Evaluating Minnesota Minimum Wage Increases Using Synthetic Control Methods

Matthew Yang

ECON 481: Advanced Econometrics

Professor Gary Krueger

Macalester College

Fall 2018

I. Introduction

Economic theory predicts that in competitive labor markets, an increase in wages will decrease employment due to the higher cost of labor for employers. Nonetheless, advocates push to increase the minimum wage so that low-wage workers will be able to earn enough money to sustain a baseline quality of life. In this view, even if the minimum wage has a small negative effect on employment, the policy is still beneficial overall because it improves earnings for low-wage workers. But if the negative effect is large, would the higher wage of some workers still outweigh the lost employment for other workers? This leads to a central question in the minimum wage debate: By how much does an increase in the minimum wage decrease employment?

In this paper, I examine the employment effects of recent minimum wage increases in Minnesota. Up until 2014, MN did not have a binding state minimum wage. That is, the state minimum wage was below the federal minimum wage, so employers were required to pay the federal minimum wage, which is currently $7.25. In August of 2014, the MN state minimum wage was raised to $8 as part of the first phase of the state's plans to adopt a $15 minimum wage. The MN state minimum wage has increased in each subsequent year, and is now $9.65.

In a 2018 policy brief from the Center for Research on the Wisconsin Economy (CROWE), Noah Williams compares employment levels in Wisconsin and Minnesota before and after the minimum wage increases in Minnesota. Williams's findings suggest that, overall, Minnesota's minimum wage increases have not been beneficial. He finds that while some workers received higher incomes after the wage increases, general employment in the state fell sharply. Moreover, Williams finds that restaurant prices increased significantly following the wage increases. This paper extends Williams' analysis by implementing different methods.

In any study of a policy intervention, one of the most important considerations is the construction of a control group. Previous studies of state-level minimum wages have chosen nearby states that kept the minimum wage constant during the time period of interest (such as Williams's choice of WI as a control for MN). This method of selecting the control group introduces arbitrariness into the analysis and can drastically affect results (Abadie et al., 2011). The synthetic control method, which is explained below, remedies this problem by using a data-driven algorithm to find the optimal control group.

The remainder of this paper is organized as follows. Section II reviews the relevant literature. Section III describes the econometric methods employed in the paper. Section IV discusses the data and implementation of the methods. Section V presents the results. Section VI provides concluding remarks.

II. Literature Review

There are a few central questions that drive research on the minimum wage. Neumark et al. (2014) summarizes them succinctly: (1) How does a minimum wage affect employment? (2) Which workers are affected by a minimum wage? (3) How can we use econometrics to isolate the effects of the minimum wage? There is no conclusive answer in the literature to any of these questions. In this section, I will examine the major contributions made in the minimum wage literature and compare them in the context of these three questions.

The vast majority of papers on minimum wage can be classified into two groups. The first group of papers uses a two-way fixed effects estimator to identify trends in employment across states with different minimum wage levels. These papers employ panel data on employment across states and time, and include state and period fixed effects in the model. This model enables the researcher to estimate the general employment effect from *all* of the minimum wage policies. The second group of papers attempts to study employment effects at a more granular level by examining individual case studies. That is, researchers match a specific place where there was variation in the minimum wage to a place where the minimum wage was constant, and then compare employment levels.

Neumark and Wascher (1992) were among the first to advocate for the use of a two-way fixed effects model. Prior to their paper, the majority of minimum wage studies used national-level time series data to identify employment effects. Neumark and Wascher add in state fixed effects to the model in order to exploit more variation in the minimum wage. The authors find employment elasticities ranging from -0.1 to -0.2 for teenagers and young adults, implying that a 10 percent increase in the minimum wage would result in a one to two percent decrease in employment. Many subsequent papers that used a two-way fixed effects estimator found similar disemployment effects. A main critique of this model is that is does not account for heterogeneous effects across states. That is, states may face different economic and social conditions overtime. Therefore, the two-way fixed effects model fails to account for time variant factors other than the minimum wage that may be confounding the employment effects.

Allegretto, Dube, and Reich (2011), hereafter ADR, remedy this critique of not accounting for spatial heterogeneity by controlling for both long-term growth differences among states and heterogeneous economic shocks. ADR control for regional differences among states by including "census division-specific time effects," which effectively compares a state only to those states within the same census division. This revised identification strategy yields employment elasticities that are statistically equal to zero. Other papers that have added spatial controls to the traditional fixed-effects model have found similar results.

Instead of comparing states within each census division, Dube, Lester, and Reich (2010), hereafter DLR, compare employment in every county pair that shares a state border. DLR find that higher minimum wages have no effect on employment, thereby confirming ADR's findings that controlling for spatial heterogeneity eliminates any significant disemployment effect. In addition to using this county matching strategy, DLR estimate employment effects using a traditional fixed effects specification and find negative employment elasticities. These findings are consistent with the pattern that has emerged in the minimum wage literature over the past several years: the traditional two-way fixed-effects model produces downwardly biased results. In his review of Neumark and Wascher's *Minimum Wages*, Dube (2010) writes, "Even simple regional controls and trends produce employment effects close to zero, as do more sophisticated approaches such as comparing contiguous counties across policy boundaries—which essentially embeds the 'case study' approach within panel data analysis." DLR's analysis demonstrates the importance of constructing adequate control groups.

Paramount to any study of an intervention on a specific group is the ability to estimate a counterfactual. In order to know the effect of a policy, we need to predict what would have happened in the absence of the policy. The case study approach attempts to estimate the unobserved counterfactual by using states or counties in close proximity to the treatment group as the control group. Case studies allow researchers to carefully choose a control group, eliminating the problem of spatial heterogeneity. Researchers often use a control group in close geographic proximity to the treatment group in order to account for regional shocks.

Card and Krueger (1994) offered one of the first case studies of the minimum wage. Much of the subsequent literature on the minimum wage in the last few decades has built on their work. The authors estimate employment effects in the fast-food industry by comparing two neighboring states: Pennsylvania and New Jersey. These two states offered a good case study because the minimum wage rose in NJ, while it stayed constant in PA. Researchers studying the minimum wage commonly focus on the fast-food or limited-service restaurant industries because many of the workers in these industries earn the minimum wage. Therefore, an analysis of employment in these sectors allows researchers to isolate the effect of the minimum wage on workers for which the minimum wage is binding. Card and Krueger use a difference-in-differences estimator to identify employment effects. They found no evidence that the rise in minimum wage reduced employment in NJ.

The synthetic control method offers a similar way to estimate the effect of an intervention on a treatment group, as compared to a control group. However, the synthetic control estimator chooses the control group in an empirical manner instead of leaving it up to the researcher's discretion, so there is less subjectivity in the construction of the control group (Abadie et al., 2007). Abadie, Diamond, and Hainmueller (2010) use the method to study the effect of a tobacco-tax program on cigarette consumption in CA. Instead of choosing a state as a direct comparison, the synthetic control method uses an optimization algorithm to compute a control group that is comprised of a weighted average of potential control states. Abadie et al. refer to these potential control states as "donor" states. The authors showed that the synthetic control method could provide a way to estimate a counterfactual as opposed to using a single control group.

A number of researchers studying the minimum wage have started turning to the synthetic control method as a way of comparing employment in states to a data-driven choice of control group. However, the application of the synthetic control method to minimum wages is not entirely straightforward. Minimum wage increases happen very frequently and in many places. As such, the states in the donor pool must be states that kept the minimum wage constant during the time period of interest, so the synthetic control does not experience any employment effects from a minimum wage increase. Neumark, Salas, and Wascher (2014), hereafter NSW, implement the synthetic control method to assess the validity of the assumptions of ADR. ADR showed that running the two-way fixed-effects model with added controls for census division and state employment patterns can produce unbiased employment elasticities. However, ADR implicitly assume that states within the same census division serve as good controls. NSW test this assumption by running the synthetic control algorithm on states in each census division and examining whether the algorithm chooses states within the same census division as controls. NSW find that, generally, it is not the case that states within the same census division provide a better control than other states do. Moreover, after comparing employment in states that increased the minimum wage to synthetic controls, NSW find evidence of disemployment effects, estimating employment elasticities for teenage workers around -0.15.

Reich, Allegretto, and Godoey (2017) implement an analogous experiment to NSW and arrive at the same conclusion. Specifically, Reich et al. employ the synthetic control method on county-level data to estimate employment effects of Seattle's minimum wage increases. The authors find that the optimal weighted control group selected by the algorithm contains counties outside Washington State, indicating that proximity does not always predict the adequacy of a control group. Indeed, the best method for constructing a control group and estimating a counterfactual continues to be a fundamental issue driving the minimum wage debate today.

Williams (2018) was the first author to examine Minnesota's recent minimum wage increases. After comparing limited-service restaurant employment in Minnesota and Wisconsin, Williams found that employment was significantly lower in Minnesota than in Wisconsin after the first minimum wage increase in 2014. Although Williams presents convincing evidence, Williams does not employ very robust statistical methods. Importantly, Williams does not present any inference for his results, so the significance of his findings is unclear. In contrast to Williams, Chinn and Johnston (2018) find no significant employment effects from the MN minimum wage increase. This paper will further explore the robustness of Williams' findings by implementing alternative methods.

III. Methods

In this paper, I employ two estimation techniques: the synthetic control method and a difference-in-differences estimator. For the former, I construct a control group that is a weighted average of states that kept the minimum wage constant. For the latter, I choose WI to serve as the control group for MN following Williams's analysis. Although

Williams directly compares employment in MN to employment in WI, he does not fit any econometric models to test his hypotheses. Therefore, a thorough investigation of Williams's hypotheses is needed to evaluate his conclusions. Each of the techniques I employ has its advantages, so I use both to provide a more complete picture of employment effects in MN.

i. Synthetic Control Estimator

The first step in the synthetic control method is establishing "donor" units. In my analysis, I define the donor units as states for which the binding state minimum wage remained constant during the period of 2010 to the present. If we were to consider all states, including those that increased the minimum wage, we would be contaminating the control group with disemployment effects, and would therefore not obtain accurate results. In the following paragraphs, I refer to the increase in the minimum wage as the "intervention," with Minnesota serving as the "treated" state.

Assuming we have J donor states and one treated state, there are $i = 1, \dots, J + 1$ states in our analysis. We define the treated state to be the first state of the total units (i.e. $i = 1$). Furthermore, we can define the number of time periods for which we observe the states as $t = 1, \dots, T$, where T_0 is the period in which the intervention occurs, and $1 \leq T_0 < T$.

We are interested in comparing employment levels in the treated state post-intervention to the employment levels that would have been observed absent the treatment. Let Y_{it}^I represent the employment level in state i at time t if the state is exposed to the intervention in periods $T_0 + 1$ to T. We define the estimated treatment effect to be $\alpha_{it} = Y_{it}^I - Y_{it}^N$ for time periods post the initial intervention. Rearranging, we get,

$$Y_{it}^{I} = Y_{it}^{N} + \alpha_{it}$$

Since only one state in the sample is exposed to the intervention, we can generalize this relationship by introducing a dummy variable that is equal to one if the state is exposed to the intervention at time t and zero otherwise. Therefore, we have:

$$Y_{it} = Y_{it}^{N} + \alpha_{it} D_{it}$$

This relationship implies that the employment level is equal to Y_{it}^{N} for all states at all times except for the treated state after $t = T_0$.

All of the quantities in the previous equation are observed except for Y_{it}^{N} for unit $i = 1$ at $t = T_0 + 1, \ldots, T$. In order to estimate the unobserved counterfactual, we need to construct a synthetic group that is a weighted average of the treated unit. We consider the following model (as described in Abadie et al. (2010)) to estimate Y_{1t}^{N} after the intervention:

$$\sum_{j=2}^{J+1} w_j Y_{jt} = \delta_t + \theta_t \sum_{j=2}^{J+1} w_j Z_j + \lambda_t \sum_{j=2}^{J+1} w_j \mu_j + \sum_{j=2}^{J+1} w_j \varepsilon_{jt}$$

In this model, the w_j's come from a $(J\ x\ 1)$ vector of weights W. All weights are nonnegative and must sum to one. Therefore, the vector W describes the weights that make up the synthetic group. The first term, δ_t, is a constant that is common among all units. Z_j is a $(r\ x\ 1)$ vector of observed explanatory variables, and θ_t is a vector containing the unknown coefficients. Furthermore, λ_t and μ_j describe an unknown common factor with different factor loadings across units, and ε_{jt} are the error terms.

A critical step in the synthetic control method is determining the vector W. As Abadie et al. (2010) describe, this is done by minimizing two quantities. First, we want to choose weights such that the synthetic group closely approximates the treated group in the pre-intervention period. If we can find a group that acts like the treatment group pre-intervention, then the simulated post intervention outcomes for the synthetic group become more believable. So, we have the following:

$$\sum_{j=2}^{J+1} w_j^* \overline{Y}_j^K = \overline{Y}_1^K,$$

where w_j^* are the optimal weights and $\overline{Y}_j^K$ is the average employment in the pre-treatment period. Rarely does this equality hold exactly, but we can find the optimal weights by minimizing the error.

Second, the synthetic group should approximate the values of the observed covariates for the treatment group. Thus, we have the following:

$$\sum_{j=2}^{J+1} w_j^* Z_j = Z_1$$

By minimizing the error in these two equations, we can find the vector W and thereby estimate the treatment effect of the intervention. The following equation gives us the estimated treatment effect:

$$\hat{\alpha}_{1t} = Y_{1t} - \sum_{j=2}^{J+1} w_j * Y_{jt}$$

ii. Difference-in-Differences

A difference-in-differences estimator is a more traditional way of estimating a treatment effect from an intervention that has been employed numerous times in minimum wage studies. This model relies on the parallel trends assumption, which states that both the control group and treatment group must have similar trajectories prior to the intervention (Powell, 2018). The synthetic control method relaxes this assumption because it finds a synthetic control that matches the treatment group very closely prior to the intervention.

The treatment effect is then calculated by taking the difference in the differences between the two groups. That is, if we expect the difference in the outcome variable to remain constant across time between the two groups, then the change in that difference post-intervention is the estimated treatment effect.

IV. Data and Implementation

In this paper, I employ data from the Quarterly Census of Employment and Wages (QCEW). Specifically, I focus on limited-service restaurant (LSR) employment, as is common in the literature. The data is from 2010-2018, which allows us to have roughly equal periods before and after the minimum wage increase. Moreover, the QCEW contains information on various labor market characteristics, such as employment, wages, and number of LSR establishments. The main outcome variable in this study is employment, as I am interested in how much employment changed due to the increase in minimum wage. Additionally, I fit a model to estimate the effect that the minimum wage increase had on total wages.

First, I use the synthetic control method to estimate the magnitude of the treatment effect for employment. I fit several models using this method, but will only present two in this paper. As discussed in the previous section, the construction of the Z matrix is important in the construction of a synthetic control. In the first model, I include average weekly wages, total taxable wages, and establishment count as covariates. Additionally, a much better synthetic control can be obtained when including specific lags of the dependent variable in the model. Therefore, I include 2010Q1, 2012Q1, and

2014Q1 employment as predictors. In order to test the sensitivity of the synthetic control algorithm, I take some predictors out of the model and compare the results.

The inference methods for the synthetic control method differ from those commonly performed with traditional regression analysis. The synthetic control method does not report standard errors, so I employ other techniques to understand the significance of the results. In this paper, I adopt the inferential methods put forth in Abadie and Gardeazabal (2003). More specifically, I fit placebo models to understand the 'randomness' of the results.

Next, I use a difference-in-differences estimator to compare the results obtained when using only WI as a control for MN, rather than a weighted average of donor states. This is implemented using the following regression model:

$$Employment = \beta_0 + \beta_1 MN + \beta_2 2014Q3 + \beta_3 MN * 2014Q3 + \varepsilon$$

where MN is a dummy variable that equals one if the state is MN and zero otherwise, and 2014Q3 is a dummy variable if the time period is after the third quarter of 2014. The interaction term between MN and 2014Q3 estimates the treatment effect.

Lastly, I measure the effect that the minimum wage increase had on total wages in MN using an analogous difference-in-differences model (with WI as the control), as shown here.

$$Total\ Quarerly\ Wages = \beta_0 + \beta_1 MN + \beta_2 2014Q3 + \beta_3 MN * 2014Q3 + \varepsilon$$

V. Results

i. Synthetic Control Method

The quality of the synthetic control can be measured by the root mean squared prediction error (RMSPE). The synthetic control model presented in Table 1 resulted in the best fit[1] (lowest RMSPE) of any of the models I ran. This model assigns positive weight to Illinois, Iowa, Maine, Missouri, and Wisconsin. Moreover, Figure 1 suggests that employment in synthetic MN becomes slightly higher than employment in actual MN after the initial minimum wage increase in August 2014, implying that the treatment effect is slightly negative. The placebo models[2] (shown in Figure 2) obtained after iteratively reassigning the treated unit show that the treatment effect for MN is not large relative to the other states. Thus, the results from this synthetic control model imply that there is not a very significant treatment effect for employment.

After taking the number of establishments out of the model, and matching only on one time period pre-intervention (as opposed to three), I obtain markedly different results. Here, the synthetic control algorithm chooses Idaho, Iowa, and North Carolina as part of the control group. The only control state in common with the previous model is Iowa, indicating that Iowa's LSR labor market exhibits many of the characteristics as that of MN. Figure 3 plots MN and synthetic MN employment. In contrast to Figure 1, this plot shows employment levels between the two groups diverging after the time of the intervention. However, the RMSPE for the second model is 1765, nearly triple the RMSPE for the first model. The large RMSPE in this model indicates that the synthetic control is a much worse approximate of MN than the corresponding synthetic control in the first model.

ii. Difference-in-Differences

Table 5 presents the difference-in-differences regression results for quarterly employment. The interaction term between MN and 2014Q, which represents the

---

[1] The RMSPE in this model is 640.

[2] Each gray line represents a placebo model, and the orange line represents the magnitude of the treatment effect.

treatment effect, is positive but insignificant. Therefore, this model implies that the treatment effect is effectively zero. This result agrees with the conclusion reached from the synthetic control models that the minimum wage increase did not have a significant effect on employment. Additionally, the MN dummy variable is insignificant, implying that there is no difference in average LSR employment between the two states, regardless of time. Figure 5 supports this finding, as a simple plot of LSR employment overtime shows almost no difference between the two states over the entire period.

Since employment in MN did not decrease after the minimum wage increase, the policy seems to be net beneficial. That is, employment did not fall, and the minimum wage increased, so total wages should increase. The results from Table 6 confirm this, as the treatment effect on total wages is positive and significant, indicating that the policy increased total wages in MN relative to WI. Figure 6, which plots total quarterly wages in MN and WI over the entire period, shows the gap in total wages between the two states increasing after the intervention.

VI. Conclusion

This paper finds no evidence that the increase in the MN minimum wage decreased LSR employment. I present two main analyses to support this conclusion. The first synthetic control model discussed in Section V estimates a counterfactual that is close to the observed employment levels in MN after the time of the intervention. Moreover, the placebo study implied that the marginal disemployment effect was not significant. Although the second synthetic control model predicts that employment in the synthetic control (i.e. in the absence of the intervention) was significantly higher, the sensitivity of the synthetic control method alone is enough to invalidate the result.

Furthermore, the synthetic control obtained in the second model had three times the error as the first model, indicating that the control group is less similar to MN.

Next, I compare LSR employment in MN and WI using a difference-in-differences model. I find that the treatment effect is slightly positive, but insignificant. Additionally, an analysis of total quarterly wages in MN and WI shows that the policy did indeed increase total wages in MN. This implies that if there was any decrease in employment, it was outweighed by the increase in the minimum wage.

There are a few main limitations of my analysis that may be affecting the results. In particular, the synthetic control method is extremely sensitive to the inputs included in the model. This could be the result of a few different factors. First, I employ non-seasonally adjusted data. The variability in employment may increase the error of the synthetic control method, as it might be harder to find a perfect fit. Next, the data contains a relatively short pre-intervention period of just under 5 years. Abadie et al. (2010) employ nearly 20 years of data prior to the intervention. Therefore, I may not have enough data to find an adequate synthetic control. The synthetic control method is difficult to implement in minimum wage studies for this reason—state minimum wage levels are constantly changing. So, when trying to isolate the effect of a minimum wage on a particular state, the number of potential control states is limited. If I were to extend the time period pre-intervention, there would be fewer states that held the minimum wage constant, which would limit the power of the synthetic control method.

As I have shown, implementing different methods than those presented in Williams's CROWE policy brief results in starkly different results. The synthetic control and difference-in-differences analyses presented in this paper provide evidence that the

MN minimum wage did not reduce LSR employment, and did increase total wages. Finding opposite or inconsistent results has become a common trend in the minimum wage literature. The variability in findings underscores the need for greater methodological refinement of control groups to be conducted. Only then will we be able to ascertain more clearly who are the winners and losers of a minimum wage policy.

Tables and Figures

Table 1

State	Weight
Alabama	0
Georgia	0
Idaho	0
Illinois	0.173
Indiana	0
Iowa	0.047
Kansas	0
Kentucky	0
Louisiana	0
Maine	0.337
Missouri	0.258
Nevada	0
New Hampshire	0
New Mexico	0
North Carolina	0
North Dakota	0
Oklahoma	0
Pennsylvania	0
South Carolina	0
Tennessee	0
Utah	0
Virginia	0
Wisconsin	0.185
Wyoming	0

Table 2

Variables	Minnesota	Synthetic Minnesota
Avg. Weekly Wage	228	246
Establishment Count	3,665	3,627
Taxable Wages	1.74e+08	1.56e+08
Employment 2010Q1	59,274	59,151
Employment 2012Q1	61,869	61,831
Employment 2014Q1	64,973	65,069

Table 3

State	Weight
Alabama	0
Georgia	0
Idaho	0.302
Illinois	0
Indiana	0
Iowa	0.347
Kansas	0
Kentucky	0
Louisiana	0
Maine	0
Missouri	0
Nevada	0
New Hampshire	0
New Mexico	0
North Carolina	0.352
North Dakota	0
Oklahoma	0
Pennsylvania	0
South Carolina	0
Tennessee	0
Utah	0
Virginia	0
Wisconsin	0
Wyoming	0

Table 4

Variables	Minnesota	Synthetic Minnesota
Avg. Weekly Wage	228	228
Taxable Wages	1.74e+08	1.73e+08
Employment 2012Q1	61,869	61,923

Table 5

Variables	
Minnesota	-554
	(874.7)
2014Q3	4,503***
	(770.8)
TreatEffect	285.8
	(1,160.2)
Constant	65,174***
	(576.2)
Observations	66
R-squared	0.50

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 6

Variables	
Minnesota	8.351e+06*
	(4.539e+06)
2014Q3	4.134e+07***
	(5.492e+06)
TreatEffect	1.818e+07**
	(8.812e+06)
Constant	1.831e+08***
	(3.010e+06)
Observations	66
R-squared	0.722

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Figure 1

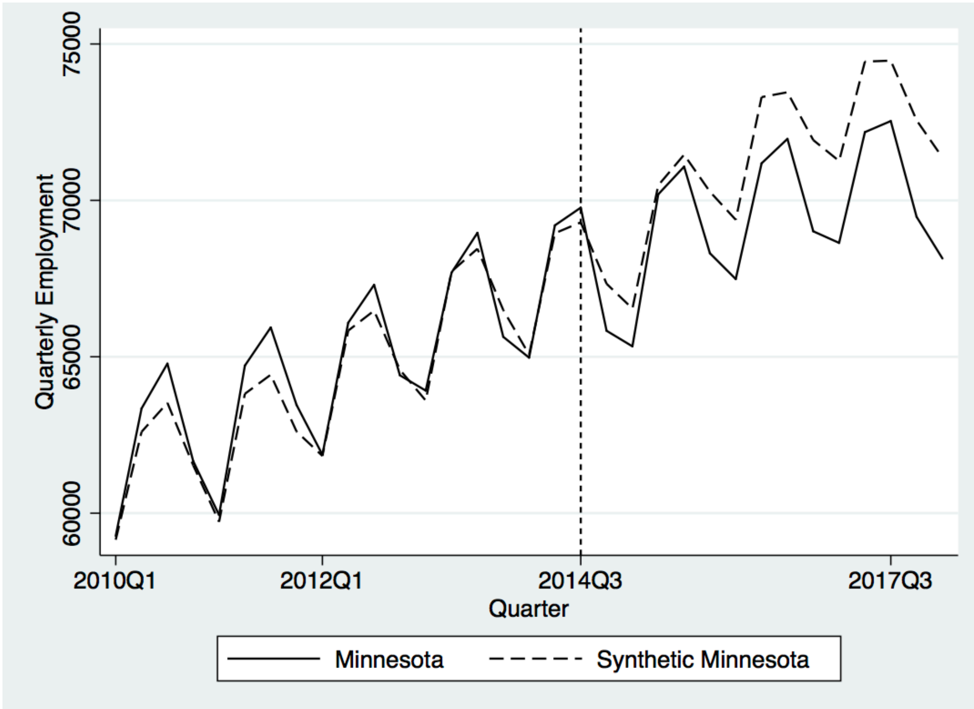

Figure 2

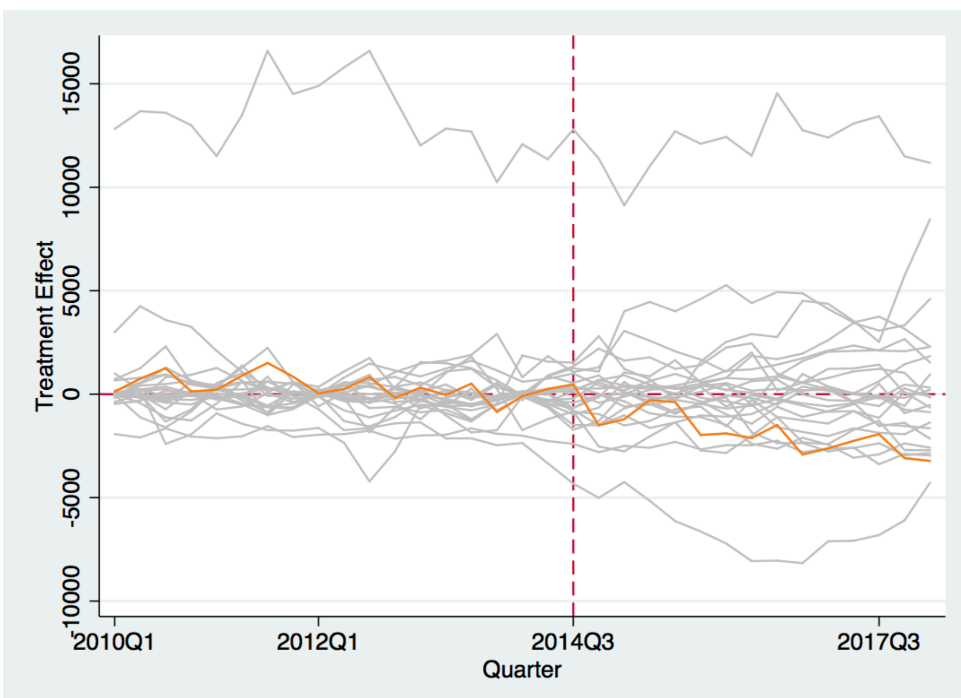

Figure 3

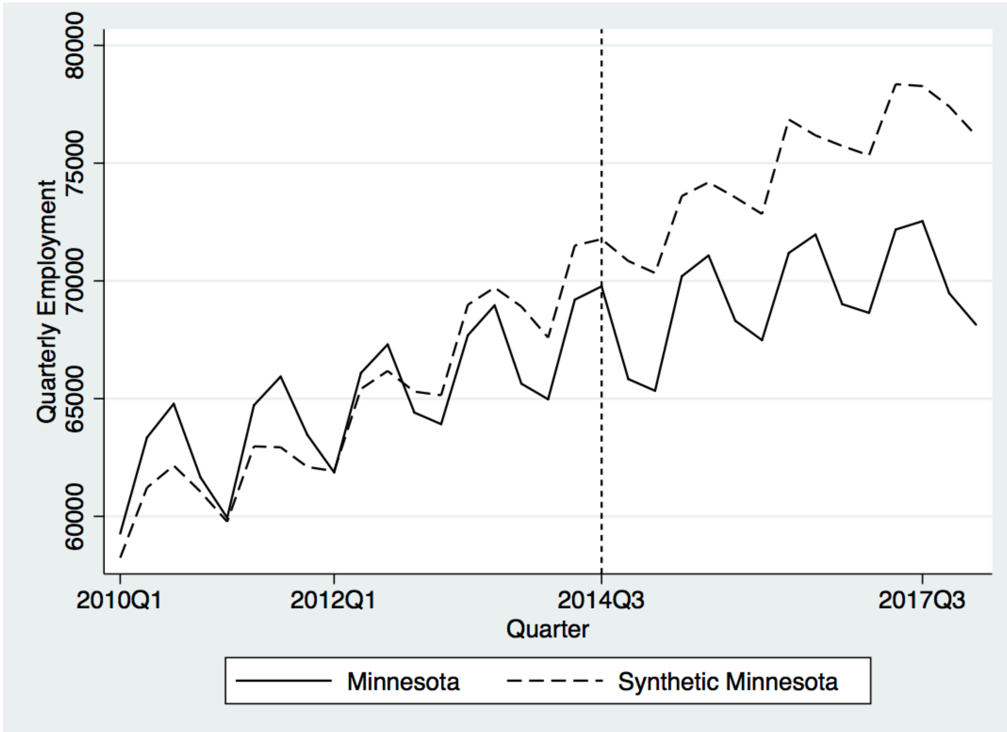

Figure 4

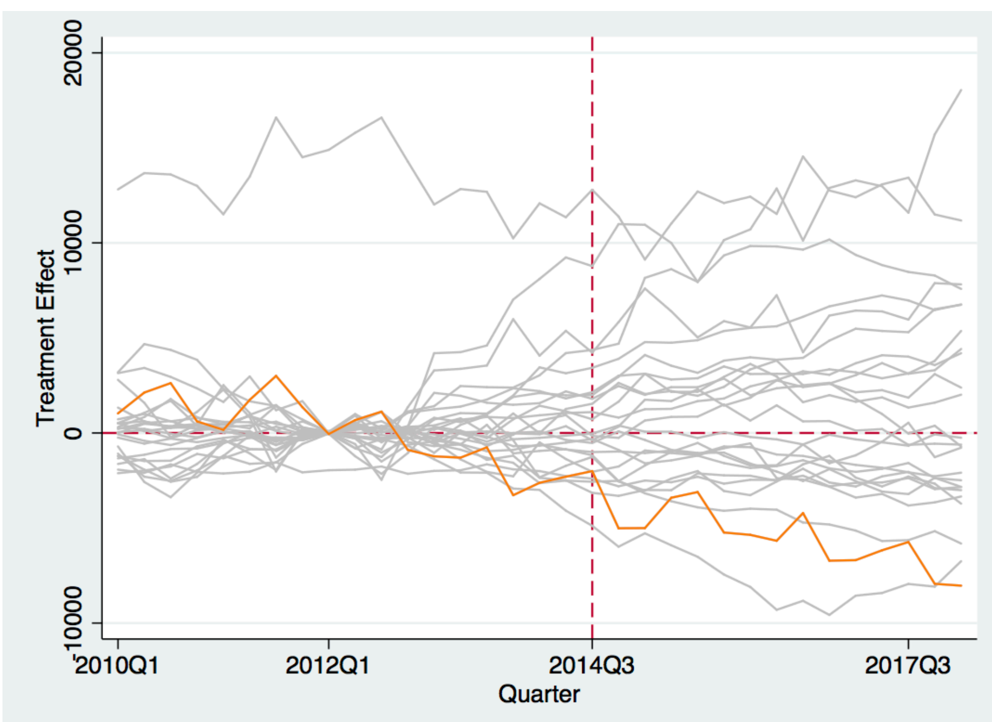

Figure 5

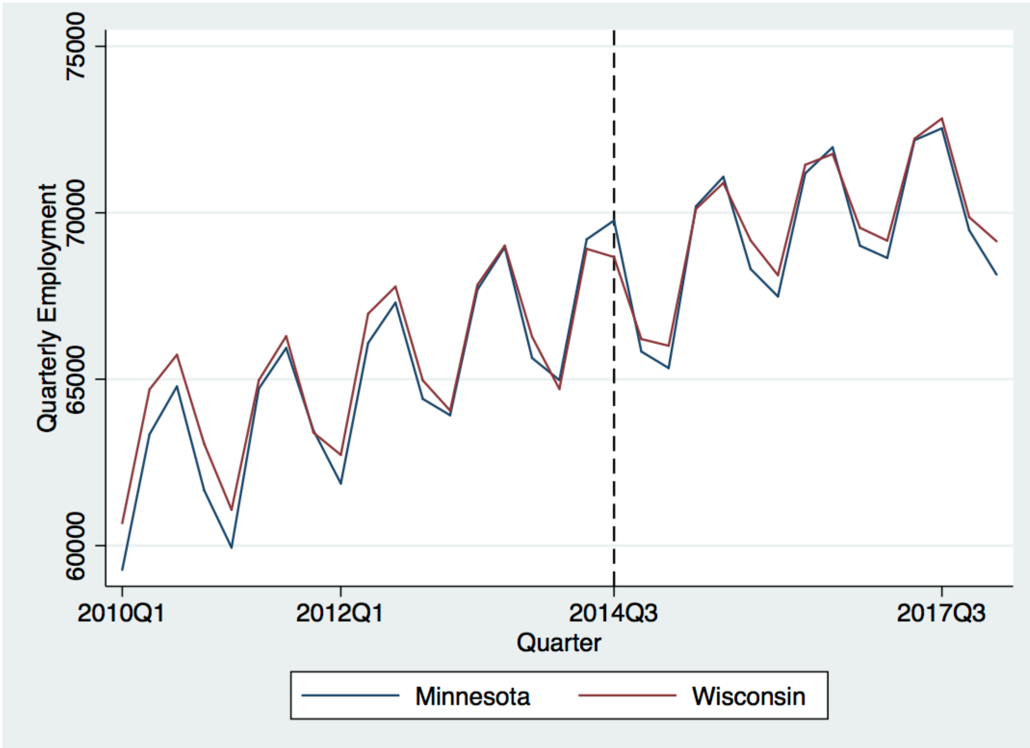

Figure 6

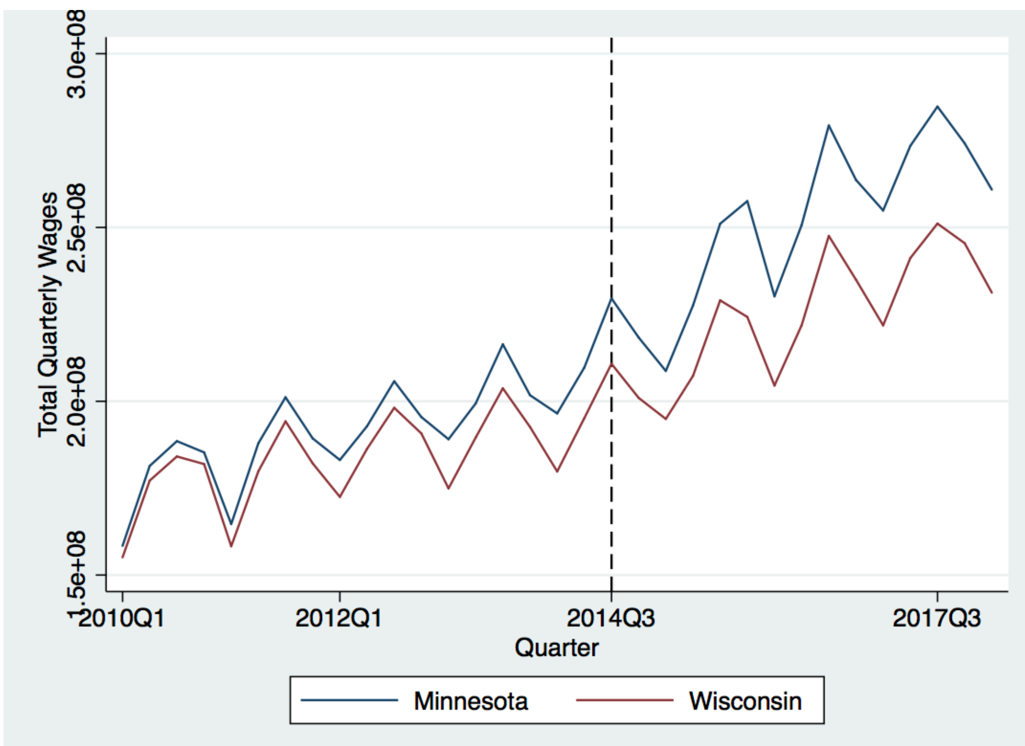

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association* 105(490): 493–505.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2011. Synth: an R package for synthetic control methods in comparative case studies. *Journal of Statistical Software.* 42(13): 1–17.

Abadie, Alberto and Javier Gardeazabal. 2003. The economic costs of conflict: a case study of the basque country. *American Economic Review* 93(1): 112-132.

Allegretto, Sylvia A., Arindrajit Dube, and Michael Reich. 2011. Do minimum wages really reduce teen employment? Accounting for heterogeneity and selectivity in state panel data. *Industrial Relations* 50(2): 205–40.

Card, David, and Alan B. Krueger. 1994. Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review* 84(4): 772– 93.

Chinn, Menzie, and Louis Johnston. 2018. The impact of minimum wage legislation in Minnesota vs. Wisconsin: contra Williams.

Chinn, Menzie. "Time Series Evidence on the Minimum Wage Impact in Minnesota vs. Wisconsin." *Econbrowser,* November 5, 2018, econbrowser.com/archives/2018/11/time-series-evidence-on-the-minimum-wage-impact-in-minnesota-vs-wisconsin.

Doucouliagos, Hristos, and Tom D. Stanley. 2009. Publication selection bias in minimum-wage research? A meta-regression analysis. *British Journal of Industrial Relations.* 47(2): 406-428.

Dube, Arindrajit. 2011. Review of *Minimum Wages* by David Neumark and William L. Wascher. *Journal of Economic Literature* 49(3): 762–66.

Dube, Arindrajit, T. William Lester, and Michael Reich. 2010. Minimum wage effects across state borders: Estimates using contiguous counties. *Review of Economics and Statistics* 92(4): 945–64.

Neumark, David, J. M. Ian Salas, and William Wascher. 2014. Revisiting the minimum wage– employment debate: Throwing out the baby with the bathwater? NBER Working Paper No. 18681. Cambridge, MA: National Bureau of Economic Research.

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

Powell, David. 2018. Imperfect Synthetic Controls: Did the Massachusetts Health Care Reform Save Lives? Santa Monica, CA: RAND Corporation.

Reich, Michael, Sylvia Allegretto, and Anna Godoey. 2017. Seattle's minimum wage experience 2015-16. Policy Brief. Center on Wage and Employment Dynamics, Institute for Research on Labor and Employment, UC Berkeley.

Williams, Noah. 2018. CROWE Policy Brief: Evidence on the Effects of Minnesota's Minimum Wage Increases.