

Effects of the Minimum Wage on Employment Dynamics

Jonathan Meer*
Texas A&M University
and NBER

Jeremy West
Massachusetts Institute of
Technology

Journal of Human Resources

August 2015

Abstract

The voluminous literature on minimum wages offers little consensus on the extent to which a wage floor impacts employment. We argue that the minimum wage will impact employment over time, through changes in growth rather than an immediate drop in relative employment levels. We show that commonly-used specifications in this literature, especially those that include state-specific time trends, will not accurately capture these effects. Using three separate state panels of administrative employment data, we find that the minimum wage reduces job growth over a period of several years. This finding is supported using several empirical specifications.

*Author emails are jmeer@tamu.edu and westj@mit.edu. We are grateful for valuable comments from Kerwin Charles and two anonymous referees, as well as from David Autor, Jeffrey Brown, Jeffrey Clemens, Jesse Cunha, Jennifer Doleac, David Figlio, Craig Garthwaite, Daniel Hamermesh, Mark Hoekstra, Scott Imberman, Joanna Lahey, Michael Lovenheim, Steven Puller, Harvey S. Rosen, Jared Rubin, Juan Carlos Saurez Serrato, Ivan Werning, William Gui Woolston, and numerous seminar participants. We benefited greatly from discussions regarding data with Ronald Davis, Bethany DeSalvo, and Jonathan Fisher at the U.S. Census Bureau, and Jean Roth at NBER. Sarah Armstrong and Kirk Reese provided invaluable research assistance. This work was supported in part by the W.E. UpJohn Institute for Employment Research. Any errors are our own.

1 Introduction

The question of how a minimum wage affects employment remains one of the most widely studied – and most controversial – topics in labor economics, with a corresponding dispute in the political sphere. Neoclassical economic theories present a clear prediction: as the price of labor increases, employers will demand less labor. However, many recent studies testing this prediction have found very small to no effects of the minimum wage on the level of employment (e.g. [Zavodny, 2000](#); [Dube et al., 2010](#); [Giuliano, 2013](#)). One possible explanation for these findings is that demand for low-wage labor is fairly inelastic; another is that more complicated dynamics cloud identification of the effect of the minimum wage on employment.¹

We argue that there is basis in theory for believing that the minimum wage may not reduce the level of employment in a discrete manner. We show that if this is indeed the case, then traditional approaches used in the literature are prone to misstating its true effects. We also demonstrate that a common practice in this literature – the inclusion of state-specific time trends as control variables – will attenuate estimates of how the minimum wage affects the employment level.

To implement our analysis, we use a number of different empirical approaches to examine effects of the minimum wage on employment growth and levels; broadly, all of our approaches leverage a difference-in-differences identification strategy using state panels. We perform numerous robustness checks to test the validity of our identification strategy, which requires that the pre-existing time-paths of outcomes for states which increase their minimum wages do not differ relative to states that do not see an increase. We evaluate this possibility by adding leads of the minimum wage into our specifications; if increases in the minimum wage showed a negative effect on employment *before* their implementation, this would suggest that the results are being driven by unobserved trends. This is not the case. Indeed, for our results to be driven by confounders, one would have to believe that increases in the minimum wage were systematically correlated with unobserved shocks to that state in the same time period, but not other states in that region, and that these shocks are not reflected in measures of state-specific demographics or business cycles. Our primary results are additionally robust to varying the specifications to account for finer spatial and temporal controls, the recent financial crisis, and inflation indexing of state minimum wages.

We use three administrative data sets in our analysis: the Business Dynamics Statis-

¹[Hirsch et al. \[2015\]](#) and [Schmitt \[2013\]](#) focus on other channels of adjustment in response to increases in the minimum wage, such as wage compression, reductions in hours worked, and investments in training.

tics (BDS), the Quarterly Census of Employment and Wages (QCEW), and the Quarterly Workforce Indicators (QWI). These data sets vary in their strengths and weaknesses, discussed at length below, but together they encompass a long (1975-2012) panel of aggregate employment metrics for the population of employers in the United States. Our findings are consistent across all three data sets, indicating that employment declines significantly in response to increases in the minimum wage over the span of several years.

If the minimum wage is to be evaluated alongside alternative policy instruments for increasing the standard of living of low-income households, a more conclusive understanding of its effects is necessary. The primary implication of our study is that the minimum wage does affect employment through a particular mechanism. This is important for normative analysis in theoretical models (e.g. [Lee and Saez, 2012](#)) and for policymakers weighing the tradeoffs between the increased wage for minimum wage earners and the potential reduction in hiring and employment. Moreover, we reconcile the tension between the expected theoretical effect of the minimum wage and the estimated null effect found by some researchers. We show that because minimum wages reduce employment levels through dynamic effects on employment growth, research designs incorporating state-specific time trends are prone to erroneously estimated null effects on employment.

This article proceeds as follows: in Section 2 we provide a brief review of the literature on the employment effects of the minimum wage and build the case for examining employment dynamics. We discuss the data in Section 3 and present our specifications and results in Section 4. We conclude in Section 5.

2 Theoretical and Econometric Framework

The economic literature on minimum wages is longstanding and vast. [Neumark and Wascher \[2008\]](#) provide an in-depth review of the field, which continues to be characterized by disagreement on how a minimum wage affects employment. The majority of recent studies, following [Card and Krueger \[1994\]](#), use difference-in-differences comparisons to evaluate the effect of these policies on employment levels. Recent papers generally focus on modifying the specification to improve the quality of the counterfactual comparisons, with disagreement on appropriate techniques and often-conflicting results (e.g. [Allegretto et al., 2011](#) and [Neumark et al., 2014](#)). Importantly, these models test whether there is a discrete change in the level of employment before and after a state changes its minimum wage, relative to the counterfactual change as measured by other states' employment.

Yet there is basis in theory for believing that the minimum wage may not reduce the level of employment in a discrete manner. While the basic analysis of the effects of the minimum wage argues for rapid adjustments to a new equilibrium employment level (e.g. [Stigler, 1946](#)), transitions to a new employment equilibrium may not be smooth [[Hamermesh, 1989](#)] or may be relatively slow ([Diamond, 1981](#); [Acemoglu, 2001](#)). In this case, the effects of the policy may be more evident in net job creation.² In worker search-and-matching models (e.g. [Van den Berg and Ridder, 1998](#); [Acemoglu, 2001](#); [Flinn, 2006, 2011](#)), summarized concisely in [Cahuc and Zylberberg \[2004\]](#), the minimum wage has opposing effects on job creation. Although it reduces demand for labor by raising the marginal cost of employing a new worker, a higher minimum wage increases the gap between the expected returns to employment relative to unemployment, inducing additional search effort from unemployed workers. By increasing the pool of searching workers (and the intensity of their searching), the minimum wage improves the quality of matches between employers and employees, generating surplus. The theory thus has ambiguous predictions for the effect of a minimum wage on job creation. If workers’ additional search effort sufficiently improves the worker-firm match quality, then job creation should not be adversely affected and may even increase. However, if the demand-side effect dominates, then increasing the minimum wage will cause declines in hiring.³

[Sorkin \[2013\]](#) builds a model that formalizes this potentially slow adjustment of labor demand, focusing on firms’ difficulties in adjusting their capital-labor ratios, and applies it to minimum wage increases. He argues that “the ability to adjust labor demand is limited in the short run” and that this “provide[s] an explanation for the small employment effects found in the minimum wage literature.” Fundamentally, this identification problem stems from the “sawtooth pattern” exhibited in states’ real minimum wages. Sorkin argues that “difference-in-difference faces challenges in measuring the treatment effect of interest, which in this case is the effect of a permanent minimum wage increase, whenever there are dynamic responses to the treatment and the treatment itself is time-varying.” [Baker et al. \[1999\]](#), using Canadian data, argue that “employment adjusts to long-run, evolutionary changes in the minimum wage,” necessitating a focus on dynamic effects; they further argue, similar to Sorkin’s

²Of course, any effect on growth does not exclude a discrete effect on the employment level. We separate these types of effects in the illustrations that follow to facilitate clearer exposition.

³With our reduced-form empirical analysis, we cannot distinguish the true mechanism driving the relationship between the minimum wage and employment. For instance, it is possible that the minimum wage would discretely affect employment, but that frictions in the labor market cause this effect to manifest over time. At a practical and policy-relevant level, these two situations are equivalent, and we are agnostic on the underlying mechanism, which limits our ability to make sweeping statements about the specific channel(s) through which the minimum wage impacts labor markets.

model, that “employment variation could result from the substitution of alternative factors of production with longer planning horizons,” such as capital. In a similar vein, [Neumark and Wascher \[1992\]](#) discuss the importance of including lagged effects of the minimum wage and find substantial impacts beyond a contemporaneous effect. To be clear, if the true effect a minimum wage is to change employment over time rather than to have an immediate effect on the employment level, then the traditional approaches used in this literature – namely, panel fixed effect estimates of the contemporaneous effects of the minimum wage on employment levels – will yield incorrect inference.⁴

2.1 Staggered Treatments and Difference-in-Differences

We illustrate the potential shortcoming of this classic difference-in-differences approach in Figure 1. This toy example depicts employment in two hypothetical jurisdictions, which initially exhibit identical growth rates. At some time t_1 , Jurisdiction A is treated; at some later time t_2 , Jurisdiction B is treated with the same intensity. In Panel (a), treatment has a *discrete* and symmetric negative effect on the employment level, whereas in Panel (b), the treatment has a symmetric negative effect on employment growth, but does not discretely alter the employment level. Consider the standard difference-in-differences (DiD) identification of the employment effect:

$$\text{Employment}_{it} = \delta_B \cdot \mathbf{I}\{\text{Jurisdiction} = \text{B}\} + \tau_t \cdot \mathbf{I}\{\text{Time} = t\} + \beta \cdot \mathbf{I}(\text{Treatment}_{it} = 1) + u_{it}$$

Because both jurisdictions are initially untreated and both are eventually treated, the only time period(s) in which the treatment effect β may be identified separately from the time fixed effects τ_t are those during which *only* Jurisdiction A is treated. During all other time periods, $\mathbf{I}(\text{Treatment}_{it} = 1)$ takes the same value for both states. Thus, the DiD model compares the average difference in employment between the jurisdictions during the time period between t_1 and t_2 to that in the time periods prior to t_1 and following t_2 .

This evaluation is obvious for the discrete employment effect in Panel (a). The difference between jurisdictions’ employment is clearly smaller during the middle time period, compared to the outer time periods, and the DiD estimate is correctly some negative number. Moreover, the duration of each of the three time periods is irrelevant for obtaining the correct inference.

⁴Several recent studies are exceptions to the focus on employment levels. [Dube et al. \[2011\]](#) examine the relationship between the minimum wage and employee turnover for teenagers and restaurant workers using a balanced panel subset (2001-2008) of the Quarterly Workforce Indicators (QWI). [Brochu and Green \[2013\]](#) assess firing, quit, and hiring rates in Canadian survey data. Both studies find a reduction in hiring rates but do not estimate the overall dynamic effect on employment levels.

If instead the treatment effect is on growth as in Panel (b), then DiD is very sensitive to the relative duration of each (outer) time period. As an example, consider first the extreme case in which there is a long pre-treatment timespan between times zero and t_1 , but a very short timespan between t_2 and T , the end of the sample period. In this situation, the average difference in employment during the outer time periods is determined nearly entirely by the pre-treatment period, and the DiD estimate for the treatment effect will be negative. Contrast this with the other extreme: a very short timespan between times zero and t_1 , but a long period following t_2 , during which both jurisdictions are treated. In this situation, the average difference in employment during the outer time periods is determined nearly entirely by the later period, and the DiD estimate for the effect of the *same* treatment will be positive. And, if T is selected such that the two outer periods have equivalent duration (i.e. $t_1 - 0 = T - t_2$), then DiD yields a zero treatment effect, visibly at odds with the plotted time paths of employment.

This toy example underscores the pitfalls in using a standard panel fixed effects model if treatments are staggered and the effects are on the growth of the outcome variable. As a state-level policy, the minimum wage clearly exhibits this type of staggered treatment: Figure 2 (along with Appendix [Online Appendix 4.](#)) shows that the effective minimum wage changed in at least one state in 33 of the 37 years from 1976 through 2012 – more than 700 changes in total – including every year after 1984.⁵ We investigate the implications of this concern more thoroughly in our Section 4. First, though, we discuss a separate but related concern.

2.2 Implications of Jurisdiction Time Trends as Controls

Many recent studies of the minimum wage include state- or county-specific time trends to control for heterogeneity in the underlying time-paths by which labor markets evolve within different areas that might be correlated with treatment intensity (e.g. [Page et al., 2005](#); [Addison et al., 2009](#); [Allegretto et al., 2011](#)). These models generally find little or no effect of the minimum wage on employment levels. However, if the policy change affects the growth

⁵Inflation is an additional consideration when evaluating the minimum wage as a policy treatment. Historically, minimum wages have been set in nominal dollars, with their value eroding substantially over time (see Appendix [Online Appendix 4.](#) for details). This means that the actual *intensity* of treatment changes over time, even in the absence of any subsequent (own or counterfactual) explicit policy change. This situation would not be problematic if the minimum wage affected employment in an abrupt, discrete manner. But if the minimum wage predominantly affects job creation, then it may take years to observe a statistically significant difference in the level of employment. In Section 4, we revisit the implications of inflation for minimum wage policy in the context of our empirical findings.

rate of the response variable, rather than its level, then specifications including jurisdiction-specific trends will mechanically attenuate estimates of the policy’s effect. The basic intuition is that including state-specific time trends as controls will adjust for two sources of variation. First, if there is any *pre-treatment* deviation in outcomes that is correlated with treatment – e.g. if states that exhibit stronger employment growth are also more likely to increase their minimum wage – then this confounding variation may be appropriately controlled for by including state-specific time trends. The potential cost of this added control is that if the actual treatment effect, the *post-treatment* employment variation, acts upon the trend itself, then inclusion of jurisdiction time trends will attenuate estimates of the treatment effect and often leads to estimating (statistical) null employment effects.⁶

A simple illustration of this is provided in Figures 3 and 4. Figure 3 depicts employment in two hypothetical jurisdictions which exhibit identical employment growth rates prior to period $t = 0$. After treatment begins in period $t = 0$, the employment growth rate in the Treated jurisdiction falls relative to the Control, and this occurs over the following four periods before leveling off after time $t = 4$. Figure 4 presents the difference in employment by time period for both levels and residuals to jurisdiction time trends. The computed employment effect is large and negative when state trends are omitted (in Panel (a)), but shrinks essentially to zero with the inclusion of jurisdiction time trends (Panel (b)). This sharp attenuation occurs despite identical pre-treatment employment trends.

We are by no means the first to make this point. In examining the effects of changes in divorce laws, Wolfers [2006] makes a general observation that a “a major difficulty in difference-in-differences analyses involves separating out trends from the dynamic effects of a policy shock.” Lee and Solon [2011] expound on this point in a discussion of Wolfers [2006], pointing out that “the sharpness of the identification strategy suffers” when jurisdiction-specific time trends are included and, “the shift in the dependent variable may vary with the length of time since the policy change.” This problem has been discussed in other contexts, including bias in estimates of the effects of desegregation (Baum-Snow and Lutz, 2011) and marijuana decriminalization (Williams, 2014).

However, this approach remains common in the minimum wage literature and, indeed, for many other important policy questions in which researchers ask “a much more nuanced question than just whether the dependent variable series showed a constant discrete shift at the moment of policy adoption” (Lee and Solon, 2011). We hope that our examples and simulations will serve as a useful guide to researchers considering how to approach estimation

⁶We are grateful to Cheng Cheng and Mark Hoekstra, as well as Justin Wolfers for this insight.

of policies whose effects may differ over time and, especially, may be reflected in changes in the growth rate of the variable of interest. We delve further into the question of how best to estimate these effects in Section 4.

3 Data

We estimate employment effects using three data sets: the *Business Dynamics Statistics* (BDS) and the *Quarterly Workforce Indicators* (QWI), both from the Bureau of the Census, and the *Quarterly Census of Employment and Wages* (QCEW) from the Bureau of Labor Statistics. The QCEW and QWI report quarterly employment for each state, while the BDS is annual. All of these data are administrative in nature; the QCEW and QWI programs collect data from county unemployment insurance commissions, while the BDS reports on employment rosters furnished to the U.S. Internal Revenue Service. As such, each of the data sets we study accounts for virtually the entire population of non-farm employment.⁷ For brevity and clarity of exposition, we report results from the BDS in the main body of the paper, with results from the full set of specifications using the QCEW and QWI in [Appendix Online Appendix 1.](#) As we note below, there is little difference in the overall results across the three data sets, which is unsurprising given that all three examine the near-population of jobs in the United States.

The BDS covers all non-agriculture private employer businesses in the U.S. that report payroll or income taxes to the IRS. The heart of the BDS is the Census Bureau’s internal Business Register, which is sourced from mandatory employer tax filings and augmented using the Economic Census and other data to compile annual linked establishment-level snapshots of employment statistics (on March 12th). The Census Bureau releases the BDS as a state-year panel (all fifty states, plus the District of Columbia), currently covering 1977 to 2011. Summary statistics from the BDS are provided in Table 1. Full descriptions of the QCEW and QWI, including their summary statistics, are located in [Appendix Online](#)

⁷The employer-sourced administrative nature of these data is important for our research question. Population-level data provide for a cleaner assessment of the overall policy impact of minimum wages by avoiding sampling error. Moreover, as discussed in Section 2, a higher minimum wage may induce additional searching effort on the part of the currently unemployed. [Mincer \[1976\]](#) shows that this positive supply elasticity often leads to an increase in the number of *unemployed* that differs substantially from the change in employment. Because employment is the policy-relevant outcome, measuring job counts using employer-sourced data provides a better identification of any disemployment effects than do surveys of individuals, such as the Current Population Survey. Finally, employment data directly reported by firms to maintain legal compliance have been shown to be more accurate than responses to individual-level surveys such as the CPS [\[Abraham et al., 2013\]](#).

3.1 State Minimum Wages

We draw historical data on state minimum wages from state-level sources.⁸ For the QCEW and QWI, we use the minimum wage value as of the first of each quarter. For the BDS, we use the value as of the previous March 12th each year, directly corresponding to the panel years in the BDS data. Some states have used a multiple-track minimum wage system, with a menu of wages that differ within a year across firms of different sizes or industries; we therefore use the maximum of the federal minimum wage and the set of possible state minimum wages for the year. To the extent that there is firm-level heterogeneity in the applicable wage level, our definition allows the minimum wage term to serve as an upper bound for the minimum wage a firm would actually face. We transform minimum wages into constant 2011 dollars using the (monthly) CPI-U from the Bureau of Labor Statistics.⁹

3.2 Other Control Variables

Although our econometric specifications include an extensive set of time period controls, precision may be gained by accounting for additional state-specific time-varying covariates. The Census Bureau’s Population Distribution Branch provides annual state-level population counts, including estimates for intercensal values. Total state population represents a determinant of both demand for (indirectly by way of demand for goods and services) and supply of employees. Because states differ non-linearly in their population changes, controlling directly for population may be important. The range in population between states and across time is enormous, so we use the natural log of state population in our specifications. We additionally include the share of this population aged 15-59, which provides a rough weight for how population might affect demand for versus supply of labor. Demographic controls such as these are commonly used in this literature (e.g. [Burkhauser et al., 2000](#); [Dube et al.,](#)

⁸Although historical state minimum wage data are available from sources such as the U.S Department of Labor (<http://www.dol.gov/whd/state/stateMinWageHis.htm>), these data suffer several limitations. For one, minimum wage values are only reported only as of January first each year, whereas the panel used in our study necessitates values as of other dates. Additionally these DOL data incompletely characterize changes to state minimum wages, especially during the early years of our panel. This DOL table is frequently used as the source of historical state minimum wage values for recent studies in this literature, and we caution future researchers to be careful not to inadvertently attribute minimum wage changes to years in which they did not occur. All data and code used in this study are available from the authors online or by request.

⁹Because we use a national-level deflator, specifying the log minimum wage term as real or nominal does not affect our results. Time period fixed effects incorporate this added variation.

2010). Following [Orrenius and Zavodny \[2008\]](#), we also include the natural log of real gross state product per capita.¹⁰ After controlling for state population, this term can be thought of as a rough proxy for average employee productivity as well as a measure of state-level fluctuations in business cycles [[Carlino and Voith, 1992](#), [Orrenius and Zavodny, 2008](#)].

4 Results

In [Section 2](#), we provide theoretical support and motivation for the hypothesis that the minimum wage affects employment over time rather than through a single, immediate drop in its level, and we illustrate several complications for attempts to quantify the magnitude of such an effect using standard approaches. In this section, we present several specifications which estimate this effect. As discussed above, we report results here using the Business Dynamics Statistics, with the same specifications using the QCEW and QWI in [Appendix Online Appendix 1.](#) There is little difference in the overall results across datasets, which is unsurprising given that all three examine the near-population of jobs in the United States. An illustrative example of the nature of these various estimates is provided in [Appendix A](#).

4.1 Standard panel fixed effect specifications

We begin with [Equation 1](#), the “classic” variant of the panel fixed effects specification in levels that has been used extensively in the literature. In this specification, emp_{it} is the level of employment in jurisdiction i at time t , α_i are jurisdiction fixed effects, τ_t are macroeconomic time period fixed effects, $\gamma_i \cdot t$ are jurisdiction-specific linear time trends, and ϵ_{it} is the idiosyncratic error term.

$$emp_{it} = \alpha_i + \tau_t + \gamma_i \cdot t + \beta_0 mw_{it} + \psi \cdot controls_{it} + \epsilon_{it} \quad (1)$$

If this model “correctly” fits the data – that is, if the effect of the minimum wage can be adequately captured via a single treatment coefficient within a panel fixed effects levels specification – then the coefficient $\hat{\beta}_0$ identifies the total causal impact of the minimum wage on employment. In light of the considerations discussed in [Section 2](#), it is questionable how well this assumption holds. We provide results using this specification in order to

¹⁰We compute the log of the real value of total GSP per capita using all industry codes, including government. Results are virtually unaffected by using $\ln(\text{real private sector GSP/capita})$ instead, but we view total GSP as the more appropriate definition given that the population term reflects total state population.

benchmark our findings within the literature, to alleviate concerns regarding differing pre-trends, to motivate that an effect could be on growth, and to underscore how dramatically the inclusion of jurisdiction time trends attenuates the classic estimates of the effect on employment.

In Table 2, we present results from estimating variations of this specification. Using national time fixed effects with no other time-specific controls in Column (1), there is a large negative effect of the minimum wage, with an elasticity of about -0.19, though the estimate is imprecise. Including state-specific time-varying controls in Column (2) does not qualitatively change the point estimate, but precision is greatly improved: this estimate is -0.15 ($p < 0.01$).¹¹ Column (3) switches to using Regional time fixed effects and the estimate is quite similar in magnitude (-0.17) and precision to that in Column (2).

If states that increased their minimum wages would have had worse employment outcomes even without these changes, then indicators for future minimum wages will show a negative effect on current employment. This would imply that the relationship between higher minimum wages and lower employment outcomes is spurious and simply reflects unobserved trends. In Columns (4) and (5), we add leading values of the minimum wage as a falsification exercise. The lead terms are statistically insignificant, while the contemporaneous effect is still large, negative, and statistically significant, albeit slightly reduced in magnitude to about -0.13. This outcome suggests that if pre-existing underlying trends are in fact different between states, they are not different by very much and are unlikely to be a key driver of the overall result. Taken as a group, the results from Columns (1) to (5) of Table 2 indicate that the minimum wage negatively affects employment and that underlying state-specific trends are not responsible for this disemployment effect. This classic specification is unable to distinguish between an effect that is immediate from one that occurs over a longer span of time, but it clearly indicates *some* negative relationship between a state's minimum wage and its employment level.

If the effect estimated in Columns (1) - (5) is indeed of a dynamic nature, as argued in Section 2, occurring over several time periods rather than discretely at the time of the policy change, then one would expect the inclusion of jurisdiction time trends to substantially attenuate the estimated treatment effect in this model. This is exactly the result when we add state-specific time trends in Column (6): in sharp contrast to the other columns, the

¹¹One approach to examine the potential endogeneity of minimum wage changes is to examine the results with different combinations of the time-varying covariates. Including (all with national time fixed effects) only log-population as a covariate yields an estimate of -0.153 (se=0.055); log-population and the share aged 15-59 yields -0.125 (0.043); with just log-GSP/capita the estimate is -0.210 (0.118).

estimated effect here is nearly zero. This sequence of results is very akin to that depicted in the illustrative example in Figure 4. Of course, the existence of this pattern does not, in itself, provide definitive evidence that the effect of the minimum wage occurs over several periods. It does, however, serve as suggestive evidence that the negative effect estimated in the classic model without jurisdiction time trends – both in this paper and elsewhere in the literature – is not simply an artifact of underlying heterogeneity in states’ time-paths of employment that merely happen to be correlated with states’ minimum wage policies.

4.2 Long-differences specifications

We turn next to long-differences specifications to better examine whether there is a dynamic effect of the minimum wage. If the true effect of the minimum wage is on growth and reflected over time, then specifications that are differenced over increasingly longer time periods should yield larger coefficients for the effect of the minimum wage on employment, with the magnitude converging towards a longer run effect.¹² Specifically, we estimate specifications of the following form, with increasing duration of time span r :

$$\Delta_r \text{emp}_{it} = \tau_t + \gamma_i \cdot t + \beta_0 \Delta_r \text{mw}_{it} + \psi \cdot \Delta_r \text{controls}_{it} + \Delta_r \epsilon_{it} \quad (2)$$

If the minimum wage has a dynamic effect on employment, then the impact of the minimum wage may be small over short durations, but it will increase in magnitude and significance as the time span is increased. Ultimately, the estimated elasticity should level off in magnitude once the full effect has taken place. Equation 2 shows such a long-differenced equation, which we estimate (without jurisdiction time trends) in Panel A of Table 3, and the results indeed follow this pattern. Beginning with a gap of a single year in the first column, there is a relatively small and statistically insignificant effect of -0.02. But as we move across columns to include longer differences, the effect increases in magnitude and becomes statistically significant. The coefficient stabilizes around -0.05 after three years and remains at that level even after differencing across eight years. These estimates provide more direct and compelling evidence of dynamic effects.

We add jurisdiction time trends to Equation 2 in Panel B of Table 3. As with the classic model, these point estimates are also substantially attenuated relative to those without trends. Moreover, as the duration of the time difference r increases across columns in Table

¹²Nichols [2009] and our Appendix A provide a graphical representation and commentary on the econometric strengths and weaknesses of the various specifications we present in this section.

3, the point estimates in Panel B become monotonically larger relative to those in Panel A. If the negative effect of the minimum wage estimated in Panel A were merely an artifact of underlying differences in time trends, then the coefficients in Panel B should remain similar to each other (and near zero) as the duration increases. The increasing degree of attenuation as the number of periods increases is indicative of the time trends soaking up variation that is in actuality due to the treatment effect.¹³

4.3 Distributed lag specifications

Finally, to provide yet a more direct and flexible view of the nature of the dynamic relationship between the minimum wage and employment, we estimate a distributed lag specification as in Equation 3, in which we not only take first differences to account for unobserved heterogeneity across states, but also include several lags of the minimum wage.¹⁴

$$\Delta \text{emp}_{it} = \theta_t + \gamma_i + \sum_{r=0}^s \beta_r \Delta \text{mw}_{it-r} + \psi \cdot \Delta \text{controls}_{it} + \Delta \epsilon_{it} \quad (3)$$

In Column (1) of Table 4, we show results from estimating Equation 3 with a contemporaneous term and three lags, with this lag structure selected based on the pattern of results seen in Table 3.¹⁵ The contemporaneous term is negative, but statistically insignificant, but the first two lags of the minimum wage are each statistically significant and similar in magnitude, around -0.03; the third lag is small and statistically insignificant, as suggested by the leveling-off of the long-differenced estimates after three years. By summing up these terms, we see that the total elasticity of employment with respect to the minimum wage is about -0.07, statistically significant, and similar in magnitude to that derived in the long-

¹³The nature of this monotonic attenuation pattern is straightforward. Consider the toy example in Figure 3 above. Long-difference estimates for this illustrative treatment effect are -0.5, -0.75, -1.25, and -1.75 for 1-4 time period differences, respectively, when time trends are not included. Including time trends yields corresponding long-difference estimates of -0.366, -0.426, -0.388, and -0.602. Thus, the estimates in this illustrative example are attenuated by 0.134, 0.324, 0.862, and 1.148 – monotonically increasing despite the clear absence of any difference in pre-treatment underlying time trends. The pattern shown via this basic illustration can be replicated more generally using simple Monte Carlo simulation.

¹⁴An alternative approach would be to estimate distributed lag coefficients in levels, rather than first-differences, but this approach may suffer from several shortcomings, as highlighted in Neumark and Wascher [1992] and Baker et al. [1999]. Estimates from distributed lags in levels using these same data are if anything *more* negative (and at least as statistically significant) as those using distributed lags in first-differences, as presented in this section.

¹⁵Additional lags do not make a qualitative difference to the sum of coefficients, and the coefficients on the first three minimum wage terms remain similar in magnitude and significance.

differenced estimates. In Columns (2) and (3), we add leads of the minimum wage to check for violations of the parallel trends assumption. The leading terms are statistically insignificant and trivial in size, both individually and summed together. The magnitude of the total minimum wage elasticity is unaffected, though some precision is lost with two leads. Even including in three leads does not greatly affect matters, with the main effect reduced slightly to -0.064 (s.e. = 0.043; $p = 0.14$), and the sum of the three leading terms equal to 0.011 (s.e. = 0.023; $p = 0.64$). Once again, this is strong evidence of a dynamic treatment effect of the minimum wage on employment, with little to indicate that differences in underlying growth trends are a relevant factor.

The remaining columns of Table 4 provide some further robustness tests. In Column (4), we allow the time effects to vary by Census Division, rather than Region; the coefficients remain similar to those in Column (1). Some precision is lost, though this to be expected – there are four Census Regions containing nine Divisions, and the median Division includes only five states. In Column (5), we assess whether states that have shifted to indexing their minimum wage for inflation affect our results by dropping these observations. Results remain similarly unchanged. Finally, in Column (6), we evaluate the role of the 2008-2009 recession. Because we include time period fixed effects, the recent recession should not unduly affect our results. However, these years of our panel experienced relatively large and high-frequency changes in real minimum wage levels, primarily resulting from the federal increases during these years (see Figure 2). As a check that these particular years are not overly influencing identification of the minimum wage term, we estimate specifications using only pre-2008 data. The estimated effects are slightly smaller, though not meaningfully different from our main results. However, the sum of the minimum wage terms is significant only at $p = 0.13$; this is somewhat unsurprising given that about fifteen percent of the observations are lost.

One possible shortcoming of the distributed lag model is that high frequency variation in treatment intensity makes it difficult to inform causal inference about the employment effects of higher-order lags of the minimum wage, because the large number of changes and potential long-run confounders make a fully-specified model fragile. Put another way: in practice the number of included lags must be fairly small in these specifications. While our results are robust to the inclusion of more lag terms, and the long-differenced specifications suggest that the effects are fairly stable after several years, a natural additional approach is to use a dynamic panel specification (e.g. [Arellano and Bond, 1991](#)). This allows us to estimate both the short- and long-run effects, at the cost of imposing a stricter assumption on the nature of this relationship. We discuss these results in Appendix [Online Appendix 3](#);

while certain specifications are somewhat fragile, the conclusions are similar to those yielded by this section.

5 Conclusion

We examine whether the minimum wage impacts employment through a discrete change in its level or if it is reflected over time. Much of the previous literature on the topic has assumed that an increase in the minimum wage would result in a relatively rapid adjustment in employment. Many have viewed a lack of such a finding as indicating that the minimum wage has minimal effects on employment, yet there are theoretical reasons to believe that this change may be slower. Using illustrative models, we show that the empirical specifications used in the prior literature will systematically err if the true effects are dynamic. Moreover, we show that the common practice in this literature of including jurisdiction-specific time trends will bias estimates towards zero in this case.

We show results from three administrative data sets that consistently indicate negative effects of the minimum wage on job growth. Our results are robust to a number of specifications and we find that the minimum wage reduces employment over a longer period of time than the literature has focused on in recent years. This phenomenon is particularly important given the evidence that minimum wage jobs often result in relatively rapid transitions to higher-paying jobs [Even and Macpherson, 2003].

This paper, of course, does not settle the debate of a contentious topic, but we do shed light on the mechanisms by which the minimum wage affects employment and provide directions for future research delving more deeply into the dynamics of this relationship.


References

- John M. Abowd and Lars Vilhuber. Statistics of jobs [lecture slides]. Retrieved from www.vrdc.cornell.edu/info7470/2013/Lecture%20Notes/5-JobStatistics.pdf, February 2013.
- Katharine G. Abraham, John C. Haltiwanger, Kristin Sandusky, and James R. Spletzer. Exploring differences in employment between household and establishment data. *Journal of Labor Economics*, 31(2):129–172, April 2013.
- Daron Acemoglu. Good jobs versus bad jobs. *Journal of Labor Economics*, 19(1):1–20, January 2001.


- John T. Addison, McKinley L. Blackburn, and Chad D. Cotti. Do minimum wages raise employment? Evidence from the U.S. retail-trade sector. *Labour Economics*, 16(4):397–408, August 2009.
- Sylvia A. Allegretto, Arindrajit Dube, and Michael Reich. Do minimum wages really reduce teen employment? Accounting for heterogeneity and selectivity in state panel data. *Industrial Relations*, 50(2):205–240, 2011.
- T.W. Anderson and Cheng Hsiao. Formulation and estimation of dynamic models using panel data. *Journal of Econometrics*, 18(1):47–82, January 1982.
- Manuel Arellano and Stephen Bond. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *The Review of Economic Studies*, 58(2):277–297, 1991.
- Michael Baker, Dwayne Benjamin, and Shuchita Stanger. The highs and lows of the minimum wage effect: A time-series cross-section study of the Canadian law. *Journal of Labor Economics*, 17(2):318–350, April 1999.
- Nathaniel Baum-Snow and Byron F. Lutz. School desegregation, school choice, and changes in residential location patterns by race. *The American Economic Review*, 101(7):3019–3046, December 2011.
- Pierre Brochu and David A. Green. The impact of minimum wages on labour market transitions. *The Economic Journal*, 123(573):1203–1235, December 2013.
- Richard V. Burkhauser, Kenneth A. Couch, and David C. Wittenburg. Who minimum wage increases bite: An analysis using monthly data from the SIPP and the CPS. *Southern Economic Journal*, 67(1):16–40, July 2000.
- Pierre Cahuc and André Zylberberg. *Labor Economics*. Cambridge: The MIT Press, 2004.
- David Card and Alan B. Krueger. Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *The American Economic Review*, 84(4):772–793, September 1994.
- Gerald A. Carlino and Richard Voith. Accounting for differences in aggregate state productivity. *Regional Science and Urban Economics*, 22(4):597–617, November 1992.
- Jeffrey Clemens and Michael Wither. The minimum wage and the great recession: Evidence of effects on the wage distributions, employment, earnings, and class mobility of low-skilled workers. NBER Working paper No. w20724, December 2014.
- Peter A. Diamond. Mobility costs, frictional unemployment, and efficiency. *Journal of Political Economy*, 89(4):798–812, August 1981.

- Arindrajit Dube, T. William Lester, and Michael Reich. Minimum wage effects across state borders: Estimates using contiguous counties. *The Review of Economics and Statistics*, 92(4):945–964, 2010.
- Arindrajit Dube, T. William Lester, and Michael Reich. Do frictions matter in the labor market? Accessions, separations and minimum wage effects. IZA Working paper 5811, June 2011.
- William E. Even and David A. Macpherson. The wage and employment dynamics of minimum wage workers. *Southern Economic Journal*, 69(3):676–690, January 2003.
- Christopher J. Flinn. Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates. *Econometrica*, 74(4):1013–1062, 2006.
- Christopher J. Flinn. *The Minimum Wage and Labor Market Outcomes*. Cambridge: The MIT Press, 2011.
- Laura Giuliano. Minimum wage effects on employment, substitution, and the teenage labor supply: Evidence from personnel data. *Journal of Labor Economics*, 31(1):155–194, January 2013.
- Daniel S. Hamermesh. Labor demand and the structure of adjustment costs. *The American Economic Review*, 79(4):674–689, September 1989.
- Barry T. Hirsch, Bruce E. Kaufman, and Tetyana Zelenska. Minimum wage channels of adjustment. *Industrial Relations*, 54(2):199–239, April 2015.
- Douglass Holtz-Eakin, Whitney Newey, and Harvey S. Rosen. Estimating vector autoregressions with panel data. *Econometrica*, 56(6):1371–1395, November 1988.
- David Lee and Emmanuel Saez. Optimal minimum wage policy in competitive labor markets. *Journal of Public Economics*, 96(9-10):739–749, October 2012.
- Jin Young Lee and Gary Solon. The fragility of estimated effects of unilateral divorce laws on divorce rates. *The B.E. Journal of Economic Analysis & Policy*, 11(1), August 2011.
- Jacob Mincer. Unemployment effects of minimum wages. *Journal of Political Economy*, 84(4):87–104, 1976.
- David Neumark and William Wascher. Employment effects of minimum and subminimum wages: panel data on state minimum wage laws. *Industrial and Labor Relations Review*, 46(1):55–81, October 1992.
- David Neumark and William Wascher. *Minimum Wages*. Cambridge: The MIT Press, December 2008.
- David Neumark, Mark Schweitzer, and William Wascher. Minimum wage effects throughout the wage distribution. *The Journal of Human Resources*, 39(2):425–450, 2004.

- David Neumark, J.M. Ian Salas, and William Wascher. Revisiting the minimum wage-employment debate: Throwing out the baby with the bathwater? *Industrial and Labor Relations Review*, 67(1):608–648, January 2014.
- Austin Nichols. Causal inference with observational data: Regression discontinuity and related methods in Stata [lecture slides]. Retrieved from www.stata.com/meeting/germany09/nichols.pdf, June 2009.
- Pia M. Orrenius and Madeline Zavodny. The effect of minimum wages on immigrants' employment and earnings. *Industrial and Labor Relations Review*, 61(4):544–563, July 2008.
- Marianne E. Page, Joanne Spetz, and Jane Millar. Does the minimum wage affect welfare caseloads? *Journal of Policy Analysis and Management*, 24(2):273–295, 2005.
- David Roodman. How to do xtabond2: An introduction to difference and system GMM in Stata. *Stata Journal*, 9(1):86–136, 2009.
- John Schmitt. Why does the minimum wage have no discernible effect on employment? Center for Economic and Policy Research Working paper, February 2013.
- Isaac Sorkin. Minimum wages and the dynamics of labor demand. Mimeo, February 2013.
- George J. Stigler. The economics of minimum wage legislation. *The American Economic Review*, 36(3):358–365, June 1946.
- Gerard J. Van den Berg and Geert Ridder. An empirical equilibrium search model of the labor market. *Econometrica*, 66(5):1183–1221, September 1998.
- Jenny Williams. Does liberalizing cannabis laws increase cannabis use? *Journal of Health Economics*, 36(1):20–32, July 2014.
- Justin Wolfers. Did unilateral divorce laws raise divorce rates? A reconciliation and new results. *The American Economic Review*, 96(5):1802–1820, December 2006.
- Madeline Zavodny. The effect of the minimum wage on employment and hours. *Labour Economics*, 7(6):729–750, 2000.



(a) Treatment effect discrete in levels



(b) Treatment effect discrete in growth

Figure 1: Illustration of two types of treatment effects with staggered treatments




Figure 2: Frequency of increases to effective state nominal minimum wages (1976-2012)






Figure 3: Simple example of disemployment effect in growth rate



(a) Levels: without trends



(b) Levels: residual to trends

Figure 4: Example difference-in-differences without versus with jurisdiction time trends

Table 1: Summary statistics for state characteristics and employment
Business Dynamics Statistics (Annual, 1977 - 2011)

	Mean	Std. Dev.	Median
State minimum wage (\$)	4.40	1.360	4.25
State minimum wage (\$ real)	7.09	0.916	6.89
Jobs (thousands)	1888.0	2103.8	1224.9
Job growth (thousands)	27.2	85.59	15.4
Job growth (log)	0.017	0.0348	0.019
Population (thousands)	5160.6	5725.6	3513.4
Share aged 15-59	0.62	0.0196	0.62
GSP/capita (\$ real)	41591.6	16309.7	38447.1
Observations	1785		

Notes: We define each state's minimum wage annually as of March 12 in the BDS, using the maximum of the federal minimum wage and the state's minimum wage each period, drawn from state-level sources. Employment statistics are computed for the aggregate population of non-agricultural employees in each state. Job growth is the annual change in each state's employment level. All real dollar amounts are indexed to \$ 2011 using the CPI-Urban.

Table 2: Classic state-panel fixed effect estimates for the effect of the minimum wage on log-employment (BDS)

	(1)	(2)	(3)	(4)	(5)	(6)
Log-MW	-0.1859 (0.1178)	-0.1515*** (0.0444)	-0.1693*** (0.0383)	-0.1376*** (0.0295)	-0.1230*** (0.0295)	-0.0125 (0.0160)
1st lead of log-MW				-0.0304 (0.0283)	-0.0065 (0.0167)	
2nd lead of log-MW					-0.0359 (0.0387)	
Observations	1785	1785	1785	1734	1683	1785
Time FE	National	National	Regional	Regional	Regional	Regional
Time-varying controls	No	Yes	Yes	Yes	Yes	Yes
Jurisdiction time trends	No	No	No	No	No	Yes

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Notes: Robust standard errors are clustered by state and reported in parentheses. All columns include state fixed effects. Where included, state-level annual time-varying controls are log-population, the share aged 15-59, and log real gross state product per capita.

Table 3: Long difference estimates for the effect of the minimum wage on log-employment (Business Dynamics Statistics)

Number of years:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: without trends								
Long difference in log-MW	-0.0202 (0.0175)	-0.0387* (0.0213)	-0.0496** (0.0235)	-0.0512** (0.0239)	-0.0463* (0.0236)	-0.0492** (0.0213)	-0.0470** (0.0201)	-0.0475** (0.0222)
Panel B: with trends								
Long difference in log-MW	-0.0196 (0.0176)	-0.0364* (0.0212)	-0.0444* (0.0234)	-0.0388 (0.0246)	-0.0223 (0.0254)	-0.0137 (0.0242)	-0.0007 (0.0211)	0.0085 (0.0201)
Observations	1734	1683	1632	1581	1530	1479	1428	1377

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Notes: Robust standard errors are clustered by state and reported in parentheses. The column numbers correspond to the number of periods over which the long difference is taken. All columns include state fixed effects, year-by-Region time fixed effects, and state-specific time-varying controls: log-population, the share aged 15-59, and log real gross state product per capita.

Table 4: Distributed lag first-differences estimates for the effect of the minimum wage on log-employment (BDS)

	Baseline (1)	Leading values (2)	(3)	Division FE (4)	Non-indexed (5)	Pre-2008 (6)
Log-MW	-0.0204 (0.0162)	-0.0206 (0.0163)	-0.0178 (0.0161)	-0.0188 (0.0163)	-0.0226 (0.0183)	-0.0219 (0.0178)
1st lag of log-MW	-0.0321** (0.0139)	-0.0336** (0.0140)	-0.0283 (0.0171)	-0.0310* (0.0175)	-0.0268 (0.0166)	0.0020 (0.0231)
2nd lag of log-MW	-0.0304** (0.0128)	-0.0317** (0.0132)	-0.0353** (0.0141)	-0.0244 (0.0152)	-0.0343** (0.0158)	-0.0467*** (0.0171)
3rd lag of log-MW	0.0093 (0.0147)	0.0084 (0.0149)	0.0109 (0.0232)	0.0107 (0.0193)	0.0074 (0.0188)	0.0162 (0.0227)
1st lead of log-MW		-0.0075 (0.0085)	-0.0073 (0.0095)			
2nd lead of log-MW			0.0057 (0.0124)			
Sum MW effects	-0.0736** (0.036)	-0.0776** (0.038)	-0.0705 (0.049)	-0.0634 (0.051)	-0.0764** (0.034)	-0.0504 (0.033)
Observations	1581	1581	1530	1581	1536	1377

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Notes: Robust standard errors are clustered by state and reported in parentheses. All columns include state fixed effects, year-by-Region time fixed effects, and state-specific time-varying controls: log-population, the share aged 15-59, and log real gross state product per capita. Columns (2) - (3) include, respectively, the leading value of the log minimum wage at time $t+1$ or both $t+1$ and $t+2$. Column (4) uses Division-by-time fixed effects, rather than Region-by-time. Column (5) drops the observations with an inflation-indexed state minimum wage, and Column (6) uses only pre-2008 data.

A Estimating growth effects for staggered treatments

This section uses a simple toy example to illustrate the nature of – and compare results from – several candidate approaches to estimating the dynamic effect of a treatment variable on the outcome of interest. In particular, it contrasts the “classic” fixed effects model with a differences model (both first-differences and longer differences) and with a distributed lag model in first-differences. Figure A.1 depicts the raw (simulated) data. There are three hypothetical jurisdictions: (1) a “control” jurisdiction, whose employment level remains at 3 units for the entire duration of the available data, (2) a jurisdiction which is treated with the binary treatment following period 2, and (3) a jurisdiction which is treated with the same-magnitude binary treatment following period 4. Figure A.1 shows the employment levels for the three hypothetical jurisdictions in each of 8 time periods.

Table A.1 presents the point estimates from six possible specifications (in rows) to estimate the treatment effect of four possible sets of jurisdictions (in columns). The columns consider the estimates that would be obtained from estimating treatment effects from comparisons of jurisdictions, respectively: 1 to 2, 1 to 3, 2 to 3, or 1 to 2 to 3. All eight time periods are included in each of these regressions.

Starting with the classic panel fixed effects difference-in-differences specification in Column (1), it is clear that this approach compares the mean difference (zero) pre-treatment, to the mean difference (-1.83) post-treatment, between jurisdictions 2 and 1. This visibly understates the long-run treatment effect here of -2, as it averages in the period 3 difference of only -1.5; this is a basic artifact of the standard difference-in-differences model, rather than an anomaly. However, moving across the columns, which compare various combinations of possible jurisdictions in this setting, one can easily see how the point estimate is sensitive to the set of included jurisdictions, varying by as much as 27% (-1.33 vs. -1.83) in this simple example in which there is a very clear actual treatment effect.

Moving down the rows to the differenced specifications, it is clear that the first-differenced specification correctly estimates the -1.0 treatment effect of the initial reduction in employment from the change in treatment, regardless of the combination of jurisdictions in this illustration. The longer differences (2, 3, 4) are, however, much more sensitive to the combination of included jurisdictions. Finally, the distributed lag model correctly sums the -2.0 treatment effect, irrespective of the included jurisdictions.

The overall point of this appendix is not to argue for any *true* model. Rather, it is meant to illustrate how sensitive results in a difference-in-differences context may be to the nature of the treatment effect and to the form of the estimator.




Figure A.1: Illustration of staggered treatments with growth effects

Table A.1: Treatment coefficients for estimates using various methods

Model	Jurisdictions included			
	1,2	1,3	2,3	1,2,3
Fixed effects (classic)	-1.83	-1.75	-1.33	-1.65
First difference (FD)	-1.00	-1.00	-1.00	-1.00
Long difference (2 period)	-1.50	-1.50	-1.25	-1.38
Long difference (3 period)	-2.00	-1.67	-1.67	-1.75
Long difference (4 period)	-2.00	-1.75	-1.50	-1.75
Distributed lag FD (2 period)	-2.00	-2.00	-2.00	-2.00

Notes: Columns include sets of two or three jurisdictions, with identities corresponding to the [1], [2], and [3] in Figure A.1. Values are point estimates from the models indicated in rows, and the linear combination of point estimates for the distributed lag specification.

Online Appendix 1. Results using additional data sets

In this appendix we provide empirical results similar to those in the main text, but using data from the Quarterly Census of Employment and Wages and the Quarterly Workforce Indicators, rather than from the Business Dynamics Statistics. The results are consonant with those in Section 4. Note that these data are quarterly rather than annual. As such, additional lags are included in the distributed lag models to cover the same temporal span as the annual specifications from the BDS.

A. Data

1. Quarterly Census of Employment and Wages (QCEW)

The *Quarterly Census of Employment and Wages* (QCEW), housed at the Bureau of Labor Statistics, is a program which originated in the 1930s to tabulate employment and wages of establishments which report to the Unemployment Insurance (UI) programs of the United States. Per the BLS, employment covered by these UI programs today represents about 99.7% of all wage and salary civilian employment in the country (including public sector employment). The BLS currently reports QCEW data by state for each quarter during 1975-2012, a span slightly longer than that of the BDS.¹⁶ The data are disaggregated by NAICS industry codes for 1990-2012.

2. Quarterly Workforce Indicators

The *Quarterly Workforce Indicators* are data provided as part of the Longitudinal Employer-Household Dynamics (LEHD) program by the Bureau of the Census. Similar to the QCEW, these data originate from county employment insurance filings.¹⁷

Yet, for our research design, a major shortcoming of the QWI is the substantially shorter – and highly unbalanced – length of the panel. At its onset in 1990, only four states participated in the QWI program, and additional states gradually joined through 2004. From 2004 on, the QWI includes forty-nine states (Massachusetts and Washington, D.C. are never included). Thus, the starting date for QWI participation varies considerably across states, and many are relatively recent. In addition to the standard concerns with unbalanced panels, this is of particular concern for the distributed lag models, as including sixteen minimum wage terms reduces the sample size by over twenty percent.

¹⁶Employment levels – and therefore also quarterly job growth rates – are not available in the QCEW for Alaska and the District of Columbia for any quarters during 1978-1980. Employment data is not missing for any other states or periods.

¹⁷In fact, the QWI and QCEW originate identically from the same county unemployment insurance records. Thus, differences in the data stem from either the periods during which each state or county is included, or differing imputation methods employed by BLS versus Census [Abowd and Vilhuber, 2013].

B. Results

We follow the same pattern of specifications as in Section 4. As with the BDS, the classic state-panel fixed effects estimates in Table OA1.2 tend to have a negative and statistically significant estimate of the impact of the minimum wage on employment. Inclusion of leads in the QCEW raises some suspicions of pre-existing trends, though the contemporaneous effect is still sizable in magnitude and remains statistically significant (and statistically equivalent to the estimate without leading terms in Column (3)). It is also worth noting that because these are quarterly data, the leading periods are much closer in time to the “treatment period” in which the minimum wage actually changes; thus, we would expect leading terms in these data sets to be more likely to detect any anticipatory action on the part of firms with respect to the future change in minimum wage. As with the BDS, inclusion of state-specific time trends in Column (6) drives the estimated effect to zero. Turning to the long-difference estimates in Table OA1.3, we increase each lag by four quarters to match the timespan in the BDS. The general pattern of effects that increase in magnitude with the length of the difference is present, particularly in the QWI. Once again, the inclusion of trends eliminates this tendency.

Finally, the distributed lag first-differences estimates in Table OA1.4 also follow the same pattern as those in Table 4. We include the contemporaneous value of the minimum wage, as well as fifteen lags, to match the same time frame as the BDS. For brevity, we present the sum of these effects rather than each individual coefficient, though full results with all coefficients are available on request (or in data and code provided by the authors online). In Column (1), we see that both the QCEW and the QWI produce a statistically significant total long-run elasticity of the minimum wage on employment of about -0.08, very similar in magnitude to that from the BDS. In Column (2), we add four lead terms and report their sum to test for pre-existing trends. This term is statistically insignificant and trivial in magnitude for both the QCEW and QWI, and the coefficients are jointly insignificantly different from zero. Moreover, the sum of the sixteen coefficients of interest is unaffected. In Column (3), we include eight lead terms and again find no evidence of pre-existing trends that would suggest that our results are being driven by confounding factors. Columns (4) through (6) follow the robustness checks in Table 4 and illustrate the stability of results.

Altogether, it is evident that our results are not driven by the choice of data set; each of the three sources produces the same conclusions.

Table OA1.1: Summary statistics for state characteristics and employment outcomes in three administrative data sets

	BDS			QCEW			QWI		
	Annual, 1977 - 2011			Quarterly, 1975 - 2012			Quarterly, varies - 2012*		
	Mean	Std. Dev.	Median	Mean	Std. Dev.	Median	Mean	Std. Dev.	Median
State minimum wage (\$)	4.40	1.360	4.25	4.53	1.535	4.25	5.86	1.094	5.15
State minimum wage (\$real)	7.09	0.916	6.89	7.28	0.975	7.05	6.89	0.729	6.85
Jobs (thousands)	1888.0	2103.8	1224.9	2167.9	2402.9	1441.7	2621.4	2794.2	1763.7
Job growth (thousands)	27.2	85.59	15.4	8.77	65.04	4.28	4.35	74.03	6.15
Job growth (log)	0.017	0.0348	0.019	0.0051	0.0256	0.0049	0.0019	0.0241	0.0061
Population (thousands)	5160.6	5725.6	3513.4	5138.0	5704.7	3502.0	6136.5	6784.5	4343.4
Share aged 15-59	0.62	0.0196	0.62	0.62	0.0199	0.62	0.62	0.0145	0.62
GSP/capita (\$real)	41,592	16,310	38,447	41,302	16,334	38,148	45,345	8384	43,969
Observations	1785			7752			3029		

Notes: We define each state's minimum wage annually as of March 12 in the BDS, and as of the first date for each quarter in the QCEW and QWI. We use the maximum of the federal minimum wage and the state's minimum wage each period, drawn from state-level sources. Employment statistics are computed for the aggregate population of non-agricultural employees in each state for each of the three listed data sets. Job growth is the change in each state's employment level from one time period to the next. We use job growth and employment outcomes annually for the BDS and quarterly for the QCEW and QWI. All real dollar amounts are indexed to \$2011 using the CPI-Urban. The QWI is a highly unbalanced panel, beginning with only four states in 1990 and gradually expanding until forty-nine states had joined by 2004. We include all available state-quarters of the QWI.

Table OA1.2: Classic state-panel fixed effect estimates for the effect of the minimum wage on log-employment

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: QCEW						
Log-MW	-0.1116 (0.1192)	-0.1344*** (0.0458)	-0.1391*** (0.0473)	-0.0976*** (0.0345)	-0.0923*** (0.0335)	0.0010 (0.0171)
1st lead of log-MW				-0.0442* (0.0248)	0.0078 (0.0131)	
2nd lead of log-MW					-0.0600*** (0.0220)	
Observations	7728	7728	7728	7677	7626	7728
Panel B: QWI						
Log-MW	-0.0384 (0.0441)	-0.0165 (0.0226)	-0.0447* (0.0231)	-0.0439*** (0.0162)	-0.0443*** (0.0163)	-0.0071 (0.0147)
1st lead of log-MW				-0.0014 (0.0167)	0.0132 (0.0175)	
2nd lead of log-MW					-0.0171 (0.0199)	
Observations	3029	3029	3029	2980	2931	3029
Time FE	National	National	Regional	Regional	Regional	Regional
Time-varying controls	No	Yes	Yes	Yes	Yes	Yes
Jurisdiction time trends	No	No	No	No	No	Yes

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Notes: Robust standard errors are clustered by state and reported in parentheses. All columns include state fixed effects. Where included, state-level annual time-varying controls are log-population, the share aged 15-59, and log real gross state product per capita.

Table OA1.3: Long difference estimates for the effect of the minimum wage on log-employment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Number of quarters:	4	8	12	16	20	24	28	32
Panel A: QCEW without trends								
Long difference in log-MW	-0.0028 (0.0069)	-0.0091 (0.0094)	-0.0152 (0.0118)	-0.0079 (0.0124)	-0.0071 (0.0136)	-0.0136 (0.0164)	-0.0276 (0.0184)	-0.0356* (0.0205)
Panel B: QCEW with trends								
Long difference in log-MW	-0.0024 (0.0067)	-0.0080 (0.0089)	-0.0127 (0.0110)	-0.0028 (0.0113)	0.0031 (0.0124)	0.0058 (0.0146)	0.0035 (0.0162)	0.0039 (0.0177)
Observations	7516	7304	7092	6896	6700	6504	6300	6096
Panel C: QWI without trends								
Long difference in log-MW	-0.0067 (0.0078)	-0.0092 (0.0105)	-0.0165 (0.0125)	-0.0208 (0.0152)	-0.0242 (0.0182)	-0.0337 (0.0209)	-0.0430* (0.0225)	-0.0469* (0.0265)
Panel D: QWI with trends								
Long difference in log-MW	-0.0044 (0.0077)	-0.0072 (0.0096)	-0.0128 (0.0105)	-0.0145 (0.0124)	-0.0133 (0.0161)	-0.0164 (0.0194)	-0.0196 (0.0193)	-0.0102 (0.0203)
Observations	2833	2637	2441	2245	2049	1853	1657	1461

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Notes: Robust standard errors are clustered by state and reported in parentheses. The “number of quarters” row corresponds to the number of periods over which the long difference is taken. All columns include state fixed effects, quarterly-by-Region time fixed effects, and state-specific time-varying controls: log-population, the share aged 15-59, and log real gross state product per capita.

Table OA1.4: Distributed lag first-differences estimates for the effect of the minimum wage on log-employment

	Baseline (1)	Leading values (2)	(3)	Division FE (4)	Non-indexed (5)	Pre-2008 (6)
Panel A: QCEW						
Current + Lags	-0.0820** (0.0313)	-0.0875*** (0.0305)	-0.0880** (0.0334)	-0.0559 (0.0345)	-0.0729** (0.0311)	-0.0869*** (0.0302)
Leads		-0.0059 (0.0070)	0.0069 (0.0138)			
Observations	6918	6714	6510	6918	6698	5898
Panel B: QWI						
Current + Lags	-0.0863** (0.0333)	-0.1114*** (0.0384)	-0.0873 (0.0532)	-0.0702** (0.0342)	-0.0612* (0.0310)	-0.1023** (0.0413)
Leads		-0.0038 (0.0152)	0.0034 (0.0251)			
Observations	2245	2049	1853	2245	2026	1366

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Notes: Robust standard errors are clustered by state and reported in parentheses. All columns include state fixed effects, quarterly-by-Region time fixed effects, and state-specific time-varying controls: log-population, the share aged 15-59, and log real gross state product per capita. Columns (2) - (3) include, respectively, the leading values of the log minimum wage during the preceding year or preceding two years (4 or 8 leading terms, respectively). Column (4) uses Division-by-time fixed effects, rather than Region-by-time. Column (5) drops the observations with an inflation-indexed state minimum wage, and Column (6) uses only pre-2008 data.

Online Appendix 2. Results by industry

In the main body of the paper, we present results for virtually the entire workforce, including workers in all industries. In this appendix, we disaggregate the effect on job growth rates by industry. The BDS does not report separate employment outcomes by state and industry, but these are disaggregated in the QCEW and QWI. In Table OA2.1, we estimate the effects of the minimum wage in different industries (two-digit NAICS code), focusing on the distributed lag model.¹⁸ Much of the literature focuses on one or several industries that are conjectured to be more responsive to changes in the minimum wage. Echoing points made in Clemens and Wither [2014] and Neumark et al. [2004], we choose to show all industries as it is not necessarily clear which particular *industry codes* ought not to be sensitive to the minimum wage. That said, industries that tend to have a higher concentration of low-wage jobs show more deleterious effects on job growth from higher minimum wages, and the results appear consistent between the QCEW and QWI.¹⁹

¹⁸See <http://www.naics.com/search.htm> for a full list of the component industries of each category.

¹⁹It may seem anomalous that professional services would be negatively affected, but firms in this category span a broad array, from lawyers' offices to direct mail advertising. The large negative effect on the offices of holding companies ("management") is perhaps stranger; note, though, that the effect is only present in the QCEW, that the estimates are quite noisy, and that this category has among the fewest firms of any industry.

Table OA2.1: Distributed lag first-differences estimates for the effect of the minimum wage on log-employment by industry

	QCEW		QWI	
	coef.	s.e.	coef.	s.e.
All: NAICS available (1990-)	-0.0815***	(0.0285)	-0.0863**	(0.0333)
11: Agriculture and wildlife	-0.1618	(0.1744)	-0.0546	(0.1266)
21: Mining	0.2011	(0.1941)	0.2638	(0.2389)
22: Utilities	-0.0043	(0.1488)	0.0625	(0.1366)
23: Construction	-0.2107	(0.1438)	-0.2003	(0.1333)
31-33: Manufacturing	-0.0957	(0.0646)	-0.0852	(0.0569)
42: Wholesale trade	-0.0073	(0.0431)	-0.0803	(0.0577)
44-45: Retail trade	-0.0253	(0.0312)	-0.0710	(0.0439)
48-49: Transportation and warehouse	-0.1010	(0.0799)	-0.0195	(0.0670)
51: Information service	0.1654	(0.2316)	-0.0086	(0.0762)
52: Finance and insurance	-0.0137	(0.0451)	-0.1410	(0.0858)
53: Real estate	-0.0639	(0.0561)	-0.0327	(0.0749)
54: Professional service	-0.2021***	(0.0614)	-0.2713***	(0.0629)
55: Management	-0.6041*	(0.3381)	-0.1628	(0.7198)
56: Administrative support	-0.1575**	(0.0595)	-0.2162**	(0.0825)
61: Education related	0.6623*	(0.3501)	0.0234	(0.0975)
62: Health care	-0.0287	(0.0320)	0.0408	(0.0689)
71: Arts and entertainment	-0.1098	(0.1452)	-0.1486	(0.0989)
72: Accommodation and food	-0.0669***	(0.0226)	-0.1098*	(0.0648)
81: Other service	-0.1235	(0.3323)	-0.0004	(0.0573)
92: Public administration	-0.0346	(0.0767)	-0.0760	(0.1075)
Observations	3825		2245	

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Notes: Robust standard errors are clustered by state and reported in parentheses. All columns include state fixed effects, quarterly-by-Region time fixed effects, and state-specific time-varying controls: log-population, the share aged 15-59, and log real gross state product per capita. Each coefficient represents a separate regression of the distributed lags in first differences model, using lags over 16 quarters.

Online Appendix 3. Dynamic panel estimates

An alternative approach to estimating the short- and long-run effects of the minimum wage on employment, at the cost of imposing a stricter assumption on the nature of this relationship, is to use a dynamic panel specification (e.g. [Arellano and Bond, 1991](#)). The specification takes the form:

$$\text{emp}_{it} = \mu \cdot \text{emp}_{it-1} + \alpha_i + \tau_t + \gamma_i \cdot t + \sum_{r=0}^s \beta_r \text{mw}_{it-r} + \psi \cdot \text{controls}_{it} + \epsilon_{it}$$

which differs from the specifications discussed in Section 4 in that the lag of employment is included on the right hand side. This can be first-differenced to eliminate the α_i jurisdiction fixed effects:

$$\Delta \text{emp}_{it} = \mu \cdot \Delta \text{emp}_{it-1} + \theta_t + \gamma_i + \sum_{r=0}^s \beta_r \Delta \text{mw}_{it-r} + \psi \cdot \Delta \text{controls}_{it} + \Delta \epsilon_{it} \quad (4)$$

In this dynamic panel model, the short run marginal effect of the minimum wage on employment is β_0 , and the effect after one year of a sustained change is captured by $\beta_1 + (1 + \mu) \cdot \beta_0$. The long run effect on employment is determined by $(\beta_0 + \beta_1)/(1 - \mu)$, following the properties of a geometric series. Importantly, this long run effect (in fact, the specific time path of the effect) can be identified using only a single lag term for the minimum wage. Thus, a dynamic panel specification skirts much – although not all – of the concern about constantly changing treatment intensities.

However, in solving one identification problem, the dynamic panel approach introduces another, as the Δemp terms are autocorrelated. The standard practice, as in [Holtz-Eakin et al. \[1988\]](#) and [Arellano and Bond \[1991\]](#), is to create “GMM-style” instruments using deeper lags of employment and substituting zeroes for the missing observations resulting from the lags. It is important to note that these instruments may be problematic as well, depending on the degree of autocorrelation. An alternative approach is to use deeper lags of the *minimum wage* rather than employment as instruments. A further alternative is to use a traditional two-stage least squares approach in which deeper lags of the minimum wage are used as instruments, without the GMM-style substitution of missing observations, similar in spirit to [Anderson and Hsiao \[1982\]](#).²⁰

In Table OA3.1, we estimate Equation 4 using both GMM-style and standard instruments. Columns (1) and (2) use [Roodman’s \(2009\)](#) Stata module, which allows for flexible estimation of dynamic panel models.²¹ In Column (1), the contemporaneous elasticity of

²⁰We are grateful to an anonymous referee for both of these suggestions. See [Roodman \[2009\]](#) for an extensive discussion of these issues.

²¹In all cases, we use deeper minimum wage lags as instruments rather than deeper lags of employment. In both specifications, results are qualitatively similar when using deeper lags of employment or when using both employment and minimum wage variables as instruments.

a minimum wage increase is -0.031 (s.e. = 0.017), with the lag term (-0.054, s.e. = 0.02) implying that the impact after one year at the same treatment intensity would be -0.10 (s.e. = 0.033) and after two years, -0.14 (s.e. = 0.049); the long-run impact of a permanent real increase in the minimum wage effect is -0.20 (s.e. = 0.088). Adding additional lags in Column (2) does not dramatically change the effect, with the impact after one year at the same treatment intensity being -0.096 and after two years, -0.13, with the long-run elasticity being -0.27 (s.e. = 0.13). In Column (3), we use standard rather than GMM instruments and find somewhat smaller effects than in Column (1), with a statistically-significant long-run elasticity of around -0.08.²² The result in Column (4) is similar; note that the magnitude here is very close to that from the distributed lag model in first differences.

Altogether, the results from the dynamic panel models also suggest that the impacts of the minimum wage on employment are dynamic rather than discrete.

²²For these specifications, we use four lags of the minimum wage as instruments, beginning with the first period not included in the primary equation. The magnitude of the overall effect of the minimum wage tends to be fairly stable based on the choice of instrument sets, though for some sets, the lagged employment coefficient is imprecise and sometimes implausible. The first stage estimates are strong, with an overall F-statistic of 142.6 in Column (3) and 241.6 in Column (4); the four instruments are jointly significant with $p = 0.009$ and $p = 0.071$ for Columns (3) and (4), respectively.

Table OA3.1: Dynamic panel estimates for the effect of the minimum wage on log-emp. (BDS)

	GMM Instruments		Standard Instruments	
	(1)	(2)	(3)	(4)
Log-MW	-0.0309* (0.0171)	-0.0390** (0.0165)	-0.0159 (0.0136)	-0.0153 (0.0150)
1st lag of log-MW	-0.0543*** (0.0204)	-0.0310* (0.0167)	-0.0379*** (0.0094)	-0.0309** (0.0153)
2nd lag of log-MW		-0.0095 (0.0121)		-0.0051 (0.0239)
3rd lag of log-MW		-0.0146 (0.0227)		0.0136 (0.0172)
1st lag of employment	0.5772*** (0.0960)	0.6539*** (0.0846)	0.3256 (0.2437)	0.5301 (0.6678)
Estimated Permanent Effect	-0.2015** (0.0884)	-0.2720** (0.1323)	-0.0799** (0.0316)	-0.0802* (0.0428)
Observations	1683	1581	1428	1326

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$ Notes: Robust standard errors are clustered by state and reported in parentheses. Columns (1) and (2) use [Roodman \[2009\]](#)'s difference GMM estimator with lags of the minimum wage as instruments. Columns (3) and (4) use two-stage least squares with four lagged minimum wage values as instruments.

Online Appendix 4. Historical Minimum Wage Increases and Erosion

Historically, minimum wages have been set in nominal dollars and not adjusted for inflation, so any nominal wage differential between two jurisdictions will become economically less meaningful over time. This appendix section presents some figures depicting the frequency and magnitude of minimum wage changes – and their subsequent erosion due to inflation. Looking first only within-state, Figure [OA4.4](#) shows that the mean real state minimum wage increase during 1976-2012 was 55 cents (the median was also 55 cents). By the time the same state next increased its real minimum wage, which took 54 months on average, the previous increase in minimum wage had eroded – via inflation – to an average cumulative real *decrease* of 11 cents (median -12 cents, see Figure [OA4.5](#)). In fact, Figure [OA4.6](#) shows that the 62 percent of state-year real minimum wage increases that were eventually fully eroded by inflation did so in, on average, twenty-two months, and the median time elapsed was only sixteen months. Turning instead to comparisons within Census Region, the mean *relative* real increase in state minimum wage was 25 cents (median 13 cents, Figure [OA4.7](#)). By the time of the next within-state increase, the prior increase had eroded – both via inflation and from other regional neighbors changing their minimum wages – to an average decrease of 1 cents (median +2 cents, Figure [OA4.8](#)). For those 47 percent of state-year increases which fully eroded relative to regional states, this took only 17 months on average (median 12 months, Figure [OA4.9](#)). This exercise demonstrates that there is a relatively short duration of time during which a state difference-in-differences estimation can identify the effects of the minimum wage on employment levels.




Figure OA4.1: Comparison of federal to state nominal minimum wages (January, 1975-2012)




Figure OA4.2: Comparison of federal to state real minimum wages (January, 1975-2012)




Figure OA4.3: Standard deviation of residual state real minimum wages (1975-2012)




Figure OA4.4: Distribution of real minimum wage increases




Figure OA4.5: Cumulative difference in real minimum wage prior to a new increase




Figure OA4.6: Erosion of real increases in minimum wage




Figure OA4.7: Distribution of relative minimum wage increases




Figure OA4.8: Cumulative difference in relative minimum wage prior to a new increase




Figure OA4.9: Erosion of relative increases in minimum wage