NBER WORKING PAPER SERIES

MINIMUM WAGES AND LABOR MARKETS IN THE TWIN CITIES

Loukas Karabarbounis
Jeremy Lise
Anusha Nath

Working Paper 30239
http://www.nber.org/papers/w30239

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2022

We thank Evan Cunningham, Katerina Gribbin, and Pedro Tanure Veloso for excellent research assistance and Oriane Casale, Mustapha Hammida, and Steve Hine for help with the administrative data used in this paper. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the Minnesota Department of Employment and Economic Development, the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the National Bureau of Economic Research. All results have been reviewed by the Minnesota Department of Employment and Economic Development to ensure that no confidential information has been revealed.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w30239.ack

Minimum Wages and Labor Markets in the Twin Cities
Loukas Karabarbounis, Jeremy Lise, and Anusha Nath
NBER Working Paper No. 30239
July 2022
JEL No. J08,J23,J38

ABSTRACT

Using merged administrative datasets from Minnesota, we bring new evidence on the labor market effects of large minimum wage increases by examining the policy changes implemented by Minneapolis and Saint Paul. We begin by using synthetic difference-in-differences methods to estimate counterfactual outcomes at the zip code level from Minnesota and at the city level from the rest of the country. The minimum wage did not affect employment in most industries but exerted a negative impact on restaurants' employment, with an elasticity of -0.8. Next, using variation in exposure to the minimum wage across establishments and workers within the Twin Cities, we find employment effects that are half as large as those from the time series. The cross-sectional estimates difference out employment effects from the pandemic or civil unrest that could confound the time series comparisons, but they do not include potential effects of the minimum wage operating through equilibrium adjustments such as entry. We quantify a model of establishment dynamics to reconcile the different estimates and argue that they plausibly reflect lower and upper bounds of employment losses. We use the model to show that our estimates are consistent with an establishment elasticity of labor demand of -1 and illustrate how they can inform deeper parameters characterizing product and labor market competition, factor substitution, and establishment dynamics.

Loukas Karabarbounis
University of Minnesota
Department of Economics
Hanson Hall
Minneapolis, MN 55455
and NBER
loukas@umn.edu

Jeremy Lise
Department of Economics
University of Minnesota
and Federal Reserve Bank of Minneapolis
jlise@umn.edu

Anusha Nath
Federal Reserve Bank of Minneapolis
90 Hennepin Avenue
Minneapolis, MN 55401
Anusha.Nath@mpls.frb.org

1 Introduction

Increasing the minimum wage has recently reemerged as one the most debated economic policies in the United States. Despite decades of academic research, both policymakers and economists are still debating the quantitative impact of minimum wages on workers, firms, and regions. Estimating the effects of the minimum wage is also important for understanding which classes of models better characterize labor markets and for estimating the deeper parameters that affect product and labor market competition, factor substitution, and establishment dynamics.

We present new evidence from two minimum wage policies instituted recently in Minneapolis and Saint Paul that bears on this debate. In 2018, Minneapolis implemented a policy that aims to raise the minimum wage to 15 dollars per hour for all workers by 2024. Saint Paul introduced a similar policy in 2018 for implementation in 2020. The changes in the minimum wage are large by historical standards. By 2020, the minimum wage in Minneapolis increased by 38 percent relative to the statewide minimum wage that would have prevailed in the absence of the increase. Just in the first year of the policy change, the corresponding increase in Saint Paul is 19 percent.

Our analysis proceeds in three steps. First, we use time series variation to compare outcomes in the Twin Cities with those of appropriate control cities within Minnesota or the rest of the country. We find wage gains in most low-wage industries in Minneapolis, with the magnitude of the gains overlapping with our estimate of the direct effect of the minimum wage increase on establishments' labor costs. We do not detect significant employment changes for most industries. The minimum wage increase, however, is associated with a roughly 30 percent decline in employment for the restaurant industry in both Minneapolis and Saint Paul. Second, using the differential exposure of establishments and workers to the minimum wage within the Twin Cities, we find employment effects that are half as large as those from the time series. Our cross-sectional estimates are consistent with an elasticity of establishment labor demand of -1. Finally, we quantify a model of establishment dynamics to reconcile estimates from the cross section with those from the time series. We argue that they plausibly reflect lower and upper bounds of employment losses, because the losses from the time series may be confounded by factors contemporaneous with the minimum wage and the losses from the cross section may omit equilibrium effects such as changes in the entry rate of establishments.

For our analyses, we use a new administrative dataset on workers and establishments from Minnesota. The dataset merges worker-level Unemployment Insurance data with establishment-level Quarterly Census of Employment and Wages data to create a quarterly dataset on workers' hours and wages, as well as the establishments where they work by industry, zip code, and city. Our dataset covers the period between 2001(1) and 2020(4) and improves measurement relative to that of previous studies in three ways. First, we provide estimates from an administrative dataset for the effects of the minimum wage increase on hours worked.[1] Second, we include in our analyses firms with multiple establishments across city borders, which account for roughly 50 percent of jobs in the Twin Cities. Finally, we leverage detailed physical location data to increase the precision of our estimates in the time series and exploit within-city variation in the cross sections of establishments and workers.

We treat Minneapolis with an increase in the minimum wage for all periods after 2018, reflecting that the entire schedule of minimum wage increases was announced at once. We also treat Saint Paul with an increase in the minimum wage after 2018, because Saint Paul committed to the policy when the Minneapolis ordinance was introduced in 2017 and its own ordinance passed in 2018. Thus, any potential effects in Saint Paul before 2020 are consistent with advance notice of future minimum wage increases.

The fundamental challenge in identifying the treatment effect of the minimum wage increase is the construction of an appropriate control group that approximates the counterfactual outcome in the absence of the increase. For our time series analysis, we adopt the synthetic control groups approach proposed by Abadie and Gardeazabal (2003) and Abadie, Diamond, and Hainmueller (2015) and augment it with fixed effects, following Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021). Unlike standard difference-in-differences approaches, this synthetic difference-in-differences methodology constructs the control group so that pre-treatment trends are aligned between treated and control groups.

Studying large minimum wage increases in a period that includes the pandemic recession

[1] Some studies examining effects on hours worked (Zavodny, 2000; Couch and Wittenburg, 2001; Neumark, Schweitzer, and Wascher, 2004; Allegretto, Dube, and Reich, 2011) have largely used reported usual weekly hours from the Current Population Survey. This measure has been documented to contain significant measurement errors (Heckman, 1993; Bound, Brown, and Mathiowetz, 2001; Barrett and Hamermesh, 2019). Our measure of hours worked is more precise than these survey measures, suggesting less attenuation bias in our estimates.

presents both a challenge and an opportunity. The pandemic offers a rare opportunity to study large increases in the minimum wage during a recession, as it impacted more severely low-wage industries facing larger exposures to the minimum wage. We document some statistically and economically significant declines in restaurants' employment before the pandemic. For restaurants in Saint Paul, the decline in employment does not accelerate significantly after the pandemic. However, for restaurants in Minneapolis, the employment effects either appear or accelerate significantly in 2020.

The post-pandemic results from the time series should be interpreted with caution, as the sensitivity to aggregate shocks may have changed for the synthetic control in 2020 relative to that of the Twin Cities. For example, the economic impact of lockdowns may have been more significant in larger, more densely populated cities than in the smaller cities or suburbs that compose our control group. We partly overcome this challenge for the interpretation of our results by using variation from other U.S. cities that are also densely populated and large and faced similar or even lengthier lockdowns than the Twin Cities. Using the sample of other U.S. cities also allows us to control for nationwide changes in economic conditions that either were induced by the pandemic or accelerated during the pandemic. Examples include demand changes such as substitution of services prone to virus transmission with online shopping and supply changes such as the rise of gig work and labor shortages in certain industries. Using the sample of other U.S. cities, we continue to find similar in magnitude jobs decline in the restaurant industries in the Twin Cities, which persist through the end of 2021. To the extent that the alternative source of variation across U.S. cities allows us to control for the effects of the recession, our estimates suggest that the recession interacted with the minimum wage increase to magnify the pre-recession job losses.

As with any research design that uses time series variation, it may still be case that the Twin Cities experienced idiosyncratic shocks or had a heterogeneous response to aggregate shocks that cannot be differenced out in the post-treatment period. An example of such an idiosyncratic shock is the civil unrest in 2020, which impacted the Twin Cities differently from other cities in Minnesota or the country. In the second part of our paper, we shift our research design away from using other cities in the control group. Instead, we use variation from the cross sections of establishments and workers within a city. The cross-sectional estimates do not suffer from the concern that other factors confound the effects of the minimum wage, to the extent that the Twin Cities shocks are differenced out either across establishments that belong to the same industry and zip code or across workers who belong to the same industry.

We demonstrate that establishments with larger exposure to the minimum wage experienced larger increases in their wage and larger declines in their jobs, hours, and wage bill. The ratio of employment to wage changes induced by exposure to the minimum wage policy is -1. Reassuringly for our research design, which attempts to difference out common factors at the zip code and industry level, we find that the response of all variables to exposure to the minimum wage is remarkably stable between 2019 and 2020. A concern about the interpretation of the results that use variation across establishments within an industry and zip code is that workers may have reallocated from exposed to non-exposed establishments or found jobs outside of their zip code.[2] For this reason, we also analyze the cross section of workers whose jobs we can track everywhere within Minnesota. Similar to establishments, workers who are more exposed to the minimum wage experience significantly larger employment and earnings losses.

To summarize our various estimates for the jobs effects, the time series analysis shows that the minimum wage increase was associated with an average jobs decline of roughly 2 percent in the Twin Cities. The great majority of job losses appear in the restaurant industry, which experienced a jobs decline of roughly 30 percent. The analysis using variation from the cross section leads to estimates of job losses around half as large as the estimates from the time series.

In the last part of the paper, we attempt to reconcile the estimates from the cross section with those from the time series. Despite our efforts to difference out other shocks, the Twin Cities may have experienced idiosyncratic shocks or had a differential response to an aggregate shock that cannot be differenced out using other cities during the post-treatment period. Time series effects of the minimum wage on employment sum up employment effects at the intensive margin, effects arising from the exit of establishments, and effects arising from the lack of entry of new establishments. By design, the estimates from the cross section do not account for the effects of entry, because they use establishments and workers that exist for at least one period.

We develop a model of establishment dynamics to formalize our intuition that it is plausible

[2] Dustmann, Lindner, Schonberg, Umkehrer, and vom Berge (2022) document the reallocation of low-wage workers from smaller, lower-paying, and less productive establishments to larger, higher-paying, and more productive establishments following the introduction of a national minimum wage in Germany in 2015.

to reconcile the time series estimates with those from the cross section by appealing to entry dynamics that are omitted from the analysis of the cross section. When we quantify the model to reproduce the establishment responses we observe in the cross section, we find that the industry-level employment decline is of similar magnitude to that we estimate in the time series of restaurants. The similarity of the responses is reassuring for the potential of endogenous entry to account for the difference between employment effects in the time series and the cross section. We calibrate the entry cost independently of our time series treatment effects by appealing to the industrial organization literature for estimates of the entry cost. To be clear, our results cannot identify whether missing entry, other equilibrium effects, or confounding factors account for the difference between the cross section and the time series. However, they point out that both types of analysis are informative and offer plausible bounds of the effects of the minimum wage increase.

Our paper contributes to various strands of literature. For most low-wage industries, we fail to detect statistically significant effects on employment. The confidence intervals we estimate for these industries overlap with estimates found in papers that examine effects on low-wage workers across industries (Allegretto, Dube, and Reich, 2011; Cengiz, Dube, Lindner, and Zipperer, 2019; Dube and Lindner, 2021; Neumark and Shirley, 2021). By contrast, our estimated employment elasticity with respect to the minimum wage for restaurants is more negative than the typical estimate found in the literature. For example, the comprehensive analysis of estimates for low-wage workers and low-wage industries in Neumark and Shirley (2021) reveals that roughly 80 percent of estimates are negative, with the average reported elasticity across studies being around -0.15. Our own analysis of published estimates for the restaurant industry finds an average elasticity of -0.1. This includes studies from the evaluation of city-wide minimum wage increases (Dube, Naidu, and Reich, 2007; Jardim, Long, Plotnick, Van Inwegen, Vigdor, and Wething, 2022).

Our estimated jobs elasticity with respect to the minimum wage for restaurants is between -0.4 and -0.8. The employment impacts we document might be larger than those in the literature because the policy change we examine is larger. The 38 percent increase in the Minneapolis minimum wage by 2020 is significantly larger than the policy changes classified as "large minimum wage increases" in Clemens and Strain (2021), which range between 20 and

25 percent. The authors estimate a -0.5 elasticity of low-skill employment with respect to the minimum wage for large policy changes and a non-negative elasticity for small policy changes.

An interpretation of our results is that the jobs elasticity is more negative when both a higher minimum wage and a recession are in effect. For example, the Seattle study by Jardim, Long, Plotnick, Van Inwegen, Vigdor, and Wething (2022) estimates an elasticity of -0.3 for restaurants. However, Seattle was booming during the implementation of their large minimum wage increase, whereas the Twin Cities were hit by a recession. The acceleration of job losses during the pandemic recession is consistent with earlier studies such as Addison, Blackburn, and Cotti (2013), who document the larger sensitivity during recessions, and more recent work such as Clemens and Wither (2019), who find an employment elasticity with respect to the minimum wage of -1 when analyzing the increase in federal minimum wage during the Great Recession.

Through the lens of the model, we show that our cross-sectional estimates are consistent with an elasticity of labor demand at the establishment level of -1. We interpret this elasticity as a long-run elasticity for establishments operating in competitive labor markets and can flexibly adjust all other inputs. Our estimate is comparable to some other estimates found the literature (Hamermesh, 1993; Beaudry, Green, and Sand, 2018). A puzzle in the labor economics literature is how to reconcile small or zero employment effects of the minimum wage with larger estimated elasticities of labor demand at the establishment level. One potential answer is that minimum wage increases have taken place in regions characterized by imperfectly competitive labor markets and that these increases are small enough relative to the competitive wage. Our results are not as puzzling as those in the literature, because we find an elastic labor demand at the establishment level and large employment effects from the introduction of the minimum wage for some low-wage industries.

We caution readers that our results cannot be used to directly identify how competitive labor markets were in the Twin Cities before the minimum wage increase. The model that we can reject is that there is labor market power and that the increase in the minimum wage was sufficiently small to induce an equilibrium wage below the competitive level. Our results are compatible with either perfectly competitive labor markets or the introduction of a minimum wage above the competitive level for most establishments that operate in a labor market with some degree of monopsony power.

2 Policy Background and Data Sources

In this section, we detail the policy background underlying the minimum wage increases in the Twin Cities and describe our sources of data.

2.1 Minimum Wage Policy

Table 1 presents the statewide minimum wage for Minnesota during the period of our study. The last policy change in Minnesota occurred in 2014. The minimum wage was set to reach 7.75 dollars for small firms and 9.50 dollars for large firms by 2016. Beginning in 2018, the statewide minimum wage has been indexed to the price deflator.

Table 2 details the minimum wage policy changes introduced by the cities of Minneapolis and Saint Paul. In 2018, Minneapolis increased the minimum wage for establishments that operate within the city. The increase was implemented in phases, with the goal of reaching 15 dollars per hour in July 2022 for large firms and in July 2024 for small firms. In 2018, Saint Paul followed Minneapolis in adopting a 15 dollar minimum wage policy. Saint Paul also enacted a phased implementation that began increasing its minimum wage in 2020, with the goal of reaching 15 dollars for all firms by 2027.[3]

These changes in the minimum wage of the Twin Cities are large relative to other minimum wage changes (even relative to the ones classified as large in the analysis of Clemens and Strain, 2021). Using the average minimum wage applicable to small and large firms, we find that in 2020 the minimum wage in Minneapolis increased by 38 percent relative to implementing the statewide minimum wage. In 2020, the minimum wage in Saint Paul increased by 19 percent relative to the statewide minimum wage. Further, in both cities the minimum wage will be indexed to inflation once the target level of 15 dollars per hour is reached. Thus, to the extent that firms and workers do not perceive further increases in the statewide minimum wage, the minimum wage increase for operating within the Twin Cities is permanent.

[3] For the statewide minimum wage policy, gratuities are excluded from the minimum wage, and employers have to pay their employees a wage that exceeds the minimum wage before gratuities are applied. The Minneapolis and Saint Paul minimum wage ordinances adopted the same policy with respect to gratuities.

Table 1: Minimum Wage Policy in the State of Minnesota

(Annual Revenue in Dollars)	Youth	Small Firms ($< 500,000$)	Large Firms ($\geq 500,000$)
2000-2005	4.25	4.90	5.15
2006-2013	4.90	5.25	6.15
2014	6.50	6.50	8.00
2015	7.25	7.25	9.00
2016	7.75	7.75	9.50
2017	7.75	7.75	9.50
2018	7.87	7.87	9.65
2019	8.04	8.04	9.86
2020	8.15	8.15	10.00
2021	8.21*	8.21*	10.08*

Notes: Symbol * denotes that the minimum wage is scheduled to increase every year according to the price deflator for personal consumption expenditures produced by the Bureau of Economic Analysis. The threshold of $500,000$ represents revenue reported to the state of Minnesota.

Table 2: Minimum Wage Policy in the Twin Cities

	Minneapolis		Saint Paul			
Firms (Employees)	Small (< 100)	Large (≥ 100)	Micro (≤ 5)	Small (6-100)	Large (101-10,000)	Macro ($> 10,000$)
2018 (Jan)		10.00				
2018 (Jul)	10.25	11.25				
2019 (Jul)	11.00	12.25				
2020 (Jan)						12.50
2020 (Jul)	11.75	13.25	9.25	10.00	11.50	
2021 (Jul)	12.50	14.25	10.00	11.00	12.50	
2022 (Jul)	13.50	15.00*	10.75	12.00	13.50	15.00*
2023 (Jul)	14.50		11.50	13.00	15.00	
2024 (Jul)	equal to large		12.25	14.00	equal to macro	
2025 (Jul)			13.25	15.00		
2026 (Jul)			14.25	equal to macro		
2027 (Jul)			15.00			
2028 (Jul)			equal to macro			

Notes: Symbol * denotes that the minimum wage is scheduled to increase every year according to the price deflator for personal consumption expenditures produced by the Bureau of Economic Analysis. The size thresholds represent total firm employment across all establishments. Franchises are considered large firms if they have more than 10 franchises nationally. For full-service restaurants, if there are fewer than 10 locations nationally, each restaurant counts as a separate business for the purpose of determining size.

2.2 Data Sources

We use two main sources of data on workers and establishments for our analysis of the effects of the minimum wage increase. Both sources are administrative and non-publicly-available data that were provided to us by Minnesota's Department of Employment and Economic Development (DEED).

The first data source is individual-level data on workers from Unemployment Insurance (UI). Minnesota requires most employers to file unemployment wage detail reports quarterly for the purpose of estimating the amount of unemployment insurance tax they owe. These reports provide us with data on quarterly earnings and hours worked for each worker. We calculate the hourly wage for each worker by dividing total quarterly earnings by quarterly hours.[4] Minnesota collects these data for each employee of a firm at the level of the establishment where they work. This feature of the data is especially important in studying the minimum wage effects, as roughly 50 percent of employment is generated in multi-establishment firms.[5]

The UI data do not contain information on the physical location of establishments, which is necessary in order to identify which establishments were affected by the minimum wage increase. To overcome this problem, the UI data is merged with establishment-level data from the Quarterly Census of Employment and Wages (QCEW). The QCEW records jobs that account for roughly 97 percent of employment in Minnesota. From these data, we observe the six-digit NAICS code for the industry that the establishment operates in, the physical location of the workplace, and the firm to which the establishment belongs. The physical location data consist of both the city and the zip code in which the establishment operates.[6]

The merged data result in a quarterly dataset between 2001(1) and 2020(4). Our geographic unit of analysis is a zip code within a city. This allows the same zip code to be affected differently

[4]We exclude roughly 1 percent of observations with jobs that reported a hourly wage below the applicable youth minimum wage for Minnesota. For calculating the wage, we exclude the roughly 5 percent of observations that reported zero hours worked. We keep these observations for calculating other outcomes.

[5]While reporting is required at the establishment level, a few firms file reports under a single account. We have to exclude from our analysis multi-establishment firms that have at least one establishment in Minneapolis or Saint Paul and at least one establishment outside of Minneapolis or Saint Paul, but report the UI wage details of all their employees under a single account. These establishments constitute roughly 3 percent of all establishments and 5 percent of all wage records.

[6]The raw data do not have physical location information for roughly 4 percent of observations. In addition, we exclude observations for which the city name and zip codes are contradictory. Such contradictions are rare and constitute roughly 0.1 percent of total establishments.

by the minimum wage policy if it belongs to two different cities. It also allows for multiple treated units within a city that faces an increase in its minimum wage. For each industry, we calculate the average wage, aggregate number of jobs (sum of full-time and part-time jobs), aggregate hours, and aggregate earnings paid within geographic units for each quarter. For the wage, we calculate the average hourly wage per worker for workers below the 90th percentile within the industry and year and treatment or control groups.[7] Finally, we aggregate all units that have fewer than 50 full-time equivalent jobs to one unit, separately for each industry and for treatment or control groups.

To summarize, by merging the worker-level UI data with the establishment-level QCEW data, we are able to create a dataset on workers' hours and wages, as well as their establishments of employment by industry, zip code, and city. Our dataset improves measurement relative to that of the typical minimum wage study along three dimensions. First, using administrative sources, we provide estimates for the effects of a minimum wage increase on hours worked.[8] Second, Minnesota is unique in that it records employee hours worked at the establishment level within firms. This feature allows us to include in our analyses firms with multiple establishments across city borders.[9] Finally, we leverage physical location data at the zip code level to increase the precision of our estimates and conduct additional analyses at the establishment level that require within-city variation.

We focus our baseline analyses on the two-digit industries in which 30 percent or more of workers earn below 15 dollars per hour in 2017 (see Online Table A.1 for the detailed estimates). The six industries that satisfy this criterion are retail trade (44); administrative services (56); health care and social assistance (62); arts, entertainment, and recreation (71); accommodation and food services (72); and other services (81), which consists of repair and maintenance shops, personal and laundry services, and various civic, professional, and religious organizations. In

[7]We examined wage results without excluding the top of the wage distribution. This measure of the wage is too noisy because of outliers at the top of the wage distribution and, in many cases, leads to statistically insignificant results for the wage. We also examined results that trimmed workers at the 75th percentile and generally find similar wage effects as those in our baseline measures, which trim at the 90th percentile.

[8]There are only four states in the U.S. that collect hours worked in the matched employer-employee administrative data, the other three being Oregon, Rhode Island, and Washington. In Minnesota, it is mandatory to report hours to UI.

[9]For example, this crucial subsample of firms, accounting for roughly 40 percent of jobs, was missing in the evaluation of the minimum wage increase in Seattle by Jardim, Long, Plotnick, Van Inwegen, Vigdor, and Wething (2022).

addition, we separately analyze full-service (722511) and limited-service (722513) restaurants, which have a high fraction of potentially impacted workers and have been studied extensively in the literature.

3 Evidence from the Time Series

We begin by laying out the econometric framework for analyzing the impact of the minimum wage increase using variation from the time series of cities. Next, we present our baseline estimates and various sensitivity checks. We conclude by discussing potential challenges in interpreting our results from the time series.

3.1 Minimum Wage Treatment

We analyze the labor market effects of the increase in the minimum wage separately in Minneapolis and Saint Paul. We exclude Saint Paul from the control group in our analysis of Minneapolis. Similarly, we exclude Minneapolis from the control group in our analysis of Saint Paul. Our choice to analyze Minneapolis and Saint Paul in parallel, as opposed to merging them in one treated unit, is appropriate because the Minneapolis ordinance was implemented in 2018, whereas the Saint Paul ordinance was implemented in 2020. We choose to treat both cities with an indicator of a minimum wage increase after 2018. Our logic for treating both cities in 2018 is that Saint Paul committed to a minimum wage policy immediately after Minneapolis passed its ordinance in 2017 and Saint Paul passed its ordinance in 2018 for implementation in 2020. Our logic for adopting a treatment indicator that covers the entire period after 2018 is that both cities announced the entire schedule of minimum wage increases in Table 2 once and not in increments. To the extent that establishments face entry costs or adjustment costs of changing their labor inputs or production techniques, we expect them to react upon the announcement of the schedule. As a result, the results we present in Saint Paul before 2020 reflect advance notice of the minimum wage increasing from 2020.

3.2 Econometric Methodology: Time Series

The key to analyzing the impact of a minimum wage increase is the credible estimation of the counterfactual in the absence of the minimum wage increase. To construct the counterfactual, we use synthetic control methods developed originally by Abadie and Gardeazabal (2003) and Abadie, Diamond, and Hainmueller (2015) and extended recently by Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021). In this section, we present the methodology from these papers and explain why it is appropriate to use it to analyze the minimum wage increase in the Twin Cities.

We have a balanced panel with N geographic units for T periods. The outcome for unit i in period t is Y_{it}. Exposure to the treatment of a minimum wage increase is $W_{it} \in \{0, 1\}$, where $W_{it} = 0$ denotes that unit i did not experience a minimum wage increase in period t and $W_{it} = 1$ denotes that it did. We order units so that the first N_{co} units are never exposed to the treatment, while the last $N_{\text{tr}} = N - N_{\text{co}}$ units are exposed to the treatment after time T_{pre}. We have multiple treated units because the unit of analysis is a zip code within a city. All zip codes in Minneapolis and Saint Paul are treated with a minimum wage increase after 2018.

Let Y_{it}^1 denote the outcome for unit i in period t if the unit has been exposed to the minimum wage increase. Let Y_{it}^0 denote the counterfactual outcome that we would have observed in the absence of the minimum wage increase. The average treatment effect in period t is $\tau_t = \frac{1}{N_{\text{tr}}} \sum_{i=N_{\text{co}}+1}^{N} \left(Y_{it}^1 - Y_{it}^0 \right)$ and the average treatment effect across all periods is $\tau = \frac{1}{T - T_{\text{pre}}} \sum_{t=T_{\text{pre}}+1}^{T} \tau_t$.

The fundamental problem in estimating the treatment effect is that the counterfactual outcome Y_{it}^0 is not observed, because unit i is exposed to the minimum wage increase at time t. Since the seminal study of Card and Krueger (1994) on the minimum wage increase in New Jersey, a popular method to overcome this problem has been to find a control group of non-treated units and use its post-treatment outcomes to estimate the counterfactual Y_{it}^0 for treated units. With multiple units and time periods in the sample, this amounts to a two-way fixed effects regression

$$Y_{it} = \alpha_i + \beta_t + \tau W_{it} + u_{it}, \tag{1}$$

where α_i is a unit fixed effect, β_t is a time fixed effect, and u_{it} is the error term.

The difference-in-differences specification in equation (1) assumes that outcomes of treated and non-treated units are equal (up to a constant) in the post-treatment period in the absence of the minimum wage increase. This "parallel trends" assumption cannot be tested in the post-treatment period, because we do not observe the counterfactual for treated units. Typically, the plausibility of parallel trends is assessed by evaluating whether trends are parallel during the pre-treatment period. If in the pre-treatment period non-treated units experience different trends than treated units, then we have low confidence that non-treated units are an appropriate control group during the post-treatment period.

The concern with the difference-in-differences specification is that there is no control group with pre-treatment outcomes that resemble those of treated units. Synthetic control methods, such as those in Abadie and Gardeazabal (2003) and Abadie, Diamond, and Hainmueller (2015), aim to overcome this problem by finding a vector of weights $\hat{\omega}$ that forces pre-treatment trends in the outcomes for the non-treated units to align with pre-treatment trends in the outcomes for the treated units. More explicitly, the goal is to find weights such that $\sum_{i=1}^{N_{\text{co}}} \hat{\omega}_i Y_{it} \approx N_{\text{tr}}^{-1} \sum_{i=N_{\text{co}+1}}^{N} Y_{it}$ for each time period before the treatment $t = 1, \ldots, T_{\text{pre}}$.

Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021) propose a synthetic difference-in-differences methodology, which uses estimating equation (1) and, additionally, weights observations with ω_i so that treated and non-treated units are as close as possible in terms of pre-treatment outcomes. The weights are estimated as

$$(\hat{\omega}_0, \hat{\omega}) = \arg \min_{\omega_0 \in \mathbb{R}, \omega \in \Omega} \sum_{t=1}^{T_{\text{pre}}} \left(\omega_0 + \sum_{i=1}^{N_{\text{co}}} \omega_i Y_{it} - \frac{1}{N_{\text{tr}}} \sum_{i=N_{\text{co}}+1}^{N} Y_{it} \right)^2 + \zeta^2 T_{\text{pre}} ||\omega||_2^2, \tag{2}$$
$$\Omega = \left\{ \omega \in \mathbb{R}_+^N : \sum_{i=1}^{N_{\text{co}}} \omega_i = 1, \omega_i = N_{\text{tr}}^{-1} \text{ for all } i = N_{\text{co}} + 1, \ldots, N \right\}.$$

Following these authors, we allow for a shifter ω_0 that aligns the pre-treatment trends for the synthetic control and the treated units up to a constant, which will be differenced out by the fixed effect specification in equation (1). The regularization parameter ζ penalizes non-zero weights to ensure the minimization problem has a unique solution.[10]

If we use the estimated $\hat{\omega}$ from equation (2) as weights in the estimating equation (1), the synthetic difference-in-differences treatment effect $\hat{\tau}$ is

$$\left(\hat{\tau}, \hat{\alpha}, \hat{\beta}\right) = \arg \min_{\tau, \alpha, \beta} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(Y_{it} - \alpha_i - \beta_t - \tau W_{it}\right)^2 \hat{\omega}_i \right\}. \tag{3}$$

[10] Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021) propose a penalty term to match the size of a typical one-period change in outcome variable Y. In practice, we find that a small penalty of $\zeta = 10^{-6}$ works well in terms of minimizing the weight on control units with dissimilar pre-trends to treated units.

Removing the estimated weights $\hat{\omega}_i$ from the least-squared problem in equation (3) leads to the standard difference-in-differences specification. Removing the unit fixed effects α_i from equation (3) and ω_0 from equation (2) leads to the standard synthetic control specification.

Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021) also propose choosing time weights λ_t to balance the pre-treatment and the post-treatment periods for the control group. A problem with using time weights is that the weights may change significantly as additional quarters of data become available. For our baseline, we settle on equally weighting all pre-treatment periods to keep the analysis as transparent as possible. However, as a robustness check, we also present analyses using estimated time weights.

To infer the statistical significance of the estimated $\hat{\tau}$, we use the "placebo method." The method takes all non-treated units and estimates the treatment effect τ in these samples, with each sample generated under a placebo treatment of a subset of non-treated units. Since we should be estimating a zero treatment effect in the absence of a treatment, the distribution of treatment effects under the placebo method gives us the distribution of noise inherent in the data.[11]

Figure 1 illustrates the use of synthetic difference-in-differences and contrasts it to standard difference-in-differences in the context of the minimum wage increase in Minneapolis. The upper panels of the figure plot quarterly time series of the average hourly wage and the total number jobs for the full-service restaurant industry between 2001(1) and 2020(4). All series are in logs and normalized to 0 in 2017(4), which is the last quarter before the minimum wage increased in Minneapolis. The solid lines in the upper panel show that full-service restaurants in Minneapolis experienced a significant increase in the wage and jobs over time.

The long-dashed blue lines show the evolution of the wage and jobs for the average of all cities in Minnesota besides Minneapolis and Saint Paul. This average represents the control group in a difference-in-differences specification that uses other cities in Minnesota in the regression. The jobs trends before 2018 in Minneapolis and in other cities in Minnesota are significantly different. The lack of parallel pre-treatment trends casts doubt on the assumptions underlying a standard

[11] See Algorithm 4 in Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021) for implementation details to construct the placebo standard errors. To confirm the robustness of our conclusions, we also examined bootstrapped standard errors, which are popular in difference-in-differences settings (Bertrand, Duflo, and Mullainathan, 2004). Our conclusions do not change when we use bootstrapped standard errors, and thus we omit them from our presentation.

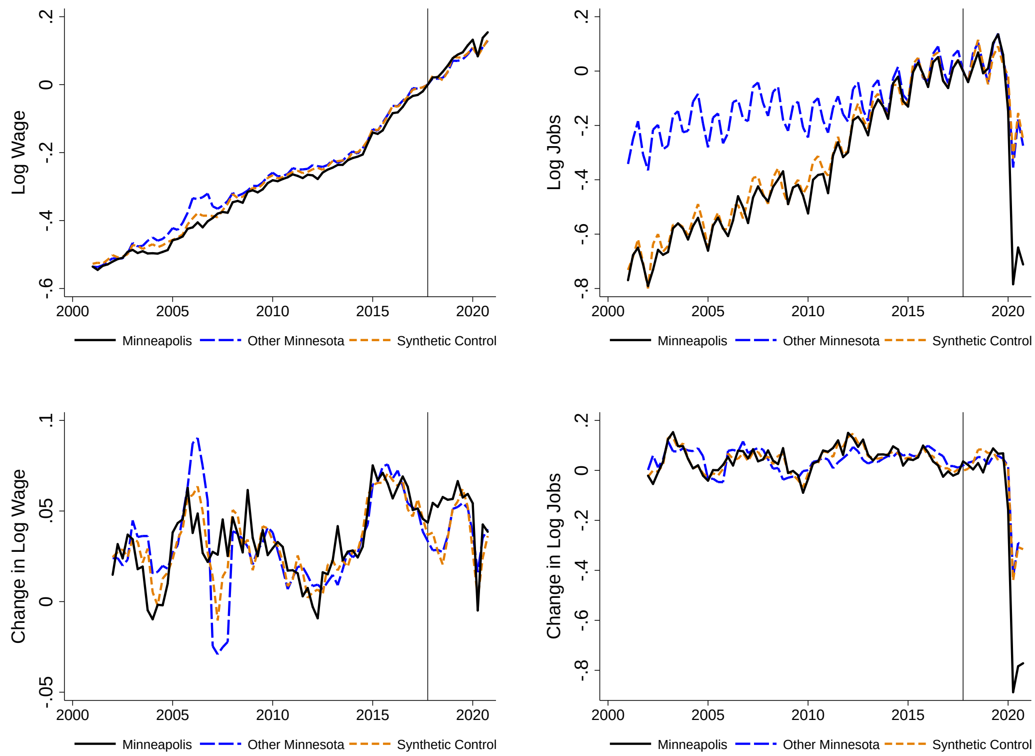

Figure 1: Time Series of Full-Service Restaurants

difference-in-differences strategy. Thus, the average of all cities in the rest of Minnesota is not a suitable control group for Minneapolis.[12]

The dashed orange line shows the evolution of the wage and jobs for the synthetic control of Minneapolis, which is the weighted average of cities in Minnesota other than Minneapolis and Saint Paul using weights $\hat{\omega}_i$. By design, the methodology weights more heavily cities with similar pre-treatment trends and less heavily cities with different pre-treatment trends. As the figure shows, the time series for the synthetic control reproduce very closely the time series of the wage and jobs in Minneapolis in the pre-treatment period.

In our analyses below, we focus on outcome variables that are expressed in yearly growth rates. The lower panels of Figure 1 demonstrate that the fit of the synthetic control during the pre-treatment period is significantly improved relative to that of the unweighted average that

[12] Another popular specification in the minimum wage literature is to add unit-specific linear time trends to equation (1). However, pre-treatment trends could be non-linear. An example of non-linearity in our context is retail trade in Minneapolis which exhibits a secular decline in the 2000s, stability in the first part of the 2010s, and an upward trend after 2015. See Online Figures A.1 to A.16 for the time series in the other low-wage industries that are included in our analyses.

underlies the difference-in-differences specification. Using synthetic difference-in-differences, we can visualize the effect of the minimum wage increase on the growth of the wage and jobs as the difference in the post-2018 period between the dashed orange lines and the solid lines.

3.3 Growth Specification of Synthetic Difference-in-Differences

We express outcome variables Y_{it} in equation (3) in growth rates. We prefer a specification in growth rates to a specification in levels for two reasons. If the Twin Cities implemented a minimum wage policy because they were growing at a different rate than other cities, that would invalidate the identifying assumption that the treatment effect is independent to other determinants of outcome variables.[13] The unit fixed effect α_i in a growth specification removes heterogeneity in average growth rates that may be correlated with the treatment of increasing the minimum wage. Additionally, using yearly growth rates allows us to remove quarterly seasonal variation, thus improving the efficiency of our estimates.

Accordingly, if y_{it} is a time series in levels, we take year-over-year differences in logs and define

$$Y_{it} \equiv \log y_{it} - \log y_{i,t-4}, \forall i = 1, ..., N_{\text{co}}, \quad Y_{it} \equiv (\log y_{it} - \log y_{i,t-4})\, \bar{\nu}_i, \forall i = N_{\text{co}} + 1, ..., N. \quad (4)$$

In equation (4), we weight zip codes of the treated cities with their share $\bar{\nu}_i$ of the corresponding variable in the three years before to the minimum wage increase. Doing so allows us to interpret the treatment effect as pertaining to the city as a whole as opposed to the average zip code within a city.[14] Holding the zip code weights $\bar{\nu}_i$ constant over time allows us to interpret the treatment effects as counterfactual outcomes that the Twin Cities would have experienced in the absence of the minimum wage increase, holding the spatial distribution of economic activity constant at the same levels observed just before the policy change.

Working with the outcome variable in equation (4) means that our treatment τ is the effect of the minimum wage increase on the average yearly growth rate of the variable over the entire

[13] Ferman and Pinto (2021) show that the synthetic control estimator is biased if treatment assignment is correlated with the factor structure underlying the dynamics of outcome variables, even when the number of pretreatment periods goes to infinity.

[14] The exception is the wage, for which we do not use any weights. The reason is that we are interested in the effects of the minimum wage increase on the wage of the average worker. For the control units, we do not weight the growth rates of zip codes, because these weights enter multiplicatively with the synthetic control weights ω_i in equation (3).

post-treatment period, $T - T_{\text{pre}}$. We transform the growth effect into a cumulative effect up to final period T on the (log) variable with the formula

$$g_T \equiv \mathbb{E}\big(\log y^1_{i,T} - \log y^0_{i,T}\big) = \frac{(T - T_{\text{pre}})\tau}{4}, \tag{5}$$

where 4 appears in the formula because τ is a yearly, as opposed to a quarterly, growth rate. Variable g_T is the log change in outcome y in the final period T due to the minimum wage increase between periods $T_{\text{pre}+1}$ and T.

Before presenting the impacts of the minimum wage increase, we pause to discuss the synthetic control method's performance in accounting for the growth of the treated units, Minneapolis and Saint Paul, in the period before the minimum wage increase. In Online Table A.2, we present R-squared coefficients from regressions of growth in Minneapolis or Saint Paul on the growth of the synthetic control calculated using the weights $\hat{\omega}_i$. The regressions are performed only during the pre-treatment period. We find that for five out of the six low-wage industries identified in Section 2 and separately for restaurants, the synthetic control accounts for a substantial fraction of the variation of growth of Minneapolis and Saint Paul before the minimum wage increase. To give an example from a key industry that we elaborate upon below, for full-service restaurants the synthetic control accounts for 86 percent of the time series variation of jobs growth in Minneapolis and 73 percent of that in Saint Paul. Despite the overall success in accounting for a substantial variation of the pre-treatment growth, the synthetic control does not perform equally well in all industries. The most notable lack of fit is for the arts, entertainment, and recreation industry. As a result, we interpret the results for this industry with more caution.

While these R-squared statistics are informative, we do not rely solely on them to assess the appropriateness of the synthetic difference-in-differences methodology. Recent research by Ferman and Pinto (2021) has documented biases when the pre-treatment fit is less than perfect. We alleviate these concerns by using a specification in growth rates with a fixed effect instead of a levels specification.[15] Additionally, in robustness checks reported below, we add time

[15] Using growth rates means that we are requiring that the synthetic control fits the high-frequency movements of the wage, jobs, hours, and earnings. While the pre-treatment fit using a levels specification would have been significantly better, we prefer to match higher-frequency variations in order to alleviate concerns about over-fitting or correlation of the incidence of the treatment with underlying structural characteristics that affect outcomes in the treated units.

weights to our specification, following Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021). Doing so allows us to balance the pre-treatment and the post-treatment periods for the control group. Finally, assuming that the data generating process is a linear factor model, we perform Monte Carlo simulations to assess the size of the bias in the presence of imperfect fit and generally conclude that the bias in our context is small.

3.4 Labor Market Effects from the Time Series

Table 3 presents results for the low-wage industries identified in Section 2 and separately for restaurants. Entries are multiplied by 100 and equal the log point change in outcomes in 2020(4) due to the minimum wage increase, $100 \cdot g_T$. The columns present different outcome variables. For example, the first row shows that the increase in the minimum wage in Minneapolis caused an 8.8 log points (roughly 9 percent) increase in the wage and a 9.4 log points (also roughly 9 percent) decrease in the number of retail jobs. Each entry in parentheses is the p-value associated with the estimated treatment effect, which is the probability of obtaining a treatment effect as extreme as the point estimate under the null hypothesis that the treatment effect is zero. Continuing the example, the placebo method produces a p-value of 0 for the wage and 12.8 percent for jobs, and thus we conclude that the wage effect is precisely estimated and that the jobs effect is not very precisely estimated and cannot be statistically distinguished from zero at conventional levels of significance.[16]

In Minneapolis, we estimate wage increases with low p-values for retail; administrative and support services; other services; and restaurants. Among industries with statistically significant increases, we document increases that range between 4 and 13 log points. For all other industries, we find either statistically insignificant wage changes or small declines. In Saint Paul, we estimate a wage increase with low p-values for retail and for accommodation and food services (and for limited-service restaurants within that industry). The wage increases in Saint Paul are generally smaller and more imprecisely estimated than the wage increases in Minneapolis.

[16] We have multiple treated units, as our geographic unit of analysis is a zip code within a city. Thus, we construct placebo estimates by assigning a treatment status to 999 random subsamples of zip codes, with each subsample having a size equal to the number of treated units in Minneapolis or Saint Paul. We use the formula $p = 2\min\{p_H, p_L\}$ to calculate the p-value for a point estimate, $g_{2020(4)}$, where p_H is the fraction of placebo samples with point estimates that are higher than the estimate of Minneapolis or Saint Paul in 2020(4) and p_L is the fraction of placebo samples with point estimates that are lower than the estimate of Minneapolis or Saint Paul in 2020(4). Similar calculations underlie our p-values that use the placebo method in other tables.

Table 3: Effects of the Minimum Wage Increase from the Time Series

Minneapolis	Wage	Jobs	Hours	Earnings
Retail Trade (44)	8.8	−9.4	−7.6	0.9
	(0.0)	(12.8)	(13.4)	(80.9)
Administration and Support (56)	8.8	−1.9	0.1	12.2
	(0.0)	(92.7)	(82.5)	(27.8)
Health Care and Social Assistance (62)	−3.3	3.5	3.2	−1.2
	(1.8)	(47.8)	(67.7)	(65.1)
Arts, Entertainment and Recreation (71)	3.9	−13.4	3.6	16.7
	(63.1)	(17.0)	(77.1)	(37.6)
Accommodation and Food Services (72)	0.3	−31.4	−46.5	−36.9
	(98.1)	(0.0)	(0.0)	(0.0)
Other Services (81)	10.7	−1.2	−6.2	9.0
	(0.0)	(91.7)	(33.0)	(9.6)
Full-Service Restaurants (722511)	4.0	−41.4	−44.5	−42.1
	(0.0)	(0.0)	(0.0)	(0.0)
Limited-Service Restaurants (722513)	13.4	−32.0	−29.7	−28.4
	(0.0)	(0.6)	(2.2)	(2.0)
Saint Paul	**Wage**	**Jobs**	**Hours**	**Earnings**
Retail Trade (44)	4.6	−2.0	−35.0	−11.5
	(0.0)	(72.7)	(0.0)	(8.2)
Administration and Support (56)	0.7	−7.3	−9.3	−69.7
	(92.1)	(60.3)	(32.6)	(0.0)
Health Care and Social Assistance (62)	−4.2	4.3	3.7	−1.9
	(0.2)	(39.2)	(60.9)	(58.5)
Arts, Entertainment and Recreation (71)	−0.0	−18.1	−6.4	−16.6
	(57.1)	(6.2)	(50.5)	(4.0)
Accommodation and Food Services (72)	7.9	−47.0	−64.8	−33.0
	(0.0)	(0.0)	(0.0)	(0.0)
Other Services (81)	2.0	15.0	1.3	11.6
	(35.4)	(0.2)	(70.7)	(3.2)
Full-Service Restaurants (722511)	1.1	−38.8	−38.4	−44.1
	(21.6)	(0.0)	(0.0)	(0.0)
Limited-Service Restaurants (722513)	4.6	−57.2	−75.2	−85.7
	(0.0)	(0.0)	(0.0)	(0.0)

Notes: The estimates are in log points, multiplied by 100. Entries in parentheses are p-values using the placebo method.

We find this result intuitive, because Saint Paul announced the minimum wage increase in 2018 but did not implement it until 2020, so there are only four quarters of data under the new minimum wage policy.

We find the wage increases in Minneapolis reasonable. The difference between the minimum wage in Minneapolis and the one in the control cities is 38 percent. However, many workers are not close to the minimum wage, even in low-wage industries, and thus the estimated effects of the minimum wage increase on the wage are expected to be smaller than the change in the minimum wage. Holding worker hours constant at their 2017 level, the average establishment in Minneapolis would have experienced an 8 percent increase in its labor cost between 2017 and 2020. Weighted with employment, the average labor cost increase is 5 percent. The direct effect of the minimum wage on labor cost falls comfortably in the range of wage gains we estimate.

Turning to the second column, we find negative jobs effects for accommodation and food in Minneapolis. Within accommodation and food, we find a 41 log points jobs decline for full-service restaurants and a 32 log points jobs decline for limited-service restaurants. We find jobs declines of 9 log points for retail and 13 log points for arts, entertainment, and recreation, but these declines are not statistically significant at conventional levels.

We also document large jobs declines for restaurants in Saint Paul, which range between 39 and 57 log points. Additionally, we find a decline of 18 log points for arts, entertainment, and recreation that is statistically significant at a 6 percent level. Another difference with Minneapolis is that in Saint Paul, we find a 15 log points increase in jobs for other services.

The third column presents results for total hours. Estimated effects for hours are generally similar to estimated effects for jobs, with three exceptions. First, in both cities, arts, entertainment, and recreation experience smaller changes in hours than in jobs. Second, in Saint Paul, we cannot detect a statistically significant change in hours for other services, whereas we found an increase for other services' jobs. Finally, in Saint Paul, we find a decline of 35 log points for retail hours, whereas we did not detect a statistically significant decline for retail jobs.

The last column of the table presents results for worker earnings. Given the modest wage gains for all industries and the significant negative effects on employment for some industries, it is not surprising that in both cities we fail to detect a statistically significant increase in earnings in any industry except for other services. We detect statistically significant declines in earnings

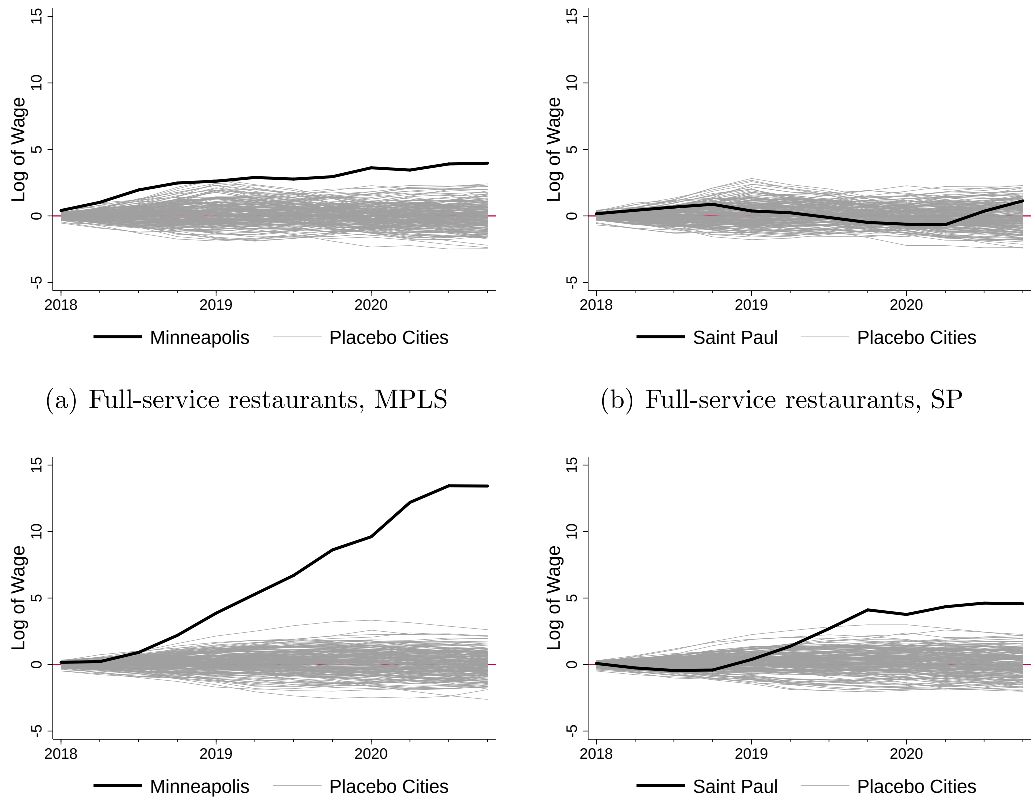

(a) Full-service restaurants, MPLS

(b) Full-service restaurants, SP

(c) Limited-service restaurants, MPLS

(d) Limited-service restaurants, SP

Figure 2: Time-Varying Wage Effects in the Twin Cities Restaurants

for retail in Saint Paul; arts, entertainment, and recreation in Saint Paul; accommodation and food services in both cities and separately for restaurants.[17]

We next examine the time variation of the estimated effects for the two industries with the most negative jobs effects, full-service and limited-service restaurants. Figure 2 plots the quarterly cumulative wage effects of the minimum wage increase for full-service and limited-service restaurants in the Twin Cities.[18] Along with our estimated effects, we plot placebo effects for 200 collections of units that were not subject to the minimum wage increase. Since

[17]We find a large decline in earnings for administration and support in Saint Paul that it is difficult to reconcile with its wage and hours changes. We investigated the time series of earnings and concluded that the estimated effect on earnings entirely reflects an extreme increase in reported earnings for some establishments just before 2018. This increase is not reflected in the wage because we trimmed the wage at the 90th percentile. When we also trim earnings at the 90th percentile, the decline in earnings becomes statistically indistinguishable from zero.

[18]We run the regression $Y_{it} = \alpha_i + \beta_t + \sum_{h=T_{\text{pre}+1}}^{T} \tau(h) W_{it}(h) + u_{it}$ using weights $\hat{\omega}_i$, where $h = T_{\text{pre}+1}, ..., T$ denotes the quarter of the treatment. For each quarter h shown in the figure, its cumulative effect equals $100 \cdot \sum_{j=T_{\text{pre}+1}}^{h} \hat{\tau}(j)/4$, where 4 appears in the formula because $\tau(h)$ is a yearly, as opposed to a quarterly, growth rate.

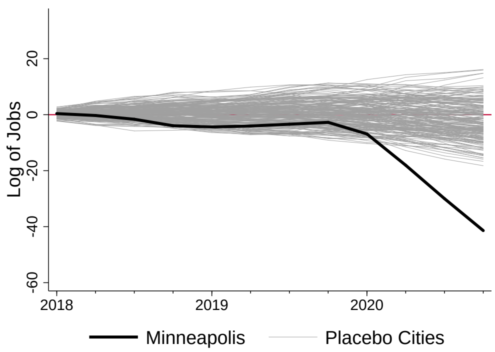

(a) Full-service restaurants, MPLS

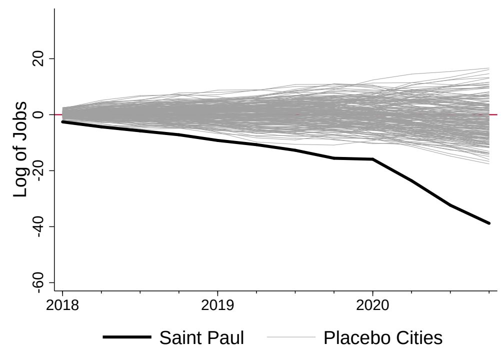

(b) Full-service restaurants, SP

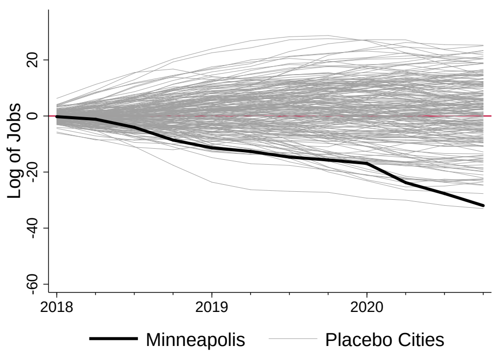

(c) Limited-service restaurants, MPLS

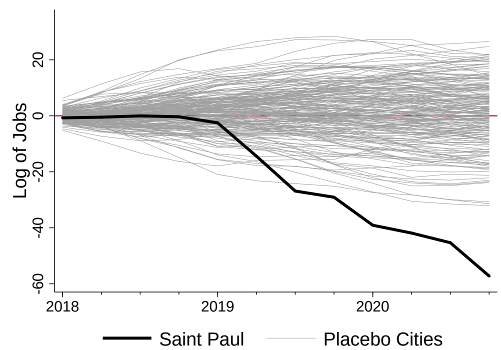

(d) Limited-service restaurants, SP

Figure 3: Time-Varying Jobs Effects in the Twin Cities Restaurants

we know that these placebo units did not experience an increase in their minimum wage, any effect we estimate for these units is due only to random noise.

The left panels of the figure show that the wage for restaurants in Minneapolis increased soon after the minimum wage ordinance went into effect. By contrast, in Saint Paul there is no statistically significant increase in the wage for full-service restaurants throughout our period. For limited-service restaurants, such an increase does not occur until the second half of 2019. We find the difference in the response of the wage between Minneapolis and Saint Paul intuitive, because the Saint Paul ordinance was implemented two years after the Minneapolis ordinance.

Figure 3 plots the quarterly cumulative jobs effects of the minimum wage increase for full-service and limited-service restaurants in the Twin Cities, from which we draw three conclusions. First, in contrast to the wage, jobs in Saint Paul declined before the implementation of the minimum wage. This evidence of advance notice is consistent with our cross-sectional results below, which also show jobs declines before 2020 in Saint Paul.[19] Second, in three out of four cases, the jobs declines for restaurants appear by the end of 2019, before the pandemic hit. Third, while the jobs declines in limited-service restaurants are relatively stable over time, the jobs declines in full-service restaurants accelerated significantly after the first quarter of 2020, after the pandemic hit.

We conclude this section by presenting two robustness checks. In Table 4, we repeat our estimates by adding time weights λ_t, following Arkhangelsky, Athey, Hirshberg, Imbens, and Wager (2021). The time weights are chosen to make the average pre-treatment growth as similar as possible to the average post-treatment growth in the control group. Thus, this exercise allows us to examine the robustness of our results when we place more weight on periods when, similar to the pandemic, the synthetic control experiences large negative growth.[20]

Table 4 shows that, even when we re-weight the data, the results are highly similar to the baseline results. There are a few differences worth emphasizing. First, when we use the λ_t weights, we find statistically significant declines in Minneapolis retail jobs and hours of 12 log points. Second, the jobs declines for arts, entertainment, and recreation in Minneapolis and for administration and support services in Saint Paul become larger and statistically significant at the 10 percent level. Finally, the previously documented changes in worker earnings for other services and for limited-service restaurants in Minneapolis become statistically insignificant.

Our second robustness check repeats our estimates in a sample of cities that excludes cities bordering Minneapolis and Saint Paul. It is conceivable that the implementation of a higher minimum wage reallocated jobs from the Twin Cities to neighboring cities. From the perspective of a city that implements a minimum wage increase, the policy-relevant statistic is its change in jobs, irrespective of whether these jobs disappeared or were reallocated to neighboring cities. Therefore, we do not merge neighboring cities with the Twin Cities in estimating the effects of the minimum wage change. However, to the extent that jobs were reallocated to neighboring

[19] We also examined advance notice in Minneapolis by backdating the treatment of the minimum wage to 2016 and 2017. We failed to detect significant effects in Minneapolis. The difference with Saint Paul is explained by the greater uncertainty during 2016 and 2017 about whether the Minneapolis minimum wage ordinance would pass, whereas in Saint Paul the ordinance was passed in 2018 but implemented in 2020.

[20] Equation (3) is replaced by $\left(\hat{\tau}, \hat{\alpha}, \hat{\beta}\right) = \arg\min_{\tau,\alpha,\beta} \left\{\sum_{i=1}^{N} \sum_{t=1}^{T} \left(Y_{it} - \alpha_i - \beta_t - \tau W_{it}\right)^2 \hat{\omega}_i \hat{\lambda}_t\right\}$. The time weights $\hat{\lambda}_t$ are chosen so that $(\hat{\lambda}_0, \hat{\lambda}_t) = \arg\min_{\lambda} \sum_{i=1}^{N_{\text{co}}} \left(\lambda_0 + \sum_{t=1}^{T_{\text{pre}}} \lambda_t Y_{it} - T_{\text{post}}^{-1} \sum_{t=T_{\text{pre}}+1}^{T} Y_{it}\right)^2$.

Table 4: Adding Time Weights in the Synthetic Difference-in-Differences Estimation

Minneapolis	Wage	Jobs	Hours	Earnings
Retail Trade (44)	7.3	−12.0	−11.9	−5.0
	(0.0)	(2.8)	(1.8)	(59.3)
Administration and Support (56)	8.9	0.6	−1.9	12.1
	(0.0)	(89.7)	(70.1)	(30.8)
Health Care and Social Assistance (62)	−1.7	5.9	−0.1	2.3
	(9.8)	(22.0)	(86.3)	(99.7)
Arts, Entertainment and Recreation (71)	5.2	−17.3	1.6	21.5
	(38.6)	(5.6)	(96.5)	(31.2)
Accommodation and Food Services (72)	0.8	−31.0	−38.9	−41.8
	(64.9)	(0.0)	(0.0)	(0.0)
Other Services (81)	9.6	−1.4	−5.2	−0.5
	(0.0)	(96.1)	(43.6)	(83.9)
Full-Service Restaurants (722511)	3.5	−42.1	−44.8	−40.1
	(0.0)	(0.0)	(0.0)	(0.0)
Limited-Service Restaurants (722513)	14.2	−26.2	−26.7	−21.2
	(0.0)	(2.6)	(6.6)	(11.4)
Saint Paul	**Wage**	**Jobs**	**Hours**	**Earnings**
Retail Trade (44)	4.8	−8.1	−32.6	−13.7
	(0.0)	(17.4)	(0.0)	(3.4)
Administration and Support (56)	−1.9	−35.1	−13.4	−76.8
	(37.0)	(1.8)	(23.8)	(0.0)
Health Care and Social Assistance (62)	−3.5	4.3	−0.0	−4.4
	(0.0)	(39.4)	(88.1)	(27.0)
Arts, Entertainment and Recreation (71)	−0.7	−15.7	−4.6	−28.6
	(53.1)	(6.8)	(64.5)	(0.0)
Accommodation and Food Services (72)	5.8	−48.6	−60.1	−41.0
	(0.0)	(0.0)	(0.0)	(0.0)
Other Services (81)	1.7	14.7	5.3	11.4
	(34.4)	(0.2)	(26.4)	(2.0)
Full-Service Restaurants (722511)	1.7	−41.7	−41.7	−38.0
	(6.8)	(0.0)	(0.0)	(0.0)
Limited-Service Restaurants (722513)	2.2	−33.4	−41.4	−55.0
	(1.6)	(0.2)	(0.4)	(0.0)

Notes: The estimates are in log points, multiplied by 100. Entries in parentheses are p-values using the placebo method.

cities and these cities are part of the synthetic control, we could be double-counting the effects of the minimum wage because cities in the control group experience jobs growth. Online Table A.3 shows that this is not the case, because our estimates do not change significantly when we exclude bordering cities from the sample of cities that form the synthetic control.

3.5 Evidence from Other U.S. Cities

While some of our estimated negative jobs effects following the minimum wage increase in the Twin Cities become apparent by the end of 2019, the largest yearly decline in jobs for full-service restaurants is observed during 2020, the year when the pandemic recession began. By design, the synthetic control aims to fit pre-treatment series of Minneapolis and Saint Paul in both expansions and downturns. However, we acknowledge that the pandemic recession is quite atypical relative to other downturns observed in our sample. A potential threat to identification would arise if in 2020 the sensitivity to aggregate shocks changed for the control group relative to that of Minneapolis and Saint Paul. For example, it may be that the enforcement and economic impact of lockdowns was larger in more densely populated cities than in smaller cities or suburbs.

To address this concern, we now extend our analysis to use variation from other U.S. cities of similar size to Minneapolis and Saint Paul. Using these cities to construct our synthetic control addresses the concern that our control from Minnesota may not be appropriate during the pandemic recession because other large, densely populated U.S. cities faced similar or more stringent lockdowns than Minneapolis or Saint Paul. Additionally, using other U.S. cities allows us to control for nationwide changes in economic conditions such as the substitution of services prone to virus transmission with online shopping, the rise of gig work, and labor shortages in low-wage industries.

For our analyses using other U.S. cities, we use publicly available data from the Quarterly Census of Employment and Wages (QCEW) produced by the U.S. Bureau of Labor Statistics.[21]

[21] Our control group includes municipalities or local government units for which data could be compiled from the publicly available files. This was possible in the following circumstances: 1) the city consists of two or more counties (for example, New York, NY); 2) the city is coterminous with a county or is governed by a consolidated city-county government (for example, San Francisco, CA; Denver, CO; Philadelphia, PA; Nashville, TN); 3) the city is independent (for example, Baltimore, MD; Saint Louis, MO; Richmond, VA); 4) the local minimum wage policy is enacted or harmonized at the county level (for example, Los Angeles County, CA; Montgomery County, MD). To further expand our control group, we also include cities that are the county seat and whose

The measure of employment refers to the number of workers who worked during or received pay for a pay period that includes the 12th of the month, as reported by establishments covered under the unemployment insurance program.

We note three differences between the research design using the QCEW data and that of our previous analyses using the DEED data. First, the QCEW does not have a measure of hours, and the wage measure differs from that in the DEED. Thus, we focus our analysis only on jobs. Second, for the QCEW data, we extend the sample by one more year to 2021(4). Doing so allows us to assess if the acceleration of the minimum wage effect during the pandemic was transitory or persistent through 2021. Finally, the unit of analysis in the QCEW data is other U.S. cities of similar size to Minneapolis and Saint Paul, whereas in the DEED data we used zip code within a city as our unit of analysis. Before the minimum wage increases, Minneapolis employment is roughly 280,000 and Saint Paul employment is roughly 150,000. We include in the control group only cities whose employment is between half and double of that of either Minneapolis or Saint Paul. This restriction results in a sample of 36 cities for the Minneapolis control group and of 42 cities for the Saint Paul control group.[22]

Table 5 presents our estimates from the QCEW until 2021(4). The first two columns present estimates using only the unit weights ω_i to construct the synthetic control. The last two columns add time weights λ_t. Despite differences in the research design, the jobs effects we estimate for the Twin Cities using variation across U.S. cities are highly similar to those we estimated previously using within Minnesota variation. As shown in the table, we find jobs declines for accommodation and food, including restaurants, in both cities and of similar magnitude to the jobs declines we documented before in the DEED data. Additionally, we find a statistically significant decline of 11 log points for retail in Saint Paul.[23]

population accounted for more than 75 percent of their county's population. In these circumstances, we use the county as a reliable proxy for the corresponding city.

[22]Online Table A.4 shows the U.S. cities included in the control groups. We have also examined results without the size restriction and find similar results when all U.S. cities are allowed to be included in the synthetic control.

[23]In Online Table A.5 we report results from Monte Carlo simulations in the QCEW data to assess potential biases due to imperfect fit of the synthetic control during the pre-treatment period. Assuming that the data generating process is a linear model with four factors, in Minneapolis we find the most negative bias for full-service restaurants (moving the estimated coefficient from -39 to -33 log points) and the most positive bias for limited-service restaurants (moving the estimated coefficient from -19 to -34 log points). In Saint Paul, we find the most negative bias for limited-service restaurants (moving the estimated coefficient from -12 to -4 log points), whereas many industries have no discernible change in their coefficients. Across all two-digit industries in the Twin Cities, the biases range from -4 to 4 log points.

Table 5: Jobs Effects of Minimum Wage Increases: Cities with Comparable Employment

	No Time Weights		With Time Weights	
	Minneapolis	Saint Paul	Minneapolis	Saint Paul
Retail Trade (44)	−3.2	−11.4	−4.1	−10.6
	(64.9)	(0.0)	(48.6)	(9.5)
Administration and Support (56)	2.5	−9.8	5.6	−25.2
	(97.3)	(38.1)	(75.7)	(14.3)
Health Care and Social Assistance (62)	−2.5	−4.5	−1.0	−7.3
	(64.9)	(41.9)	(70.3)	(41.9)
Arts, Entertainment, and Recreation (71)	−12.4	−21.6	−19.0	−26.4
	(32.4)	(18.6)	(32.4)	(18.6)
Accommodation and Food Services (72)	−25.2	−21.9	−23.6	−23.0
	(5.4)	(9.3)	(16.2)	(4.7)
Other Services (81)	−9.6	−1.9	−10.2	2.8
	(34.5)	(61.1)	(34.5)	(77.8)
Full-Service Restaurants (722511)	−38.5	−22.1	−32.7	−17.6
	(5.4)	(9.3)	(5.4)	(9.3)
Limited-Service Restaurants (722513)	−19.0	−12.3	−7.9	−18.4
	(5.4)	(4.8)	(16.2)	(4.8)

Notes: The estimates are in log points, multiplied by 100. Entries in parentheses are p-values using the placebo method.

To summarize, using additional variation from outside of Minnesota, we conclude that our results are not driven exclusively by the pandemic recession, which may have impacted control regions within Minnesota differently than the Twin Cities. This is because other U.S. cities are also large and densely populated and faced similar negative economic impacts from lockdowns and other pandemic restrictions. Rather, given the acceleration in 2020 of the negative employment impacts of the minimum wage in the Twin Cities, our conclusion is that the minimum wage increase interacted with the pandemic recession to more negatively impact employment. Further, the results from the QCEW show that the negative effects on jobs in restaurants and retail persisted through 2021, despite the partial lifting of lockdowns and other pandemic restrictions.

3.6 Challenges in Interpreting the Time Series Results

The time series methods attribute any differences after 2018 between outcomes in the Twin Cities and those in the control group to the minimum wage increase. Our starting point is to use the synthetic control from the state of Minnesota to difference out any effects unrelated to the minimum wage. Despite the synthetic control fitting well the time series of the treated units in the pre-treatment period, it may be that the sensitivity to aggregate shocks changed for the Twin Cities relative to that of the synthetic control in the post-treatment period. As we argued, using time weights or a synthetic control from other U.S. cities alleviates somewhat this concern because these strategies make it more plausible that we are differencing out pandemic effects in the post-treatment period that are unrelated to the minimum wage increase.

As with any research design that uses time series variation, it may still be case that the Twin Cities experienced idiosyncratic shocks, such as civil unrest, that cannot be differenced out in the post-treatment period. Using the QCEW sample of other U.S. cities we find a persistent jobs decline through the end of 2021 when it is reasonable to assume that civil unrest was no longer impacting the Twin Cities differently from other cities. However, perhaps the negative impacts of civil unrest propagated in 2021 through mechanisms other than the minimum wage. For this reason, we shift our research design away from using other cities in the control group and now use variation from the cross sections of establishments and workers within a city.

4 Evidence from the Cross Section

We use variation from the cross sections of establishments and workers within a city, a strategy that allows us to absorb any common shock shared by these establishments and workers to the intercept of our regressions. We begin by laying out the econometric framework and then present our estimates.

4.1 Econometric Methodology: Cross Section

Our starting point is the statistical model

$$Y_{jszt} = \gamma_{szt} + \sum_{t=2018}^{2020} \tau_t \left(\mathrm{GAP}_{jszt-3} \cdot d_t\right) + u_{jszt}, \tag{6}$$

where Y_{jszt} denotes an outcome for establishment j in industry s, zip code z, and period t. The outcome variables are the arc percent change of y_{jszt} over three years

$$Y_{jszt} = \frac{y_{jszt} - y_{jszt-3}}{(1/2)(y_{jszt} + y_{jszt-3})},$$

where y_{jszt} is the level of the wage, jobs, hours, and worker earnings for an establishment. We adopt the arc percent change transformation of growth rates to capture potential changes in the propensity of establishments to exit in response to the minimum wage increase. The lowest value of Y_{jszt} is -2, which we obtain for jobs, hours, and earnings when an establishment exists in period $t-3$ and exits in period t. The establishments we include in this regression are located only within Minneapolis or Saint Paul and have to exist in the sample in period $t-3$.

In regression (6), the fixed effect γ_{szt} absorbs the common growth in period t of all establishments that belong to the same industry s and zip code z of the Twin Cities. For example, among other things, the fixed effect could capture the common effect for each industry within zip code arising from the pandemic recession or civil unrest in 2020(3).

The key variable of interest in regression (6) is the gap in labor costs over three years

$$\text{GAP}_{jszt} = \frac{\sum_{i \in j} \max(15/(1+\pi_{t,2017}) - w_{ijszt}, 0) h_{ijszt}}{\sum_{i \in j} w_{ijszt} h_{ijszt}}. \tag{7}$$

The numerator of the GAP variable is the additional costs incurred by establishment j when its workers i earn wages that are below the projected level of the minimum wage. The denominator of the GAP variable denotes the wage bill of the establishment. Therefore, the GAP variable captures the exposure of an establishment to the minimum wage increase, where exposure is expressed as the fraction of the wage bill accruing to additional labor costs.[24] In equation (7), we adjust the projected level of the minimum wage in each period with the metro-level CPI deflator $\pi_{t,2017}$, where $\pi_{2017,2017} = 1$. As an example, if an establishment pays all of its workers above 15 dollars per hour in 2017, its GAP measure equals zero.

One might be tempted to interpret the coefficients τ_{2018}, τ_{2019}, and τ_{2020} as the difference in establishment outcomes arising from differences in their exposure to the minimum wage increase in 2018, 2019, and 2020, after differencing out any common time effect that these

[24]Previous studies that also used the GAP measure of exposure to the minimum wage include Card and Krueger (1994), Draca, Machin, and Van Reenen (2011), Harasztosi and Lindner (2019), and Dustmann, Lindner, Schonberg, Umkehrer, and vom Berge (2022).

establishments share with other establishments in the same zip code and industry.[25] These coefficients, however, do not only capture differences in exposure to the minimum wage increase because typical establishment dynamics unrelated to exposure introduce a spurious correlation between exposure and various outcomes. Smaller establishments pay lower wages and thus have larger gaps. At the same time, smaller establishments tend to exit at faster rate, which tends to generate a negative τ_t for jobs, hours, and earnings. The wage regressions include only establishments that exist in both period t and period $t-3$. We expect smaller establishments that survived to experience higher wage growth, generating a positive τ_t for the wage.

We augment our regression to include three more years before the minimum wage increase. The final specification is

$$Y_{jszt} = \gamma_{szt} + \sum_{t=2018}^{2020} \tau_t \left(\text{GAP}_{jszt-3} \cdot d_t\right) + \tau_0 \text{GAP}_{jszt-3} + u_{jszt}, \tag{8}$$

where τ_0 controls for any correlation between GAP and outcomes due to typical establishment dynamics unrelated to the minimum wage increase. Using this specification, we now interpret the coefficients τ_{2018}, τ_{2019}, and τ_{2020} as the difference in establishment outcomes due to differential exposure to the minimum wage increase.

4.2 Labor Market Effects from the Cross Section of Establishments

The upper panel of Table 6 presents estimates of the coefficients τ_t from specification (8) applied to the wage, jobs, hours, and earnings, separately by each city. The entries are multiplied by 100 and are interpreted as the percent change in establishments' outcomes when the GAP changes from 0, which is the value for an establishment that is not exposed to the minimum wage, to 1, which is the value for an establishment that experiences 100 percent increase in its wage bill due to the minimum wage.[26] The maximum GAP is around 100 percent and the average GAP is around 6 percent. We will later use these moments of the GAP to translate the coefficient estimates from the cross section into most extreme and average labor market effects

[25] We run regression (6) with quarterly data but estimate one coefficient common to all quarters within a year. To improve the readability, we have suppressed the notation of the quarters from regression (6).

[26] Our sample includes many establishments with a zero GAP. The average outcome of these establishments estimates the constant γ_{szt}. We believe it is appropriate to include non-exposed establishments in the regression, because they are a valid control group for exposed establishments within a zip code and industry. To examine how sensitive our results are to the linear specification adopted in regression (8), we have repeated our regressions by excluding establishments with a zero GAP. We find no significant differences in our results.

Table 6: Labor Market Effects of Minimum Wage Increases: Cross Section of Establishments

	Minneapolis				Saint Paul			
Baseline	Wage	Jobs	Hours	Earnings	Wage	Jobs	Hours	Earnings
2018	11.5	−10.9	−12.7	−7.9	4.9	−11.8	−12.9	−13.4
	(0.0)	(1.3)	(0.5)	(13.1)	(15.4)	(6.0)	(4.2)	(6.0)
2019	13.7	−15.7	−16.2	−11.4	5.7	−24.2	−22.4	−25.1
	(0.0)	(0.5)	(0.4)	(8.0)	(21.4)	(0.1)	(0.4)	(0.4)
2020	15.2	−14.3	−13.1	−12.8	4.7	−24.4	−23.8	−23.3
	(0.0)	(1.5)	(2.6)	(6.3)	(32.2)	(0.1)	(0.2)	(0.8)
Add Lagged Growth	Wage	Jobs	Hours	Earnings	Wage	Jobs	Hours	Earnings
2018	7.5	−10.7	−12.5	−8.0	5.5	−10.8	−12.3	−12.7
	(1.1)	(1.6)	(0.6)	(11.9)	(14.0)	(8.5)	(5.3)	(7.2)
2019	8.6	−15.1	−16.0	−11.3	3.3	−23.6	−22.0	−23.4
	(1.8)	(0.7)	(0.4)	(7.7)	(53.3)	(0.2)	(0.4)	(0.7)
2020	6.4	−14.1	−13.1	−13.3	2.3	−24.5	−23.7	−22.8
	(9.3)	(1.7)	(2.7)	(5.1)	(67.3)	(0.1)	(0.2)	(0.9)
Pre-sample to Six Years	Wage	Jobs	Hours	Earnings	Wage	Jobs	Hours	Earnings
2018	10.0	−9.0	−10.8	−8.0	4.4	−11.8	−13.3	−12.3
	(0.1)	(5.3)	(2.3)	(13.8)	(23.3)	(6.3)	(4.0)	(9.1)
2019	12.1	−13.7	−14.3	−11.6	5.2	−24.2	−22.8	−24.0
	(0.0)	(1.3)	(1.0)	(7.0)	(26.1)	(0.1)	(0.2)	(0.4)
2020	13.7	−12.3	−11.2	−12.9	4.3	−24.4	−24.2	−22.2
	(0.0)	(3.1)	(4.8)	(5.0)	(36.7)	(0.1)	(0.1)	(0.8)

Notes: The estimates are in percent, multiplied by 100. Entries in parentheses are p-values using standard errors clustered at the establishment level.

arising from the minimum wage increase. Entries in parentheses are p-values associated with each coefficient. We cluster standard errors at the establishment level.

Beginning with the wage effects, we estimate wage growth between 12 and 15 percent in Minneapolis establishments. We fail to detect statistically significant wage increases in Saint Paul establishments. This result echoes the result using time series variation, which showed stronger wage responses in Minneapolis implementing a minimum wage ordinance in 2018 than in Saint Paul passing a minimum wage ordinance in 2018 and implementing it in 2020.

Turning to the employment responses, we estimate declines of jobs and hours that range between 11 and 16 percent in Minneapolis establishments. We estimate even larger employment declines in Saint Paul establishments in 2019 and 2020. The employment declines in Saint Paul establishments, without a corresponding increase in the wage, are consistent with advance notice from the increase in the minimum wage, which shows up in quantities rather than prices. Finally, in both cities, we estimate negative relationships between exposure to the minimum wage and earnings of workers at the establishment level. In Minneapolis, earnings coefficients are generally smaller in absolute value and less statistically significant than employment coefficients, reflecting the positive effects we estimate for the wage.

The responses of the wage, jobs, hours, and earnings are above and beyond those generated by typical establishment dynamics because regression (8) includes the GAP measure in the period before the minimum wage increase. However, it could still be the case that there is a trend in establishment dynamics that makes these coefficients larger in absolute value over time, irrespective of the minimum wage increase. To examine this possibility, we allow the τ_0 coefficients to vary by period and perform a test of the null hypothesis that they are stable over time. For each city and variable, we cannot reject the null hypothesis that the coefficients before the minimum wage policy are all equal to each other. Thus, we conclude that the responses of all variables to the GAP measure display no trend before the minimum wage policy and become larger in absolute value only after the minimum wage policy.

A reasonable concern about our cross-sectional results in 2020 is whether our strategy identifies establishments' sensitivity to the minimum wage or whether it identifies the sensitivity of smaller establishments with larger GAP exposure to the pandemic recession or civil unrest. However, with the exception of the estimated coefficients in Saint Paul in 2018, the estimated coefficients on all variables are quite stable over time. We find the stability of the estimated coefficients reassuring and conclude that our identification strategy from the cross section of establishments isolates differential exposure to the minimum wage rather than other forces contemporaneous with the minimum wage change.

The middle panel of Table 6 presents estimated coefficients τ_t in a specification in which we add lags of the dependent variable into the regression. Our estimated coefficients do not change much, with the exception of the wage effects in Minneapolis establishments, which decrease between 4 and 9 percentage points. In the lower panel, we include six years of data before the minimum wage increase, as opposed to the three year period in the baseline specification. Our logic for including three years of data in the baseline specification is that the minimum wage increase is implemented over three years and that we wish to control for typical establishment dynamics during a period close to the minimum wage increase. However, our results are not sensitive to expanding the sample to include the last six years before the minimum wage increase.

A challenge in interpreting the results that use variation from the cross section of establishments is that there may be spillovers from high to low GAP establishments. These spillovers may be important, given that we use within-zip-code and within-industry variation in establishments' outcomes. As an example, if workers reallocated from high to low GAP establishments, then we would be double-counting the effects of the minimum wage increase on establishments' employment. Another challenge arises from reallocation outside of the Twin Cities, because our estimates could reveal negative employment effects from the minimum wage even if all affected workers find jobs outside of the Twin Cities. Thus, while our estimates directly speak to the outcomes of establishments that were located in the Twin Cities before the minimum wage increase, they may not be informative about workers' labor market outcomes and welfare. To address these challenges, we now turn to specifications from the cross section of workers.

4.3 Labor Market Effects from the Cross Section of Workers

In this section, we use variation in wage gaps across workers and track workers' outcomes over time. Our first specification is

$$Y_{it} = \sum_{s} \gamma_{st} X_{ist} + \sum_{t=2018}^{2020} \tau_t \left(\mathrm{GAP}_{it-3} \cdot d_t\right) + \tau_0 \mathrm{GAP}_{it-3} + \rho Y_{it-1} + u_{it}, \tag{9}$$

where the dependent variable for worker i in period t, Y_{it}, is defined as the arc percent change over three years and the wage gap over three years is calculated at the worker level,

$$\mathrm{GAP}_{it} = \frac{\max(15/(1+\pi_{t,2017}) - w_{it}, 0)}{w_{it}}. \tag{10}$$

Our specification is again agnostic about the intercept γ_{st}, which absorbs all time effects common to workers belonging to industry s. Workers may work in more than one industry in a year, so the variable X_{ist} denotes the share of worker i's employment in industry s. The

Table 7: Labor Market Effects of Minimum Wage Increases: Cross Section of Workers

	Minneapolis			Saint Paul		
Workers' Own GAP	Wage	Hours	Earnings	Wage	Hours	Earnings
2018	3.6	0.2	1.8	1.6	−9.1	−7.3
	(0.0)	(89.3)	(23.8)	(2.2)	(0.0)	(0.0)
2019	11.4	−5.0	−2.6	4.5	−10.4	−7.6
	(0.0)	(0.2)	(11.8)	(0.0)	(0.0)	(0.0)
2020	14.2	−15.1	−10.5	7.3	−16.4	−12.1
	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
Workers' Establishment GAP	Wage	Hours	Earnings	Wage	Hours	Earnings
2018	−1.4	−0.8	−2.7	−0.7	−8.8	−8.0
	(18.7)	(76.9)	(30.4)	(53.1)	(0.7)	(1.8)
2019	7.0	−15.2	−12.4	−2.1	−13.0	−11.9
	(0.0)	(0.0)	(0.0)	(15.2)	(0.0)	(0.1)
2020	7.9	−14.5	−8.8	3.7	−14.6	−9.2
	(0.0)	(0.0)	(0.4)	(1.3)	(0.0)	(1.5)

Notes: The estimates are in percent, multiplied by 100. Entries in parentheses are p-values using standard errors clustered at the worker level.

other difference relative to our specification for establishments is that now we include in the regression the lagged outcome Y_{it-1} for workers. Thus, we interpret the τ coefficients as the percent change in worker outcomes resulting from a higher wage gap for workers with the same growth rate in the period immediately preceding the wage increase and after differencing out the common effect that workers in the same industry experience, γ_{st}, and any effects we would detect due to typical worker dynamics, τ_0.

The upper panel of Table 7 presents estimates of the coefficients τ_t from specification (9) applied to the wage, hours, and earnings, separately by each city.[27] The entries are multiplied by 100 and equal the percent change in an outcome when GAP changes from 0, which is the value for a worker who is not exposed to the minimum wage increase, to 1, which is the value for a worker who experiences 100 percent increase in their wage due to the minimum wage

[27] As in the time series analysis, we exclude workers with wage below the youth minimum wage for Minnesota. For the worker-level analysis, we include only workers with wage below 45 dollars per hour and run the regression at the yearly frequency.

increase.

In Saint Paul, the results from the cross section of workers are similar to those from the establishments. We generally find small wage effects, with the exception of the wage in 2020, the year Saint Paul implemented its minimum wage ordinance. By contrast, we find significant declines in worker hours and earnings. In Minneapolis, in 2018 we do not detect statistically significant responses of hours and earnings and we find only a modest increase in workers' wage. In 2019, we find larger increases in workers' wage and a moderate decline in their hours. The 2020 results for workers in Minneapolis are comparable to the results for establishments.

A benefit of using the cross section of workers to estimate the labor market effects of the minimum wage increase is that it allows us to track directly worker outcomes, irrespective of whether workers reallocated to other establishments in or outside of the Twin Cities. The worker results are different from the establishment results, for which the responses are stable over time. For example, between 2019 and 2020, the coefficients on hours increase in absolute value by 10 percentage points in Minneapolis and 6 percentage points in Saint Paul, raising the concern that roughly half of the employment effects capture low-wage workers' difficulty finding jobs due to the pandemic or civil unrest, rather than their difficulty finding jobs due to the minimum wage.

To address this concern, we consider a final specification

$$Y_{it} = \sum_{s} \gamma_{st} X_{ist} + \sum_{t=2018}^{2020} \tau_t \left(\overline{\text{GAP}}_{it-3} \cdot d_t \right) + \tau_0 \overline{\text{GAP}}_{it-3} + \rho Y_{it-1} + u_{it}, \tag{11}$$

where the wage gap over three years now becomes

$$\overline{\text{GAP}}_{it} = \frac{1}{\# J_t(i)} \sum_{j \in J_t(i)} \text{GAP}_{jt}, \tag{12}$$

$\# J_t(i)$ denotes the number of establishments that worker i worked in during period t, and GAP_{jt} is establishment's j gap in labor costs defined in equation (7). The only difference relative to our previous worker-level specification in equation (9) is that we treat workers with their establishments' gaps instead with their own gaps. This specification combines elements from both the establishment-level regressions and the worker-level regressions we ran previously. It allows us to track workers' outcomes over time, as workers reallocate to other establishments both in and outside of the Twin Cities. However, the treatment is defined at the establishment

Table 8: Jobs Effects of Minimum Wage Increases: Summary of Estimates

Jobs (2020, percent)	Time Series	Cross Section	Ratio
Minneapolis Average	−2.1	−0.7	0.33
Minneapolis Most Negative	−30.1	−13.9	0.46
Saint Paul Average	−3.3	−1.4	0.42
Saint Paul Most Negative	−27.6	−19.2	0.70
Average			**0.48**

Notes: Average from the time series includes only industries with statistically significant changes in jobs. Most Negative from the time series uses the estimates for the restaurant industries. The estimates for the cross section multiplies the 2020 coefficients from the establishments' and workers' regressions with the weighted average and maximum GAP. The ratio of 0.48 is the average across the four ratios.

level, thus alleviating the concern that low-wage workers' difficulty finding jobs in 2020 is because of the pandemic or civil unrest.

The bottom panel of Table 7 presents the estimated coefficients from this specification. We continue to find stable results in Saint Paul, with small changes in the wage and large declines in hours and earnings. In Minneapolis, the 2018 estimated coefficients are similar to those in the upper panel of Table 7 for workers. However, the 2019 estimated coefficients for hours and earnings are more similar to those from the establishment regressions in Table 6. Additionally, the 2020 estimated coefficients for the wage and hours in Minneapolis are stable relative to those in 2019.

5 Summary of Estimates

In this section, we compare our estimates from the time series with those from the cross section. We then compare our estimates of employment losses with previous estimates found in the minimum wage literature. We conclude by comparing our estimates of the elasticity of labor demand at the establishment level with other estimates in the literature.

5.1 Comparison of the Time Series with the Cross Section

Table 8 summarizes our estimates using variation from the time series of cities and the cross sections of establishments and workers. In the first row, we present the average job losses in

Minneapolis in 2020. The time series estimate of the job losses is 2.1 percent. We calculate this number as the average job losses across all two-digit industries, where losses are weighted with the employment of the corresponding industry in total Minneapolis employment before the minimum wage increase. We include only industries with statistically significant changes in jobs and average our estimates between the DEED and the QCEW data sources. Because only accommodation and food experienced a statistically significant decline, 2.1 percent represents the decline in jobs in that industry weighted with its employment share. The estimate of job losses using variation from the cross section is 0.7 percent. We calculate this number by multiplying the 2020 average of the coefficients from establishments and workers (roughly −15 percent) with the employment-weighted average GAP (roughly 5 percent) in Minneapolis in 2020. Similar calculations in Saint Paul lead to estimated job losses of 3.3 percent from the time series and 1.4 percent from the cross section.

The second and fourth rows summarize our most negative jobs estimates. For the time series, we use the estimates for restaurants and conclude that the most negative jobs effects are 30 percent in Minneapolis and 28 percent in Saint Paul. For the cross section, we multiply the 2020 average of the coefficients from the establishments' and workers' regressions with the maximum GAP for each of the cities in 2020. We use the maximum GAP so that we can get a comparable estimate of the most negative jobs effects. This yields estimated job losses of 14 percent in Minneapolis and 19 percent in Saint Paul.

The last column of Table 8 shows the ratio of estimates from the cross section and the time series. This ratio ranges between 0.3 and 0.7, with an average value of 0.48. In Section 6, we discuss reasons why the time series estimates of job losses differ from estimates that use the cross section and argue that they plausibly reflect upper and lower bounds.

5.2 Comparison of Minimum Wage Estimates with Other Studies

In this section, we compare our estimates to estimates of other papers that use data from the U.S. restaurant industry. To compare our findings with those of the literature, we find it useful to transform the estimated employment effects from the time series into elasticities. The

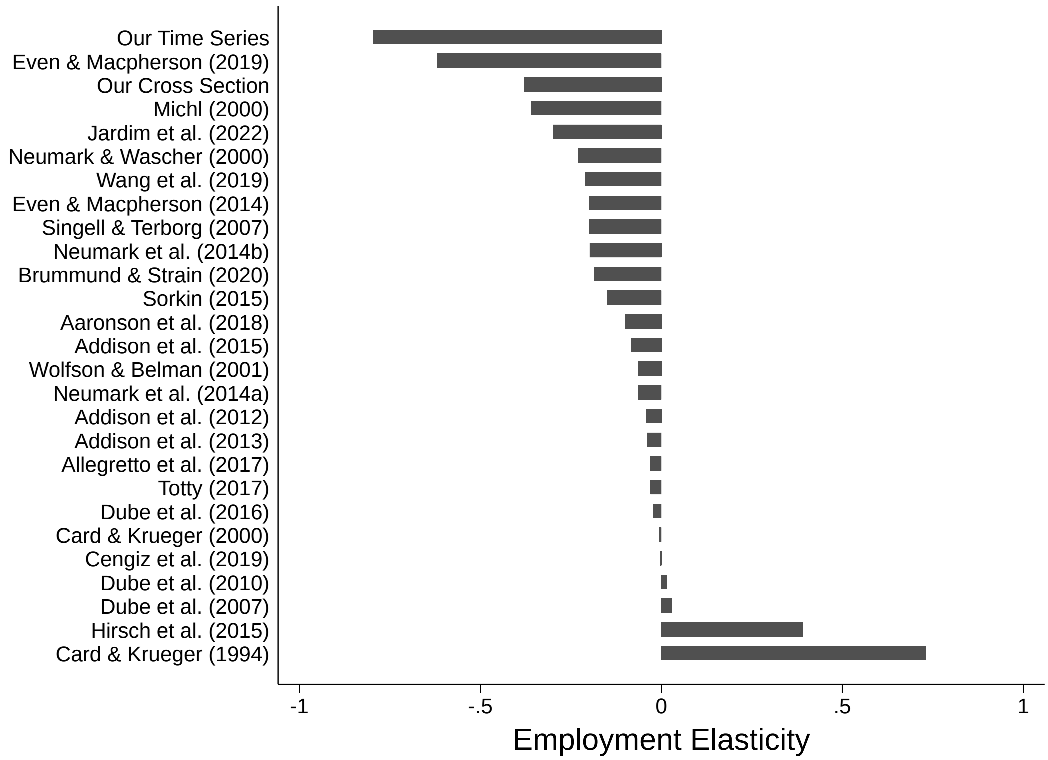

Figure 4: Comparison of Restaurant Employment Elasticity to the Literature

elasticity of an outcome in period T with respect to the change in the minimum wage is

$$\varepsilon_T = \frac{\exp\left((T - T_{\text{pre}})\hat{\tau}/4\right) - 1}{w^1_{\min,T}/w^0_{\min,T} - 1}, \tag{13}$$

where $w^1_{\min,T}$ is the minimum wage in effect at time T for the treated units and $w^0_{\min,T}$ is the minimum wage in effect at time T for the untreated units. Our time series elasticity of employment with respect to the minimum wage is -0.8. This estimate uses only the Minneapolis results because the Saint Paul effects before the implementation of the minimum wage increase in 2020 reflect advance notice. For the elasticity using variation from the cross sections of establishments and workers, we multiply the elasticity from the time series with 0.48, which Table 8 shows to be the average ratio of estimates from the cross section to those from the time series. This yields an elasticity of employment with respect to the minimum wage of -0.38.

Figure 4 presents our estimated elasticity of restaurant employment with respect to the minimum wage and elasticities found in other studies.[28] As seen in the figure, our employment

[28]We include papers published after the study of Card and Krueger (1994) and for which we could obtain an estimated elasticity with respect to the minimum wage for the U.S. restaurant industry.

elasticities are more negative than those found in the literature, which average around −0.1. Our analysis for the estimates in the restaurant industry from other studies corroborates the more comprehensive analysis of estimates for all low-wage workers and low-wage industries in Neumark and Shirley (2021), which reveals that while almost 80 percent of estimates are negative, reported elasticities average −0.15 across studies.

There are two reasons why our employment elasticities with respect to the minimum wage might be larger than those found in the literature. First, the policy variation we examine is larger than the typical variation found in the literature. Clemens and Strain (2021) estimate an elasticity of −0.5 for low-skilled employment during "large minimum wage increases," and the 38 percent increase in the Minneapolis minimum wage by 2020 is larger than the large policy changes examined by these authors.[29] The Seattle study by Jardim, Long, Plotnick, Van Inwegen, Vigdor, and Wething (2022) examines a policy change as large as the one in the Twin Cities and estimates a −0.3 elasticity of employment with respect to the minimum wage for restaurants. The difference in their elasticity with the one we estimate for the Twin Cities could reflect that Seattle was booming during the implementation of its large minimum wage increase. It is conceivable that the sensitivity of employment with respect to the minimum wage becomes larger during recessions, like the one induced by the pandemic. The more negative elasticity during the pandemic that we estimate is consistent with previous studies documenting larger elasticities during recessions (Addison, Blackburn, and Cotti, 2013; Clemens and Wither, 2019).

5.3 Comparison of Labor Demand Elasticity with Other Studies

Dividing the cross-sectional coefficients for jobs or hours, τ_ℓ, in the upper panel of Table 6 with the cross-sectional coefficients for the wage, τ_w, we find a cross-sectional own-wage employment elasticity of around −1. This elasticity is estimated using employment and wage variation induced by heterogeneous exposure to the minimum wage increase. As shown in the table, this ratio is stable across time. In Section 6, we show more formally that this ratio is informative about the own-wage elasticity of labor demand at the establishment level, which we interpret as

[29] Appendix Figure A1 of Clemens and Strain (2021) shows that the 95 percent confidence interval of large minimum wage increases ranges between 20 and 25 percent after three years. The five-year range is 30 to 40 percent, whereas the Minneapolis minimum wage is expected to increase by more than 40 percent by 2023.

a long-run elasticity, as we allow establishments to substitute between labor and other inputs.

How does our estimate of a -1 labor demand elasticity at the establishment level compare with the estimates in the literature? Hamermesh (1993) summarizes various early estimates of labor demand elasticities. While many studies estimate elasticities below one in absolute value, some studies from British plants find an elasticity of around -1. Lichter, Peichl, and Siegloch (2015) update the analysis of Hamermesh (1993) with newer studies and find a mean own-wage elasticity of around -0.6, with roughly 80 percent of the estimates ranging between 0 and -1. However, as discussed by Hamermesh (1993) and Lichter, Peichl, and Siegloch (2015) several of the studies in the literature hold output constant and thus are more appropriately interpreted as estimating factor substitution rather than the total elasticity of labor demand. Beaudry, Green, and Sand (2018) estimate a labor demand elasticity of -1 at the city-industry level. Using a model with search frictions and scarcity in new firm creation, they estimate that the increase in the minimum wage observed in Seattle is associated with a 2.1 percent decline in employment. This estimate matches very closely our estimates in Table 8 for the Twin Cities minimum wage increase, which average 1.9 percent between cities and estimation methods.

6 Reconciling the Time Series with the Cross Section

There are three reasons why time series estimates of employment losses differ from those that use variation from the cross section. First, despite our efforts to difference out other shocks, the Twin Cities may have experienced idiosyncratic shocks or had a differential response to an aggregate shock that cannot be differenced out using other cities during the post-treatment period. The cross-sectional estimates do not suffer from this concern, to the extent that the Twin Cities shocks are differenced out across establishments in the same industry and zip code or across workers in the same industry. Second, the time series effects of the minimum wage on employment are at the industry level and sum up employment effects at the intensive margin, effects arising from the exit of establishments, and effects arising from the lack of entry of new establishments. By design, the estimates from the cross section do not account for the effects of entry, because they use establishments and workers that exist for at least one period. Finally, any other equilibrium adjustment at the industry level affecting the average establishment or worker is included in the time series estimates but not in those from the cross section.[30]

We expect that the first two differences in research design generate more negative employment effects in the time series than in the cross section. We formalize this claim by developing a model of establishment dynamics. We use the model to answer two questions. First, is it plausible to reconcile the time series estimates with those from the cross section by appealing to entry dynamics that are omitted from the analysis of the cross section? Second, what do our estimates imply about the deeper determinants of labor demand, such as parameters that characterize the degree of product and labor market competition, factor substitution within establishments, and establishment dynamics?

6.1 Model Environment

Our model economy consists of establishments that are heterogeneous along four dimensions: productivity z, technology of production ϕ, labor supply to the establishment $\bar{w}$, and fixed cost of entry κ. Establishments choose prices, factors of production, entry, and exit. We consider two periods and denote with a subscript 0 the period before the minimum wage change and with a subscript 1 the period after the minimum wage change. The model is partial equilibrium and the only market clearing condition is that an establishment's labor demand equals workers' labor supply at the establishment level, which, in turn, pins down the wage that each establishment faces.

Period before the minimum wage increase. Establishment $(z, \phi, \bar{w}, \kappa)$ solves the following static problem.[31] Conditional on having entered, the establishment chooses price p_0, labor demand ℓ_0, and all other inputs m_0 (for example, materials and capital) to maximize flow profits

$$\max_{p_0, \ell_0, m_0} \pi_0 = \max\{p(y_0) \cdot y_0(m_0, \ell_0) - w(\ell_0)\ell_0 - m_0 - f, 0\}, \tag{14}$$

where the price of all other inputs is normalized to one. Establishments face operating fixed cost

[30] Examples of such equilibrium effects are wage spillovers to establishments not directly exposed to the minimum wage, a shift of product demand away from an industry, or a shift of labor supply toward an industry. We addressed the concern that non-exposed establishments changed their employment because of worker reallocation by using the cross section of workers to infer the effects of the minimum wage on employment.

[31] Establishment type $(z, \phi, \bar{w}, \kappa)$ is fixed over time, with the only exception of an aggregate shock that we consider later. The model can be extended to accommodate changes in establishment productivity and labor supply over time. If changes in establishment heterogeneity occur with a perceived probability of zero, then we can still focus on a sequence of static profit maximization problems.

f that is constant over time and establishments. In equation (14), profits cannot be negative, because establishments choose to exit when total costs exceed revenues.

We allow establishments to potentially have market power in both product and labor markets. In the product market, establishments internalize that their product demand is downward sloping

$$p_0 = y_0^{-\frac{1}{\varepsilon}}, \tag{15}$$

where the elasticity of product demand $\varepsilon > 1$ is constant over time and across establishments. When $\varepsilon \to \infty$, we obtain the limiting case of perfect competition in the product market.[32] In the labor market, establishments internalize that their labor supply is upward sloping

$$w_0 = \bar{w}\ell_0^{\frac{1}{\theta}}, \tag{16}$$

where the labor supply elasticity $\theta > 0$ is also constant over time and across establishments. When $\theta \to \infty$, we obtain the limiting case of perfect competition in the labor market.[33]

Establishments operate a CES technology

$$y_0 = z\left(\phi m_0^{\frac{\sigma-1}{\sigma}} + (1-\phi)\ell_0^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}, \tag{17}$$

where $\phi \in [0,1]$ is a distribution parameter characterizing the technology of production and $\sigma \geq 0$ is the elasticity of substitution between all other inputs m and labor ℓ, which is constant over time and across establishments.

Establishments do not know part of their type, $(z, \phi, \bar{w})$, before they enter. As a result, the decision to enter compares expected profits to the fixed cost of entry. Establishments enter if

$$\mathbb{E}\pi_0 \geq \kappa. \tag{18}$$

Establishments have different entry outcomes because they differ in their cost of entry κ.

[32] We have normalized the shifter of product demand to one for all establishments. Given this normalization, productivity z is better understood as a convolution of physical productivity and demand for an establishment's product. For brevity, we do not present the household side of the model, which gives rise to establishments' demand functions. Establishments face a downward sloping demand for their product when households view their products as imperfect substitutes, with ε governing the substitutability of product demand across different establishments.

[33] As we did for product demand, for brevity we do not present the worker side of the model. Establishments face an upward sloping labor supply when workers view them as imperfect substitutes, with θ governing the substitutability of labor supply across different establishments and $\bar{w}$ capturing parameters related to the disutility of work across establishments.

Period after the minimum wage increase. Fraction $\delta \in [0, 1]$ of establishments are exogenously destroyed in the end of the first period after exit, pricing, and production decisions are made. A fraction $1 - \delta$ of establishments continues in the next period. For a continuing establishment, the maximization problem is similar to that of the first period, with two differences. First, continuing establishments do not pay the entry cost κ, because they know their type. Second, establishments potentially face a binding minimum wage. The minimum wage increase is unexpected as of the first period, which justifies focusing on a sequence of static problems for establishments. For continuing establishments, profits equal

$$\max_{p_1,\ell_1,m_1} \pi_1^c = \max\{p(y_1) \cdot y_1(m_1, \ell_1) - \max\{w(\ell_1), w_{\min}\}\ell_1 - m_1 - f, 0\}, \tag{19}$$

where the max operator inside the brackets denotes that establishments cannot pay a wage that falls below the minimum wage. The establishments face the same functional forms for product demand in equation (15), labor supply in equation (16) if the minimum wage does not bind, and production function in equation (17).

There is a mass δM of potential entrants in the second period, where M is the mass of potential entrants in the first period. Potential entrants in the second period behave similarly to potential entrants in the first period. Conditional on having entered, they maximize

$$\max_{p_1,\ell_1,m_1} \pi_1^n = \max\{p(y_1) \cdot y_1(m_1, \ell_1) - \max\{w(\ell_1), w_{\min}\}\ell_1 - m_1 - f, 0\}, \tag{20}$$

where product demand is given by equation (15), labor supply by equation (16) if the minimum wage does not bind, and the production function by equation (17). Potential entrants in the second period enter if

$$\mathbb{E}\pi_1^n \geq \kappa, \tag{21}$$

where the expectation of profits uses the distribution of establishment characteristics of potential entrants, as opposed to the characteristics of continuing establishments.

6.2 Effects of the Minimum Wage on Establishment Labor Demand

Labor demand of operating establishments in the period before the minimum wage increase is

$$\ell_0 = (1-\phi)^{\varepsilon} z^{\varepsilon-1} \left(\phi \left(\frac{\phi \mu_w w_0}{1-\phi} \right)^{\sigma-1} + (1-\phi) \right)^{\frac{\varepsilon-\sigma}{\sigma-1}} (\mu_p \mu_w w_0)^{-\varepsilon}, \tag{22}$$

where $\mu_p = \frac{\varepsilon}{\varepsilon-1} \geq 1$ denotes the product market markup and $\mu_w = \frac{\theta+1}{\theta} \geq 1$ denotes the labor market markup. Both markups reduce labor demand for a given level of the wage.

As in Hamermesh (1993), we calculate the total elasticity of labor demand

$$\eta_0 = (1-\alpha_0)\varepsilon + \alpha_0\sigma, \tag{23}$$

which is a weighted average of the two elasticities ε and σ. The elasticity of product demand ε appears in the elasticity of labor demand because a higher wage increases the marginal cost and prices and thus lowers demand for an establishment's product. This is the scale effect that allows output y to adjust. The elasticity of factor substitution σ appears in the elasticity of labor demand because a higher wage induces establishments to lower labor and increase all other inputs along a stable isoquant. This is the substitution effect that holds output y constant.[34] Weights $\alpha_0 = \frac{\left(\frac{\phi}{1-\phi}\right)^{\sigma}(\mu_w w_0)^{\sigma-1}}{1+\left(\frac{\phi}{1-\phi}\right)^{\sigma}(\mu_w w_0)^{\sigma-1}}$ equal ϕ in the limiting case of a Cobb-Douglas production function with $\sigma = 1$. With an elasticity of substitution different than one, weights α_0 vary as a function of the wage. Weights α_0 do not equal the input share of variable costs unless there is a perfectly competitive labor market, $\mu_w = 1$.

Labor demand of operating establishments in the period after the introduction of the minimum wage is

$$\ell_1 = \begin{cases} (1-\phi)^{\varepsilon} z^{\varepsilon-1}\left(\phi\left(\frac{\phi\mu_w w_1}{1-\phi}\right)^{\sigma-1} + (1-\phi)\right)^{\frac{\varepsilon-\sigma}{\sigma-1}} (\mu_p\mu_w w_1)^{-\varepsilon}, & \text{if } w_1 > w_{\min}, \\ (1-\phi)^{\varepsilon} z^{\varepsilon-1}\left(\phi\left(\frac{\phi w_{\min}}{1-\phi}\right)^{\sigma-1} + (1-\phi)\right)^{\frac{\varepsilon-\sigma}{\sigma-1}} (\mu_p w_{\min})^{-\varepsilon}, & \text{if } w_1 \leq w_{\min}. \end{cases} \tag{24}$$

The effects of the minimum wage on establishments' labor demand can be understood by comparing labor demand before the minimum wage in equation (22) to labor demand with a binding minimum wage in the second line of equation (24). In the absence of market power in the labor market, $\mu_w = 1$, labor demand unambiguously falls for establishments for which the minimum wage $w_{\min}$ exceeds their competitive wage w_0. This is a movement along the labor demand curve, with the elasticity of labor demand η governing the magnitude of reduction in labor. With labor market power, $\mu_w > 1$, the employment effect of the minimum wage is theoretically

[34] Our elasticity of labor demand is a long-run elasticity because we allow other inputs to adjust. If we keep other inputs m fixed at the level before the minimum wage increase, the labor demand elasticity equals $\eta^m = \frac{\varepsilon\sigma}{\sigma+\tilde{\alpha}(\varepsilon-\sigma)}$, where $\tilde{\alpha} = \phi m^{\frac{\sigma-1}{\sigma}} / \left(\phi m^{\frac{\sigma-1}{\sigma}} + (1-\phi)\ell^{\frac{\sigma-1}{\sigma}}\right)$. We show that $\eta > \eta^m$ in the limiting case $\tilde{\alpha} \to \alpha$, which one obtains when $\sigma = 1$ for example.

Table 9: Parameters and Data Moments

Parameter	Value	Data Moments	Value
Elasticity of Factor Substitution, σ	0.19	Cross-Sectional Elasticity, τ_ℓ/τ_w	-1.01
Mean Intensity of Other Inputs, μ_ϕ	0.52	Labor Share of Costs, Mean	0.30
Share of Affected Establishments, q	0.23	GAP, Employment-Weighted	0.14
Mean Entry Cost, μ_κ	0.06	Entry Cost to Profit, Median	0.36
Fixed Operating Cost, f	0.08	Operating Cost to Profit, Median	0.59

Notes: The first two columns present the value of calibrated parameters. The last two columns present the moments in the data most closely associated with the chosen value of each parameter.

ambiguous. A binding minimum wage removes labor market power because establishments perceive an infinite elasticity of labor supply for all labor units below the minimum wage. This effect represents an upward shift of labor demand. Depending on the values of $\mu_w w_0$ and $w_{\min}$, the minimum wage may result in an increase or decrease in equilibrium labor.

6.3 Effects of Minimum Wage on Aggregate Labor Demand

We now parameterize the model economy and evaluate how aggregate outcomes at the industry level change when the minimum wage is introduced.

Parameterization. We set the exogenous destruction rate to $\delta = 0.28$ to match the average exit rate of the Twin Cities restaurants in the DEED data during the years before the minimum wage increase. We assume that $(\log z, \phi, \log w, \kappa)$ are normally distributed and, as a baseline, uncorrelated with each other.[35] From the DEED data, we calculate that the standard deviation of the log wage across restaurants equals 0.46 and thus we set $\sigma_w = \sigma_z = 0.46$.[36] For the parameters characterizing market power, we begin our analysis with $\varepsilon = 2$ and $\theta \to \infty$.

[35] To simplify the exposition, we invert the maximization problem and represent establishments as directly drawing a wage w. We then find ℓ from the labor demand functions, equations (22) and (24), and use the labor supply function to solve for the primitive source of heterogeneity characterizing each establishment, $\bar{w} = w\ell^{-1/\theta}$. The only subtlety is how we treat aggregate productivity shocks that directly change the equilibrium wage. For that case, we have opted for adjusting labor supply parameters $\bar{w}$ so that the equilibrium wage changes only in response to the increase in the minimum wage, because our time series and cross-sectional evidence attempt to difference out aggregate shocks and therefore show an increase in the wage.

[36] We do not have information that allows us to estimate the standard deviation of technology, σ_ϕ, and entry cost, σ_κ. To be conservative, we have opted for a small dispersion of these sources of heterogeneity and set the two standard deviations to 10 percent of their means.

Our choice reflects two observations. First, under a higher product market elasticity ε, the model produces higher labor demand elasticities than those implied from the cross-sectional relationship between employment and the wage that we estimate in Section 4. Second, we can rationalize the negative employment effects from both the time series and the cross section as outcomes from a model with perfectly competitive labor markets. Given the data that we have access to, we are not able to identify θ from other parameters and thus we will examine the sensitivity of our results to this parameter.

Table 9 presents the remaining five parameters, which we calibrate so that our model perfectly matches five moments from the data. An elasticity of factor substitution $\sigma = 0.19$ allows the model to replicate the cross-sectional elasticity of employment with respect to the wage, $\tau_\ell/\tau_w = -1.01$, found by dividing the jobs column with the wage column of the upper panel of Table 6. The mean technology parameter is $\mu_\phi = 0.52$, so that the model generates a labor share of operating costs for restaurants equal to 0.30, as calculated by Aaronson and French (2007). Next, we calibrate the share of the establishments q for which the minimum wage exceeds their wage in the absence of a minimum wage. We find that the employment-weighted average GAP in our model equals its analog from the data for $q = 0.23$.[37] Finally, we calibrate the mean fixed cost of entry μ_κ and the fixed cost of operation f to match estimates of these costs for restaurants by Aguirregabiria and Mira (2007). The entry cost represents 36 percent of profits and 19 percent of revenues. The operating cost is larger than the entry cost, representing 59 percent of profits and 30 percent of revenues.

Baseline results. The first row of Table 10 presents outcomes from our baseline calibrated economy. By design, the model reproduces the elasticity of employment with respect to the wage, $\tau_\ell/\tau_w = -1.01$. Following the introduction of the minimum wage, the average wage in the economy grows by 9 percent. Had we used τ_ℓ/τ_w to infer the employment effects of the minimum wage, we would have concluded that employment also declines by 9 percent. However, aggregate employment in the calibrated economy declines by 25 percent. The logic is that the increase in the minimum wage puts a downward pressure on expected profits and thus limits the

[37]We cannot directly target the increase in the minimum wage, because our model does not have a binding minimum wage in the first period. The minimum wage is 40 percent higher than the 5th percentile of the wage without the minimum wage, an increase which is consistent with the Minneapolis minimum wage implemented in 2020.

Table 10: Aggregate Effects of the Minimum Wage Increase in the Model

	τ_ℓ/τ_w	η	W_1/W_0-1	L_1/L_0-1	Ratio of Entry
Baseline	−1.01	−1.07	0.09	−0.25	2.28
$\mu_\kappa = 0$	−1.01	−1.07	0.09	−0.08	1.00
$\delta = 0$	−1.01	−1.07	0.09	−0.08	∞
$f = 0.08$	−0.98	−1.07	0.09	−0.21	1.75
$f = \bar{f}w^{1/2}$	−0.95	−1.04	0.08	−0.29	3.50
$\mu_\phi = 0.80$	−0.88	−0.96	0.08	−0.16	1.40
$\sigma = 1$	−1.57	−1.48	0.14	−0.36	2.37
$\mathrm{corr}(\log z, \log w) = 0.50$	−1.46	−1.10	0.07	−0.29	3.43
$\mathrm{corr}(\phi, \log w) = 0.50$	−1.03	−1.07	0.09	−0.26	2.35
$\theta = 5$	−0.77	−1.12	0.07	−0.20	2.33
$\theta = 2$	−0.45	−1.19	0.06	−0.12	2.46
$\Delta \log Z = -0.11, \mu_\kappa = 0$	−1.01	−1.26	0.09	−0.25	1.00

Notes: The table presents values of the cross-sectional elasticity of employment with respect to the wage, τ_ℓ/τ_w, the elasticity of labor demand at the establishment level, η, the percent change in the aggregate wage between the period with the minimum wage and the period without the minimum wage, $W_1/W_0 - 1$, the percent change in aggregate employment between the two periods, $L_1/L_0 - 1$, and the ratio of entry in the period with the minimum wage to entry in the period without the minimum wage. The first row presents these values for the baseline parameters and the other rows present these values for alternative model parameterizations.

entry of establishments. The effect of reduced entry on aggregate employment is not reflected in the cross-sectional elasticity τ_ℓ/τ_w, because τ_ℓ/τ_w uses information only from establishments which have entered.

The wage increase of 9 percent in the model matches perfectly the 9 percent wage increase for restaurants that we estimated from the time series. The employment decline of 25 percent in the model is close to the 30 percent decline that we estimated for restaurants. The similarity of the employment and wage responses between the model and the data is reassuring for the potential of endogenous entry to account for the difference between employment effects in the time series and in the cross section. We calibrated the entry cost independently of our time series treatment effects by appealing to the industrial organization literature for estimates of this cost.

Our model generates 17 log points decline in the number of establishments operating after the minimum wage increase. We find some evidence that declines in the number of operating establishments are observed for full-service restaurants in the data.[38] Partly because of the decline in the number of establishments and partly because we find negative employment effects using variation from both the time series and the cross section, our model of establishment dynamics is closer in spirit to the model of Hopenhayn (1992), which allows for flexible adjustment of labor, than putty-clay models such as those in Sorkin (2015) and Aaronson, French, Sorkin, and To (2018), which highlight the importance of higher entry of capital-intensive establishments. Consistent with our mechanism, Draca, Machin, and Van Reenen (2011) document a decline in firm profitability induced by the introduction of a UK national minimum wage in 1999. Harasztosi and Lindner (2019) also document a decline in firm profitability for the case of Hungary, but find that consumers bear the largest share of the increased cost from the minimum wage partly by paying higher product prices.

The column labeled η presents the elasticity of labor demand at the establishment level, averaging formula (23) across the two periods. The model-generated elasticity of labor demand is -1.07, which is close to the cross-sectional elasticity of employment with respect to the wage, $\tau_\ell/\tau_w = -1.01$, induced by the change in the minimum wage in the data. We conclude that our results from the cross section are informative about the elasticity of labor demand at the establishment level.[39]

Understanding model mechanisms. The second row shows the importance of endogenous entry by setting the entry cost to zero for all establishments. In the absence of endogenous entry, the decline in aggregate employment is only 8 percent and can be inferred from the cross-sectional elasticity of employment with respect to the wage. Similarly, when the exogenous rate of destruction δ is zero, the decline in aggregate employment is only 8 percent because there is

[38] We define an establishment as active in a quarter if it reports positive employment in that quarter. Using DEED data and the same synthetic difference-in-differences approach as in our other time series estimates, we find 16 log points decline in Minneapolis and 6.3 log points decline in Saint Paul in the number of establishments.

[39] The two elasticities need not be equal, partly because there is an omitted variable bias when projecting employment on the wage in the cross section in order to infer the elasticity of labor demand. The bias arises because the weight α and productivity z are correlated with the wage owing to selection into entry and exit and because α varies endogenously with the wage whenever the elasticity of factor substitution $\sigma \neq 1$. In practice, the gap between the two elasticities is small under the baseline parameterization, because using the arc growth of employment removes persistent sources of heterogeneity correlated with the wage and allows us to include exiting establishments.

no entry in the period with the minimum wage.

The next row shows that lowering the fixed operating cost f leads to a smaller sensitivity of entry to the minimum wage and thus smaller employment losses at the aggregate. Denominating the fixed cost as a geometric mean of input prices, $f = \bar{f}w^{1/2}$, increases the sensitivity of entry and the employment losses in response to the increase in the minimum wage, because establishments with a binding minimum wage face an increasing operating cost over time. Lowering the labor share of operating costs by setting $\mu_\phi = 0.80$ weakens the cross-sectional elasticity of employment with respect to the wage, τ_ℓ/τ_w, and also makes entry less sensitive to the minimum wage increase. By contrast, raising the elasticity of factor substitution to $\sigma = 1$ increases the labor demand elasticity η and leads to larger employment losses from the introduction of the minimum wage.

Next, we consider how of our results change if we allow the sources of heterogeneity to be correlated within establishments. Introducing a correlation of 0.5 between productivity and the wage strengthens the cross-sectional elasticity of employment with respect to the wage, τ_ℓ/τ_w, and also makes entry more sensitive to the minimum wage increase. By contrast, introducing a correlation between the intensity of all other inputs in production and the wage does not have a discernible effect on the outcomes of the model economy.

The next experiment is to introduce labor market power by allowing parameter θ to be lower than infinity. In an economy with labor market power, the cross-sectional elasticity of employment with respect to the wage, τ_ℓ/τ_w, is less informative about establishments' elasticity of labor demand η. The logic is that employment and the wage in the cross section of establishments reflect both movements along a stable labor demand and shifts of the labor demand, because some establishments with a binding minimum wage increase their employment. The employment growth for some establishments also makes aggregate employment decline by less than in an economy without labor market power.[40]

It is worth highlighting that our results do not allow us to directly identify how competitive

[40]The cross-sectional elasticity of employment with respect to the wage is more informative about η if establishments with a binding minimum wage have a low labor market power. Berger, Herkenhoff, and Mongey (2022) develop a general equilibrium model with labor market power that varies with establishment size to evaluate the efficiency and redistributive effects of a minimum wage. They find small total factor productivity and employment gains from introducing a minimum wage because establishments with a binding minimum wage tend to be unproductive and have a low labor market power.

labor markets were in the Twin Cities before the introduction of the minimum wage. However, the model we reject is that labor markets are imperfectly competitive and that the increase in the minimum wage is sufficiently small to induce an equilibrium wage below the competitive level. This model generates positive employment effects in both the cross section and the time series, which are inconsistent with our estimates. We conclude that our results are compatible either with perfectly competitive labor markets or the introduction of a minimum wage above the competitive level for most establishments that may operate in labor markets with some degree of monopsony power.

Aggregate shock contemporaneous with the minimum wage increase. The final row of the table shuts off endogenous entry and introduces an aggregate shock to productivity, $\Delta \log Z = -0.11$.[41] The outcomes in this economy are identical to the outcomes in the baseline economy with entry but no aggregate shocks. We conclude that the divergence between the employment effects in the time series and the employment effects in the cross section can reflect either endogenous entry that is missing from the cross section or an aggregate shock that is contemporaneous with the minimum wage change and differenced out in the cross section. Thus, both estimation strategies are informative and provide plausible bounds for the employment effects of the minimum wage increase.

7 Conclusion

We use high-quality administrative data from the state of Minnesota to analyze the labor market effects of two large increases in the minimum wage from Minneapolis and Saint Paul. Our analysis proceeds in three steps. Leveraging recent advances in synthetic difference-in-differences approaches, we estimate counterfactual outcomes in the absence of the minimum wage using variation at the zip code within Minnesota or at the city level from the rest of the country. Using variation from the cross sections of establishments and workers within the Twin Cities, we estimate the labor market effects of differential exposure to the minimum

[41]The economy is recalibrated to match the targets in Table 9. This leads to $\sigma = 0.57$ and $q = 0.20$, whereas all other parameters are held constant at their previous values. In this economy, we adjust the labor supply shifters $\bar{w}$ so that the wage is constant when aggregate Z declines. Thus, the 9 percent growth of the wage reflects only the introduction of the minimum wage.

wage increase. Finally, using a model of establishment dynamics, we attempt to reconcile the results from the time series with those from the cross section and discuss how our estimates are informative for deeper determinants of establishments' labor demand.

We reach several substantial conclusions. The minimum wage increase is associated with wage gains in most low-wage industries that accord well with estimates of the direct effect of the minimum wage on establishments' labor costs. The time series analysis shows that in the Twin Cities the minimum wage increase is associated with an average jobs decline of roughly 2 percent. The great majority of job losses appear in the restaurant industry, which experiences a jobs decline of roughly 30 percent. The analysis using variation from the cross section leads to estimates of job losses that are half as large as the estimates from the time series. Our cross-sectional estimates are consistent with an elasticity of labor demand at the establishment level equal to -1, for establishments that operate in competitive labor markets and can adjust flexibly their inputs.

We offer two ways to reconcile the results from the time series with those from the cross section. One possibility is that the Twin Cities may have experienced idiosyncratic shocks or had a differential response to an aggregate shock that cannot be differenced out using other cities during the post-treatment period. Another possibility is that estimates from the cross section do not account for the effects of entry or other equilibrium effects. Using our model, we illustrate that either endogenous entry or confounding factors contemporaneous with the minimum wage can account for the divergence of the results from the two research designs. We argue that both research designs are informative and plausibly bound the effects of the minimum wage increase.

References

Aaronson, D., and E. French (2007): "Product Market Evidence on the Employment Effects of the Minimum Wage," *Journal of Labor Economics*, 25(1), 167–200.

Aaronson, D., E. French, I. Sorkin, and T. To (2018): "Industry Dynamics and the Minimum Wage: A Putty-Clay Approach," *International Economic Review*, 59(1), 51–84.

Abadie, A., A. Diamond, and J. Hainmueller (2015): "Comparative Politics and the Synthetic Control Method," *American Journal of Political Science*, 59(2), 495–510.

Abadie, A., and J. Gardeazabal (2003): "The Economic Costs of Conflict: A Case Study of the Basque Country," *American Economic Review*, 93(1), 113–132.

Addison, J., M. Blackburn, and C. Cotti (2012): "The Effect of Minimum Wages on Labour Market Outcomes: County-Level Estimates from the Restaurant-and-Bar Sector," *British Journal of Industrial Relations*, 50(3), 412–435.

——— (2013): "Minimum Wage Increases in a Recessionary Environment," *Labour Economics*, 23, 30–39.

——— (2015): "On the Robustness of Minimum Wage Effects: Geographically-Disparate Trends and Job Growth Equations," *IZA Journal of Labor Economics*, 4(1), 1–16.

Aguirregabiria, V., and P. Mira (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75(1), 1–53.

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

Allegretto, S., A. Dube, M. Reich, and B. Zipperer (2017): "Credible Research Designs for Minimum Wage Studies: A Response to Neumark, Salas, and Wascher," *ILR Review*, 70(3), 559–592.

Arkhangelsky, D., S. Athey, D. Hirshberg, G. Imbens, and S. Wager (2021): "Synthetic Difference in Differences," *American Economic Review*, 111(12), 4088–4118.

Barrett, G. F., and D. S. Hamermesh (2019): "Labor Supply Elasticities: Overcoming Nonclassical Measurement Error Using More Accurate Hours Data," *Journal of Human Resources*, 54(1), 255–265.

Beaudry, P., D. Green, and B. Sand (2018): "In Search of Labor Demand," *American Economic Review*, 108(9), 2714–2757.

Berger, D., K. Herkenhoff, and S. Mongey (2022): "Minimum Wages, Efficiency, and Welfare," University of Minnesota Working Paper.

Bertrand, M., E. Duflo, and S. Mullainathan (2004): "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 119(1), 249–275.

Bound, J., C. Brown, and N. Mathiowetz (2001): "Measurement Error in Survey Data," in *Handbook of Econometrics*, vol. 5, pp. 3705–3843. Elsevier.

Brummund, P., and M. Strain (2020): "Does Employment Respond Differently to Minimum Wage Increases in the Presence of Inflation Indexing?," *Journal of Human Resources*, 55(3), 999–1024.

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

——— (2000): "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Reply," *American Economic Review*, 90(5), 1397–1420.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019): "The Effect of Minimum Wages on Low-Wage Jobs," *Quarterly Journal of Economics*, 134(3), 1405–1454.

Clemens, J., and M. Strain (2021): "The Heterogeneous Effects of Large and Small Minimum Wage Changes: Evidence over the Short and Medium Run Using a Pre-Analysis Plan," NBER Working Paper No. 29264.

Clemens, J., and M. Wither (2019): "The Minimum Wage and the Great Recession: Evidence of Effects on the Employment and Income Trajectories of Low-Skilled Workers," *Journal of Public Economics*, 170, 53–67.

Couch, K., and D. Wittenburg (2001): "The Response of Hours of Work to Increases in the Minimum Wage," *Southern Economic Journal*, 68(1), 171–177.

Draca, M., S. Machin, and J. Van Reenen (2011): "Minimum Wages and Firm Profitability," *American Economic Journal: Applied Economics*, 3(1), 129–151.

Dube, A., T. W. Lester, and M. Reich (2010): "Minimum Wage Effects across State Borders: Estimates Using Contiguous Counties," *Review of Economics and Statistics*, 92(4), 945–964.

——— (2016): "Minimum Wage Shocks, Employment Flows, and Labor Market Frictions," *Journal of Labor Economics*, 34(3), 663–704.

Dube, A., and A. Lindner (2021): "City Limits: What Do Local-Area Minimum Wages Do?," *Journal of Economic Perspectives*, 35(1), 27–50.

Dube, A., S. Naidu, and M. Reich (2007): "The Economic Effects of a Citywide Minimum Wage," *ILR Review*, 60(4), 522–543.

Dustmann, C., A. Lindner, U. Schonberg, M. Umkehrer, and P. vom Berge (2022): "Reallocation Effects of the Minimum Wage," *Quarterly Journal of Economics*, 137(1), 267–328.

Even, W., and D. Macpherson (2014): "The Effect of the Tipped Minimum Wage on Employees in the US Restaurant Industry," *Southern Economic Journal*, 80(3), 633–655.

Even, W., and D. Macpherson (2019): "Where Does the Minimum Wage Bite Hardest in California?," *Journal of Labor Research*, 40, 1–23.

Ferman, B., and C. Pinto (2021): "Synthetic Controls with Imperfect Pretreatment Fit," *Quantitative Economics*, 12(4), 1197–1221.

Hamermesh, D. (1993): *Labor Demand*. Princeton University Press.

Harasztosi, P., and A. Lindner (2019): "Who Pays for the Minimum Wage?," *American Economic Review*, 109(8), 2693–2727.

Heckman, J. (1993): "What Has Been Learned about Labor Supply in the Past Twenty Years?," *American Economic Review*, 83(2), 116–121.

Hirsch, B., B. Kaufman, and T. Zelenska (2015): "Minimum Wage Channels of Adjustment," *Industrial Relations*, 54(2), 199–239.

Hopenhayn, H. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60(5), 1127–1150.

Jardim, E., M. Long, R. Plotnick, E. Van Inwegen, J. Vigdor, and H. Wething (2022): "Minimum Wage Increases and Low-Wage Employment: Evidence from Seattle," *American Economic Journal: Economic Policy*, 14(2), 263–314.

Lichter, A., A. Peichl, and S. Siegloch (2015): "The Own-Wage Elasticity of Labor Demand: A Meta-Regression Analysis," *European Economic Review*, 80, 94–119.

Michl, T. (2000): "Can Rescheduling Explain the New Jersey Minimum Wage Studies?," *Eastern Economic Journal*, 26(3), 265–276.

Neumark, D., J. Salas, and W. Wascher (2014a): "More on Recent Evidence on the Effects of Minimum Wages in the United States," *IZA Journal of Labor Policy*, 24(3), 1–26.

——— (2014b): "Revisiting the Minimum Wage-Employment Debate: Throwing Out the Baby with the Bathwater?," *ILR Review*, 67(3), 608–648.

Neumark, D., M. Schweitzer, and W. Wascher (2004): "Minimum Wage Effects throughout the Wage Distribution," *Journal of Human Resources*, 39(2), 425–450.

Neumark, D., and P. Shirley (2021): "Myth or Measurement: What Does the New Minimum Wage Research Say about Minimum Wages and Job Loss in the United States?," NBER Working Paper No. 28388.

NEUMARK, D., AND W. WASCHER (2000): "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Comment," *American Economic Review*, 90(5), 1362–1396.

SINGELL, L., AND J. TERBORG (2007): "Employment Effects of Two Northwest Minimum Wage Initiatives," *Economic Inquiry*, 45(1), 40–55.

SORKIN, I. (2015): "Are There Long-Run Effects of the Minimum Wage?," *Review of Economic Dynamics*, 18(2), 306–333.

TOTTY, E. (2017): "The Effect of Minimum Wages on Employment: A Factor Model Approach," *Economic Inquiry*, 55(4), 1712–1737.

WANG, W., P. PHILLIPS, AND L. SU (2019): "The Heterogeneous Effects of the Minimum Wage on Employment across States," *Economics Letters*, 174, 179–185.

WOLFSON, P., AND D. BELMAN (2001): "The Minimum Wage, Employment, and the AS-IF Methodology: A Forecasting Approach to Evaluating the Minimum Wage," *Empirical Economics*, 26(3), 487–514.

ZAVODNY, M. (2000): "The Effect of the Minimum Wage on Employment and Hours," *Labour Economics*, 7(6), 729–750.

Minimum Wages and Labor Markets in the Twin Cities

Online Appendix

Loukas Karabarbounis Jeremy Lise Anusha Nath

This Appendix reports additional results and analyses.

- Table A.1 reports the industry distribution of employment shares and the fraction of workers earning below 15 dollars in 2017 by industry. The shares of employment do not add up to 100 percent, as some industries have been excluded because of confidentiality concerns based on the presence of few establishments. The excluded industries are Agriculture, Forestry, Fishing, and Hunting (11); Mining, Quarrying, and Oil and Gas Extraction (21); Construction (23); Information (51); Real Estate and Rental and Leasing (53); and Public Administration (92). The fraction of workers earning below 15 dollars reported in Table A.1 for the restaurant industries is a lower bound for the fraction of workers who are affected by the minimum wage increase. This is because earnings reported to DEED include tips and the minimum wage ordinance excludes tips.

- Figures A.1 and A.2 show the time series of the wage and jobs in Minneapolis and Saint Paul, as well as for the Minnesota average of other cities and for the synthetic control, for retail; Figures A.3 and A.4 for administration and support; Figures A.5 and A.6 for health care and social assistance; Figures A.7 and A.8 for arts, entertainment, and recreation; Figures A.9 and A.10 for accommodation and food services; Figures A.11 and A.12 for other services; Figures A.13 and A.14 for full-service restaurants; and Figures A.15 and A.16 for limited-service restaurants.

- Table A.2 presents R-squared coefficients from regressions of outcome variables in Minneapolis or Saint Paul on the outcome variables of the synthetic control calculated using the weights $\hat{\omega}_i$. To set a baseline of comparison, we also present the R-squared coefficients when using the outcome variables of the unweighted average of all other zip codes within cities in Minnesota. The regressions are performed only during the pre-treatment period.

- Table A.3 repeats our estimates from the DEED data when we exclude bordering cities from the sample of cities that form the synthetic control.
- Table A.4 presents the list of cities included in the control group of Minneapolis and Saint Paul for the analyses using QCEW data from other U.S. cities of similar size.
- Table A.5 presents Monte Carlo simulations to assess the size and sources of bias of the synthetic difference-in-differences method when the true data generating process is a factor model.

Table A.1: Employment Shares and Fraction of Workers Earning below 15 Dollars

(2017)	Share of Employment (percent)			Fraction of Workers Earning Below 15 Dollars		
	MPLS	SP	Other MN	MPLS	SP	Other MN
Manufacturing (31)	4	4	12	14	18	17
Wholesale Trade (42)	3	3	4	11	16	15
Retail Trade (44)	5	7	12	59	63	65
Transportation (48)	2	2	3	20	21	23
Finance and Insurance (52)	11	5	4	5	6	13
Professional Services (54)	11	4	4	5	12	12
Management of Companies (55)	5	4	3	15	29	12
Administration and Support (56)	6	6	5	58	66	48
Educational Services (61)	13	13	8	22	23	23
Health Care and Social Assistance (62)	17	18	17	30	42	34
Arts, Entertainment, and Recreation (71)	2	2	2	42	45	61
Accommodation and Food Services (72)	8	10	9	54	63	71
Other Services (81)	3	4	3	40	34	49
Restaurant Industries						
Full-Service Restaurants (722511)	4	4	3	46	51	56
Limited-Service Restaurants (722513)	2	4	3	80	82	90

Note: MPLS denotes Minneapolis, SP denotes Saint Paul and Other MN denotes the sum of all other cities in Minnesota.

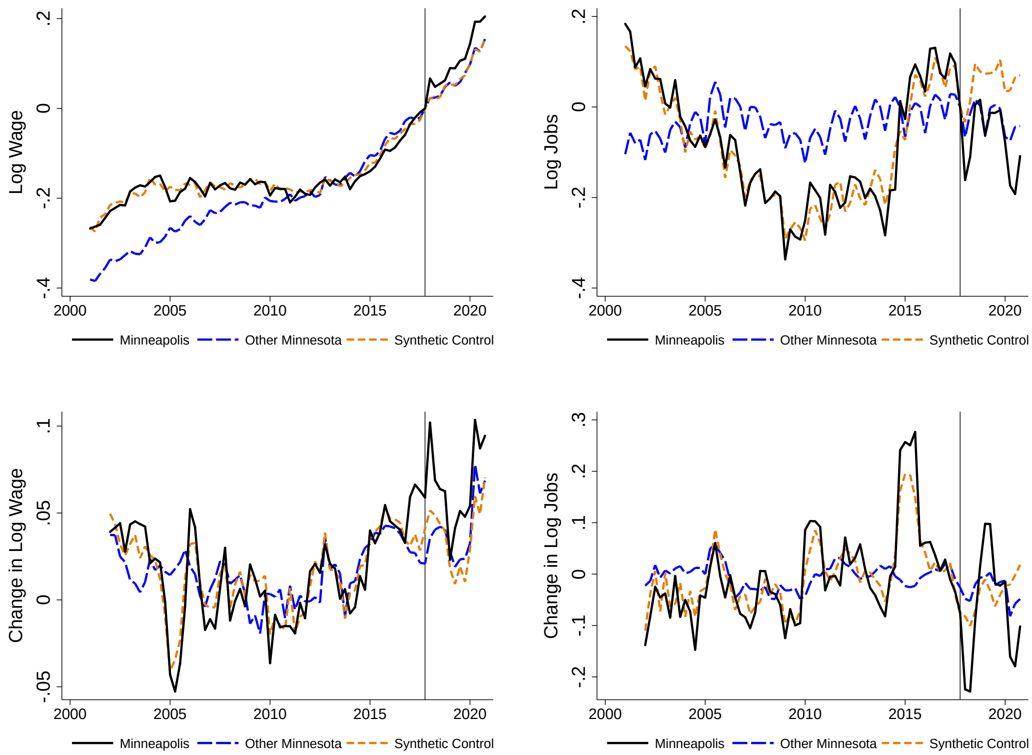

Figure A.1: Time Series of Retail Trade in Minneapolis

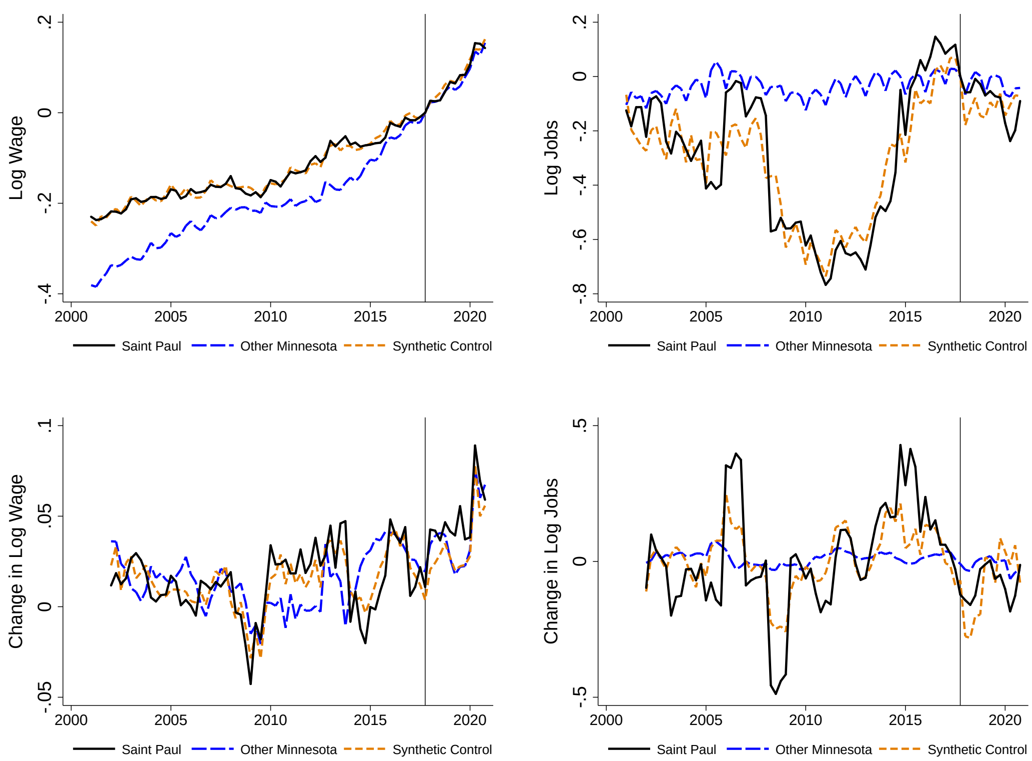

Figure A.2: Time Series of Retail Trade in Saint Paul

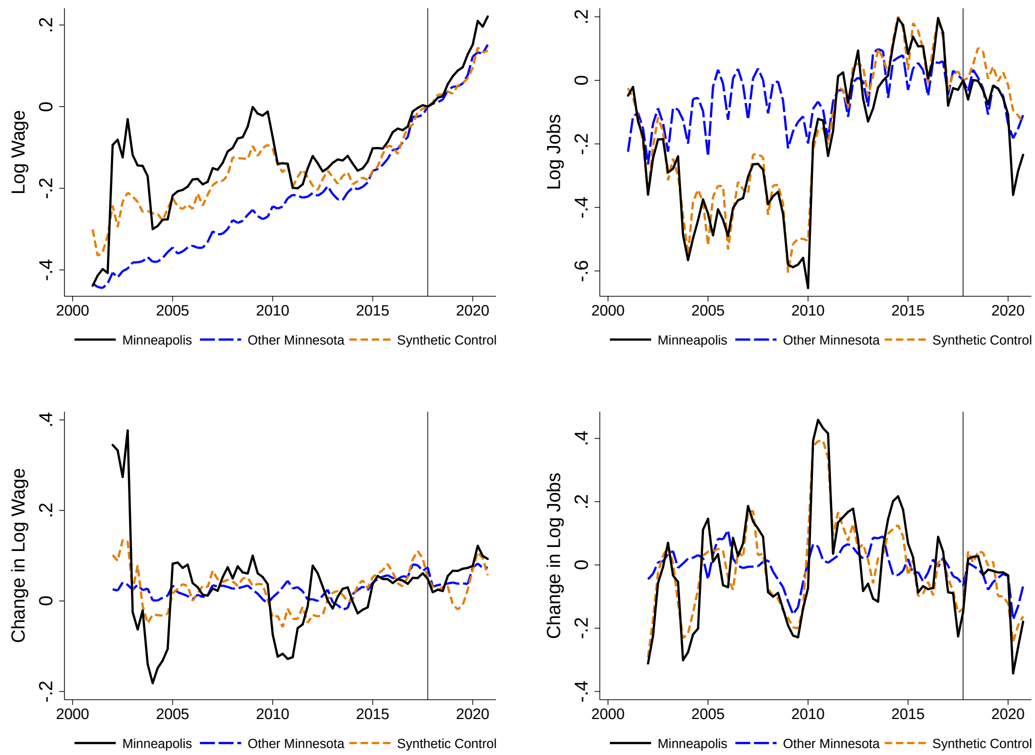

Figure A.3: Time Series of Administration and Support in Minneapolis

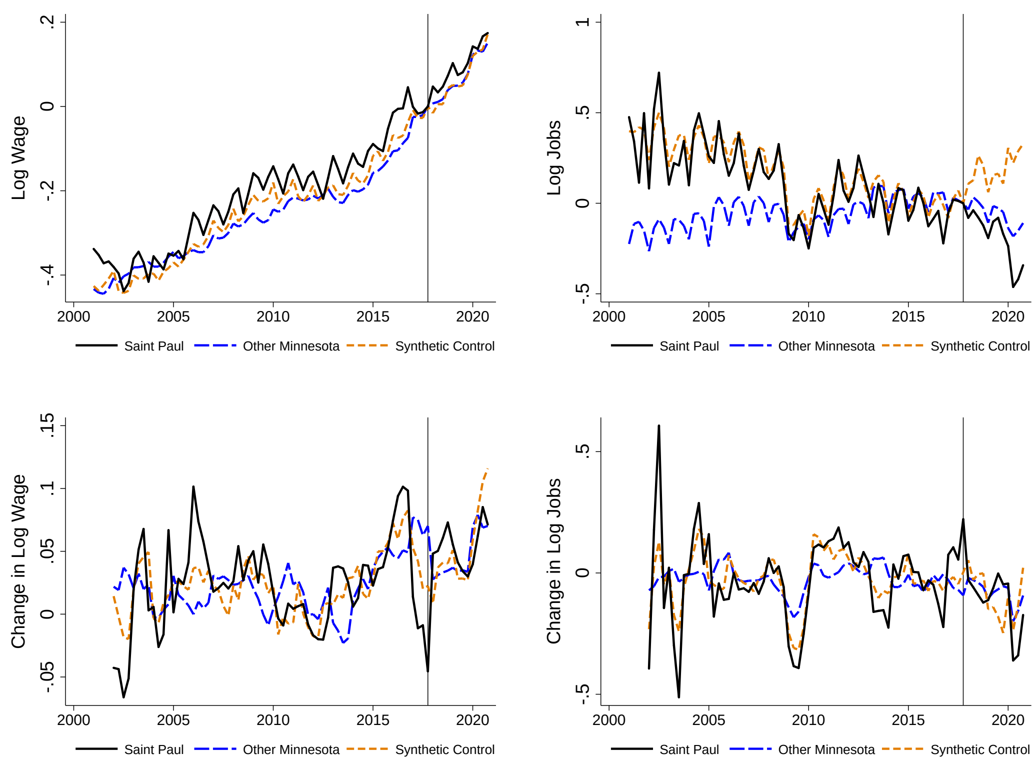

Figure A.4: Time Series of Administration and Support in Saint Paul

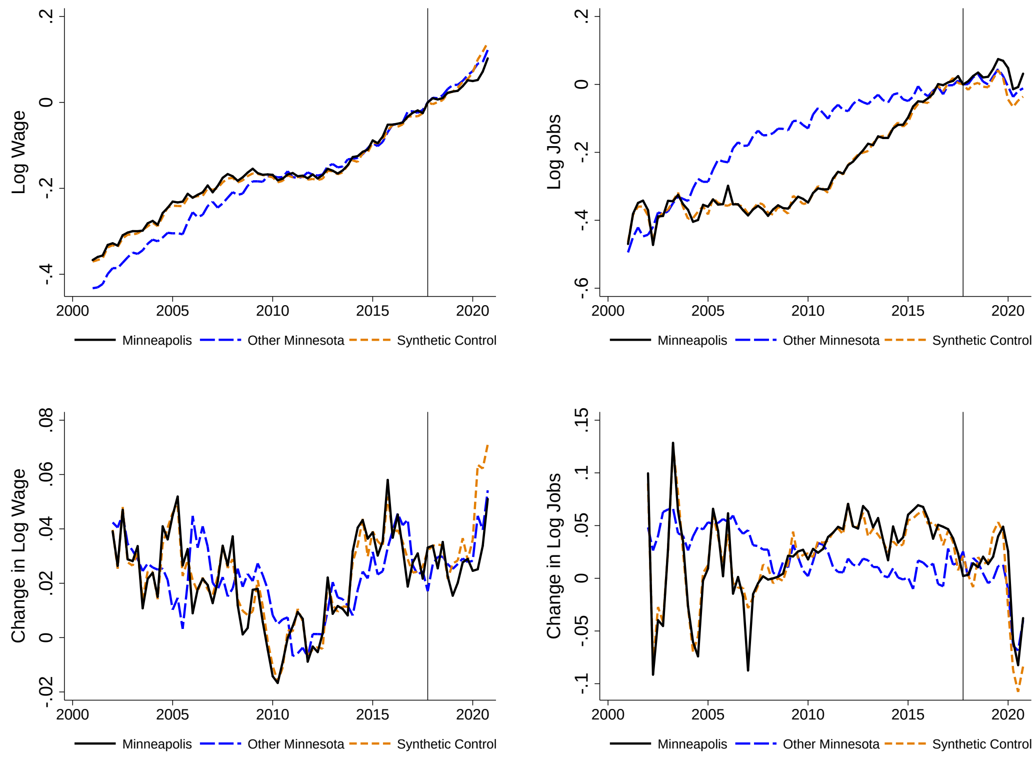

Figure A.5: Time Series of Health Care and Social Assistance in Minneapolis

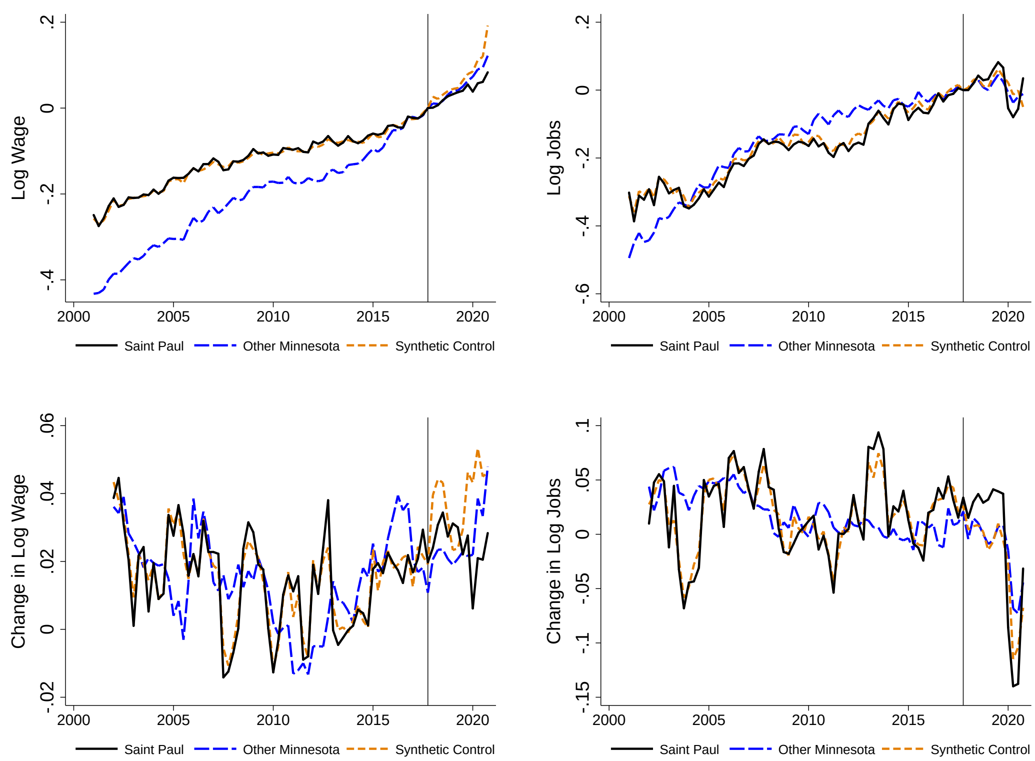

Figure A.6: Time Series of Health Care and Social Assistance in Saint Paul

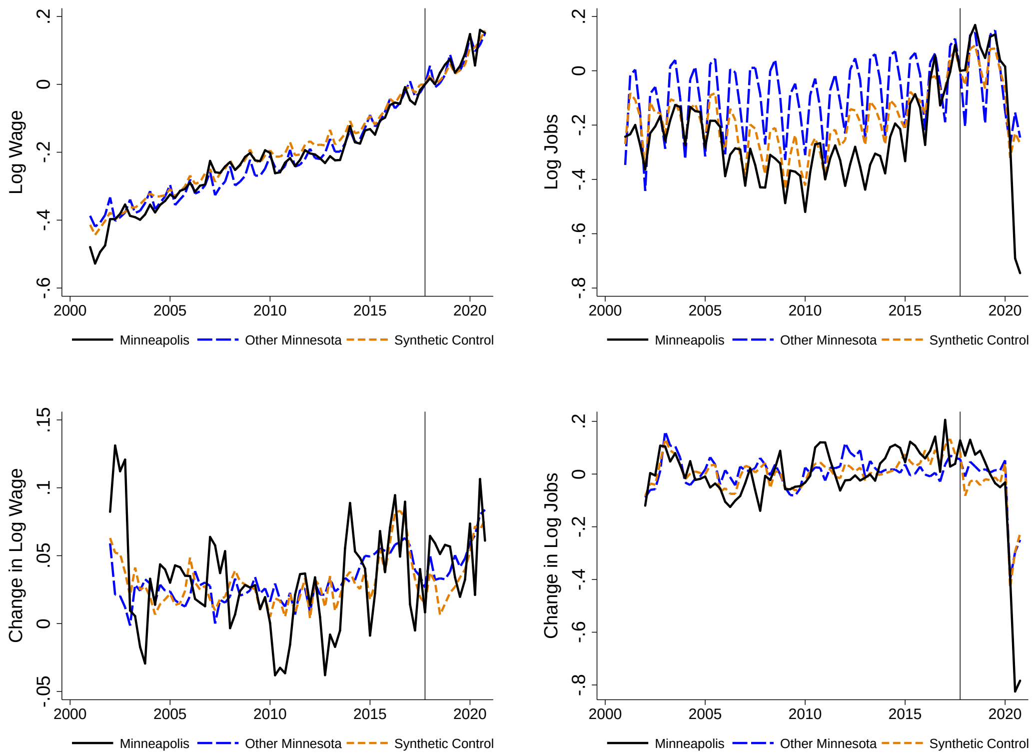

Figure A.7: Time Series of Arts, Entertainment, and Recreation in Minneapolis

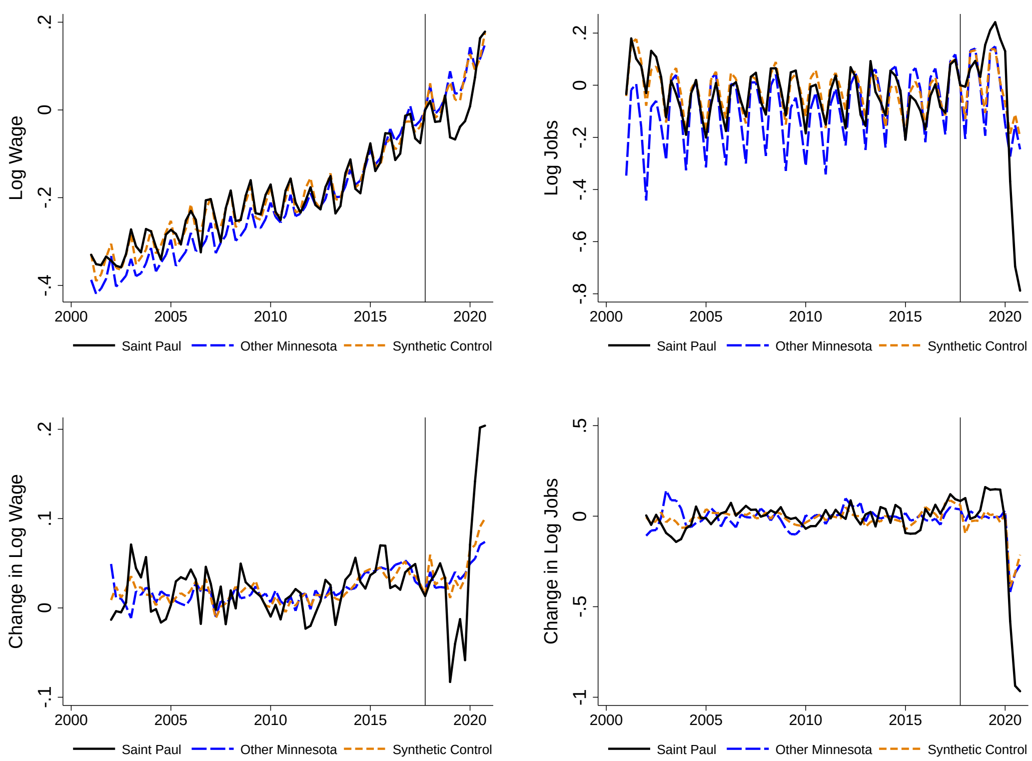

Figure A.8: Time Series of Arts, Entertainment, and Recreation in Saint Paul

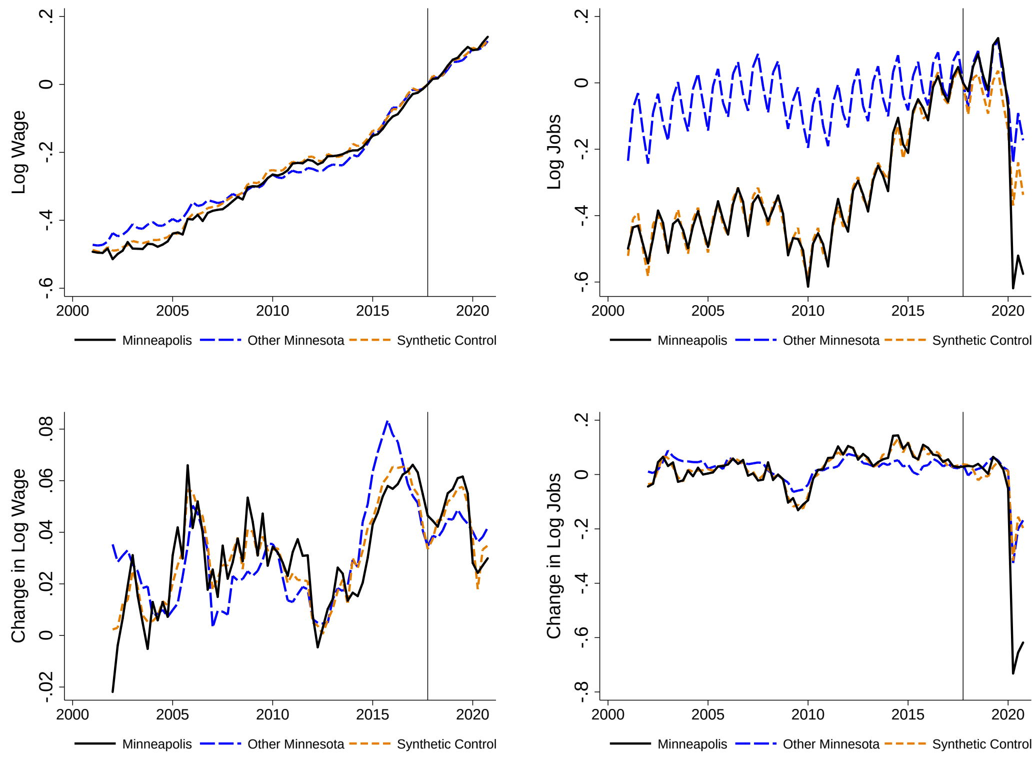

Figure A.9: Time Series of Accommodation and Food Services in Minneapolis

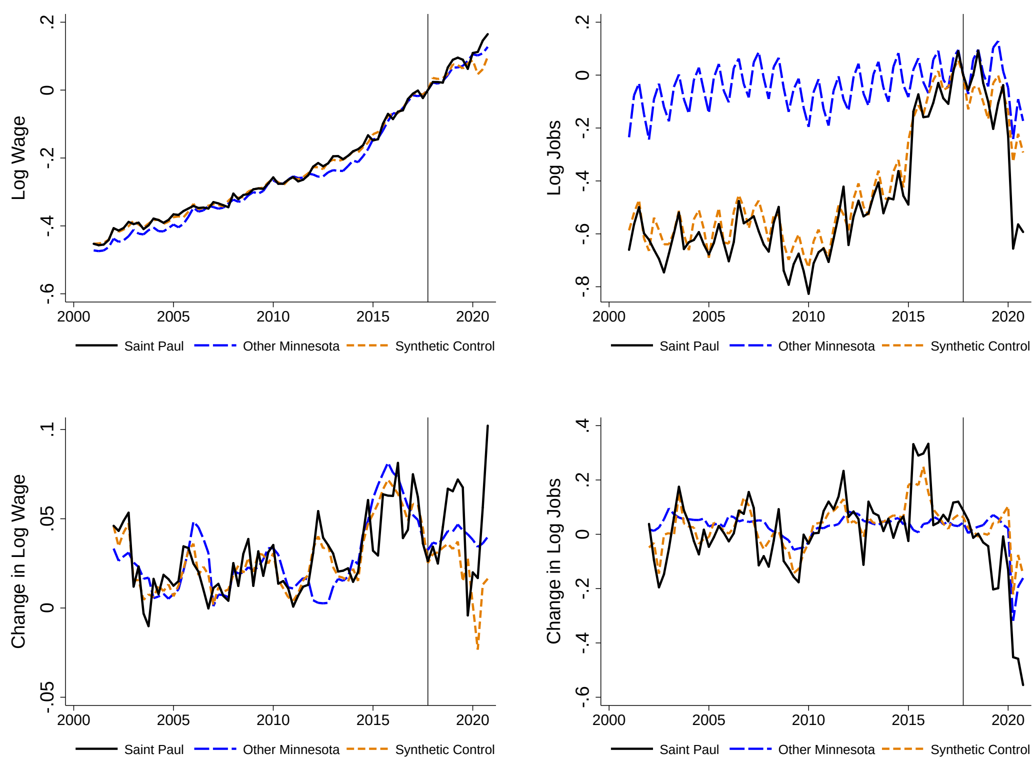

Figure A.10: Time Series of Accommodation and Food Services in Saint Paul

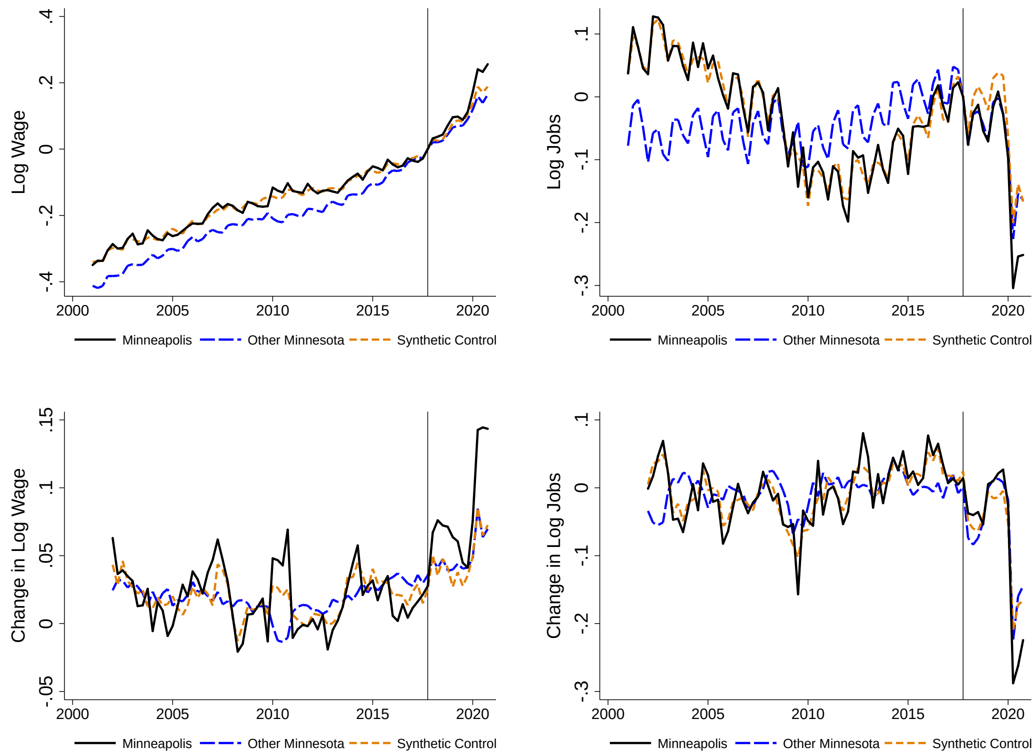

Figure A.11: Time Series of Other Services in Minneapolis

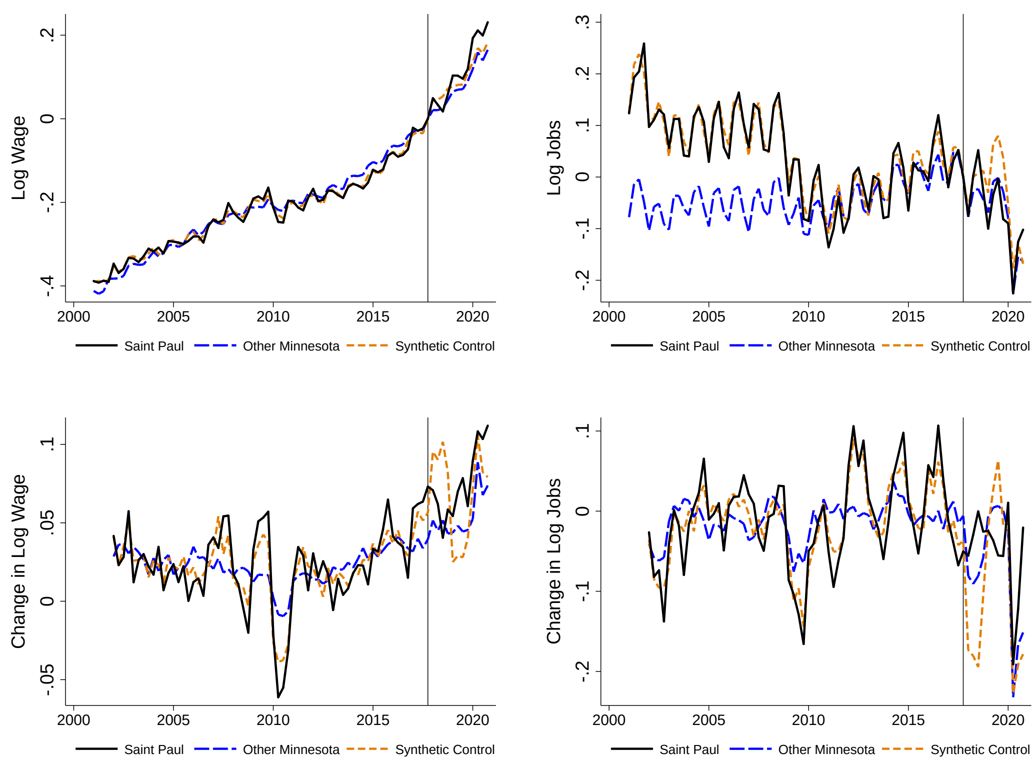

Figure A.12: Time Series of Other Services in Saint Paul

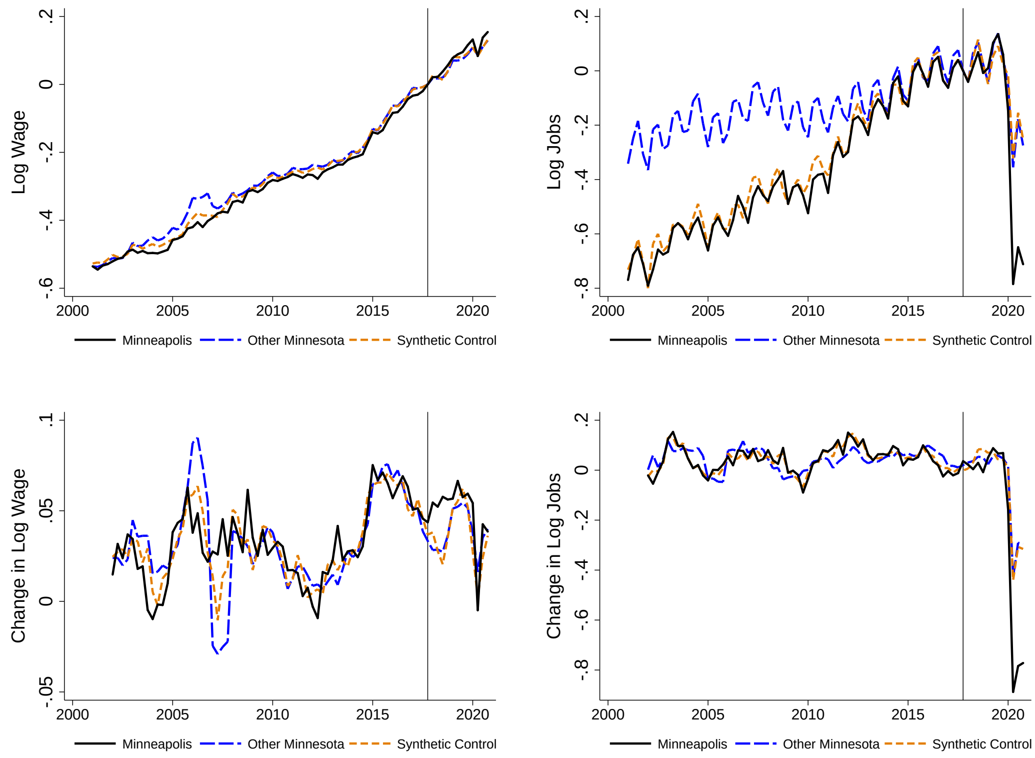

Figure A.13: Time Series of Full-Service Restaurants in Minneapolis

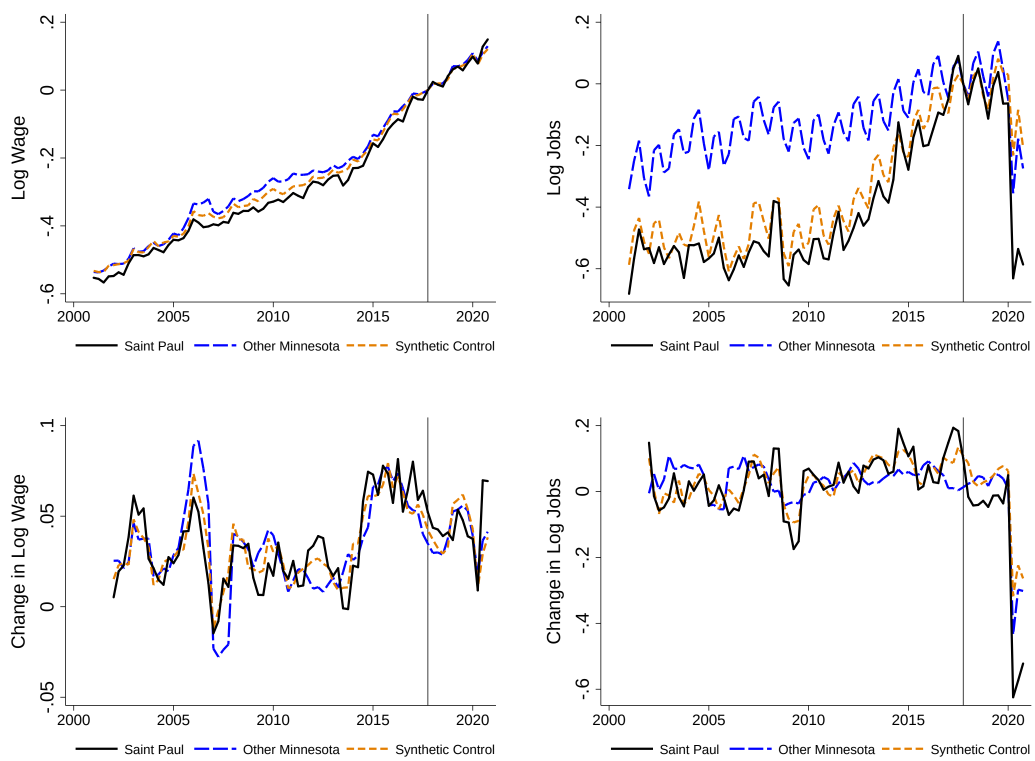

Figure A.14: Time Series of Full-Service Restaurants in Saint Paul

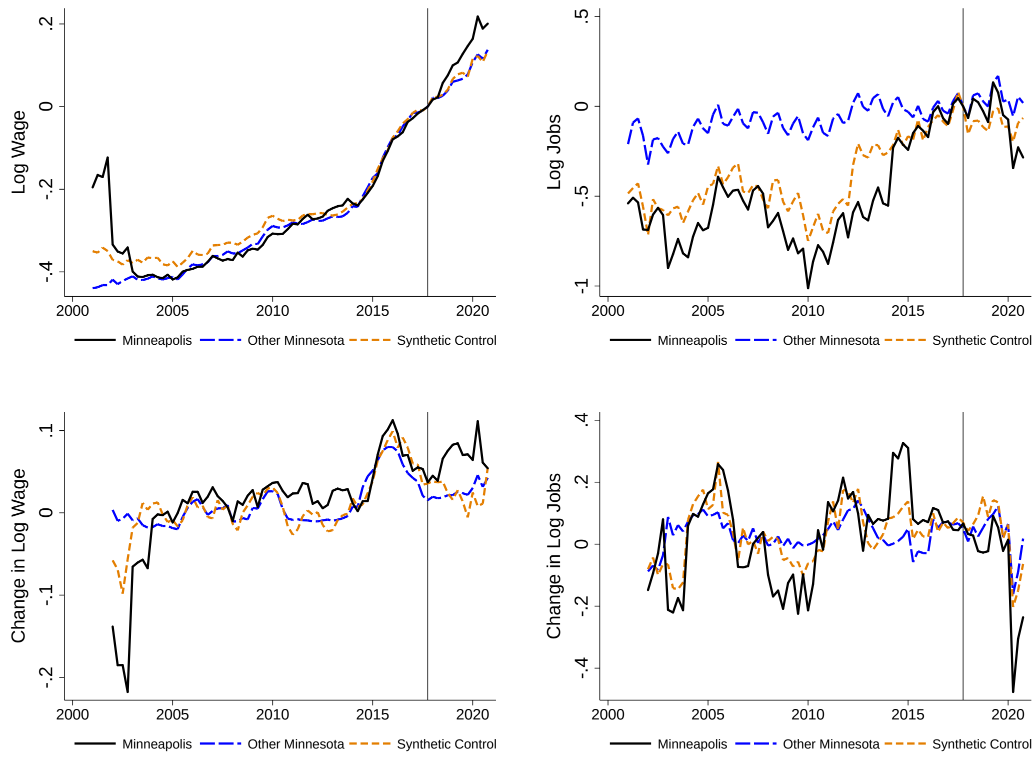

Figure A.15: Time Series of Limited-Service Restaurants in Minneapolis

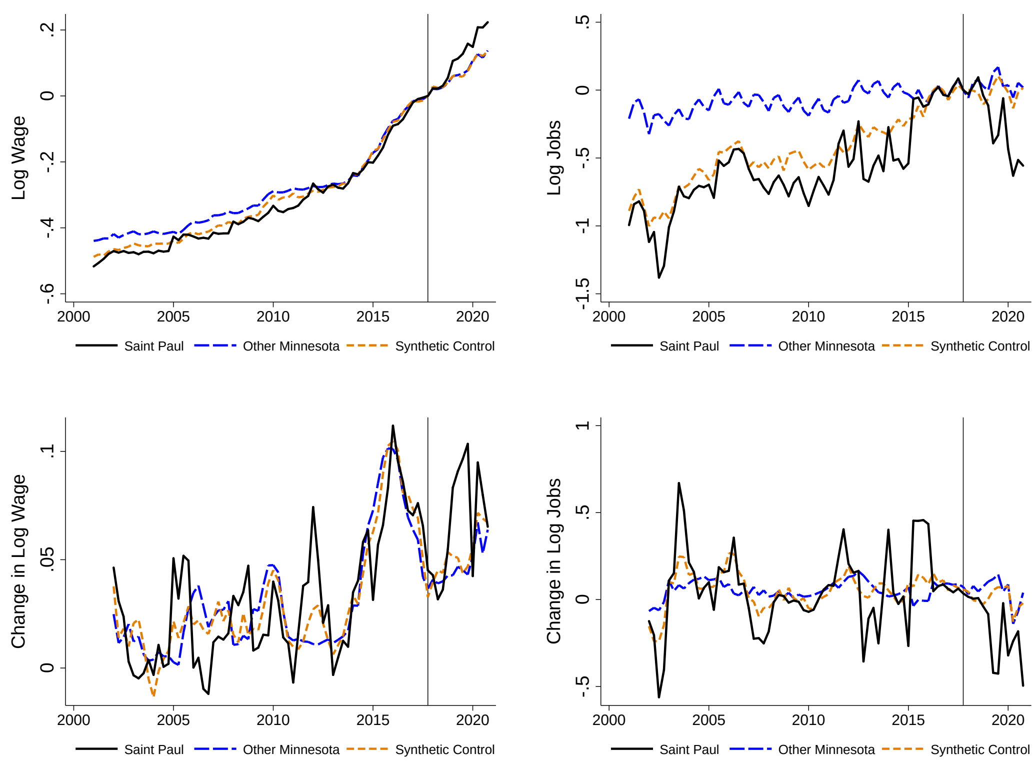

Figure A.16: Time Series of Limited-Service Restaurants in Saint Paul

Table A.2: Pre-treatment Fit: Synthetic Control versus Difference-in-Differences

	Wage		Jobs		Hours		Earnings	
(R-squared, percent)	SC	DD	SC	DD	SC	DD	SC	DD
				Minneapolis				
Retail Trade (44)	84	25	84	0	77	5	72	3
Administration and Support (56)	53	3	87	12	72	13	80	19
Health Care and Social Assistance (62)	94	27	92	7	79	15	91	7
Arts, Entertainment and Recreation (71)	30	5	46	5	45	14	21	5
Accommodation and Food Services (72)	82	41	94	45	93	36	95	58
Other Services (81)	61	0	79	2	78	2	85	12
Full-Service Restaurants (722511)	65	33	86	25	84	38	84	26
Limited-Service Restaurants (722513)	64	29	62	13	59	4	55	5
				Saint Paul				
Retail Trade (44)	70	2	65	1	60	0	67	3
Administration and Support (56)	47	0	69	5	80	12	75	3
Health Care and Social Assistance (62)	90	14	91	1	95	15	97	34
Arts, Entertainment and Recreation (71)	35	9	32	2	53	3	31	0
Accommodation and Food Services (72)	77	40	70	6	57	0	68	13
Other Services (81)	84	38	85	28	87	14	92	19
Full-Service Restaurants (722511)	79	51	73	2	65	3	67	5
Limited-Service Restaurants (722513)	66	47	53	3	47	0	59	2

Notes: SC: synthetic control. DD: difference-in-differences.

Table A.3: Excluding Neighboring Cities from the Control Group

Minneapolis	Wage	Jobs	Hours	Earnings
Retail Trade (44)	9.0	−9.6	−6.1	0.5
	(0.0)	(13.6)	(18.0)	(83.1)
Administration and Support (56)	8.8	4.6	2.0	7.2
	(0.2)	(63.7)	(99.1)	(58.9)
Health Care and Social Assistance (62)	−4.0	3.1	0.8	−0.0
	(0.4)	(54.9)	(94.9)	(78.3)
Arts, Entertainment and Recreation (71)	3.6	−15.1	5.2	27.8
	(69.9)	(13.2)	(53.3)	(4.4)
Accommodation and Food Services (72)	1.3	−31.4	−32.5	−40.0
	(43.0)	(0.0)	(0.0)	(0.0)
Other Services (81)	10.7	0.7	−5.6	10.0
	(0.0)	(70.3)	(35.6)	(8.2)
Full-Service Restaurants (722511)	4.0	−41.4	−43.4	−42.7
	(0.0)	(0.0)	(0.0)	(0.0)
Limited-Service Restaurants (722513)	13.4	−28.9	−28.9	−27.6
	(0.0)	(1.0)	(3.0)	(3.8)
Saint Paul	**Wage**	**Jobs**	**Hours**	**Earnings**
Retail Trade (44)	4.5	−2.0	−34.7	−11.4
	(0.0)	(78.5)	(0.0)	(9.0)
Administration and Support (56)	0.7	−8.4	−9.5	−73.3
	(99.3)	(57.9)	(36.2)	(0.0)
Health Care and Social Assistance (62)	−3.8	4.0	2.1	−2.1
	(1.0)	(43.8)	(82.3)	(55.1)
Arts, Entertainment and Recreation (71)	−0.1	−18.3	−6.4	−16.6
	(61.7)	(3.8)	(46.2)	(5.0)
Accommodation and Food Services (72)	8.0	−43.4	−64.8	−34.2
	(0.0)	(0.0)	(0.0)	(0.0)
Other Services (81)	1.9	14.6	3.1	11.1
	(29.8)	(0.4)	(51.5)	(3.8)
Full-Service Restaurants (722511)	1.1	−37.3	−39.5	−45.6
	(32.4)	(0.0)	(0.0)	(0.0)
Limited-Service Restaurants (722513)	3.6	−57.2	−75.2	−85.8
	(0.0)	(0.0)	(0.0)	(0.0)

Notes: The estimates are in log points, multiplied by 100. Entries in parentheses are p-values using the placebo method.

Table A.4: Cities of Similar Size to Minneapolis and Saint Paul

City	Jobs (000's)	City	Jobs (000's)
Washington, DC	533	Baltimore, MD	276
Indianapolis, IN	527	Albuquerque, NM	264
Jacksonville, FL	461	Greensboro, NC	251
Denver, CO	444	El Paso, TX	236
Nashville, TN	440	Prince George's County, MD	232
Memphis, TN	438	Colorado Springs, CO	225
Milwaukee, WI	434	Baton Rouge, LA	222
Portland, OR	433	Wichita, KS	220
Louisville, KY	425	Little Rock, AR	201
Montgomery County, MD	380	St. Louis, MO	197
Honolulu, HI	380	Reno, NV	193
Oklahoma City, OK	374	New Orleans, LA	170
Tulsa, OK	322	Fort Wayne, IN	169
Kansas City, MO	314	Winston-Salem, NC	167
Fresno, CA	310	Lexington, KY	159
Omaha, NE	301	Huntsville, AL	155
Tucson, AZ	299	Virginia Beach, VA	149
Aurora, CO	295	Springfield, MO	147
Minneapolis, MN	**280**		
Aurora, CO	295	Corpus Christi, TX	135
Baltimore, MD	276	Salem, OR	132
Albuquerque, NM	264	Anchorage, AK	120
Greensboro, NC	251	Sioux Falls, SD	115
El Paso, TX	236	Rockford, IL	114
Prince George's County, MD	232	Richmond, VA	114
Colorado Springs, CO	225	Lubbock, TX	111
Baton Rouge, LA	222	Norfolk, VA	104
Wichita, KS	220	Tallahassee, FL	102
Little Rock, AR	201	Montgomery, AL	97
St. Louis, MO	197	Shreveport, LA	95
Reno, NV	193	Amarillo, TX	90
New Orleans, LA	170	Jackson, MS	86
Fort Wayne, IN	169	Midland, TX	85
Winston-Salem, NC	167	Chesapeake, VA	85
Lexington, KY	159	Newport News, VA	83
Huntsville, AL	155	Fayetteville, NC	83
Saint Paul, MN	**149**	Augusta, GA	81
Virginia Beach, VA	149	Laredo, TX	79
Springfield, MO	147	Kansas City, KS	78
Lincoln, NE	137	Birmingham, AL	77
Savannah, GA	136		

Table A.5: Assessing the Size and Sources of the Bias

Minneapolis	Jobs Effect	Bias μ	Bias u	Bias Total
Retail Trade (44)	−3	−2	0	−2
Administration and Support (56)	2	0	0	0
Health Care and Social Assistance (62)	−2	0	0	0
Arts, Entertainment and Recreation (71)	−12	7	−3	4
Accommodation and Food Services (72)	−25	3	−6	−3
Other Services (81)	−10	1	−1	0
Full-Service Restaurants (722511)	−39	0	−6	−6
Limited-Service Restaurants (722513)	−19	22	−7	15
Saint Paul	**Jobs Effect**	**Bias μ**	**Bias u**	**Bias Total**
Retail Trade (44)	−11	0	−2	−2
Administration and Support (56)	−10	−1	−1	−2
Health Care and Social Assistance (62)	−4	−1	−1	−2
Arts, Entertainment and Recreation (71)	−22	2	−4	−2
Accommodation and Food Services(72)	−22	1	−4	−3
Other Services (81)	−2	−4	0	−4
Full-Service Restaurants (722511)	−22	1	−5	−4
Limited-Service Restaurants (722513)	−12	−7	−1	−8

Notes: Estimates are in log points, multiplied by 100. The first column repeats our jobs estimates using synthetic difference-in-differences from the QCEW dataset. The data generating process is a factor model $Y_{it} = \alpha_i + \beta_t + \sum_{k=1}^{4} \mu_i^k \gamma_t^k + u_{it} + \sum_{s=T_{\text{pre}}+1}^{T} \tau_s W_{is}$. After generating data from this model, we perform our synthetic difference-in-differences model to assess the bias of the estimates relative to those in the factor model. The bias μ column assumes $u_{Nt} = \sum_{i=1}^{N_{\text{co}}} \omega_i u_{it}, \forall t = 1, ..., T$ and gives the bias due to not fitting the underlying factor structure. The bias u column assumes $\sum_{i=1}^{N_{\text{co}}} \omega_i \mu_i^k = \mu_N^k$ and gives the bias due to the scale of errors relative to the length of the pre-treatment period. The total bias column is the sum of the two columns.