-2012), and 79 events (Pre-2015). The second header indicates whether the estimation is weighted (by constant state population) or unweighted. Rows indicate the sector (Total or Restaurant) and sample used; the Pre-2015 sample extends through 2019Q4. These earnings elasticities serve as the denominator in the own-wage elasticity calculations in Table 6 (baseline columns) and its bundled-events counterpart Table A4 (bundled columns). Standard errors (in parentheses) are clustered at the event-state level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

3 Minimum Wages, Employment, and Population

In Section 2 we argued that the employment-to-population ratio with a time-varying denominator gives a biased measure of the minimum wage effect on employment: if both numerator and denominator fall after a minimum wage increase, the ratio falls by less than employment does (and may even increase), so the event-study coefficients understate the percentage change in employment. The previous tables and figures confirm this upward bias by showing, for both the total economy and the restaurant industry, a clear gap between the (weighted or unweighted) estimates using a varying-population denominator and those using a constant-population denominator or log employment. This section explains that gap through two related channels: a negative relationship between minimum wages and state population, and an additional bias in CDLZ's *level* varying-population specification relative to its log counterpart, $\ln(\text{Emp}/\text{Pop})$.

As discussed in Section 2.1, the advantage of directly using the level of the employment-to-population ratio is that we can cleanly split each wage bin's contribution to the implied percentage change in the total employment-to-population ratio, which CDLZ directly interpret as the percentage change in total employment under the implicit assumption that population changes are negligible. If we are only interested in total employment effects of minimum wages, we could instead use log employment as the dependent variable, as we did above. Here, we turn instead to the log of the employment-to-population ratio—the same dependent variable used by Dube and Lindner (2024)—which is useful because it admits an exact decomposition of the minimum wage effect into employment and population contributions: since $\ln(\text{Emp}/\text{Pop}) = \ln(\text{Emp}) - \ln(\text{Pop})$, estimating equation (1) with log population as the dependent variable and combining the result with our log employment estimates from Section 2.2 yields a log ratio effect that equals the log employment effect minus the log population effect.

Table 8 reports weighted and unweighted average post-period estimates of minimum wage effects on $\widehat{\%\Delta(\text{Emp}/\text{Pop})}^{\text{log}}_{\text{post}}$, the percentage change in the employment-to-population ratio from the log specification; on $\widehat{\%\Delta\text{Emp}}^{\text{log}}_{\text{post}}$ and $\widehat{\%\Delta\text{Pop}}^{\text{log}}_{\text{post}}$, the percentage changes in employment and population from the log specifications; and on $\widehat{\%\Delta(\text{Emp}/\text{Pop})}^{\text{level}}_{\text{post}}$, the percentage change in the employment-to-population ratio from the level specification used by CDLZ. Estimates are reported for the all-events, pre-2012, and pre-2015 samples in our baseline (Panel A) and bundled-events (Panel B) datasets. The level employment-to-population and log employment estimates reproduce, respectively, the first and second columns (CDLZ's Emp/Pop Varying specification) and the fifth and sixth columns (Log Employment) of Tables 3 and 4, Panel A; the log employment-to-

Table 8: Exact Log Decomposition of the Varying-Population Employment Effect

	All Events		**Pre-2012 Events**		**Pre-2015 Events**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Panel A: Baseline						
$\widehat{\%\Delta(\text{Emp/Pop})}^{\text{log}}_{\text{post}}$	-0.334	-0.564**	-0.604	-0.713**	-0.405	-0.567**
	(0.397)	(0.283)	(0.410)	(0.311)	(0.389)	(0.275)
$\widehat{\%\Delta\text{Emp}}^{\text{log}}_{\text{post}}$	-0.518*	-1.376***	-0.642**	-1.562***	-0.570**	-1.326***
	(0.274)	(0.350)	(0.300)	(0.380)	(0.273)	(0.342)
$\widehat{\%\Delta\text{Pop}}^{\text{log}}_{\text{post}}$	-0.184	-0.813***	-0.038	-0.849***	-0.165	-0.758***
	(0.383)	(0.278)	(0.411)	(0.301)	(0.378)	(0.270)
$\widehat{\%\Delta(\text{Emp/Pop})}^{\text{level}}_{\text{post}}$	-0.134	0.113	-0.368	0.030	-0.203	0.092
	(0.387)	(0.313)	(0.398)	(0.350)	(0.379)	(0.307)
Gap (level − log)	0.200	0.677	0.235	0.743	0.202	0.659
Panel B: Bundled Events						
$\widehat{\%\Delta(\text{Emp/Pop})}^{\text{log}}_{\text{post}}$	-0.815	-0.657	-1.150*	-0.821*	-0.796	-0.575
	(0.594)	(0.419)	(0.616)	(0.465)	(0.581)	(0.389)
$\widehat{\%\Delta\text{Emp}}^{\text{log}}_{\text{post}}$	-0.818**	-1.589***	-0.962**	-1.805***	-0.812**	-1.430***
	(0.346)	(0.539)	(0.378)	(0.586)	(0.337)	(0.505)
$\widehat{\%\Delta\text{Pop}}^{\text{log}}_{\text{post}}$	-0.003	-0.932**	0.187	-0.984**	-0.016	-0.855**
	(0.612)	(0.431)	(0.675)	(0.471)	(0.593)	(0.402)
$\widehat{\%\Delta(\text{Emp/Pop})}^{\text{level}}_{\text{post}}$	-0.661	0.098	-0.957	0.025	-0.644	0.125
	(0.579)	(0.478)	(0.599)	(0.536)	(0.567)	(0.439)
Gap (level − log)	0.153	0.756	0.193	0.846	0.152	0.700

Notes: In each panel, the first four rows report the estimated average post-treatment percentage change in the employment-to-population ratio from the log specification (row 1), the percentage changes in employment (row 2) and population (row 3) from log specifications, and the percentage change in the employment-to-population ratio from the level specification (row 4). The last row shows the gap between row 4 and row 1. Because all three log regressions share the same regressors and weights, the identity $\ln(\text{Emp/Pop}) = \ln(\text{Emp}) - \ln(\text{Pop})$ implies that row 1 equals row 2 minus row 3 exactly. Panel A uses the baseline dataset, with the top header indicating the sample used: All Events (138 events, 109,326 observations), Pre-2012 (98 events with full post-period windows, 93,176 observations), and Pre-2015 (119 events with full post-period windows, 106,520 observations). Panel B uses the bundled dataset, with the top header indicating the sample used: All Bundled Events (80 events, 65,184 observations), Pre-2012 Bundled Events (60 events with full post-period windows, 56,451 observations), and Pre-2015 Bundled Events (79 events with full post-period windows, 68,515 observations). The Pre-2015 sample extends through 2019Q4. The second header indicates whether the estimation of equation (1) is weighted (by constant state population) or unweighted. All regressions include event-state and event-quarter fixed effects, and fixed effects accounting for federal minimum wage changes, overlapping minimum wage changes, and other minimum wage events. Standard errors (in parentheses) are clustered at the event-state level. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

population and log population rows are new. The table also reports the gap between the level and log employment-to-population estimates.

There is a negative relationship between minimum wages and population. Across the three samples and in both panels, the unweighted estimates show that a minimum wage event is associated with an average post-period decline in population in the $[-0.98\%, -0.76\%]$ range, corresponding to an average post-period elasticity in the $[-0.08, -0.07]$ range. Moreover, as with employment, the unweighted event-study paths in Appendix Figure A5 show the relationship between minimum wages and population becoming more negative over time; these paths also show no evidence of pre-trends. The weighted population estimates, by contrast, are insignificant in every cell and the event-study plots (Appendix Figure A5) show no clear dynamic path. This points to a small number of large states—such as California, Texas, New York, and Florida—driving the weighted estimates.

To assess the size of the bias arising from the negative relationship between minimum wages and population, we can compare the population estimates against the employment estimates and the log ratio estimates. In both panels, the population estimates are between 54% and 60% of the size of the employment estimates in the unweighted specifications, but the employment effects are so large that even after the population channel halves them, the log ratio estimates remain sizable (in the $[-0.82\%, -0.56\%]$ range) and significant in four out of six cases. In the weighted specifications, the population estimates are between 6% and 36% of the size of the employment estimates in Panel A, and even smaller in Panel B, where the population estimates are close to zero in two cases and have the opposite sign in pre-2012. Even if the varying-population upward bias is thus smaller in the weighted specifications, the log ratio estimates show that adding population dramatically reduces the precision of the estimates. Take, for example, the bundled all-events estimates in the first column: the population bias is zero and the log ratio estimate is about the same as the log employment estimate at about -0.82%; however, the standard error for the log ratio estimate is 1.7 times larger than the standard error for the log employment estimate.

The comparison between the $\widehat{\%\Delta(\text{Emp}/\text{Pop})}^{\text{level}}_{\text{post}}$ and $\widehat{\%\Delta(\text{Emp}/\text{Pop})}^{\text{log}}_{\text{post}}$ rows in each panel of Table 8 highlights another source of bias in CDLZ's dependent variable choice. Across both panels, CDLZ's dependent variable (the level of the employment-to-population ratio) causes an upward bias ranging between 0.15 and 0.24 percentage points in the weighted specifications, and between 0.66 and 0.85 percentage points in the unweighted specifications. This bias is so large in the unweighted specifications that it flips the sign in CDLZ's level specifications, even when the log ratio specifications yield negative and sizable estimates. Had CDLZ used a log ratio specification for total employment—as two of the authors later do in Dube and Lindner (2024)—their null story would no longer stand.

But what is the source of this large bias? It is a covariance: the level functional form, using a denominator that itself is related to the policy, makes the dependent variable's level and its log change move together. To see this, start from the first-order approximation $\Delta \ln(\text{Emp}_{hst}/\text{Pop}_{st}) \approx$

$\Delta(\text{Emp}_{hst}/\text{Pop}_{st})/(\text{Emp}_{hst}/\text{Pop}_{st})$, which can be rearranged as

$$\Delta(\text{Emp}_{hst}/\text{Pop}_{st}) \;\approx\; \text{EPOP}_{hst} \cdot \Delta \ln(\text{Emp}_{hst}/\text{Pop}_{st}), \tag{6}$$

where $\text{EPOP}_{hst} = \text{Emp}_{hst}/\text{Pop}_{st}$, and Δ denotes the change induced by a minimum wage event. Let α^{level} and $\alpha^{\log}$ denote, respectively, the post-period minimum wage effect on the level and log of the employment-to-population ratio; that is, $\alpha^{\text{level}} \equiv E\big[\Delta(\text{Emp}_{hst}/\text{Pop}_{st})\big]$ and $\alpha^{\log} \equiv E\big[\Delta \ln(\text{Emp}_{hst}/\text{Pop}_{st})\big]$, with expectations taken over the post-period. Taking the expectation of equation (6) yields

$$\alpha^{\text{level}} \;\approx\; E[\text{EPOP}_{hst}] \cdot \alpha^{\log} \;+\; \text{Cov}(\text{EPOP}_{hst}, \Delta \ln(\text{Emp}_{hst}/\text{Pop}_{st})). \tag{7}$$

Using the approximation $E[\text{EPOP}_{hst}] \approx \overline{\text{EPOP}}_{-1}$—used in the percentage-change rescaling of (4)—and dividing both sides of equation (7) by $\overline{\text{EPOP}}_{-1}$ and multiplying by 100, it then follows that

$$\widehat{\%\Delta(\text{Emp}/\text{Pop})}^{\text{level}}_{\text{post}} \approx \widehat{\%\Delta(\text{Emp}/\text{Pop})}^{\log}_{\text{post}} + \underbrace{\frac{\widehat{\text{Cov}}(\text{EPOP}_{hst}, \Delta \ln(\text{Emp}_{hst}/\text{Pop}_{st}))}{\overline{\text{EPOP}}_{-1}} \times 100}_{\text{Gap}}. \tag{8}$$

Thus, the gap observed between the level and log rows of Table 8 is given by the covariance term in equation (8).

The gap term need not be signed in general, but the varying-population denominator gives it a structural reason to be positive. A minimum wage shock that is negatively related to Pop_{st} raises both EPOP_{hst} and $\Delta \ln(\text{Emp}_{hst}/\text{Pop}_{st})$—since Pop_{st} enters the denominator of both—producing a systematically positive gap. In contrast, with a constant-population denominator, $\Delta \ln(\text{Pop}_{st}) = 0$ by construction; since $\Delta \ln(\text{Emp}_{hst}/\text{Pop}_{st}) = \Delta \ln(\text{Emp}_{hst}) - \Delta \ln(\text{Pop}_{st})$, the covariance reduces to $\widehat{\text{Cov}}(\text{EPOP}_{hst}, \Delta \ln(\text{Emp}_{hst}))$, which has no obvious sign. Moreover, the evidence indicates that with a constant-population denominator the bias is minor; in particular, looking at Tables 3 and 4, comparing the constant-population estimates (Panel A, third and fourth columns) and the log employment estimates (fifth and sixth columns), the "level minus log" gap between those columns is small and stays within ±0.1 percentage points across all comparisons.

4 The Dube and Lindner (2024) Pooled Estimates

As Solon et al. (2015) note, "[I]t often is good practice to report both weighted and unweighted estimates." In our previous sections, we presented weighted and unweighted estimates side-by-side and highlighted when they diverge. This practice becomes critical in this section. In particular, this section shows that the null employment effects of minimum wages obtained by Dube and Lindner (2024) do not survive a modest switch from population-weighted to unweighted estimation.

4.1 The LP-DiD Approach

In their review of the minimum wage literature, Dube and Lindner (2024)—DL hereafter—present in their Table 1, columns 1–2, a state-level event-study design in the spirit of CDLZ, with a few important differences. As in our Section 2.4, they focus on QCEW total and restaurant employment rather than wage bins.[14] They use yearly rather than quarterly data, and extend the post-period to 6 years (instead of 5). To avoid the overlapping-events problem in CDLZ, discussed in Section 2.3, their event pool consists of 60 bundled events spanning 1986 to 2019.[15] Crucially, each event's post-period would extend from the event year through $\tau = 5$, but DL truncate the post-period at the year before the next federal minimum wage increase, or at 2019 (the last year in the sample); their 4 events from 2006, for example, contribute only the event year ($\tau = 0$) because the federal minimum wage rose in 2007.

DL implement the local projections difference-in-differences (LP-DiD) estimator of Dube et al. (2025), which for their application is equivalent to CDLZ's stacked design. For state s at year t, the LP-DiD specification is given by

$$\overline{\Delta y_{st}} = \alpha_{\text{post}} D_{st} + \delta_t + \nu_{st}, \tag{9}$$

where $\overline{\Delta y_{st}} = \frac{1}{6}\sum_{\tau=0}^{5}(y_{s,t+\tau} - y_{s,t-1})$ is the average post-period difference (with respect to $\tau = -1$) for the outcome variable y, $D_{st} = 1$ if state s had a minimum wage event in year t (and zero otherwise), δ_t is a year fixed effect, and ν_{st} is the error term. For each year t, the sample is restricted to include only observations with treated or control states. We can also obtain the treatment effect at each horizon τ by estimating

$$y_{s,t+\tau} - y_{s,t-1} = \alpha_\tau D_{st} + \delta_t + \nu_{st}. \tag{10}$$

[14] The specifications we review in this section are the ones DL emphasize (Dube and Lindner, 2024, p. 283). These specifications are also the closest in empirical design, sample period, and data source to our analysis in Section 2.4. (Their Table 1 also includes four additional columns using CPS samples, as well as three other panels presenting a 46-event 1998–2019 subsample and TWFE estimates.) DL do not provide a replication package, but we were able to match their results almost exactly; for example, in their columns 1–2, Panel A, they found OWEs of −0.026 (0.151) and 0.331 (0.217), while we found −0.028 (0.155) and 0.329 (0.223), respectively.

[15] Out of the 60 DL events, 38 match one of our 80 bundled events from Section 2.3 by state-year. The remaining 22 DL events have no state-year bundled counterpart: 5 are outside CDLZ's event-time window (2017–2019), and 17 are inside the window but excluded by CDLZ's event-definition rules. Conversely, 42 of our 80 bundled events have no state-year match with DL's events: 2 are outside DL's event-time window (DC 1980 and Connecticut 1981), and the remaining 40 are inside the window but excluded by DL's event-definition rules (of those 40, 21 fall in a federal minimum wage year).

If all events have identical post-period window lengths, the pooled estimate equals the simple average of the horizon-specific effects, $\widehat{\alpha}_{\text{post}} = \frac{1}{6}\sum_{\tau=0}^{5}\widehat{\alpha}_{\tau}$.

However, this equivalence does not hold in DL's application because their post-period lengths differ across events; thus, $\overline{\Delta y_{st}}$ is calculated differently from event to event.[16] The consequence is not harmless: with truncated post-period windows, the LP-DiD pooled estimate places disproportionate weight on the earliest horizons, biasing the estimate towards zero and undermining its interpretation as an average post-period effect.[17] In our reassessment of DL's Table 1 we therefore also present the event-study paths, reporting estimates at every horizon $\tau \in -3, -2, \ldots, 5$, with the understanding that the $\tau = 0$ estimates use all 60 events while the $\tau = 5$ estimates use only the 18 events with a full six-year post-period.

4.2 Results

Table 9 reports estimates of a varying-windows version of (9) for three outcomes: log average earnings, log employment-to-population (with time-varying population in the denominator), and log employment; the first two are DL's outcomes. We also report OWEs, computed via DL's two-stage approach but equivalent to the ratio of the log employment-to-population (or log employment) estimate to the log average earnings estimate.[18] As in Tables 1 and 3, all estimates except the OWEs are scaled as percentage changes, $\widehat{\alpha}_{\text{post}} \times 100$. In contrast to DL, we report both weighted and unweighted estimates. Figures 5 and 6 show the corresponding event-study paths for the restaurant industry and the total economy, respectively.

The shaded cells replicate the results in DL's Table 1, Panel A, columns 1–2. Using weighted (by state population in the event year) specifications, DL find positive and significant earnings effects of minimum wages in both restaurants (3.12%) and the total economy (1.45%), but no significant employment-to-population effects or OWEs: negative and close to zero for restaurants, and positive but imprecisely estimated for the total economy. In contrast to our findings in Section 2.4, switching to log employment (which avoids the varying-population bias) does not make much difference. This

[16] As discussed in Section 2.2, CDLZ worry about events with incomplete post-period windows and present results with the pre-2012 sample, which includes only full post-period events. In Section 2.4 we extend the QCEW data to 2019, which allows us to expand the full post-window sample (the pre-2015 sample). DL's design implicitly prioritizes the potential contamination of estimates from federal minimum wage changes—which CDLZ control for with fixed effects—over the consequences of varying post-period lengths for the LP-DiD estimator.

[17] The intuition for this truncation bias is clearest in the unweighted estimation case. In the unweighted pooled regression each event contributes roughly equally to $\widehat{\alpha}_{\text{post}}$, regardless of its post-window length. A short-window event contributes only to the impact period ($\tau = 0$), while a full-window event contributes equally to each of the six post-period horizons. Because most DL events have short post-windows, $\widehat{\alpha}_{\text{post}}$ ends up dominated by the impact-period effect, while the long-horizon effects enter only through the small set of events with full post-windows.

[18] DL interpret the log employment-to-population coefficients as predicted percentage changes in employment, implicitly assuming no population response to minimum wages.

Table 9: Pooled Event-Study Estimates Following Dube and Lindner (2024)

	Full Sample		Excluding CA/NY	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*
Panel A: Restaurants sample				
% Change in avg. earnings	3.122***	2.615***	2.441***	2.491***
	(0.568)	(0.537)	(0.614)	(0.572)
% Change in emp-to-pop ratio	-0.086	-1.000**	-1.054***	-1.267***
	(0.479)	(0.433)	(0.333)	(0.419)
OWE (emp/pop)	-0.028	-0.383**	-0.432***	-0.509***
	(0.155)	(0.156)	(0.137)	(0.154)
% Change in employment	-0.243	-0.954*	-1.257**	-1.213**
	(0.683)	(0.477)	(0.544)	(0.463)
OWE (emp)	-0.078	-0.365*	-0.515**	-0.487**
	(0.224)	(0.184)	(0.249)	(0.188)
Panel B: Total sample				
% Change in avg. earnings	1.445***	0.903**	0.883***	0.787**
	(0.429)	(0.352)	(0.316)	(0.365)
% Change in emp-to-pop ratio	0.476	-0.324	-0.332**	-0.525***
	(0.407)	(0.219)	(0.138)	(0.178)
OWE (emp/pop)	0.329	-0.360	-0.376*	-0.667
	(0.223)	(0.310)	(0.198)	(0.404)
% Change in employment	0.319	-0.279	-0.536	-0.471
	(0.651)	(0.397)	(0.429)	(0.376)
OWE (emp)	0.221	-0.309	-0.607	-0.599
	(0.404)	(0.494)	(0.547)	(0.603)

Notes: Each cell reports $\widehat{\alpha}_{\text{post}} \times 100$ from the estimation of equation (9) for the outcome in the row label, except for the OWE rows (defined below). The top header indicates the sample: "Full Sample" uses 60 state-level minimum wage event states and their clean controls (745 observations); "Excluding CA/NY" drops California and New York entirely (55 events, 727 observations). The second header indicates whether the estimation of equation (9) is weighted (by state population in the event year) or unweighted. Panel A reports estimates for restaurant employment; Panel B reports estimates for total employment. All regressions use annual QCEW data over 1985–2019 and include year fixed effects. OWE is the own-wage elasticity of employment, estimated by IV using the minimum wage event indicator as an instrument for log wage; each OWE row uses the employment outcome indicated in parentheses. Standard errors (in parentheses) are clustered at the state level. Shaded cells highlight DL's results. $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

is surprising, given that we use QCEW data over roughly the same period; the only data difference is frequency (quarterly in Section 2, yearly in DL's analysis). Moreover, the weighted full-sample event-study paths are flat for restaurants (see Figure 5, Panels (c) and (e)); for the total economy, the employment-to-population path turns positive and significant after $\tau = 3$ (see Figure 6, Panel (c)). This points to differences in event definitions between CDLZ and DL as a major driver of estimated minimum wage effects.

The key finding from Table 9 is how DL's story changes if we look instead at unweighted estimates.[19] Using the full sample, the estimated percentage change in average earnings continues to be positive and significant (though smaller in magnitude), but the most striking result is that the estimated percentage change in the employment-to-population ratio is -1%, which is significant and more than 11 times larger in magnitude than the weighted estimate. The implied OWE is -0.38, also significant. We obtain similar estimates if we use log employment instead. The event-study paths in Figure 5 show negative restaurant employment effects that become stronger through $\tau = 3$ (reaching about -2%) and then partially recover by $\tau = 5$. For the total sample (Panel B, column 2), the employment and OWE estimates turn negative, though insignificant, with the event-study paths in Figure 6, Panels (c) and (e), showing insignificant estimates at every horizon.

To further explore the source of the gap between weighted and unweighted estimates, columns 3 and 4 of Table 9 present results for a subsample that excludes California and New York (18 observations, including 5 events). The dramatic similarity between the weighted and unweighted columns across all cells, once California and New York are excluded, shows that they single-handedly drive DL's null employment results. For restaurants, every weighted and unweighted estimate is significant, with OWEs in the $[-0.52, -0.43]$ range. Panels (d) and (f) of Figure 5 reinforce this: the weighted and unweighted event-study paths track each other closely and are significant through most of the post-period. These panels also reveal a non-trivial truncation bias in DL's pooled specification: the pooled weighted estimates of -1.05% for the employment-to-population ratio and -1.26% for employment are smaller in magnitude than the post-period averages of -1.77% and -2.54% from Panels (d) and (f); the unweighted comparison is similar, with pooled estimates of -1.27% and -1.21% against averages of -1.61% and -1.73%.

Panel B of Table 9, third and fourth columns, also shows negative and significant estimates of the percentage change in the employment-to-population ratio in the total sample excluding California and New York, with OWE estimates in the $[-0.67, -0.38]$ range (one significant). There are, however, truncation issues worth highlighting. The weighted pooled estimate for employment-to-population is significant at -0.33%, but Panel (d) of Figure 6 shows a rather flat weighted path (post-period average of -0.32%) that is significant only at $\tau = 0$; this shows how the impact year drives not only the magnitude but also the significance of the pooled estimate. The unweighted path in Panel (d), by contrast, is mostly significant after $\tau = 1$, with a post-period average of -0.71% (larger in magnitude than the pooled -0.53% estimate). Panel (f) shows similar weighted and unweighted paths for employment, with the weighted estimates significant in the long term; the

[19] In their Table A3, DL provide five columns of robustness checks for the restaurant sample; none re-estimates without weights.

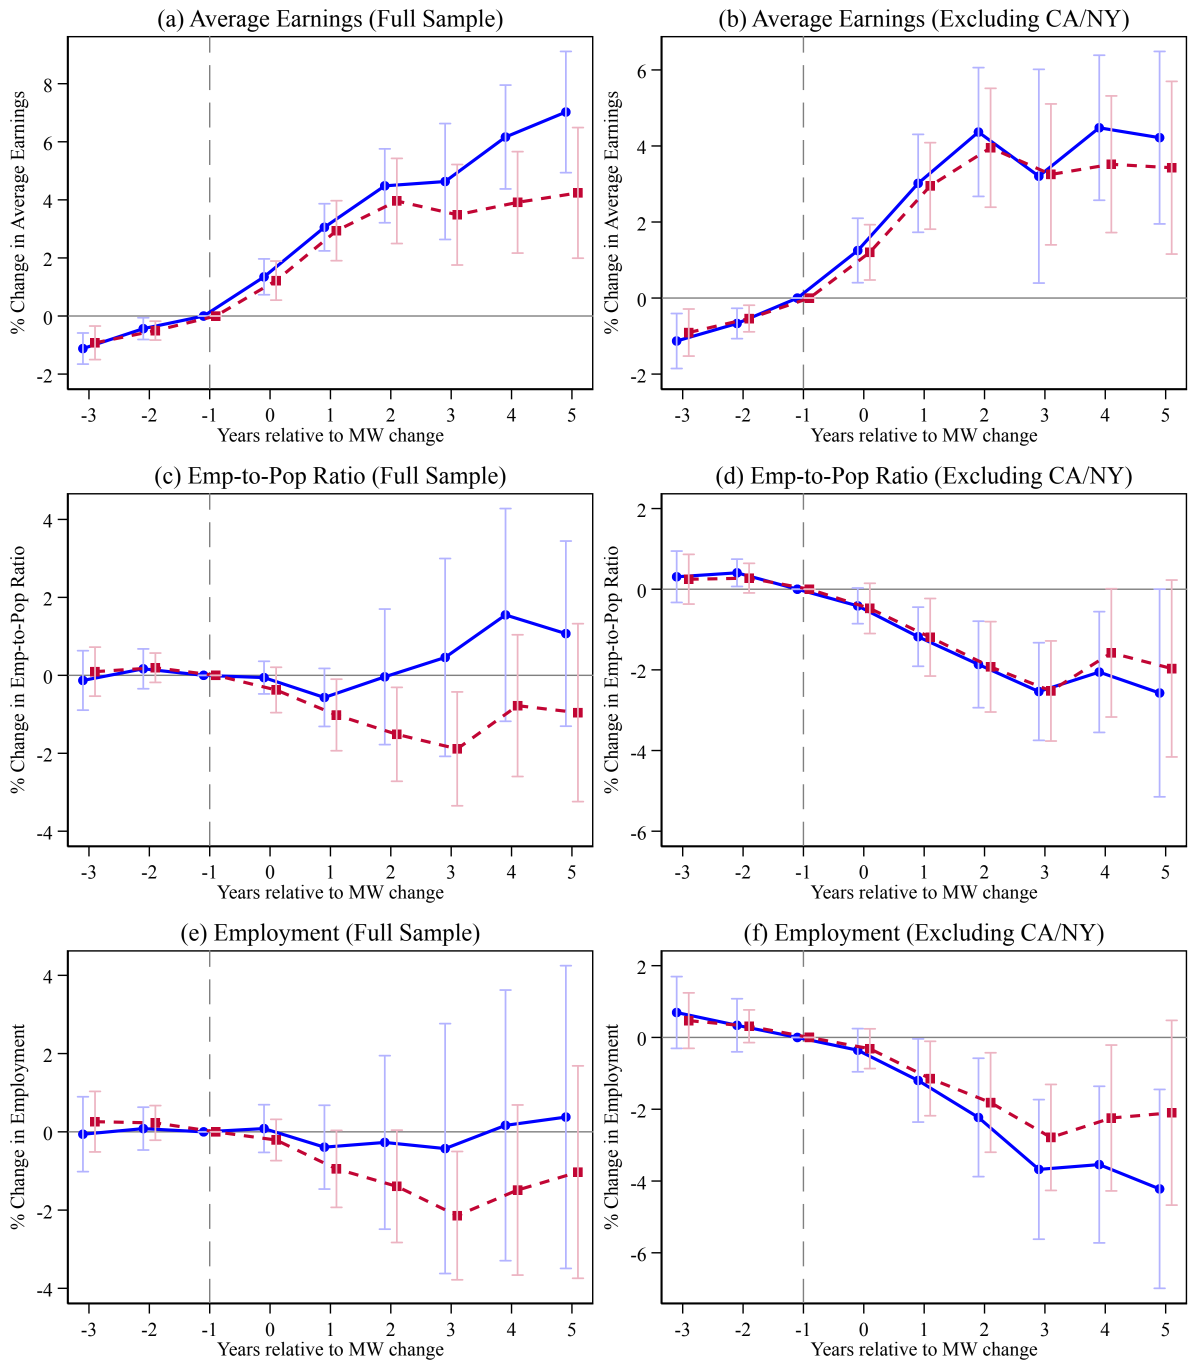

Figure 5: Event-Study Paths for the Restaurants Sample with 95% CIs:
Weighted (solid, •) and Unweighted (dash, ■)

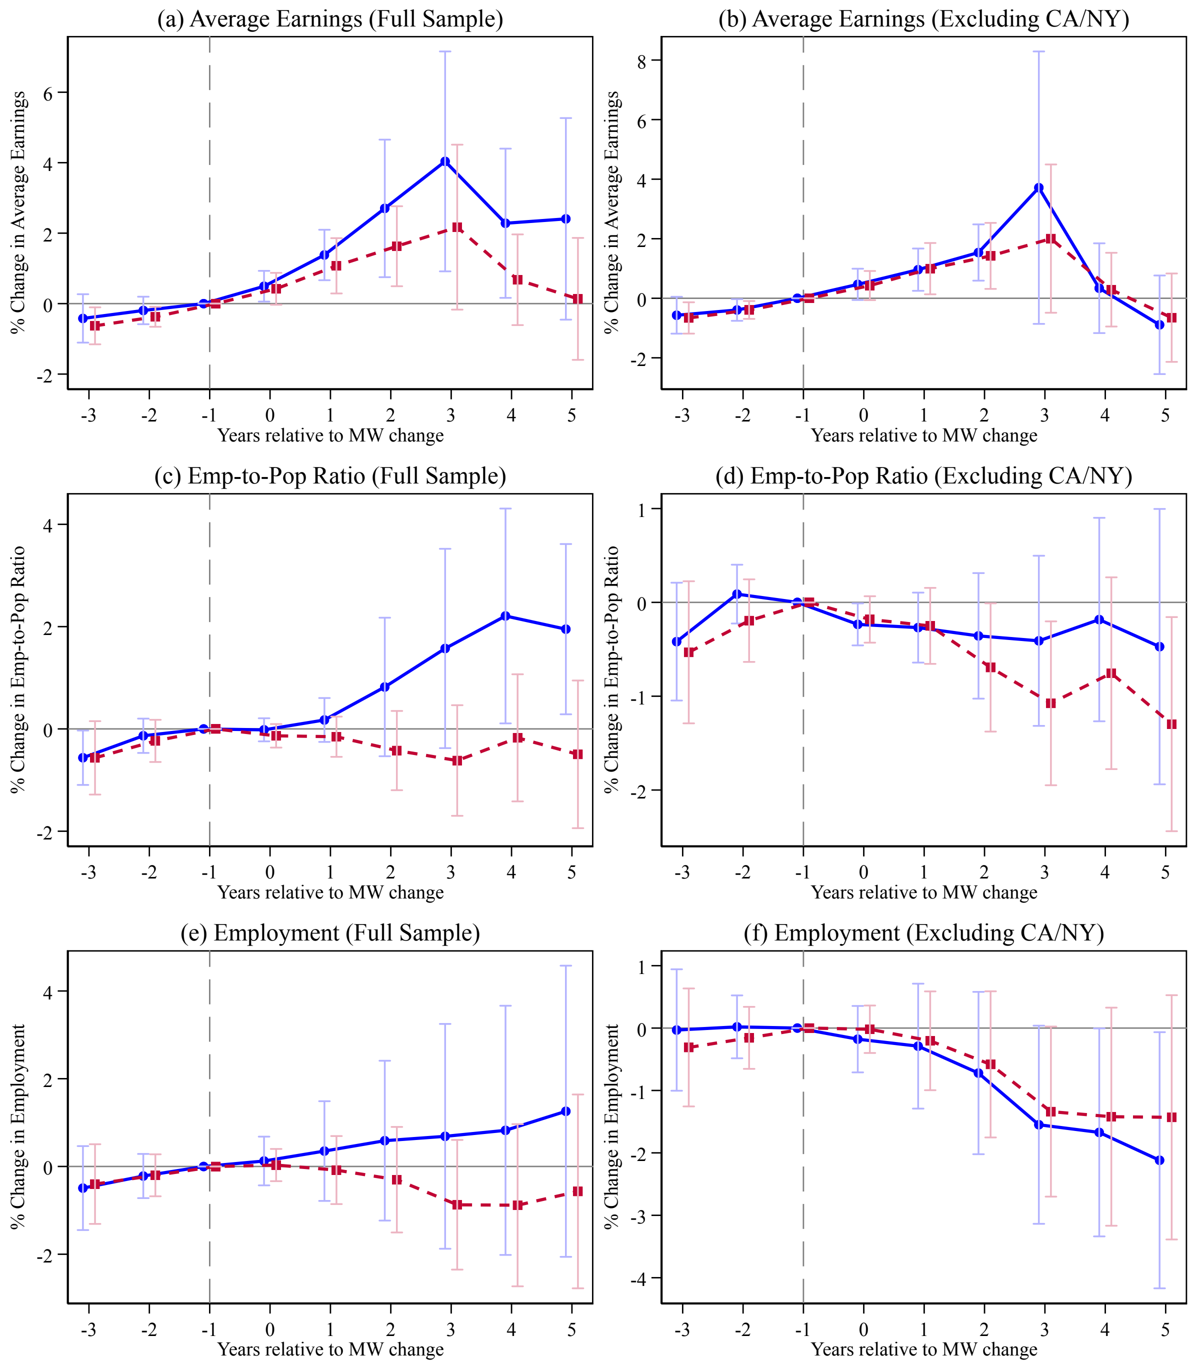

Figure 6: Event-Study Paths for the Total Sample with 95% CIs:
Weighted (solid, •) and Unweighted (dash, ■)

truncation bias is more severe here, with an average post-period weighted estimate of −1.09% that is twice the −0.54% pooled estimate, and an unweighted average of −0.83%, larger in magnitude than the pooled −0.47% estimate.

4.3 Discussion

The data provide the rationale for excluding California and New York. To see why they drive DL's null result, we focus on DL's weighted estimate of the percentage change in the employment-to-population ratio in the restaurant sample, $\widehat{\alpha}^{w}_{\text{post}} \times 100 = -0.086\%$. By the Frisch–Waugh–Lovell theorem, we can write the weighted LP-DiD pooled estimate as a sum of state-level contributions,

$$\widehat{\alpha}^{w}_{\text{post}} = \sum_{s} I_s, \qquad I_s = \frac{\sum_t \text{Pop}_{st}\, \widetilde{D}_{st}\, \tilde{y}_{st}}{\sum_{s',t'} \text{Pop}_{s't'}\, \widetilde{D}^{2}_{s't'}},$$

where Pop_{st} is the population of state s in year t, and $\widetilde{D}_{st}$ and $\tilde{y}_{st}$ are the event indicator D_{st} and the outcome $\overline{\Delta y_{st}}$ residualized on the year fixed effects δ_t. The term I_s aggregates the contributions of state s across all of its appearances in the sample, either as the treated state or as a control, with its sign indicating whether the state increases or reduces $\widehat{\alpha}^{w}_{\text{post}}$.

Figure 7 plots $I_s \times 100$ (percentage point contributions) for the top 25 contributors to the −0.086% estimate. California and New York have, by far, the two largest bars, both positive and of similar size. Compared to Washington (the third-largest state by absolute contribution, and the largest negative bar), California's contribution is 1.81 times as large, and New York's is 1.71 times as large. Together they offset most of the negative contributions from the remaining states, leaving the pooled estimate near zero.[20] Dropping CA and NY together—the sample used in columns 3 and 4 of Table 9—is the natural endpoint of this exercise.

In one of their robustness checks for their TWFE specification, CDLZ (Appendix A, p. 8) mention—citing Solon et al. (2015)—"The similarity of the weighted and unweighted estimates is reassuring, since a substantial difference between the two could reflect potential misspecification." One might have thought that Dube and Lindner—two of the authors of CDLZ—would have sought the same reassurance.

Given the difference between the weighted and unweighted estimates in the first and second columns of Table 9, does it mean that the LP-DiD model in (9) is misspecified? When looking at the full DL sample, the answer to this question is likely yes, and the reason, also mentioned

[20] Appendix Table A5 confirms the same point. Each row drops a single state from the regression and re-estimates the model; the table lists the ten states with the largest absolute leave-one-out shift. Dropping California shifts $\widehat{\alpha}^{w}_{\text{post}}$ from −0.09% to −0.54%, a swing of −0.45 percentage points; dropping New York yields a swing of −0.33 percentage points. The next-largest shift, from Washington, moves the estimate in the opposite direction by 0.16 percentage points.

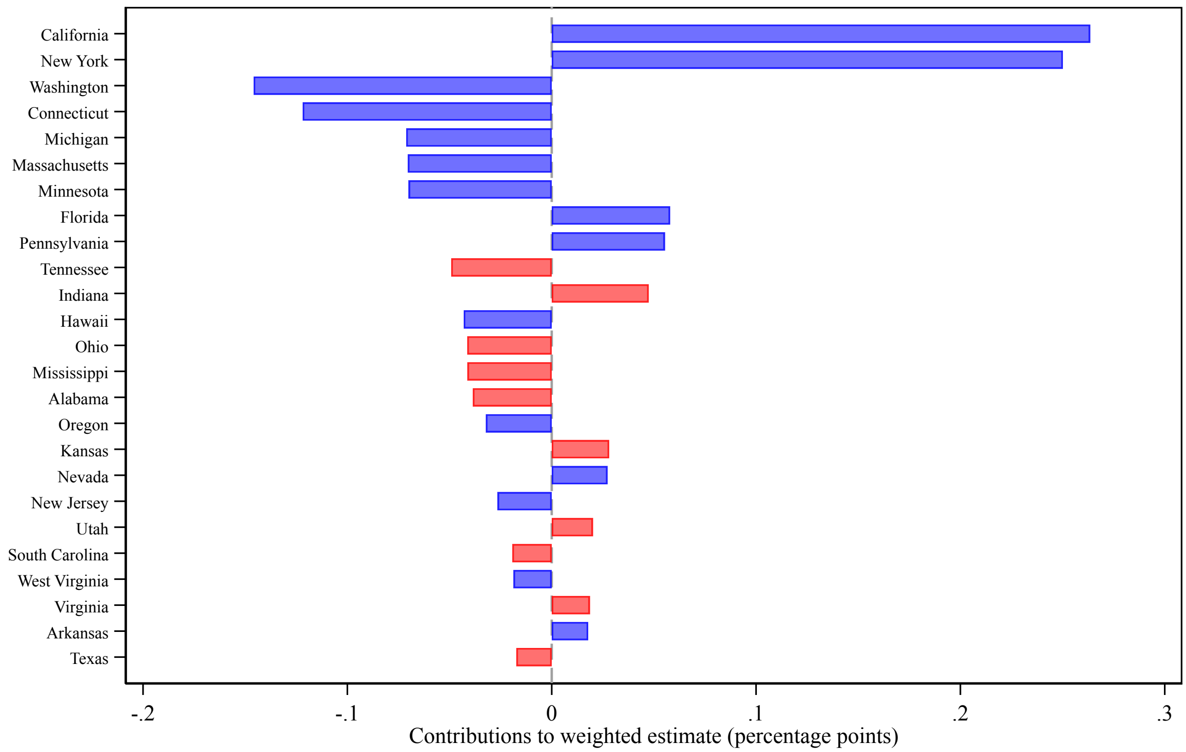

Figure 7: State-Level Contributions to the Weighted Employment-to-Population Estimate for Restaurants, Full Sample: Ever-Treated States (blue), Always-Control States (red)

in Solon et al. (2015), is the failure to model heterogeneous effects. This section points to one likely source of heterogeneity in minimum wage effects. California and New York are the 1st and 4th most populous U.S. states, so they carry large weights in the weighted estimation. And they also have large urban areas, where minimum wages are likely to bind the least.[21] Using city-level data, for example, Jha et al. (2026) show that minimum wage effects in the restaurant industry are negative in most cities, but closer to zero in larger places. Although removing both states aligns weighted and unweighted coefficients, the correct next step is "to study the heterogeneity, not just try to average it out" (Solon et al., 2015).

5 Conclusion

Modern difference-in-differences event-study designs are the latest tool in the arsenal of applied economists studying the effects of minimum wages on employment. They are still subject, however, to discretionary decisions of the researcher concerning the definition of events, variables used, and

[21] According to Census data, the two most populous urban areas in the U.S. are New York–Jersey City–Newark (NY–NJ) and Los Angeles–Long Beach–Anaheim (CA). The Census also reports that 94.2% of California's population is urban (the largest share in the country).

what results to present (e.g., weighted and/or unweighted). In this paper, we show that two of the most prominent minimum wage studies using these designs yield null results due to a set of decisions that tend to favor that result.

CDLZ is among the most influential minimum wage papers of the last decade. Its impact can be explained along two dimensions. First, CDLZ asks a novel question in the minimum wage literature: rather than following the common approach of asking how minimum wages affect employment among a low-skilled group, such as teens or the restaurant industry, CDLZ estimates job changes across segments of the wage distribution in the entire economy. Second, CDLZ introduces a stacked event-study design with desirable econometric properties—such as avoiding negative-weighting bias—that can be widely applied in economics research (Dube, Girardi, Jordà and Taylor, 2025). We show, however, that their perfect-reallocation story across wage bins—and the overall null effect of minimum wages on employment that follows from it—is fragile, an artifact of their choice of dependent variable: the employment-to-population ratio with a varying-population denominator.

We use instead an employment-to-population ratio with a constant-population denominator, as well as log employment, and obtain negative and significant overall minimum wage effects. The reason is simple: population is negatively related to the minimum wage. With population also declining after a minimum wage event, using the varying-population ratio as the dependent variable yields upward-biased estimates of the employment effects.

Moreover, we show that minimum wage effects extend beyond the narrow $[-4, 5)$ wage distribution segment (around the new minimum wage) that CDLZ consider, obtaining negative and significant effects for employment more than \$4 below and more than \$5 above the new minimum wage. From CDLZ's view, their finding of negative employment effects for higher wages when using TWFE can only be the result of a spurious relationship, with two of its authors noting "... beyond a narrow range above the new minimum, it [the minimum wage] is unlikely to have a major causal impact on the upper tail of the frequency distribution. Therefore the lack of such upper-tail effects serves as a useful check on the research design" (Dube and Lindner, 2024, p. 293). But this limited view fully ignores potential minimum wage effects on entry and exit of firms. If minimum wages induce firm exit (as documented, for example, in Luca and Luca, 2019, Aaronson, French, Sorkin and To, 2018, Chava, Oettl and Singh, 2019, and Jha and Rodriguez-Lopez, 2021) or reduce firm entry, the employment effects extend along the entire wage distribution—higher-wage workers, of course, work alongside minimum-wage workers in the same firms.

In addition, using CDLZ's exact stacked design, events, and sample period, we find that restaurant disemployment is universal across specifications and samples—even in the varying-population specifications—with average post-period elasticities and own-wage elasticities in the $[-0.19, -0.09]$ and $[-1.44, -0.57]$ ranges, respectively. This contrasts sharply with the pooled restaurant OWE of -0.03 headlined by DL and with the "indistinguishable from zero" framing of that literature. DL's finding, however, is itself fragile: a remarkably modest change—switching from population-weighted to unweighted estimation—moves the OWE to a significant -0.38. We then show that California and New York drive their weighted null; removing both states yields negative and significant OWEs of -0.43 (weighted) and -0.51 (unweighted).

The authors of these studies have argued that their evidence shows that using clean event studies is necessary to estimate the employment effects of minimum wages, and that doing so demonstrates that minimum wage increases do not reduce employment. However, their evidence is very fragile. More defensible decisions and analyses using modern difference-in-differences yield robust evidence of the same old result: minimum wages reduce jobs.

References

Aaronson, Daniel, Eric French, Isaac Sorkin, and Ted To, "Industry Dynamics and the Minimum Wage: A Putty-Clay Approach," *International Economic Review*, 2018, *59* (1), 51–84.

Allegretto, Sylvia, Arindrajit Dube, and Michael Reich, "Do Minimum Wages Really Reduce Teen Employment? Accounting for Heterogeneity and Selectivity in State Panel Data," *Industrial Relations: A Journal of Economy and Society*, 2011, *50* (2), 205–240.

Brown, Charles, Curtis Gilroy, and Andrew Kohen, "The Effect of the Minimum Wage on Employment and Unemployment," *Journal of Economic Literature*, 1982, *20* (2), 487–528.

Card, David, "Using Regional Variation in Wages to Measure the Effects of the Federal Minimum Wage," *Industrial and Labor Relations Review*, 1992, *46* (1), 22–37.

_ , "Do Minimum Wages Reduce Employment? A Case Study of California, 1987–89," *Industrial and Labor Relations Review*, 1992, *46* (1), 38–54.

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

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

_ , _ , _ , and _ , "Replication Data for: 'The Effect of Minimum Wages on Low-Wage Jobs'," *Harvard Dataverse*, 2019, *https://doi.org/10.7910/DVN/TJCTC7.*

Chava, Sudheer, Alexander Oettl, and Manpreet Singh, "Does a One-Size-Fits-All Minimum Wage Cause Financial Stress for Small Businesses?," Working Paper 26523, National Bureau of Economic Research December 2019.

Douty, H. M., "Minimum Wage Regulation in the Seamless Hosiery Industry," *Southern Economic Journal*, 1941, *8* (2), 176–190.

Dube, Arindrajit and Attila Lindner, "Chapter 4 - Minimum Wages in the 21st Century," in Christian Dustmann and Thomas Lemieux, eds., *Handbook of Labor Economics*, Vol. 5, Elsevier, 2024, pp. 261–383.

_ , Daniele Girardi, Òscar Jordà, and Alan M. Taylor, "A Local Projections Approach to Difference-in-Differences," *Journal of Applied Econometrics*, 2025, *40* (7), 741–758.

_ , Michael Reich, Akash Bhatt, and Denis Sosinskiy, "Restaurant Employment, Minimum Wages, and Border Discontinuities," *Journal of Political Economy Microeconomics*, forthcoming.

_ , T. William Lester, and Michael Reich, "Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties," *Review of Economics and Statistics*, 2010, *92* (4), 945–964.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277. Themed Issue: Treatment Effect 1.

Jha, Priyaranjan and Antonio Rodriguez-Lopez, "Minimum Wage and Employer Variety," Technical Report 9312, CESifo Munich 2021.

— , **Jyotsana Kala, David Neumark, and Antonio Rodriguez-Lopez**, "JUE Insight: City Size, Monopsony, and the Employment Effects of Minimum Wages," *Journal of Urban Economics*, 2026, *154*, 103868.

Karabarbounis, Loukas, Jeremy Lise, and Anusha Nath, "Minimum Wages and Labor Markets in the Twin Cities," Working Paper 30239, National Bureau of Economic Research July 2022.

Lordan, Grace and David Neumark, "People versus Machines: The Impact of Minimum Wages on Automatable Jobs," *Labour Economics*, 2018, *52*, 40–53.

Luca, Dara Lee and Michael Luca, "Survival of the Fittest: The Impact of the Minimum Wage on Firm Exit," Working Paper 25806, National Bureau of Economic Research May 2019.

Meer, Jonathan and Jeremy West, "Effects of the Minimum Wage on Employment Dynamics," *Journal of Human Resources*, 2016, *51* (2), 500–522.

Neumark, David and William Wascher, "Employment Effects of Minimum and Subminimum Wages: Panel Data on State Minimum Wage Laws," *ILR Review*, 1992, *46* (1), 55–81.

Solon, Gary, Steven J. Haider, and Jeffrey M. Wooldridge, "What Are We Weighting For?," *Journal of Human Resources*, 2015, *50* (2), 301–316.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199. Themed Issue: Treatment Effect 1.

Vaghul, Kavya and Ben Zipperer, "Historical State and Sub-State Minimum Wage Data," *Washington Center for Equitable Growth*, 2016.

Appendix

A Supporting Figures and Tables

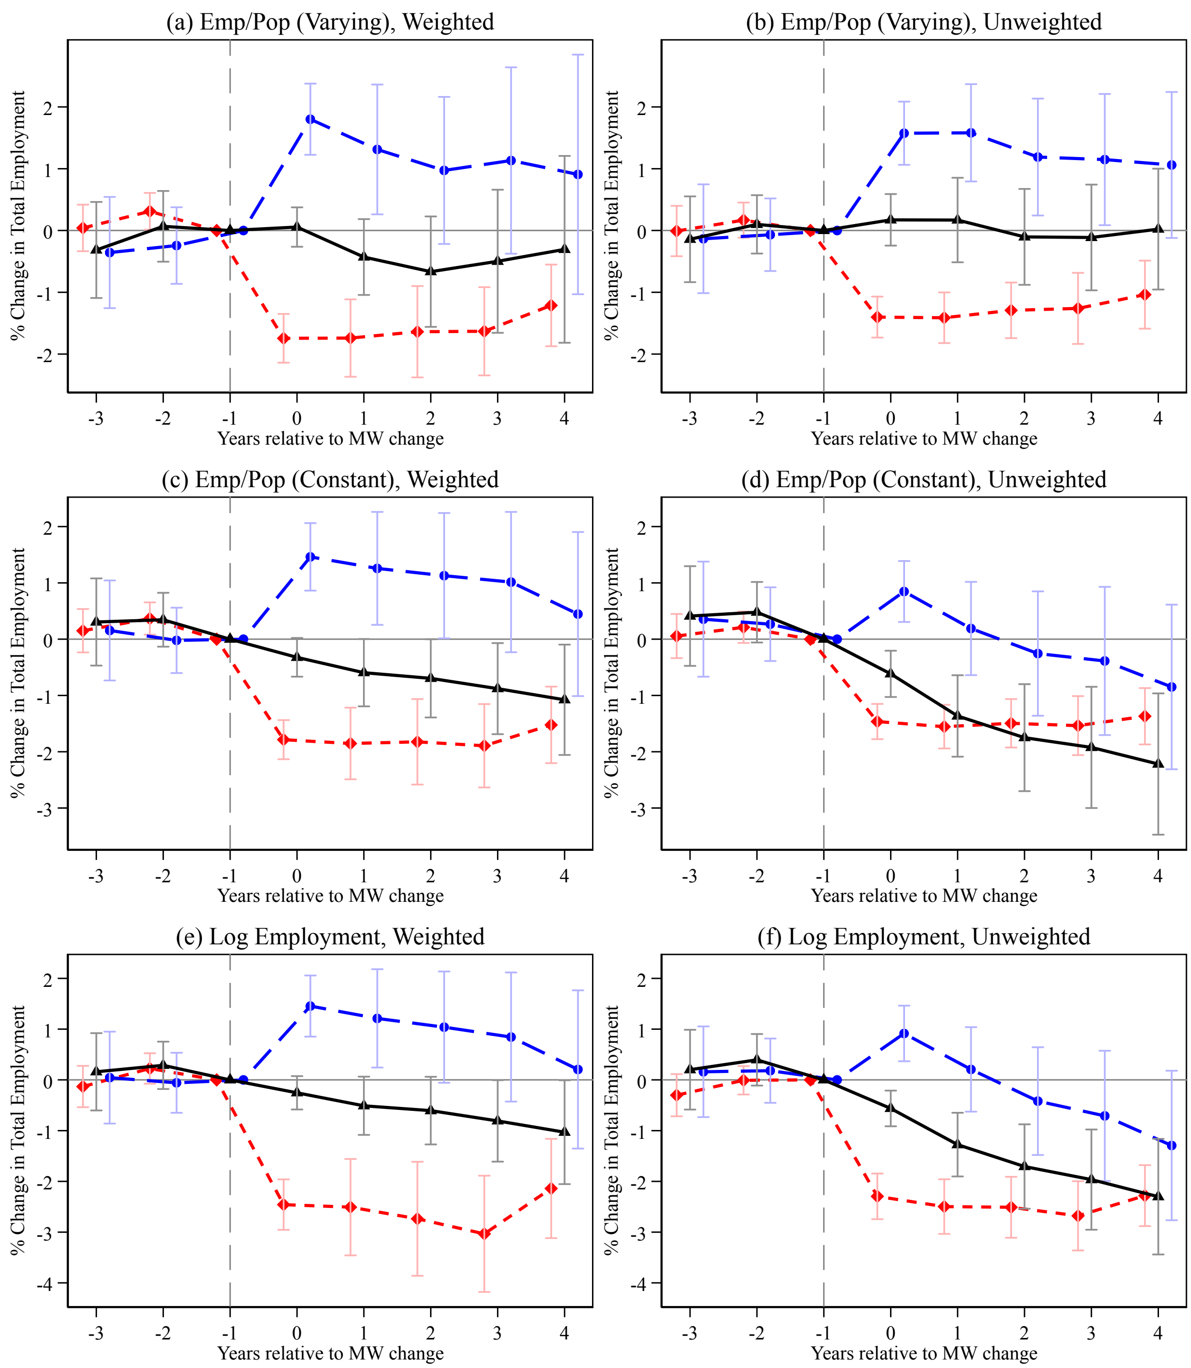

Figure A1: Event-Study Paths for % Change in Total Employment with 95% CIs (Pre-2012 Events): Below (dash, ◆), Above (long dash, ●), and Total (solid, ▲)

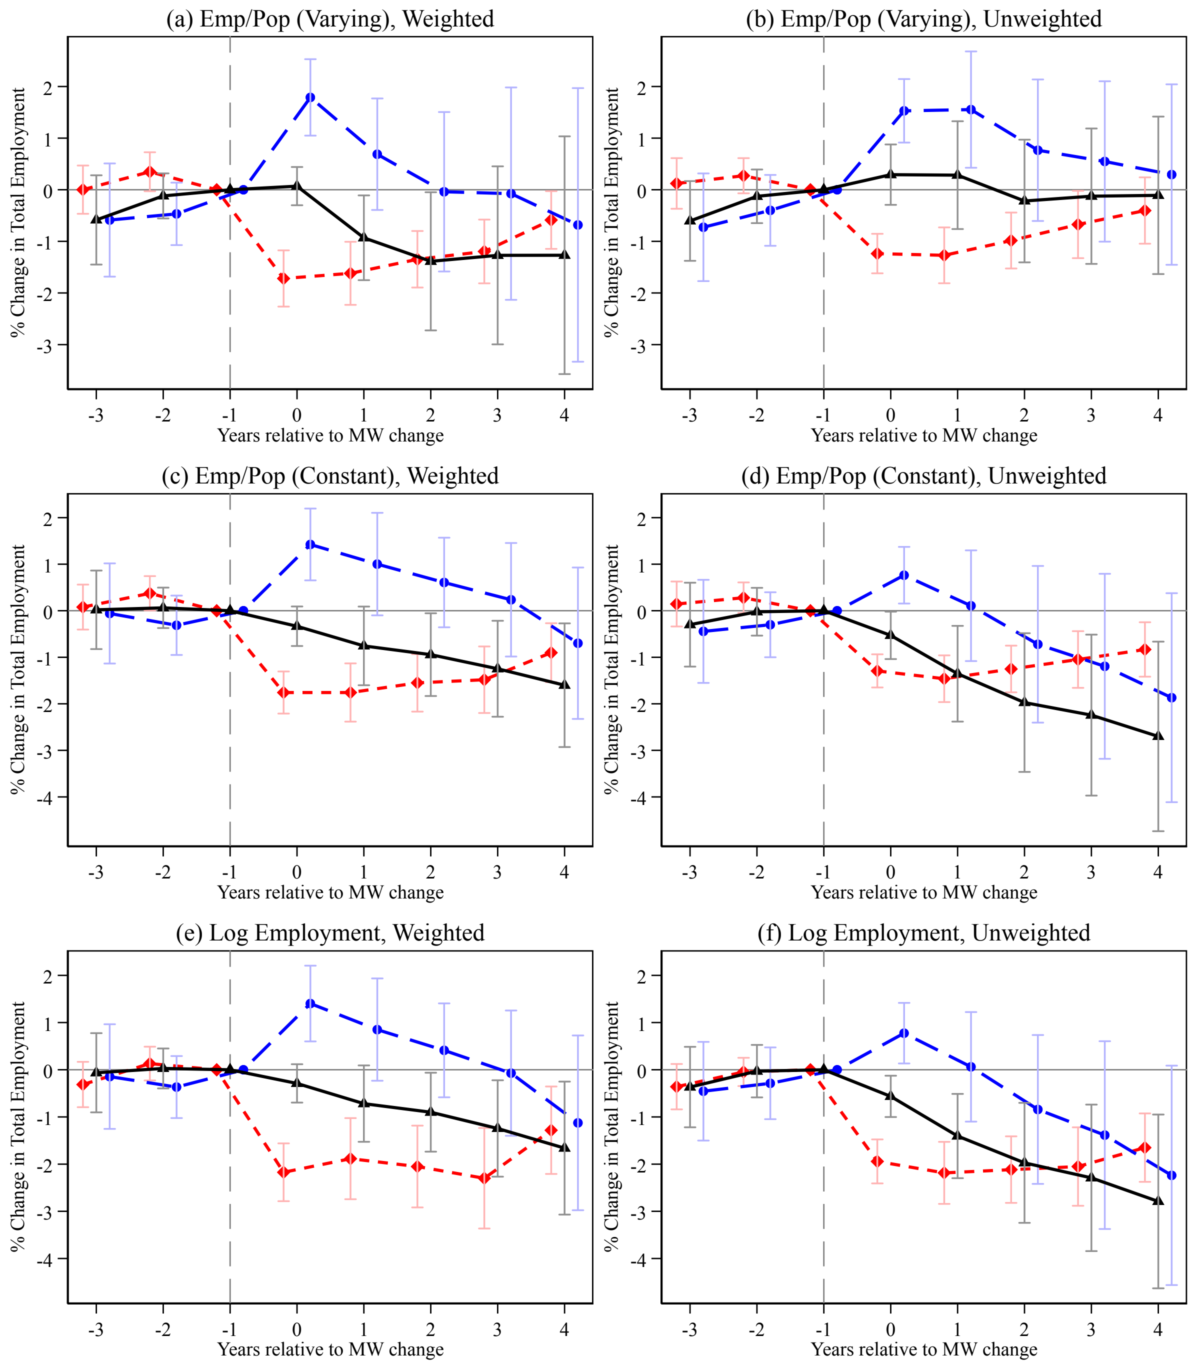

Figure A2: Event-Study Paths for % Change in Total Employment with 95% CIs (Pre-2012 Bundled Events): Below (dash, ♦), Above (long dash, •), and Total (solid, ▲)

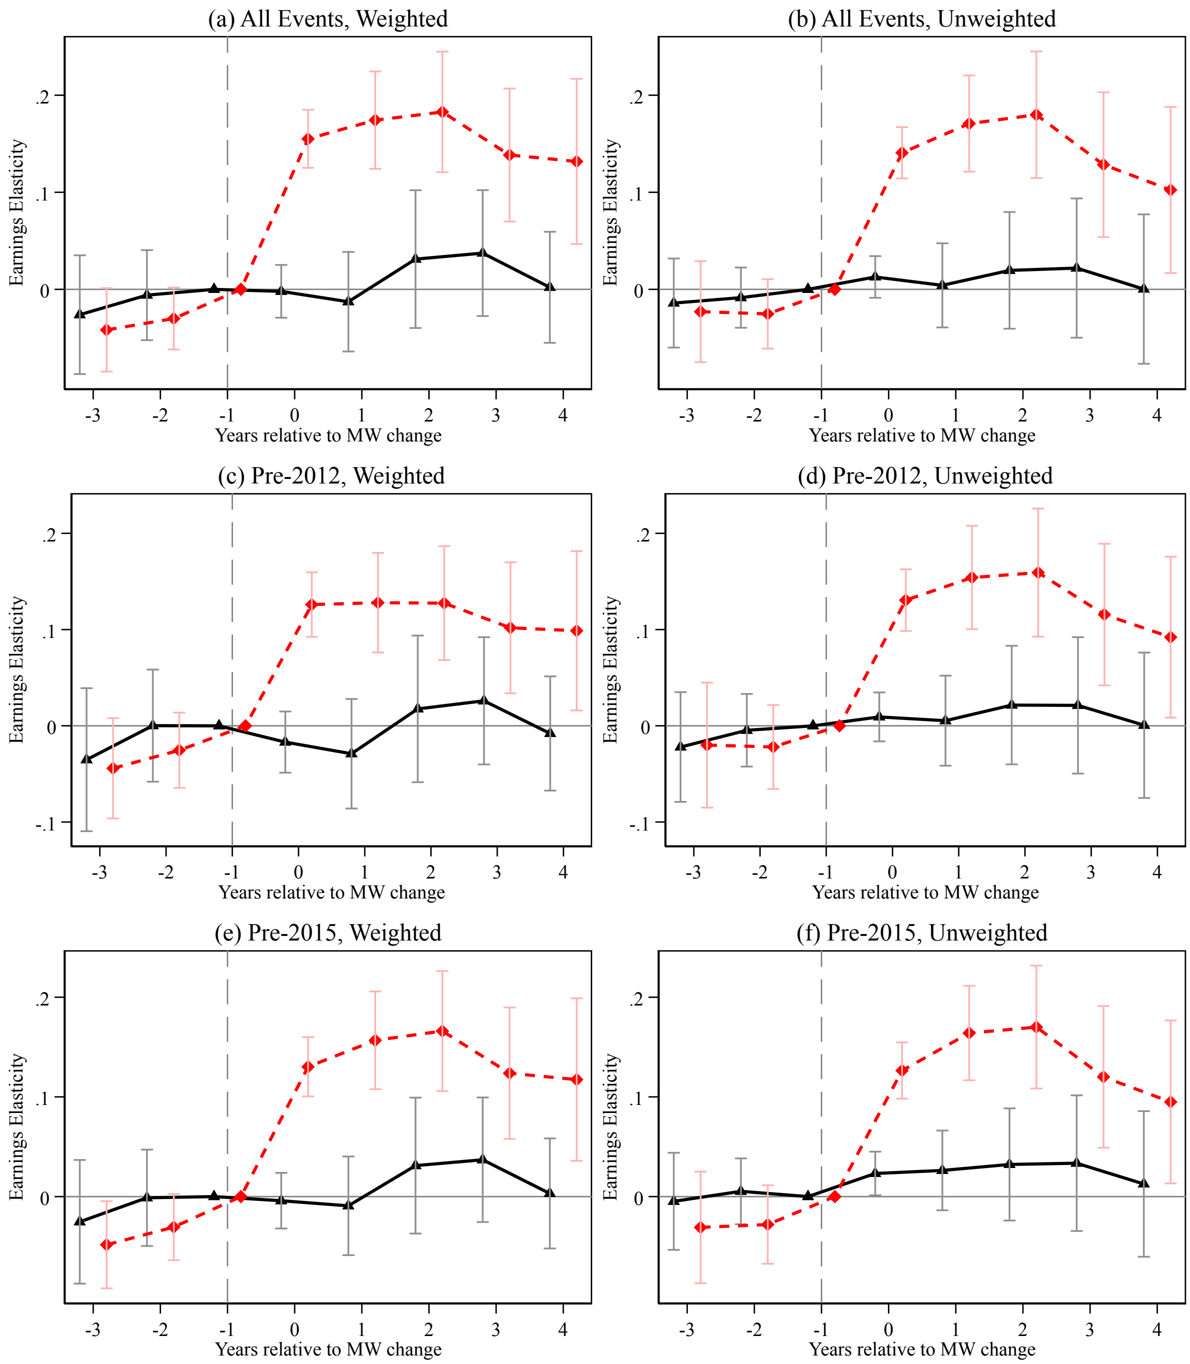

Figure A3: Event-Study Paths for the Estimated Minimum Wage Elasticity of Average Earnings with 95% CIs: Total (solid, ▲), Restaurant (dash, ♦)

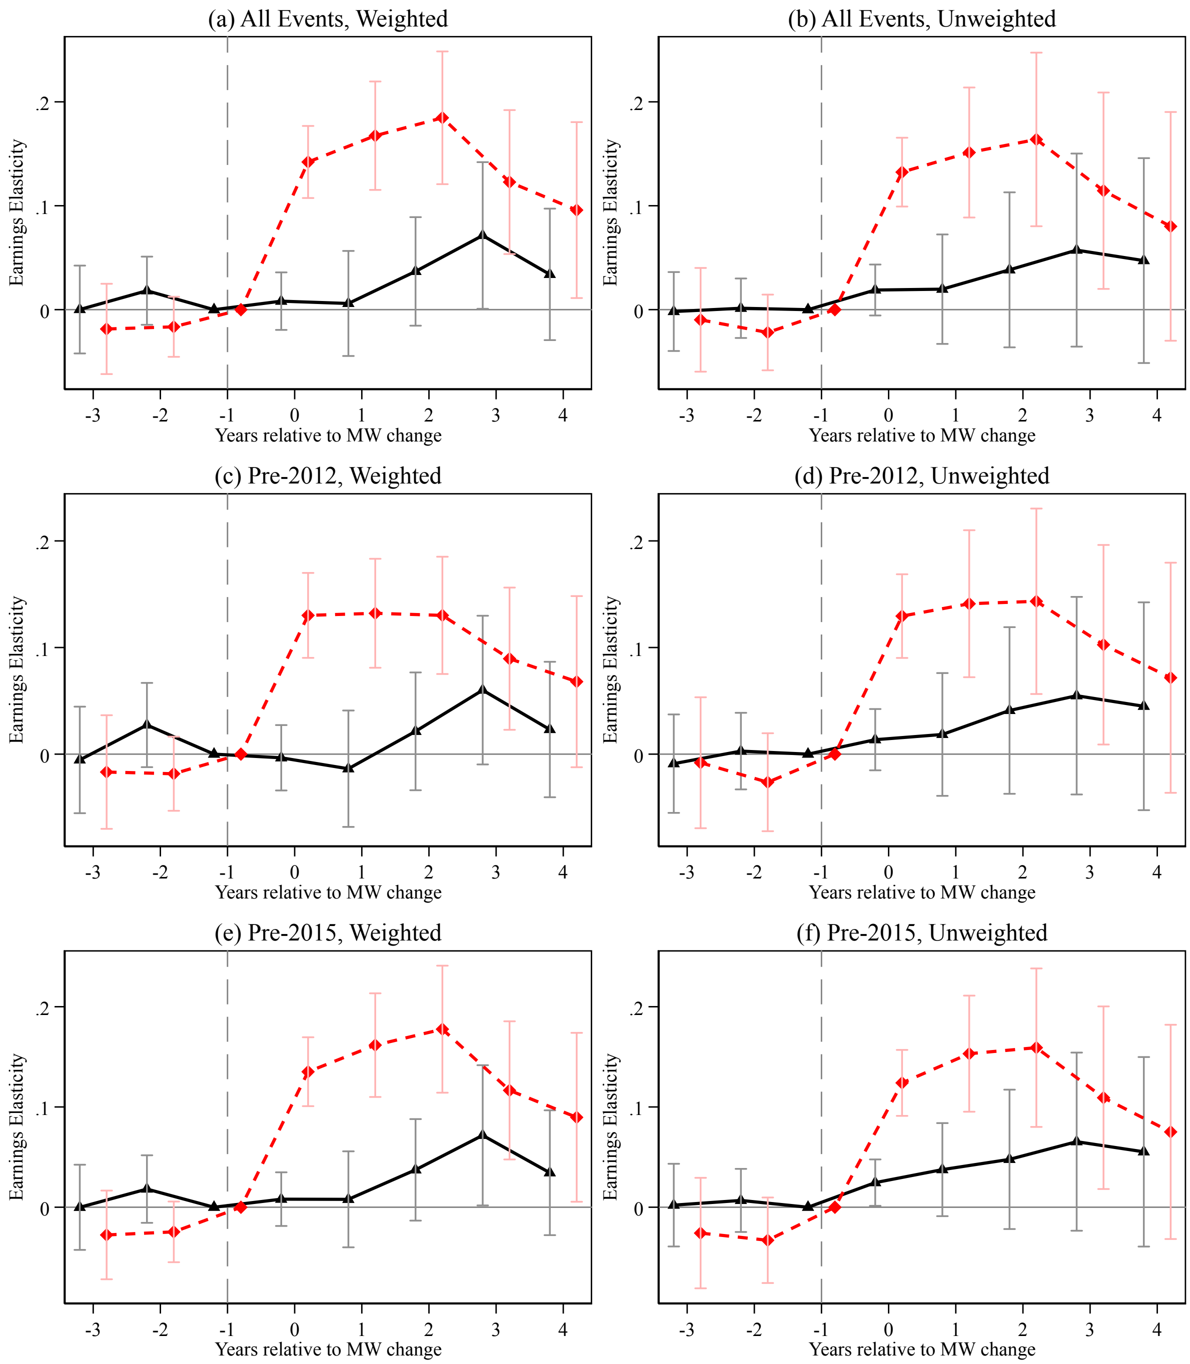

Figure A4: Event-Study Paths for the Estimated Minimum Wage Elasticity of Average Earnings with 95% CIs (Bundled Events): Total (solid, ▲), Restaurant (dash, ♦)

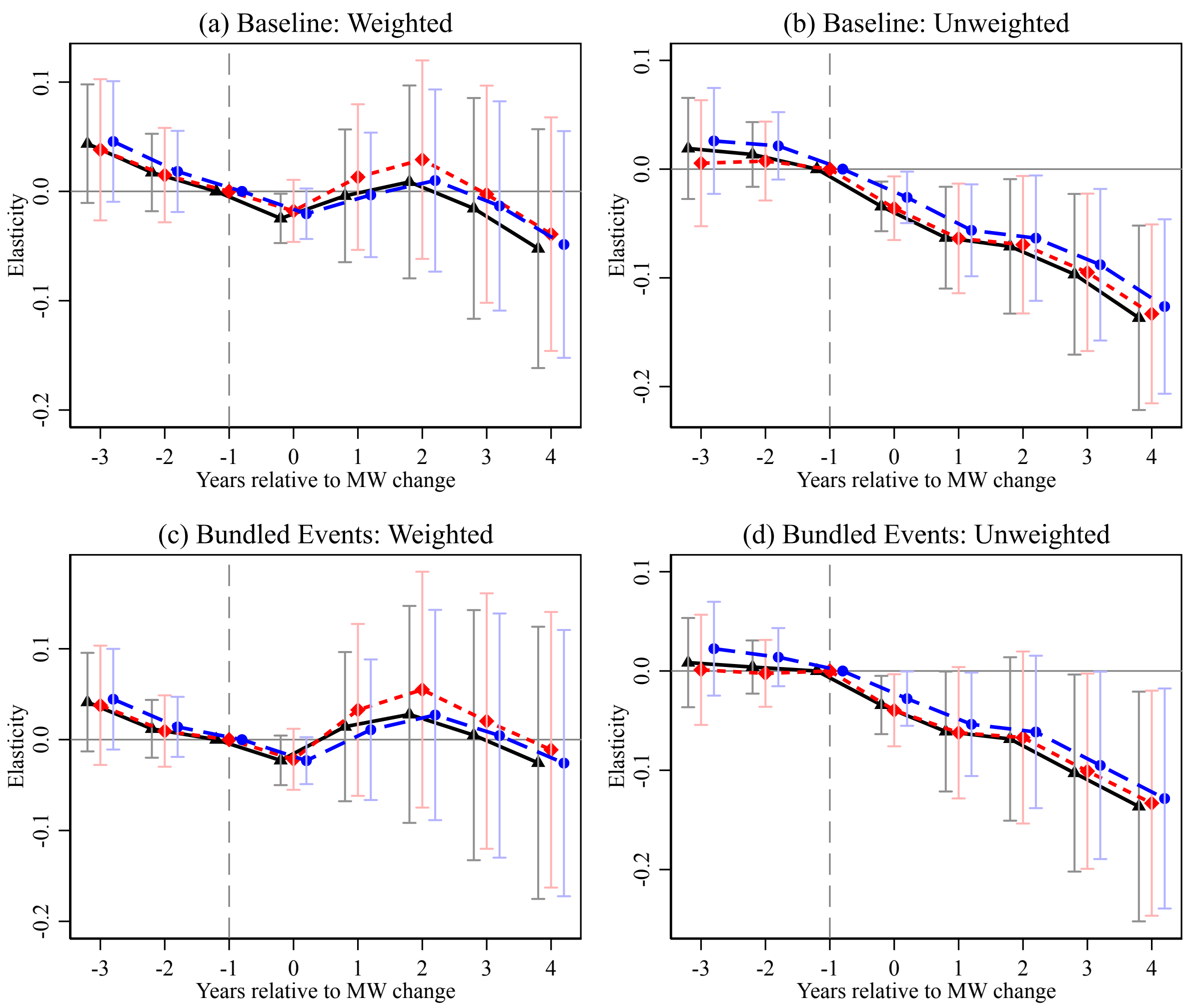

Figure A5: Event-Study Paths for the Estimated Minimum Wage Elasticity of Population with 95% CIs: All Events (solid, ▲), Pre-2012 (dash, ♦), Pre-2015 (long dash, •)

Table A1: Denominator Components for Estimated Percentage Employment Changes ($\%\widehat{\Delta\text{Emp}}$) and Elasticities ($\widehat{\varepsilon}$)

	Emp/Pop Varying		**Emp/Pop Constant**		**Log Employment**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Affected (M+E), All Events	0.573	0.623	0.615	0.670	—	—
	10.44	10.11	10.44	10.11	10.44	10.11
Affected (M+E), Pre-2012	0.576	0.620	0.602	0.646	—	—
	10.95	10.68	10.95	10.68	10.95	10.68
Total, All Events	0.573	0.623	0.615	0.670	—	—
	10.44	10.11	10.44	10.11	10.44	10.11
Total, Pre-2012	0.576	0.620	0.602	0.646	—	—
	10.95	10.68	10.95	10.68	10.95	10.68
Total, Pre-2015	0.573	0.617	0.608	0.655	—	—
	10.89	10.53	10.89	10.53	10.89	10.53
Restaurant, All Events	0.039	0.043	0.042	0.046	—	—
	10.44	10.11	10.44	10.11	10.44	10.11
Restaurant, Pre-2012	0.038	0.041	0.040	0.043	—	—
	10.95	10.68	10.95	10.68	10.95	10.68
Restaurant, Pre-2015	0.039	0.042	0.041	0.045	—	—
	10.89	10.53	10.89	10.53	10.89	10.53

Notes: Each cell shows two values used in the calculation of the percentage employment changes and elasticities in Tables 1, 3, and 5: the first line is the baseline employment-to-population ratio ($\overline{EPOP}_{-1}$, replaced by $\overline{EPOP}^{R}_{-1}$ in the Restaurant rows), and the second line is $\overline{\%\Delta MW}$ expressed in percent. The first line is "—" in the log-employment columns, where $\overline{EPOP}_{-1}$ does not enter the percentage-change or elasticity formulas. The top header indicates the dependent variable: employment to varying population, employment to constant population, and log employment. The second header indicates whether the estimation is weighted (by constant state population) or unweighted. Rows indicate the employment group (CDLZ's "Affected", Total, and Restaurant) and sample used: All Events (138 events), Pre-2012 (98 events), and Pre-2015 (119 events).

Table A2: Estimated Minimum Wage Elasticities of Employment (Bundled Events)

	Emp/Pop Varying		**Emp/Pop Constant**		**Log Employment**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Panel A: Average Post-Period Elasticities (τ=0 through τ=4)						
Affected (M+E), All Events	0.002	0.012	-0.013	-0.023	-0.031	-0.127
	(0.027)	(0.026)	(0.031)	(0.027)	(0.095)	(0.093)
Affected (M+E), Pre-2012	0.006	0.016	-0.004	-0.019	0.014	-0.105
	(0.030)	(0.027)	(0.035)	(0.028)	(0.105)	(0.096)
Total, All Events	-0.055	0.009	-0.073**	-0.132**	-0.068**	-0.137***
	(0.048)	(0.041)	(0.029)	(0.052)	(0.029)	(0.047)
Total, Pre-2012	-0.077	0.002	-0.078**	-0.144***	-0.077**	-0.148***
	(0.048)	(0.044)	(0.031)	(0.055)	(0.030)	(0.048)
Total, Pre-2015	-0.053	0.011	-0.072**	-0.117**	-0.067**	-0.123***
	(0.047)	(0.038)	(0.028)	(0.048)	(0.028)	(0.043)
Restaurant, All Events	-0.128**	-0.114***	-0.153***	-0.189***	-0.126***	-0.198***
	(0.056)	(0.039)	(0.049)	(0.045)	(0.043)	(0.047)
Restaurant, Pre-2012	-0.153***	-0.119***	-0.153***	-0.197***	-0.140***	-0.203***
	(0.059)	(0.042)	(0.053)	(0.048)	(0.045)	(0.048)
Restaurant, Pre-2015	-0.127**	-0.117***	-0.153***	-0.186***	-0.126***	-0.191***
	(0.055)	(0.036)	(0.048)	(0.044)	(0.042)	(0.045)
Panel B: Long-Term Elasticities (τ=4)						
Affected (M+E), All Events	-0.007	0.003	-0.041	-0.058	-0.089	-0.218*
	(0.038)	(0.038)	(0.042)	(0.042)	(0.117)	(0.131)
Affected (M+E), Pre-2012	-0.003	0.008	-0.032	-0.053	-0.046	-0.192
	(0.039)	(0.037)	(0.045)	(0.042)	(0.124)	(0.129)
Total, All Events	-0.084	-0.001	-0.125**	-0.213**	-0.128**	-0.224***
	(0.098)	(0.064)	(0.054)	(0.085)	(0.059)	(0.078)
Total, Pre-2012	-0.102	-0.009	-0.129**	-0.221***	-0.133**	-0.229***
	(0.094)	(0.064)	(0.054)	(0.085)	(0.058)	(0.077)
Total, Pre-2015	-0.082	0.001	-0.122**	-0.198**	-0.126**	-0.209***
	(0.097)	(0.061)	(0.054)	(0.081)	(0.058)	(0.075)
Restaurant, All Events	-0.179*	-0.173***	-0.233***	-0.287***	-0.210**	-0.316***
	(0.101)	(0.063)	(0.083)	(0.076)	(0.082)	(0.082)
Restaurant, Pre-2012	-0.198**	-0.172***	-0.233***	-0.292***	-0.216***	-0.311***
	(0.101)	(0.064)	(0.087)	(0.078)	(0.082)	(0.080)
Restaurant, Pre-2015	-0.177*	-0.173***	-0.230***	-0.281***	-0.209**	-0.307***
	(0.099)	(0.061)	(0.082)	(0.075)	(0.081)	(0.080)

Notes: Each cell reports the estimated elasticity of employment with respect to the minimum wage, using the bundled-events samples that retain only the first event per state and any subsequent event occurring at least 20 quarters after the previous retained event in that state. The top header indicates the dependent variable: employment to varying population, employment to constant population, and log employment. The second header indicates whether the estimation of equation (1) is weighted (by constant state population) or unweighted. Rows in each panel indicate the sector (CDLZ's "Affected" bins, Total, and Restaurants) and sample used: All Events (80 bundled events), Pre-2012 (60 bundled events), and Pre-2015 (79 bundled events). The Pre-2015 sample extends through 2019Q4. Standard errors (in parentheses) are clustered at the event-state level. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table A3: Denominator Components for Estimated Percentage Employment Changes ($\%\widehat{\Delta\text{Emp}}$) and Elasticities ($\widehat{\varepsilon}$) — Bundled Events

	Emp/Pop Varying		**Emp/Pop Constant**		**Log Employment**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Affected (M+E), All Events	0.570	0.615	0.607	0.661	—	—
	12.03	11.56	12.03	11.56	12.03	11.56
Affected (M+E), Pre-2012	0.574	0.617	0.597	0.646	—	—
	12.44	12.21	12.44	12.21	12.44	12.21
Total, All Events	0.570	0.615	0.607	0.661	—	—
	12.03	11.56	12.03	11.56	12.03	11.56
Total, Pre-2012	0.574	0.617	0.597	0.646	—	—
	12.44	12.21	12.44	12.21	12.44	12.21
Total, Pre-2015	0.570	0.616	0.606	0.661	—	—
	12.09	11.67	12.09	11.67	12.09	11.67
Restaurant, All Events	0.039	0.042	0.041	0.045	—	—
	12.03	11.56	12.03	11.56	12.03	11.56
Restaurant, Pre-2012	0.038	0.041	0.040	0.043	—	—
	12.44	12.21	12.44	12.21	12.44	12.21
Restaurant, Pre-2015	0.038	0.042	0.041	0.045	—	—
	12.09	11.67	12.09	11.67	12.09	11.67

Notes: Each cell shows two values used in the calculation of the percentage employment changes and elasticities in the bundled-events counterparts of Tables 1, 3, and 5: the first line is the baseline employment-to-population ratio ($\overline{EPOP}_{-1}$, replaced by $\overline{EPOP}^{R}_{-1}$ in the Restaurant rows), and the second line is $\overline{\%\Delta MW}$ expressed in percent. The first line is "—" in the log-employment columns, where $\overline{EPOP}_{-1}$ does not enter the percentage-change or elasticity formulas. The top header indicates the dependent variable: employment to varying population, employment to constant population, and log employment. The second header indicates whether the estimation is weighted (by constant state population) or unweighted. Rows indicate the employment group (CDLZ's "Affected", Total, and Restaurant) and sample used: All Bundled Events (80 events), Pre-2012 Bundled (60 events), and Pre-2015 Bundled (79 events).

Table A4: Own-Wage Elasticities of Employment (Bundled Events)

	Emp/Pop Varying		**Emp/Pop Constant**		**Log Employment**	
	Weighted	*Unweighted*	*Weighted*	*Unweighted*	*Weighted*	*Unweighted*
Panel A: Average Post-Period OWEs (τ=0 through τ=4)						
Total, All Events	-1.759	0.235	-2.329	-3.632	-2.176	-3.788
	(1.973)	(1.160)	(1.879)	(3.577)	(1.781)	(3.654)
Total, Pre-2012	-4.399	0.059	-4.485	-4.173	-4.422	-4.280
	(6.370)	(1.272)	(6.124)	(4.402)	(6.031)	(4.436)
Total, Pre-2015	-1.672	0.233	-2.246	-2.543	-2.109	-2.660
	(1.843)	(0.831)	(1.734)	(1.977)	(1.651)	(1.992)
Restaurant, All Events	-0.899**	-0.890**	-1.077***	-1.472***	-0.885***	-1.544***
	(0.425)	(0.390)	(0.391)	(0.534)	(0.343)	(0.561)
Restaurant, Pre-2012	-1.395**	-1.008**	-1.388**	-1.672**	-1.271**	-1.729**
	(0.617)	(0.476)	(0.568)	(0.660)	(0.499)	(0.675)
Restaurant, Pre-2015	-0.932**	-0.941**	-1.123***	-1.498***	-0.923***	-1.540***
	(0.436)	(0.387)	(0.408)	(0.533)	(0.354)	(0.547)
Panel B: Long-Term OWEs (τ=4)						
Total, All Events	-2.474	-0.029	-3.671	-4.507	-3.767	-4.750
	(3.722)	(1.356)	(3.844)	(5.126)	(3.986)	(5.326)
Total, Pre-2012	-4.416	-0.198	-5.571	-4.926	-5.781	-5.090
	(7.429)	(1.436)	(8.171)	(5.773)	(8.495)	(5.890)
Total, Pre-2015	-2.389	0.012	-3.560	-3.581	-3.668	-3.786
	(3.573)	(1.100)	(3.642)	(3.454)	(3.792)	(3.574)
Restaurant, All Events	-1.865	-2.155	-2.428*	-3.578	-2.196*	-3.949
	(1.345)	(1.701)	(1.397)	(2.683)	(1.311)	(2.951)
Restaurant, Pre-2012	-2.917	-2.400	-3.424	-4.075	-3.185	-4.344
	(2.301)	(2.048)	(2.430)	(3.316)	(2.265)	(3.519)
Restaurant, Pre-2015	-1.973	-2.308	-2.569*	-3.734	-2.327	-4.086
	(1.455)	(1.859)	(1.536)	(2.888)	(1.437)	(3.149)

Notes: Each cell reports the estimated own-wage elasticity of employment (OWE), defined as the ratio of the estimated minimum wage elasticity of employment to the minimum wage elasticity of average earnings, using the bundled-events samples that retain only the first event per state and any subsequent event occurring at least 20 quarters after the previous retained event in that state. The top header indicates the dependent variable: employment to varying population, employment to constant population, and log employment. The second header indicates whether the estimation of equation (1) is weighted (by constant state population) or unweighted. Rows in each panel indicate the sector (Total or Restaurants) and sample used: All Events (80 bundled events), Pre-2012 (60 bundled events), and Pre-2015 (79 bundled events). The Pre-2015 sample extends through 2019Q4. Standard errors (in parentheses) are clustered at the event-state level. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table A5: Leave-One-State-Out Sensitivity of the Weighted Employment-to-Population Estimate for Restaurants

	$\widehat{\alpha}^{w}_{\text{post}} \times 100$	
State dropped	Estimate	Shift
California	−0.536	−0.450
New York	−0.413	−0.327
Washington	+0.069	+0.155
Connecticut	+0.039	+0.125
Michigan	−0.012	+0.074
Massachusetts	−0.013	+0.073
Minnesota	−0.015	+0.071
Florida	−0.154	−0.068
Pennsylvania	−0.153	−0.067
Hawaii	−0.043	+0.043
None (full sample)	−0.086	

Notes: Each row drops one state from the full DL sample and re-estimates equation (9) with the restaurant employment-to-population ratio as outcome. "Estimate" is the value after the state is dropped, and "Shift" is the change from the full-sample baseline in the last row. Both are in percentage points ($\widehat{\alpha}^{w}_{\text{post}} \times 100$). The ten states shown are those with the largest absolute shifts.