ESSAYS IN APPLIED MICROECONOMICS

by

Robert Tucker Omberg
A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
In Partial fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Economics

Committee:

______________________________ Director

____

____

______________________________ Department Chairperson

______________________________ Dean, College of Humanities and Social Sciences

Date: _________________________ Spring Semester 2022
George Mason University
Fairfax, VA

Essays in Applied Microeconomics

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at George Mason University

By

Robert Tucker Omberg
Master of Arts
George Mason University, 2020
Bachelor of Science
George Mason University, 2017

Director: Alex Tabarrok, Professor
Department of Economics

Spring Semester 2022
George Mason University
Fairfax, VA

Copyright © 2022 by Robert Tucker Omberg
All Rights Reserved

Dedication

I dedicate this dissertation to my wife Bridget and my daughters Julia and Allison. I couldn't have done this without you, and I love all three of you very much.

Acknowledgments

I would like to thank my advisor, Dr. Alex Tabarrok, and my committee members, Dr. Thomas Stratmann and Dr. Sita Slavov, for their advice and guidance. I am also incredibly thankful for the support of the Mercatus Center, The Institute for Humane Studies, and the Bradley Foundation.

Table of Contents

List of Tables

List of Figures

Abstract

ESSAYS IN APPLIED MICROECONOMICS

Robert Tucker Omberg, PhD

George Mason University, 2022

Dissertation Director: Alex Tabarrok

This dissertation consists of three evaluations of "natural experiments" using modern methods of causal inference. The first chapter, The Labor Market Impacts of Ridesharing, examines the impact of the emergence of the âgig economyâ on the broader labor market by exploiting the staggered introduction of the ridesharing service Uber to American Cities between 2013 and 2018. Using difference-in-differences methods, Chaisemartin and DâHaultoeuilleâs time-corrected Wald estimator, Callaway and SantâAnnaâs doubly robust difference-in-differences estimator, and Abadie et alâs synthetic control method, I estimate that Uberâs arrival to a city resulted in decline in the unemployment rate by between a fifth and a half of a percentage point. This suggests that Uber allowed many workers to supplement their earnings during periods of unemployment, framing the ridesharing service as a complement to, rather than a substitute for, traditional employment. I also find some evidence that Uber had a very small positive effect on wages at the lower end of the wage distribution, suggesting that Uber may have altered worker search behavior or affected bargaining power. The second chapter, which has been published in the *IZA Journal of Labor Policy*, explores the impact of the minimum wage in Puerto Rico. Revisiting research from the 1990s from Castillo-Freeman and Krueger, I use the synthetic control method of Abadie et al. to estimate the impact of the most recent increase in the federal minimum wage

on employment in Puerto Rico. I estimate that the employment/population ratio of various groups in Puerto Rico was significantly lower than that of a data-constructed synthetic Puerto Rico which did not raise its minimum wage. Placebo tests on other donor units, time periods, and population groups suggest that a significant portion of this gap is a result of the minimum wage. Groups with greater exposure to the minimum wage, such as teens and restaurant workers, experienced proportionally greater declines in employment. My results suggest an own-wage elasticity of employment in Puerto Rico of â0.68, higher than estimates from the mainland, which suggests that the employment response to minimum wages may be more dramatic at higher relative minimum wages. The final chapter explores the impact that removing monuments to the confederacy has had on race based hate crime. Beginning in the late 2010s, many municipalities have begun removing Confederate monuments, memorials, and/or flags in response to both public outcry and recent white supremacist acts of terrorism. While the stated goal of these removals is to move past a troubled episode in our nation's history, what actual effect does monument removal have on acts of racially-biased violence? Using data from the FBI and the Southern Poverty Law Center and a difference-in-differences research design, I find that the number of hate crimes based on race occurring in a city increases by between 20 and 40% following the city's removal of a confederate monument, even after controlling for a variety of covariates. Decomposition of these results shows that this increase only occurs when monuments are removed in states that were formally part of the Confederacy; although there is still no evidence that removal decreases hate crime even in non-Confederate areas.

Chapter 1: The Labor Market Impacts of Ridesharing

1.1 Introduction and Background

1.1.1 Introduction

How will the "gig economy" change the face of work in the 2020s and beyond? Since the early 2010s, platforms like Uber, Amazon Mechanical Turk, and TaskRabbit have allowed workers to engage in temporary, flexible jobs at their own discretion as opposed to working a pre-determined number of hours for a single employer. The digitization of communication now allows platforms to hire workers on a task-by-task basis, leaving workers to work as many or as few tasks as they please, or to take "gigs" from multiple different platforms. Some (Ravenelle (2019)[1], Munger (2019)[2]) see the the growth in these platforms as foretelling a future where traditional work is replaced entirely by gig work. Others (Oyer (2020)[3]) believe that the rise in gig work has not supplanted traditional employment arrangements thus far, and that many individuals who engage in gig work often do so to supplement earnings from another more traditional job. In other words, there is disagreement as to whether gig economy work arrangements are a substitute for or a complement to traditional jobs.

Up until now, this debate has mostly revolved around quantifying the rise in independent and contract workers in recent decades. This paper takes an alternative approach, by investigating how the arrival of gig economy firms to a metropolitan area affects that area's labor market. To do this, I exploit a natural experiment created by the behavior of one of the most iconic gig economy firms: the ridesharing service Uber. In April 2013, residents of only 26 U.S. cities had the opportunity to become driver-partners using their own vehicles, but over the next five years Uber would gradually expand until it was present in 310 of the 384 Metropolitan Statistical Areas (MSAs) in the U.S. in May 2018. This staggered adoption design allows the use of difference-in-differences econometric methods to identify the causal effect of the gig economy on a variety of economic outcomes.

Using two-way fixed effects regressions, Chaisemartin and D'Haultoeuille's (2019)[4] Wald-TC estimator , Callaway and Sant'Anna's (2021)[5]doubly-robust difference-in-difference estimator, and synthetic controls (Abadie et al. (2010)[6]), I consistently find that Uber's presence lowers the unemployment rate in a city by between a fifth and a half of a percentage point. I also use the city of Austin, TX, as a case study. Austin had a unique experience with Uber: the company operated in Austin between June 2014 and May 2016, before a local ordinance requiring driver-partners to submit to city background checks and fingerprinting caused Uber to leave in reaction. One year later, in June 2017, a state law superseded Austin's regulation, whereupon Uber returned to the city. The fact that Uber left from and returned to Austin for political reasons unrelated to local labor market conditions creates a cleaner natural experiment with fewer concerns regarding endogeniety. Uber's effect in Austin is consistent with those found when examining the nation as a whole: the unemployment rate was 0.38 percentage points above trend during the period of Uber's absence. These results suggest that Uber plays a complementary role to traditional employment by allowing frictionally unemployed workers to supplement their earnings while searching for another job. I also use a similar empirical strategy to examine Uber's effect on the average wages in a city.

The rest of the paper is laid out as follows. The remainder of Part I summarizes research on the gig economy to date and gives background on Uber's development. Part II describes my empirical strategy and summarizes the data used. Part III presents my estimates of Uber's effect on both unemployment and wages. Part IV discusses my results, their limitations, and concludes.

1.1.2 Prior Research

A January 2020 cartoon in the New York Times, attached to an article by E. Tammy Kim entitled "The Gig Economy is Coming for Your Job", depicts a smartphone full of apps with names like "SRGN", "Cheff", "PiLOT", and "Vetify". This cartoon embodies a popular view of how the burgeoning "gig economy", typified by firms like AirBnB, TaskRabbit, and Uber, will shape the future of work. Rather than report to work at a logging firm, a lumberjack might instead register as an independent contractor on an app called "Timbr", which would connect him directly with clients who need a tree removed, eliminating the firm entirely. Some view this as a positive development for the Lumberjack: he likely enjoys more flexibility in setting his schedule and hours as he reaps the benefits of "being his own boss". Others would see "Timbr" as a net negative for the Lumberjack profession since, for example, as independent contractors they may not be subject to the same legal protections as traditional employees. What these groups agree on is the fundamental role that the sharing or "gig" economy will play in the future of work: that of a replacement or *substitute* for traditional employment. An alternative, and sometimes neglected, perspective is that the gig economy functions, or will function, as a *complement* for firm-based work. The flexibility of gig work might allow people to supplement income earned from more traditional work arrangements during periods of low earnings, or provide an alternative source of income during a period of frictional unemployment.

In her piece, Kim paints a gloomy picture of how Uber and other "gig economy" firms may influence the future of work, noting that current trends "...represent a subtle, sneaky form of technological displacement, care of the gig economy. They're not robots stepping in for humans on a factory floor, but rather smartphone-based independent contractors and supplemental "cobots" (a portmanteau of "co-worker" and "robot") chipping away at the careers of full-time and in some cases unionized employees".

UNC Chapel Hill Sociologist Alexandrea Ravenelle provides a similarly bleak view in the conclusion of her 2019 monograph <u>Hustle and Gig: Struggling and Surviving in the Sharing Economy</u>:

> "...the disruption offered by the sharing economy isn't about moving forward. Instead of offering a way out, the sharing economy has simply increased economic insecurity and worker vulnerability. Workers go from gig to gig, ostensibly as their own bosses, but subject to the whims of platform pivots and deactivations....As sharing economy services have grown and proliferated, they've successfully subverted generations of financial gains and workplace protections...Hard-won victories for workers' rights and protections are being hacked and disrupted in the name of a "cheaper, poorer quality" progress that is eviscerating a hundred years of workers' rights. The disruption offered by the sharing economy is simply a hustle" (p. 207)[1].

However, Paul Oyer gives an alternative vision of the role that the sharing economy has played in the evolution of work:

> "Non-traditional work relationships (independent work') has grown steadily enveloped economies as it has become easier to break work into discrete blocks. The app-based 'gig economy' has increased that growth but has not brought (and will not soon bring) fundamental change in most people's work and in the economic centrality of the employment model" (Oyer, 2020)[3].

This encapsulates the view that "gig work" primarily functions as a complement for traditional employment. Supporting this view, Oyer notes that the growth of the prevalence of independent/contract work changes depending on how it is measured. When people are asked about "*any* independent work (including occasional work and work in addition to a traditional job", the numbers "suggest that the fraction of US workers who do gig or independent work in a given year is a quarter to a third of all workers" (p. 7). However, when people were instead asked only about their "main job", researchers concluded:

> "that non-traditional work had not grown substantially from 2005 to 2017 and had, in fact, declined slightly. This was partially a matter of labor market conditions (the extremely tight labor market in 2017 meant very few people

> worked independently if they did not want to) and partially differences in samples across the two surveys. But the bottom line is that, when looking strictly at 'main job', the rise of non-traditional work has not been nearly as dramatic as media coverage of the gig economy might suggest" (p. 7)

In their 2016 profile of Uber's Driver-Partners, Jonathan Hall and Alan Krueger come to the conclusion that "...most driver-partners do not appear to turn to Uber out of desperation or because they face an absence of other opportunities in the job market...but rather because the nature of the work, the flexibility, and the compensation appeals to them compared with other available option". Later in the paper, Hall and Krueger note that "another aspect of Uber's flexibility is that spending time on the platform can help smooth the transition to another job, as driver-partners can take off time to prepare for and search for another job at their discretion".

The idea of using Uber's staggered introduction as a natural experiment did not originate with this paper. Hall, Palsson, & Price (2018)[7] use a difference-in-difference design to identify the relationship between Uber entry and public transit ridership, finding that public transit ridership increases by roughly five percent after Uber's entry. More recently, Barrios, Hochberg, and Yi (2022)[8] exploit Uber's staggered entry to examine the effect of the gig economy on new business formation, concluding that new business registrations increase by roughly 5% following the entry of Uber and Lyft. Beyond using new response variables, this paper is the first to use identifcation strategies beyond two-way fixed effects, which recent scholarship has suggested may be a biased estimator (Callaway & Sant'Anna (2021)[5]).

1.1.3 Uber's Development

Uber first began operation in San Francisco in May 2010 as "Ubercab", which allowed users to hire a luxury black vehicle from their mobile devices, costing roughly 1.5 times as much as a cab. After rebranding to "Uber" in October of that same year, the service began expanding outside of San Francisco, launching in New York City, Chicago, and Paris by the end of 2011. In July 2012, the company launched the "UberX" service, which used lower cost hybrid vehicles instead of the more expensive luxury vehicles, which were termed "UberBlack". While UberX offered fairs that were 35% cheaper than the previous black car rates, commentators at the time of the announcement still thought of Uber as something closer to "a taxi dispatch service like GetTaxi or MyTaxi" (Tsotsis, 2012)[9].

Discussion of Uber's business model and the broader future of work began in April 2013, when Uber published a policy white paper on its blog where it noted that it had thusfar "work[ed] almost exclusively with commercially licensed, insured, and regulated entities". However, the paper goes on to announce Uber's intentions to begin "offering transportation services without traditional commercial insurance or licensing- known as ridesharing". The introduction of ridesharing allowed drivers to use their personal vehicles as part of the UberX service, which not only lowered prices for riders, but also dramatically expanded the number of potential Uber employees. Prior to the introduction of ridesharing, becoming an Uber driver required commercial licensing similar to what would be required of a taxi driver, but suddenly anyone with a vehicle, and clean driving record, and willingness to submit to a background check was a potential employee. Figure 1.1, taken from Hall and Krueger (2016)[10], demonstrates how critical the introduction of ridesharing was to the exponential growth of the number of driver-partners working with Uber.

While the number of luxury UberBlack drivers grew only modestly between Uber's founding and 2016, the number of UberX drivers only began to explode with the introduction of ridesharing in mid-2013. This increase in the number of active driver-partners coincided with a rapid expansion of Uber into new cities across the United States. When it introduced ridesharing in April 2013, Uber was active in 16 U.S. cities. One year later, in April 2014, that number had grown to 49. After another year, riders could order an Uber in 169 cities. By May 2018, Uber confirmed that is was operating in 310 of the 384 metropolitan statistical areas in the United States. The staggered introduction of Uber across U.S. cities constitutes a useful natural experiment: by comparing labor market outcomes in cities where Uber arrived early to those where Uber arrived later, we can more confidently estimate the causal effect of Uber's presence on our variable of interest: the unemployment rate.

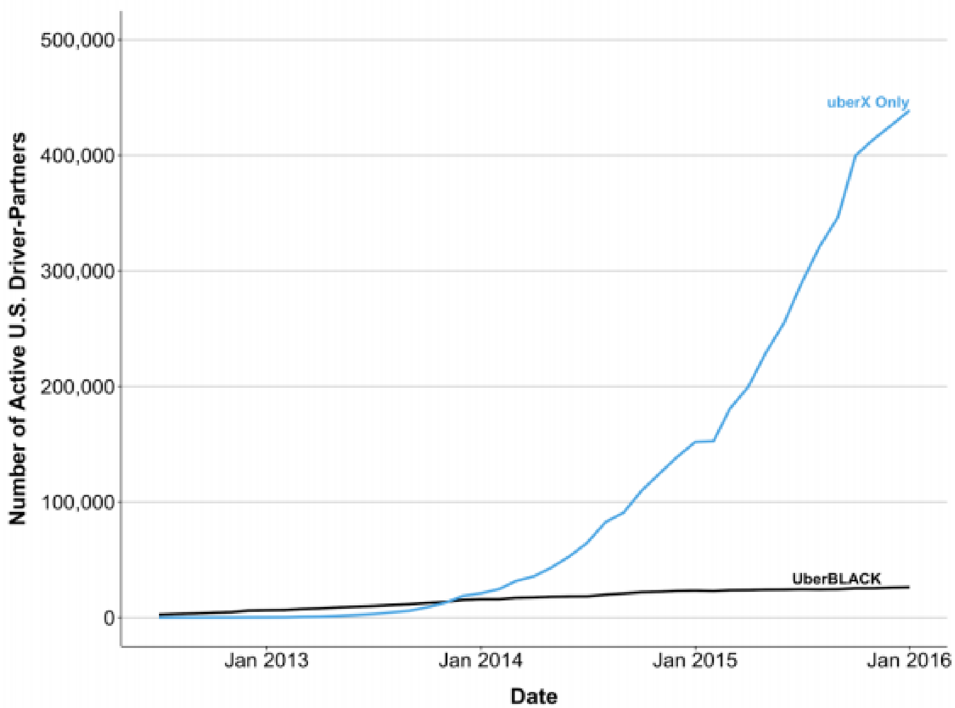

Note: Sample consists of all U.S. UberBLACK and uberX driver-partners making at least four trips in any month (1,149,024 individuals).

Figure 1.1: Number of Driver-Partners through the Uber X and Uber Black Services (From Hall & Krueger (2016), Figure 4)

1.1.4 Why the Unemployment Rate?

How can the unemployment rate be informative as to Uber's impact on the structure of the labor market? If Uber functions as a complement to traditional employment with a firm, as posited by Oyer and Krueger, then we would expect to find a negative effect of ridesharing's introduction on the unemployment rate in a given metropolitan area. Any city will have some fraction of its workforce "between jobs", or fictionally unemployed, at any given time. Prior to the introduction of ridesharing, these individuals see their incomes fluctuate while they search for another job. Once Uber begins operating, however, some of these individuals choose to drive with Uber in between scrolling through job boards and sending out resumes, allowing them to cushion the transitory dip in their income. What's more, these people are no longer considered unemployed by official statistics, leading to a drop in the unemployment rate. An alternative mechanism by which Uber may complement traditional employment and decrease the unemployment rate is by solving problems of spatial mismatch between homes and jobs. Existing research, such as Brandtner et al. (2017)[11] and Hernandez et al. (2020)[12], has already indentified a relationship between increased access to public transportation services and unemployment. The introduction of a ridesharing service like Uber could have an identical effect: expanding the pool of job opportunities available to individuals with limited transportation options.

Conversely, if Uber primarily functions as a substitute for traditional employment, then we shouldn't expect to observe a relationship between ridesharing's introduction and the level of unemployment, since driver-partners are drawn from the ranks of the previous employed. Indeed, if Uber's introduction causes jobs to be destroyed, then it's arrival may be associated with a temporary *increase* in the unemployment rate. For example, if a local taxicab company cuts staff or goes out of business due to pressure from Uber, then the former Taxi drivers will become unemployed for a short time before finding other work or joining the ranks of Uber's driver-partners.

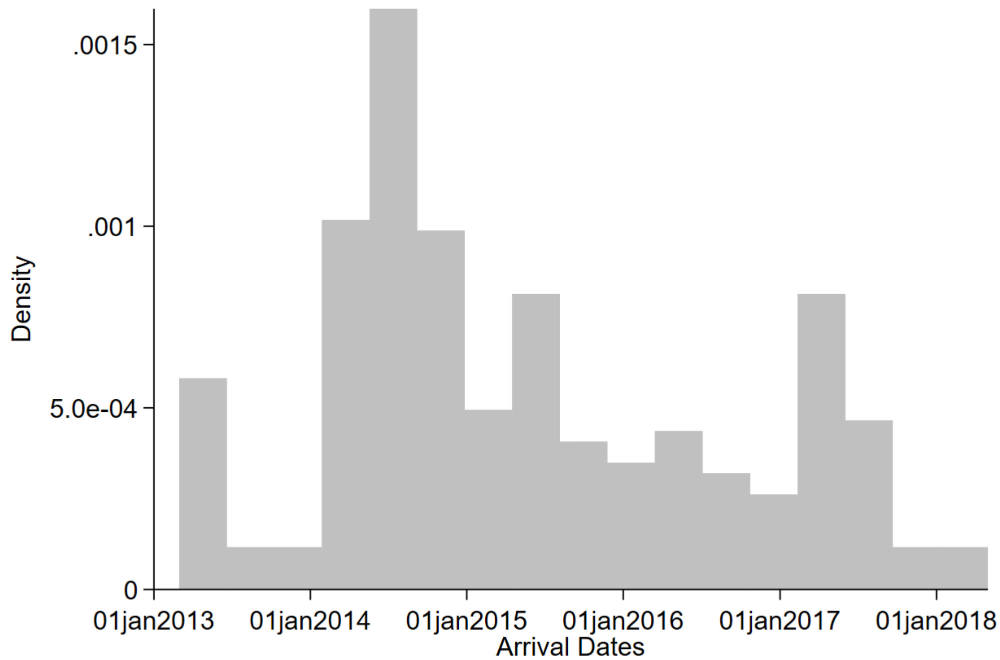

Figure 1.2: Histogram of Uber's Arrival to MSAs

1.2 Data

1.2.1 Data

Data on the labor force, employment, unemployment, and unemployment rate were obtained for 389 metropolitan areas ranging from January 2012 until July 2018 from the Bureau of Labor Statistics. For each area, the date at which Uber began operation was found using press releases from Uber's own website, or, if none were available, from local news articles announcing Uber's arrival, and an indicator was coded indicated whether Uber was operating in an area during a given month. Since Uber only began offering the UberX service, which allowed driver-partners to use their own vehicles, on April 1st 2013, I chose this as the start date for the few cities that Uber launched in between 2010 and 2013. If no conclusive information could be found about whether or not Uber operated in an area, it was dropped from the regression, leaving the total number of areas at 323. The histogram in Figure 1.2 plots the distribution of entry dates for MSAs, while Table 1.1 reports summary statistics.

Table 1.1: Summary Statistics

Variable		Mean	Std. Dev.	Min	Max	Observations
unemp	overall	5.7	2.2	1.4	28.3	N=25517
	between		1.6	2.6	21.1	n=323
	within		1.6	-2.0	16.5	T=79
uber	overall	.46	.49	0	1	N=25517
	between		.21	0	.82	n=323
	within		.45	-.35	1.4	T=79
laborforce	overall	387662	747412	24601	7275687	N=25517
	between		748302	25353	7083602	n=323
	within		19509	192474	625924	T=79

1.3 Results

1.3.1 Ridesharing's Impact on Employment

Two-Way Fixed Effects Estimation

To estimate the effect of ridesharing's introduction on the unemployment rate, I use a two-way fixed effects regression of the following functional form:

$$unemp_{it} = \beta_0 + \beta_1 U_{it} + \beta_2 labforce_{it} + \alpha_i + \lambda_t + \varepsilon_{it} \tag{1.1}$$

where $unemp_{it}$ is the unemployment rate in city i in time period t, U_{it} is a binary variable equal to 1 if Uber is currently operating in city i in time period t, $labforce_{it}$ is the size of city i's labor force in time period t, α_i is a vector of city fixed effects, λ_t is a vector of time-period fixed effects, and ε_{it} is the error term. Additionally, since the conventional standard errors tend to understate the standard deviation of difference-in-differences estimators as discussed by Bertrand, Duflo, & Mullainathan (2004)[13], I choose to utilize cluster robust standard errors. The results of this specification are found in Table 1.2.

I also estimate a specification of the above regression with the natural logarithm of the

Table 1.2: Uber's Effect on the Unemployment Rate

	(1)	(2)
Uber Oberating?	-0.206**	-0.211**
	(-3.06)	(-3.14)
Labor Force Size		-0.0000020i2
		(-1.58)
City Fixed Effects	Yes	Yes
Time Fixed Effects	Yes	Yes
N	25517	25517

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.3: Uber's Effect on the Natural Log of the Unemployment Rate

	(1)	(2)
Uber Oberating?	-0.0324**	-0.0343**
	(-3.06)	(-3.28)
Labor Force Size		-0.000000695**
		(-3.22)
City Fixed Effects	Yes	Yes
Time Fixed Effects	Yes	Yes
N	25517	25517

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

unemployment rate as the dependent variable:

$$\ln(unemp_{it}) = \beta_0 + \beta_1 U_{it} + \beta_2 labforce_{it} + \alpha_i + \lambda_t + \varepsilon_{it} \tag{1.2}$$

which yields the results found in Table 1.3.

Wald-TC Estimation and Parallel Trends Test

While linear regressions with period and group fixed effects are commonly used to estimate treatment effects in cases such as these where adoption of the treatment is staggered across groups and time, recently concerns have been raised about their validity. For example, Chaisemartin and D'Haultoeuille (2020)[14] note that two-way fixed effects regressions "estimate weighted sums of the average treatment effects (ATE) in each group and period, with weights that may be negative. Due to the negative weights, the linear regression coefficient may for instance be negative while all the ATEs are positive". The authors propose using an alternative estimator, the time-corrected Wald ratio, or Wald-TC, that is valid even if the treatment effect is heterogeneous over time or across groups:

$$W_{TC} = \frac{E(Y_{11}) - E(Y_{10} + \delta_{D_{10}})}{E(D_{11}) - E(D_{10})} \tag{1.3}$$

Where $\delta_d = E(Y_{d01}) - E(Y_{d00})$ denotes the change in the mean outcome between period 0 and 1 for control groups with treatment status d. The estimate of W_{TC}, $\hat{W}_{TC}$, can be calculated for each time period using equation (4).

$$\hat{W}_{TC} = \frac{\frac{1}{n_{11}} \sum_{i \in I_{11}} Y_i - \frac{1}{n_{10}} \sum_{i \in I_{10}} [Y_i + \hat{\delta} D_i]}{\frac{1}{n_{11}} \sum_{i \in I_{11}} D_i - \frac{1}{n_{10}} \sum_{i \in I_{10}} D_i} \tag{1.4}$$

However, like the two-way fixed effects estimator used above, the Wald-TC estimator is reliant on the assumption of common trends between the treatment and control groups. To test this assumption, we can examine whether there is a difference between the cities where Uber did and did not eventually operate in the months *before* Uber began operating in the treated cities. If unemployment was lower in treated cities before Uber ever began operating, then perhaps the above results are being driven by an omitted variable or simultaneity bias. Calculating the Wald-TC estimator for the 50 months before and after Uber's arrival for each city using Chaisemartin et al. (2020)'s method yields the results in Figure 1.3 and Table A.1.

Fortunately, the common trends assumption appears to hold in our data; there is no

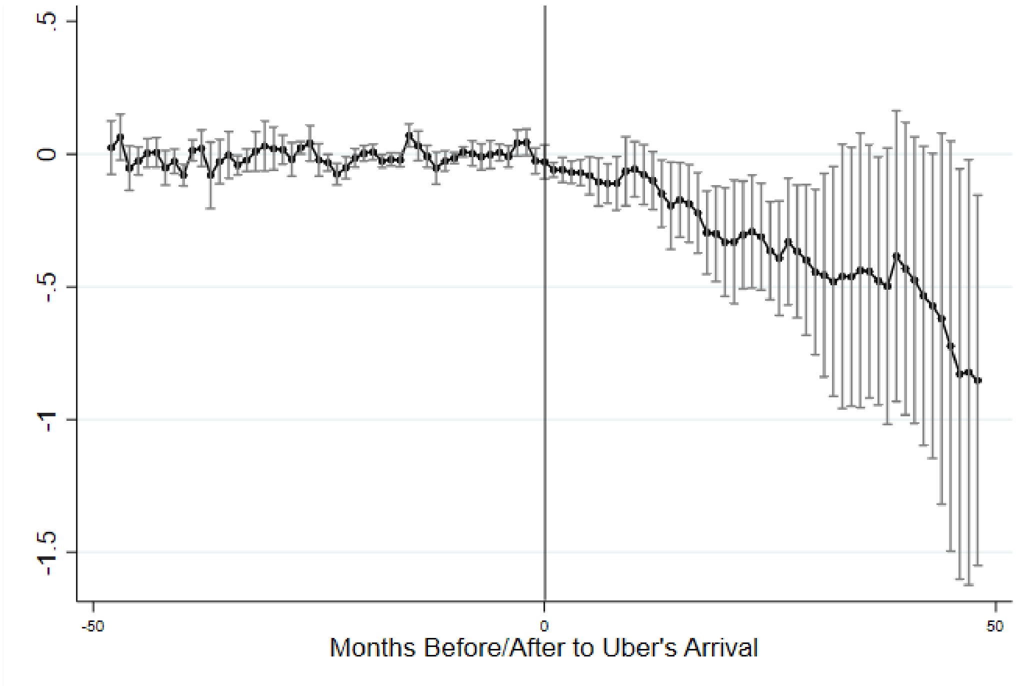

Figure 1.3: Time-Corrected Wald Estimates of Uber's Effect on the Unemployment Rate

noticeable difference between treated and untreated cities prior to Uber's arrival. The average treatment effect during the 50 "placebo" periods was -0.007, while the average standard error was 0.02. Additionally, the Wald-TC estimator yields an estimate of the effect of Uber's arrival that is consistent with the above OLS results: Uber tends to lower the unemployment rate in cities where it begins operating. The average treatment effect during the 50 post-treatment periods was -0.36, with an average bootstrapped standard error of 0.09, slightly higher in magnitude then the estimates from two-way fixed effects. Also noteworthy is the gradual increase in the treatment effect with the amount of time Uber operates in a metropolitan area. The effect on the unemployment rate does not become statistically significant until 6 months after Uber's introduction.

Callaway and Sant'Anna Estimation

Anpther alternative procedure for adressing the potential bias in two-way-fixed-effects regression models comes from Callaway and Sant'Anna (2021)[5]. The authors point out that two-way-fixed-effects estimators can be decomposed into a weighted average of three different comparisons:

- Comparing treated units to *never* treated units
- Comparing treated units to *not yet* treated units
- Comparing treated units to *previously* treated units

The final comparison can be problematic, as earlier treated units are likely a poor counterfactual for the outcomes of later treated units. Callaway and Sant'Anna reccomend an alternative procedure where the outcomes for each group treated units that undergo treatment at the same time are compared to those of either *never treated* units*:*

$$ATT^{nev}_{unc}(g,t) = \mathbb{E}(Y_t - Y_{g-1}|G_g = 1) - E(Y_t - Y_{g-1}|C = 1)$$

or *not yet treated units:*

$$ATT^{ny}_{unc}(g,t) = \mathbb{E}(Y_t - Y_{g-1}|G_g = 1) - E(Y_t - Y_{g-1}|D_t = 0, \mathrm{G_g} = 0)$$

Where $G_{i,g} = 1$ if unit i is first treated at time g, and zero otherwise, $C = 1$ if the group is never treated. Estimated treatement effects for each

"cohort" for each time period can then be aggregated to yield an estimate of the average effect of participating in the treatment for exactly e time periods:

$$\theta^{event}_D(e) = \sum_{g-2}^{\tau} 1\{g + e \leq \tau\} ATT(g, g+e) P(G = g|G + e \leq \tau, C \neq 1)$$

Using the authors' *csdid* package for Stata, the treatment effect of Uber's arrival can be estimated using both never-treated cities and not-yet treated cities as comparison groups. Table 1.4 shows that Callaway and Sant'Anna's method yields similar results to those used in sections 6.1 and 6.2: Uber's arrival is associated with a decline in the unemployment

Table 1.4: Callaway and Sant'Anna Estimates of Uber's Effect on the Unemployment Rate

	(1)	(2)
ATT	-0.711***	-0.324**
	(-3.51)	(-2.80)
Average Pre-Treatment Effect	-0.020**	-0.002
	(-3.37)	(-0.64)
Control	Never Treated	Not Yet Treated

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

rate, regardless of whether never treated cities or not yet treated cities are chosen as the comparison group. Figure 1.4 depicts the event study estimates for the specification using the not yet treated control group.

Since Callaway and Sant'Anna's method estimates the treatment effect for each "cohort" of units that undergo treatment at the same time, we can also compare the estimated treatment effect for cities where Uber arrived earlier to cities where Uber arrived later. As seen in Figure 1.5, cities where Uber arrived earlier, particularly the earliest treated cities with arrival dates in 2013, saw relative changes in unemployment that were smaller, and sometimes even of a different sign, than those experienced by later treated cities.

Synthetic Control Estimation

One additional estimation strategy is the the synthetic control methodology first introduced by Abadie et al. (2010)[6]. The intuition behind synthetic control methods are simple: by constructing a synthetic version of the treated units consisting of a weighted average of donor units which minimizes the pre-treatment root mean squared error in the dependent variable and other predictor variables, the post-treatment behavior of the synthetic unit can analyzed as the counterfactual of the treated united had the treatment not taken place. Thus the effect of the treatment on treated unit 1 in post treatment period t on dependent variable δ be written as:

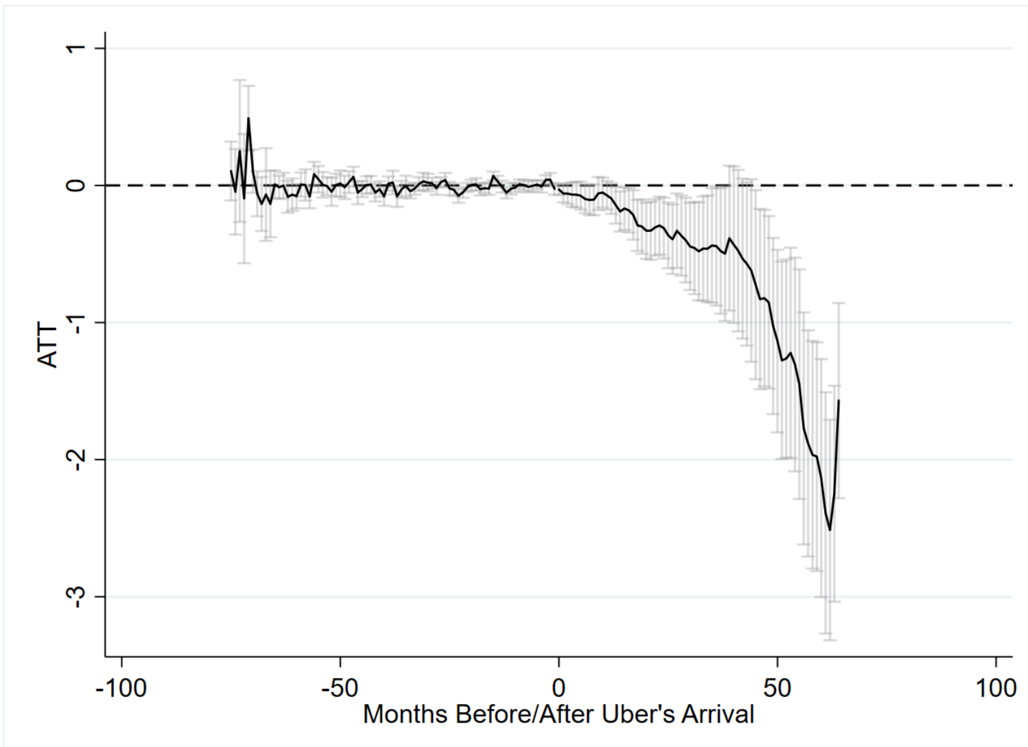

Figure 1.4: Callaway and Sant'Anna Estimates of Uber's Effect on the Unemployment Rate

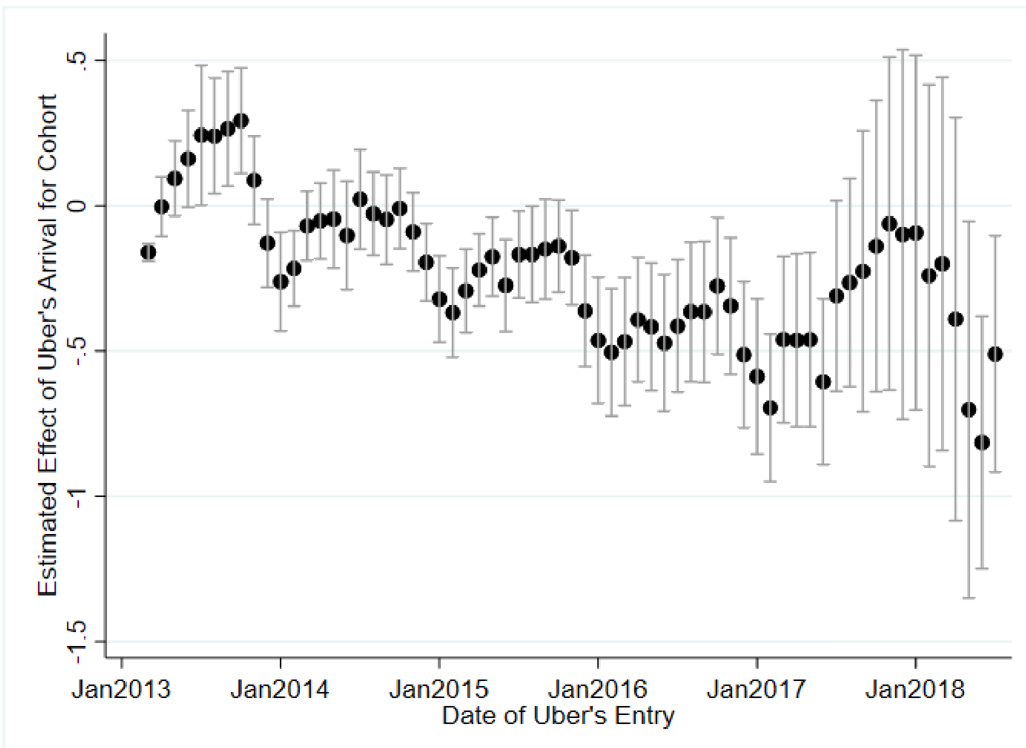

Figure 1.5: Callaway and Sant'Anna Results by Entry Date Cohorts

$$\hat{\alpha}_{1t} = \delta_{1at} - \delta_{1st} \tag{1.5}$$

Where δ_{at} and δ_{st} are the levels of the dependent variable in the actual and synthetic treated unit at time t.

After constructing the treatment effects using the synthetic control for the treated units, statistical significance can be determined by running placebo tests. By estimating the same model on each untreated donor unit, while disallowing the treated unit to be used as a donor, one can generate a distribution of effect sizes for the placebo unit. If the size of the treatment effect for the treated unit is much larger than those generated for the untreated units, than it is unlikely that the estimated effect was the result of chance. If the distribution of placebo effects at time t is $\hat{\alpha}_{1t}^{PL} = \{\hat{\alpha}_{jt} : j \neq 1\}$, then the two-sided p-values for period t are:

$$p - value_{2s} = Prob(|\hat{\alpha}_{1t}^{PL}| \geq |\hat{\alpha_{1t}}|) \tag{1.6}$$

If some placebo units have poor matches than the p-values may be too conservative. Galiani and Quistorff (2016)[15] recommend dividing all effects by the pre-treatment fit to generate "pre-treatment adjusted" p-values.

For an example of how the Synthetic Control method can be used to assess the impact of Uber's arrival, we can use the case study of Louisville, KY. Uber began operation in Louisville in May, 2014, a date towards the center of the distribution of Uber's arrival dates, which allows a large portion of the sample to serve as potential untreated donors. I construct a synthetic Louisville from donor cities with arrival dates after April 2015, matching on the average size of the labor force during the pre-treatment period, and the unemployment rate in January, May, and September in 2012 and 2013.The weights used to construct the counterfactual synthetic Louisville are found in Table 1.5. When the trajectory of Louisville's unemployment rate is compared to that of the untreated Louisville, the effect of Uber's arrival is clear: Louisville's unemployment rate noticeably dips below that of

Table 1.5: Weights used to Construct Synthetic Louisville

City	Weight
Farmington, New Mexico	0.219
Las Vegas, Nevada	0.034
New Orleans, Louisiana	0.017
Odgen, Utah	0.095
Pine Bluff, Arkansas	0.201
Pocatello, Idaho	0.250
Portland, Oregon	0.184

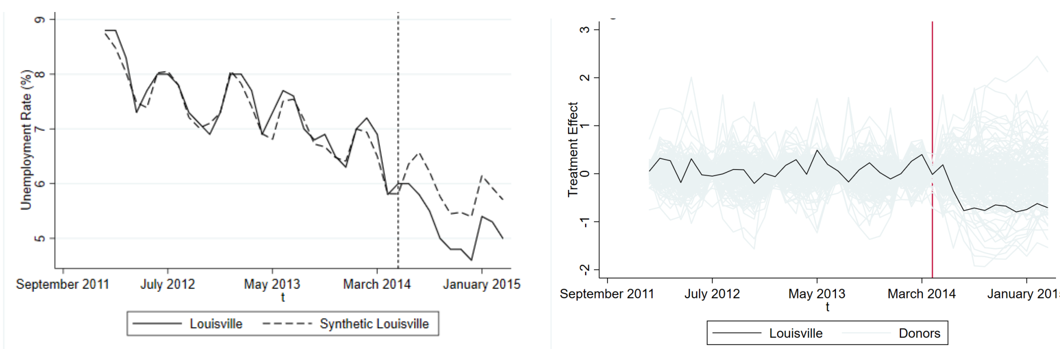

(a) Unemployment Rate in Actual and Synthetic Louisville

(b) Treatment Effects for Louisville and Donor Cities

Figure 1.6: Treatment Effects and Placebo Test for Louisville, KY

Synthetic Louisville following the date of Uber's arrival (Figure 1.6a).

The results of the placebo test suggest that the decline in Louisville's unemployment rate relative to its synthetic counterpart following Uber's arrival are likely not the result of chance, as the adjusted p-values approach statistical significance for all but the earliest post-treatment periods (Figure 1.6b, Table 1.6). Of course, Louisville is just a single case study: what would be optimal would be to construct a synthetic control for each and every city in our sample, so that we could calculate an average effect across *all* cities.

Cavallo et al. (2013)[16] develop an extension of the synthetic control methodology which accommodate such designs. If treated units are indexed $g \in \{1...G\}$, and J are those units

Table 1.6: Synthetic Control Estimation of Uber's Effect on the Unemployment Rate in Louisville, KY

Period	Estimated Effect	P-Value	Pre-Treatment Adjusted P-Value
June, 2014	0.1867	0.563	0.455
July, 2014	-0.355	0.335	0.259
August, 2014	-0.769	0.088	0.018
September, 2014	-0.713	0.158	0.056
October, 2014	-0.765	0.151	0.069
November, 2014	-0.649	0.227	0.101
December, 2014	-0.674	0.196	0.132
January, 2015	-0.795	0.151	0.044
February, 2015	-0.745	0.126	0.082
March, 2015	-0.621	0.177	0.113
April, 2015	-0.706	0.170	0.044

that never undergo treatment, then for each treated unit g, one can estimate treatment effects using the methodology described above, for example, the first post-treated period effect $\hat{\alpha}_g$, with the t subscript omitted as treatment dates differ across events. Over all of the treated units, the average treatment effect is $\bar{\alpha} = G^{-1} \sum_{g=1}^{G} \hat{\alpha}_g$.

Treated units that share the same treatment period share a set of placebo effects $\hat{\alpha}_g^{PL}$. Just as we averaged the treatment effects across treated units in a given post-treatment period to obtain $\bar{\alpha}$, we can construct $\bar{\alpha}^{PL}$, the average placebo treatment effect for each treatment period. With the placebo averages calculated, conducting inference can be done with:

$$p - value = Prob(|\bar{\alpha}^{PL}| \geq |\bar{\alpha}|) \tag{1.7}$$

To construct a subsample for synthetic control estimation, I drop from the analysis those cities with between one and nine post-treatment observations, as well as those cities in which Uber began operating during or before March 2014. The properties of the subsample are found in Table 1.7. Like the example with Louisville, the synthetic control for each unit

Table 1.7: Summary Statistics: Synthetic Control Subsample

Variable		Mean	Std. Dev.	Min	Max	Observations
unemp	overall	5.7	2.3	1.4	28.3	N=21804
	between		1.4	2.7	21.1	n=276
	within		1.5	-2.0	16.5	T=79
uber	overall	.43	.49	0	1	N=21804
	between		.18	0	.65	n=276
	within		.45	-.22	1.3	T=79
laborforce	overall	235722	287647	24601	2046192	N=21804
	between		287987	25353	1939291	n=276
	within		10045	139142	368478	T=79

is generated by matching on the average size of the labor force during the pre-treatment period, as well as the unemployment rate in January, May, and September in 2012 and 2013.

Estimating a synthetic control for each treated city using the *synth_runner* Stata package from Galiani and Quistorff (2016)[15] yields the results in Table 1.8. Similar to the other identification strategies, I estimate that Uber's introduction had a negative and statistically significant effect on the unemployment rate in a city following its introduction. The magnitudes of the estimated treatment effects are larger in magnitude when synthetic controls are used, which is likely a result of the subsample selection process. Cities with fewer post-treatment observations will likely have smaller estimated effects on average, due to the "phase-in" effects observed in the previous section. For example, if the phase-in period is 6 months, then a city with 7 post-treatment observations will have only 1 observation where the effect has reached it's full magnitude, while a city with 14 post-treatment observations will have 8 such observations. Since I purposefully exclude cities with few post-treatment observations in order to better construct synthetic controls for each unit, we should expect the estimated treatment effects to be larger.

Table 1.8: Synthetic Control Estimation of Uber's Effect on the Unemployment Rate in All Cities

Period	$\bar{\alpha}$	p-value	Adjusted p-value
1	-0.473	0	0
2	-0.542	0	0
3	-0.539	0	0
4	-0.551	0	0
5	-0.514	0	0
6	-0.503	0	0
7	-0.548	0	0
8	-0.541	0	0
9	-0.574	0	0
10	-0.526	0	.000012

Austin's Experience: When Uber Leaves and Returns

Even if multiple methods suggest that Uber's arrival was associated with a reduction in the unemployment rate, some may still be cautious to view these results as a causal effect. After all, Uber's decision to begin operation in an area could be partially influenced by the health of that area's labor market. If this was the case, then causation would run in the other direction: a lower unemployment rate *causes* Uber to begin operating, rather than the other way around. The confirmation of parallel trends in the previous section partially assuages this fear: a city's unemployment rate does not go down relative to other cities until *after* Uber arrives. Still, it could be the case that Uber chooses new markets where it forecasts that unemployment will soon drop, which would lead to both reverse causality and parallel trends. The best way to identify Uber's effect would be to take the decision of where to operate out of the company's hands entirely, for example, by randomly choosing a city each month for Uber to begin operation. We could then be confident that any effect we observed on the unemployment rate was being *caused* by Uber's arrival.

The next best alternative to a random experiment is to examine the labor market in Austin, Texas specifically, which had a unique experience with Uber. While the ridesharing app began operating in Austin in June, 2014, Uber suspended its operations in Austin in

Table 1.9: Uber's Effect on the Unemployment Rate in Austin, Texas

	(1)	(2)	
	2012-2018	2014-2018	
Uber	-0.3838***	-0.176*	-0.3967***
	(-3.57)	(-1.95)	(-4.12)
Linear Time Trend	-0.0365***	-0.0238***	-0.2543***
	(-15.92)	(-9.31)	(-6.00)
Quadratic Time Trend			0.001***
			(5.51)
N	80	56	56

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

May, 2016 after voters rejected Proposition 1, which would have overridden city ordinances requiring driver-partners to submit to fingerprinting and background checks. Uber returned to Austin nearly one year later in June, 2017, after statewide regulations passed superseding Austin's regulations on the service. This case study provides a neater natural experiment with less endogeniety, albeit with far fewer data points. I examine Uber's impact on employment in Austin for both the full 2012-2018 sample used above, as well as a truncated 2014-2018 sample, beginning when Uber arrived in Austin in order to isolate the effects only of Uber's departure and eventual return. I use a regression of the following functional form:

$$unemp_t = \beta_0 + \beta_1 U_t + \beta_2 t + \beta_3 t^2 + \varepsilon_t \tag{1.8}$$

The results shown in Figure 1.7 and Table 1.9 show that Austin's experience seems to mirror the experience of the nation at large. The effect of Uber's absence on the unemployment rate was roughly -0.4 percentage points.

Effects of Driver-Partner Employment

My analysis thus far has used Uber's entry as the independent variable of interest, which cannot capture differences in Uber's activity *within* a city over time. As an alternative, I

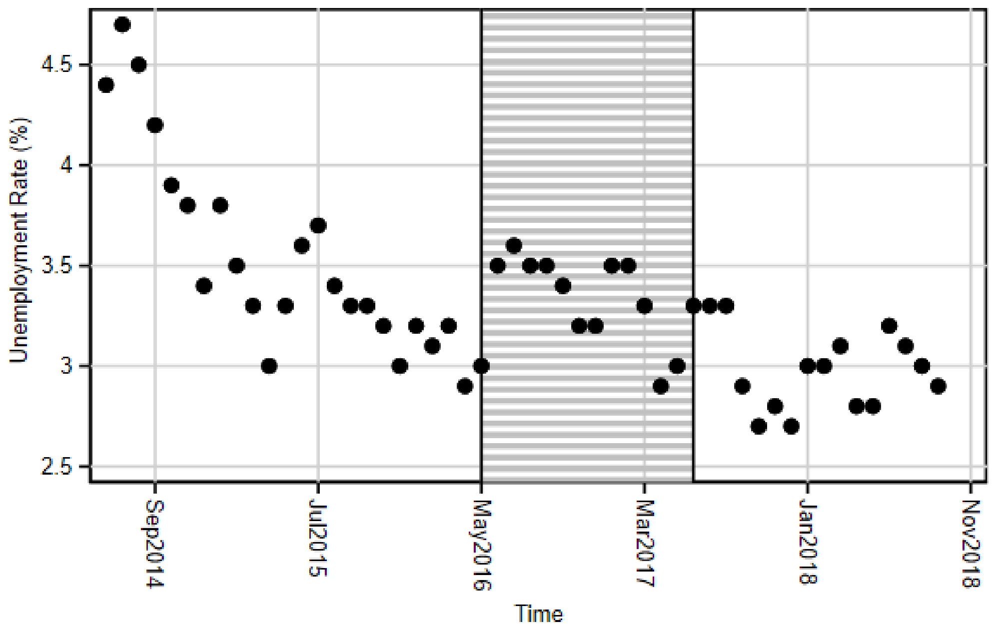

Figure 1.7: Unemployment Rate in Austin TX, Period of Uber's Absence Shaded

Table 1.10: Summary Statistics for Driver-Partner Data

Variable		Mean	Std. Dev.	Min	Max	Observations
drivers	overall	6.33	8.93	0.034	56.67	N=640
	between		4.79	1.13	15.78	n=16
	within		7.60	-9.35	47.22	T=40

obtained monthly data on the number of Uber driver-partners active in 16 large US cities from Hall et al. (2018). Uber considers a driver-partner "active" if they have completed at least four trips in a given month. Summary statistics for the driver-partner data are reported in Table 1.10, and Figure 1.8 plots the change in driver-partner employment by city.

I use the following model to estimate the relationship between active driver-partners in a city and that city's unemployment rate:

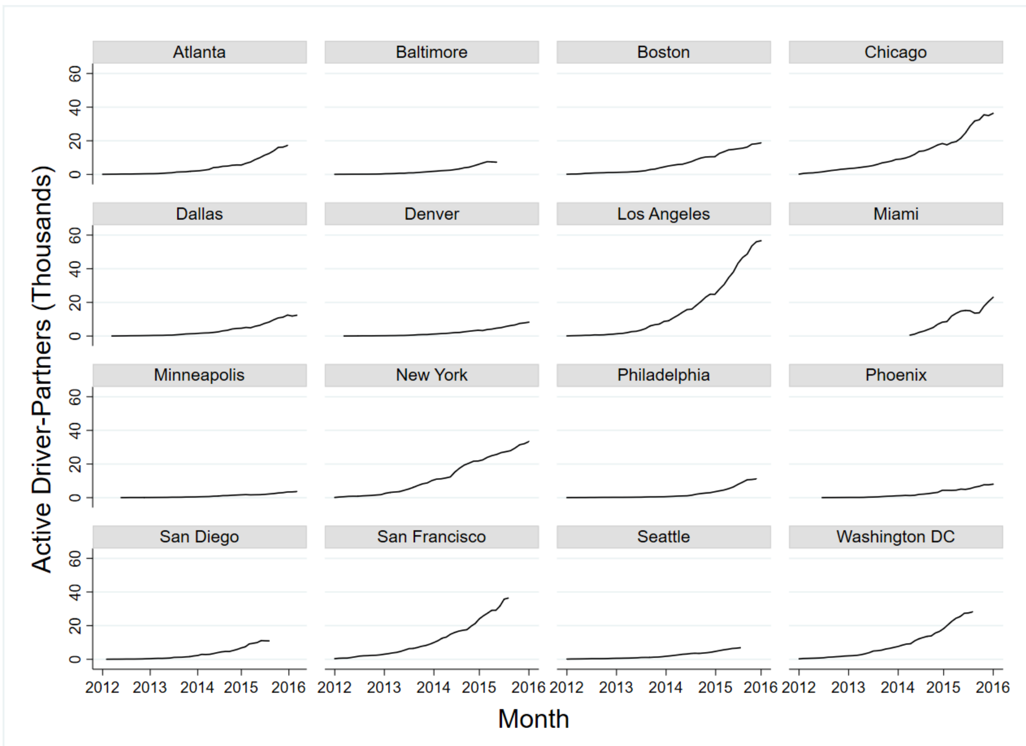

Figure 1.8: Driver-Partner Employment By City

Table 1.11: Effects of Driver-Partner Employment on the Unemployment Rate

	(1)	(2)	(3)
Active Driver-Partners (Thousands)	-0.0275	-0.0284*	-0.0192
	(-2.02)	(-2.15)	(-1.56)
Labor Force Size		0.00000211	0.00000179**
		(1.02)	(3.05)
City Fixed Effects	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes
City-Specific Linear Time Trends	No	No	Yes
N	640	640	640

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

$$unemp_{it} = \beta_0 + \beta_1 drivers_{it} + \beta_2 labforce_{it} + \alpha_i + \lambda_t + X + \varepsilon_{it} \tag{1.9}$$

Where $drivers_{it}$ is the number of active Uber driver-partners in city i in time period t in thousands, X is a set of city-specific linear time trends, and all other variables are the same as in equation (1). Note that the sample for this model only consists of the 640 city/month observations with driver-partner data.

As seen in Table 1.11, even with the addition of city-specific linear time trends, growth in Uber employment is associated with declines in the unemployment rate, as shown in 2011. Each additional 1,000 driver-partners in a city is estimated to decrease the unemployment rate by between -0.019 and -0.028 percentage points. Since the average number of driver-partners after the introduction of UberX ride-sharing in April 2013 was roughly 8,000, the effect of Uber entry based on this model should be between -0.152 and -0.224, consistent with the effect of -0.211 found in section 6.1.

Table 1.12: Uber's Effect on the Unemployment Rate: Effects by Car Ownership

	(1) Highest Car Ownership	(2) Lowest Car Ownership
ATT	-0.303*	-0.142
	(-2.36)	(-1.32)
Average Pre-Treatment Effect	-0.002**	0.010
	(-0.64)	(1.72)
Control	Not Yet Treated	Not Yet Treated

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

1.3.2 Subgroup Analysis

There are two plausible mechanisms by which Uber's arrival decreases the unemployment rate. The first is by serving as an "employer of last resort" for individuals who are frictionally unemployed, while the second is by providing individuals an opportunity to commute to jobs they would otherwise be unable to reach, expanding their pool of potential work opportunities. Both mechanisms imply that the city-level treatment effect of Uber on the unemployment rate will be heterogenous depending on certain city-level characteristics.

The first mechanism occurs when individuals drive for Uber during periods of frictional unemployment between two more traditional jobs. Because driving with Uber requires that an individual use their own vehicle, this effect should be stronger in cities with higher levels of car ownership. Using data from the 2016 American Community Survey, I calculate the fraction of households that own a car in 139 cities. Then, using Callaway and Sant'Anna's (2021)[5] method, I estimate Uber's effect on the unemployment rate for only those cities in the top and bottom quartile of this variable. As reported in Table 1.12, Uber's effect is negative and statistically significant in the cities with the highest rate of car ownership among households, while smaller and not statistically significant in the cities with the lowest rate of car ownership. Figure 1.9 shows the event study results for high and low ownership cities.

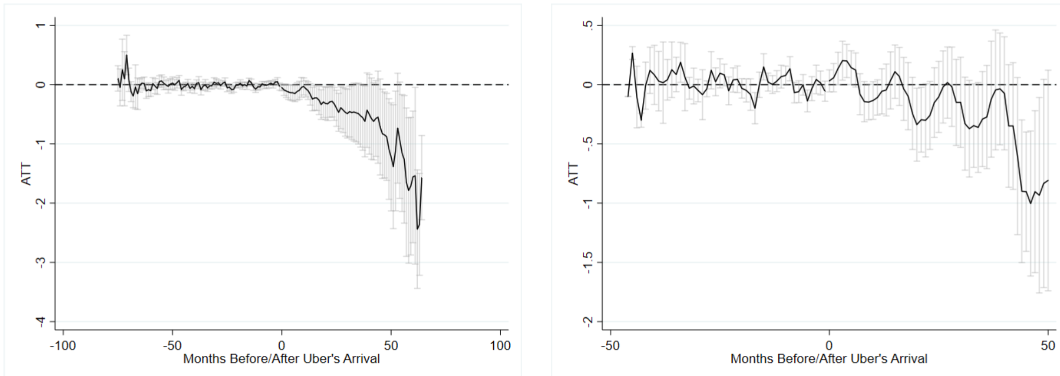

(a) Highest Car Ownership (b) Lowest Car Ownership

Figure 1.9: *csdid* Results by Car Ownership

The second mechanism whereby people use ridesharing to commute to jobs they would otherwise be unable to reach implies treatment effect heterogeniety across cities with different initial unemployment rates. Job seekers in cties with tighter labor markets characterized by more job vaccancies may be more likely to find an opening that Uber can help them commute to. To examine this possibility, I calculate Uber's ATT for cities in the highest and lowest quartile of the unemployment rate in January 2012, and record the results in Table 1.13 and Figure 1.10. Cities with the most slack in their labor markets experienced no statistically significant decline in unemployment following Uber's arrival, while Uber's arrival to cities with the tightest labor markets was associated nearly a half of a percentage point reduction in the unemployment rate.

1.3.3 Ridesharing's Impact on Wages

Two-Way Fixed Effects Estimation

Uber's arrival in a metropolitan area may also impact the wages that workers earn, though theory is ambiguous on the direction of the effect. Critics like Ravenelle (2019)[1] fear that growth in less-regulated and non-union work in the sharing economy comes at the expense of higher-paying "traditional" jobs. If this is the case, then Uber's arrival in a city should cause

Table 1.13: Uber's Effect on the Unemployment Rate: Effects by Initial Unemployment Rate

	(1)	(2)
	Highest Initial Unemployment	Lowest Initial Unemployment
ATT	-0.211	-0.487**
	(-1.49)	(-3.00)
Average Pre-Treatment Effect	0.0008	0.001
	(0.23)	(0.42)
Control	Not Yet Treated	Not Yet Treated

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 1.10: *csdid* Results by Initial Unemployment Rate

(a) Highest Initial Unemployment Rate

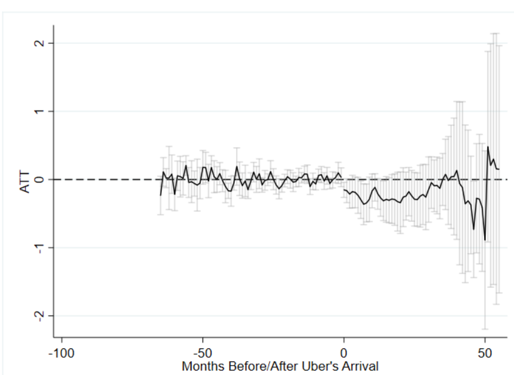

(b) Lowest Initial Unemployment Rate

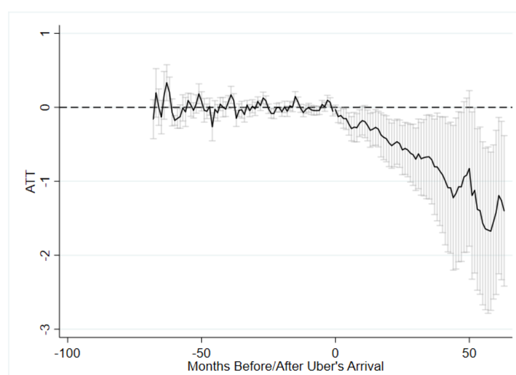

Table 1.14: Uber's Effect on Wages

	(1) 10th Percentile	(2) 25th Percentile	(3) Median	(4) Mean
uber	-0.00417*	0.00445**	0.00183	0.000158
	(-1.99)	(2.82)	(1.17)	(0.10)
City Fixed Effects	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes
N	2058	2058	2058	2058

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

a decrease in wages. Conversely, there are other mechanisms by which Uber could *raise* average wages in a city. Uber could serve a similar function to unemployment insurance by cushioning earnings during periods of frictional unemployment, which would in turn raise the reservation wages of unemployed workers and lengthen job search times. Uber may also put upward pressure on wages if some workers find driving for Uber full-time more attractive than their current occupation, either because they earn more as a driver-partner or because they value the flexibility. If this is the case, then existing employers must increase the wage they pay in order to retain and attract employees away from Uber.

To test these possibilities, I gathered data on 25th percentile, Median, and Mean wages for all occupations from the Bureau of Labor Statistics' Occupational Employment Statistics program for Metropolitan Statistical Areas for 2012 through 2017. To identify this potential effect, I use a similar Difference-in-Differences specification to the one previously used to examine Uber's effect on employment:

$$ln(wage_{it}) = \beta_0 + \beta_1 U_{it} + \alpha_i + \lambda_t + \varepsilon_{it} \tag{1.10}$$

Which gives the results contained in Table 1.14.

While the Bureau of Labor statistics reports employment data monthly, wage data is only published on a yearly basis. As a result, a city where Uber has been operating for 11 months is statistically identical to one where Uber has been operating for only 2. This

Table 1.15: Uber's Effect on Wages (Alternate Treatment Variable)

	(1)	(2)	(3)	(4)
	10th Percentile	25th Percentile	Median	Mean
sixmos	-0.00175	0.00282	0.000260	-0.00102
	(-0.73)	(1.45)	(0.14)	(-0.57)
City Fixed Effects	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes
N	2058	2058	2058	2058

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

is potentially concerning given that the effects on employment did not appear until several months after Uber commenced operation. To account for this, I estimate the following model with an alternative treatment variable:

$$ln(wage_{it}) = \beta_0 + \beta_1 U6mos_{it} + \alpha_i + \lambda_t + \varepsilon_{it} \tag{1.11}$$

Where $U6mos_{it}$ is a binary variable equal to 1 if Uber had been operating for at least four months in city i in May (when BLS data is collected) of year t. The results of these models are found in Table 1.15.

Callaway and Sant'Anna Estimation

The results in Tables 1.14 and 1.15 indicate that Uber's arrival may have had a small positive effect on workers' wages, but only at the lower end of the wage distribution. As an additional robustness test, I apply estimator from section 6.3 for the four periods before and the 2 periods after the date of Uber's arrival, as seen in Table 1.16 and Figure 1.11.

csdid estimates of Uber's effect on wages are similar to those found using OLS. In addition, there does not appear to be any significant differences between treated and untreated cities prior to Uber's arrival, which is indicative of parallel trends.

Table 1.16: Callaway and Sant'Anna Estimates of Uber's Effect on Log 25th Percentile Wages

	(1)
	25th Percentile
ATT	0.088***
	(3.86)
Average Pre-Treatment Effect	0.008
	(0.88)
Control	Not Yet Treated

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

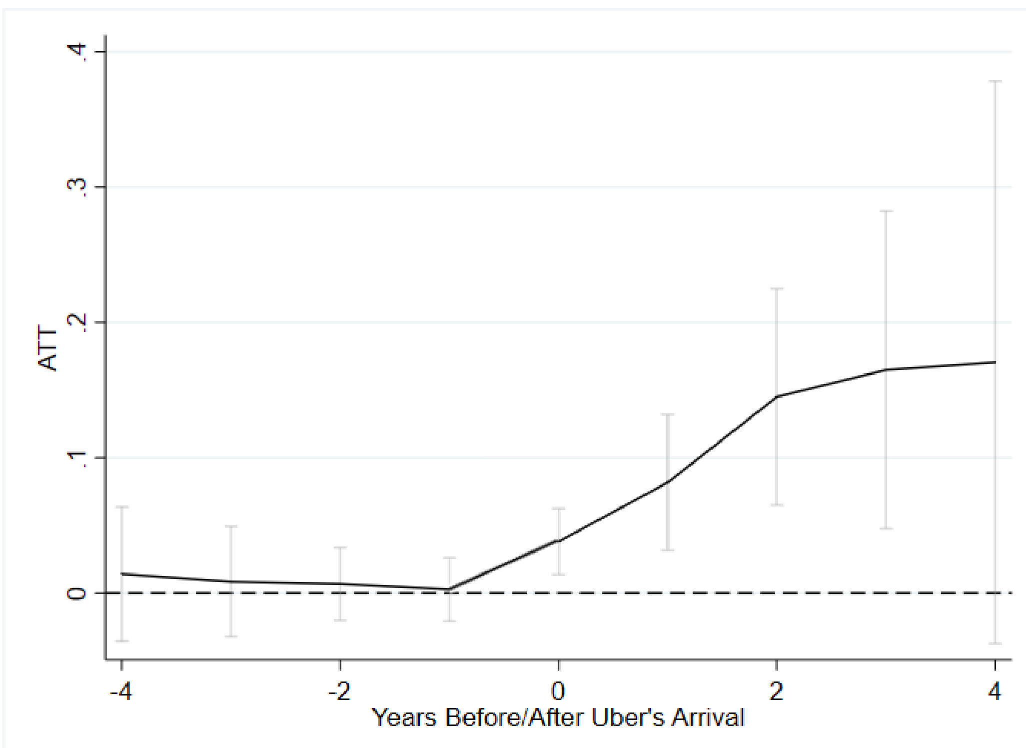

Figure 1.11: Uber's Effect on Log 25th Percentile Wages

1.4 Conclusion

There are two main views of how the sharing economy will impact the future of work. The first is that, as a substitute for traditional employment, all "jobs" will gradually be replaced by "gigs", for better or for worse. The second is that the sharing economy serves primarily as a complement for traditional employment, enabling people to more deftly navigate the space *between* conventional jobs and smooth fluctuations in their incomes. By examining the impact of the introduction of one of the most famous sharing economy firms, Uber, I conclude that the "complement view" is likely more relevant. Uber's introduction in a metropolitan area appeared to cause a decrease in the unemployment rate in that area, implying that Uber was utilized by those who were fictionally unemployed. The experience of Austin, TX, which experienced a cleaner "natural experiment" with Uber confirms that the ridesharing service lubricated the frictions between jobs. Finally, Uber's arrival had, if anything, a very small positive effect on workers' wages, casting doubt on both the view that Uber reduces worker welfare by destroying higher-paying jobs.

This analysis has limitations. Most evidently, it is limited by the choice to examine Uber's effect by means of an indicator variable, rather than via a continuous measurement such as the number of driver-partners currently active within a metropolitan area. It's also possible that the external validity of my results are limited, and the effects of Uber are unlike the effects of other gig economy firms like TaskRabbit or even unlike other ridesharing firms like Lyft. Finally, like all analysis using difference-in-differences style methods, it is impossible to verify the presence of parallel counterfactual trends directly. My results are indicative of parallel *pre-treatment* trends, but there is always the possibility of a confounder effecting unemployment and wages simultaneously with Uber's arrival.

Chapter 2: Puerto Rico's Minimum Wage

2.1 Introduction

The employment effect, or lack thereof, of the minimum wage remains a contentious issue in empirical economics, but a significant amount of progress on the issue has been made since the dueling studies of Card & Krueger (1993)[17] and Neumark & Wascher (1995)[18]. Modern research on the minimum wage generally finds an employment response that is negative, small, and mostly localized to groups such as teenagers (Belman & Wolfson (2019)[19], Neumark & Shirley (2021)[20]). However, most studies use variation in state and local minimum wages to identify the effect of the minimum, and states and localities in the United States have generally chosen minimum wages that are modest relative to their median wages. In Cengiz et al.'s (2019)[21] comprehensive study of state level minimum wage changes, for example, the highest relative minimum wage was 59% of the median wage.

With current proposals to dramatically increase the federal minimum wage to $15/hr being considered by policymakers, the current literature on the employment effects of the minimum wage may be ill-equipped to forecast the effects of much larger minimum wage increases. A federal minimum wage of $15/hr would be 95% of the median wage in Mississippi and 87% of the median wage in West Virginia if implemented immediately, for example. Such an increase is far beyond those studied by the majority of the minimum wage literature.The employment effects of such a high minimum may be better gauged by looking to the US territory of Puerto Rico, which has been bound by the Federal Minimum wage since 1983. While Puerto Rico is wealthier than other Caribbean countries, with GDP per capita in 2020 equal to roughly $31,000, compared to $9,000 in Jamaica, $18,000 in the Dominican Republic, and $25,000 in St. Kitts and Nevis, it is simultaneously poorer than the U.S. states with which it shares Federal Laws. GDP per capita for all states was \$65,281, nearly double Puerto Rico's, and even the poorest states like Mississippi (\$40,000) and West Virginia (\$43,000) are still much wealthier than Puerto Rico. Average hourly earnings in Puerto Rico were equal to \$12.21 in 2010, compared to \$21.92 for the mainland United States, \$17.74 for Mississippi, and \$17.65 in West Virginia. This results in Puerto Rico having a *relative* minimum wage that is much higher than any other state. Most recently, in 2007, the Federal Minimum Wage was increased from \$5.15/hr, which was 59% of Puerto Rico's median wage at the time, to \$7.25/hr, 76% of the median wage. Estimating the effect of this increase on Puerto Rican employment might provide insight into whether the minimum wage has a modest disemployment effect even at higher relative levels.

This paper revisits the question of the minimum wage in Puerto Rico using the most recent increase in the federal minimum wage as a case study. Since there is no *a priori* sensible control group to use for Puerto Rico, such as a neighboring state with a higher minimum wage, I utilize Abadie et al.'s (2010)[22] synthetic control method, constructing a plausible counterfactual for the path of Puerto Rican employment without the minimum wage increase using the labor markets of other nations, Puerto Rican industries less affected by the minimum wage increase, and cities on the mainland United States. Identification is threatened by the Great Recession of the late 2000s, which coincided with the minimum wage increase and was particularly pronounced in Puerto Rico. To combat this potential source of bias, I use "triple differences" strategies comparing the estimated employment effects of groups with differential exposure to the wage floor and find that groups with lower wage levels and thus more exposure to the minimum experienced proportionally larger relative decreases in employment. I also construct a synthetic Puerto Rico matching on fluctuations in pre- and post-treatment GDP per capita, rather than just pre-treatment employment, creating a counterfactual which experienced similar macroeconomic fluctuations during the post-treatment period. I find that the increase in the minimum wage lead to substantial reductions in Puerto Rican employment across all specifications. On average, my results suggest an own-wage elasticity of employment of -0.68, larger than estimates from previous studies of the minimum wage in the United States, which are generally in the -0.3 to -0.5 range. I discuss several reasons why the employment response to the minimum wage may have a non-constant elasticity. Alternative avenues by which employers may choose to adjust, such as cuts to hours or fringe benefits, have a limited capacity to absorb higher labor costs, leaving cuts to employment as the only remaining option at higher relative minimum wages. A higher minimum wage is also more likely to bind in the tradable goods sector, where employment is demonstrated to be much more sensitive to minimum wage increases due to more elastic product demand.

The remainder of the paper is organized as follows. Part II gives a brief summary of the state of the literature on the employment effects of minimum wages and previous investigations into Puerto Rico more specifically. Part III provides background on the minimum wage in Puerto Rico and provides evidence that the 2007 increase in the minimum lead to a large increase in hourly wage. Part IV explains the synthetic control method, and Part V presents my results. Finally, Part VI concludes.

2.2 Prior Research

Research on the employment effects of the minimum wage is as voluminous as it is controversial. As Neumark & Shirley (2021)[20] summarize: "depending on what one reads about how economists summarize the evidence, one might conclude that: (1) it is not well-established that higher minimum wages do not reduce employment, (2) the evidence is very mixed with effects centered on zero with no basis for a strong conclusion one way or the other, or (3) most evidence points to adverse employment effects" (p. 2). In general, however, the evidence seems to be synthesized into supporting an employment effect of the minimum wage that is negative, small, and localized within subgroups such as teens and high school dropouts. In their meta-analysis summarizing the past 15 years of research on the employment effects of the minimum wage, Belman & Wolfson (2019)[19] summarize the "consensus range" of elasticities of teen employment to the minimum wage as being -0.13 to -0.07, implying that a 10% increase in the minimum wage decreases employment among teenagers by between 1.3% and 0.7%. Neumark & Shirley's (2021)[20] meta analysis reaches similar conclusions, with a median elasticity of -0.11. The majority of this research relies on evidence from the United States, where cross-state differentials in the minimum wage generate "natural experiments" useful for identifying the employment effect. The average US minimum is 39% of the median wage and the largest minima are just above 50% of the median (Cengiz et al., 2019)[21].

The idea of using Puerto Rico to examine the effects of a high relative minimum wage is not a new one. Santiago (1985)[23] sought to examine the employment and unemployment effects of the minimum wage in Puerto Rico soon after the gap between the continental and Puerto Rican minima closed around 1983. Using multivariate time series techniques and transfer functions, Santiago concluded that "the empirical findings suggest that both disemployment and unemployment effects resulted from the post-1974 minimum wage policy...consistent with theoretical hypothesis" (p. 308). Soon after the revival in interest in minimum wage research in the 1990s, Castillo-Freeman and Freeman (1992)[24] published research on the minimum wage's effect in Puerto Rico, primarily utilizing time-series data for their analysis. The authors found significant impacts on employment, concluding that

> Imposing the U.S.-level minimum reduced total island employment by 8-10 percent compared to the level that would have prevailed had the minimum been the same proportion of average wages as in the United States. In addition, it reallocated labor across industries, greatly reducing jobs in low-wage sectors that had to raise minima substantially to reach federal levels (p. 178).

Castillo-Freeman and Freeman's average estimate of the elasticity of employment to the minimum wage was -0.41. Two years later, Alan Krueger reexamined the impacts of the minimum wage in Puerto Rico, and reached different conclusions from Castillo-Freeman and Freeman, stating:

> The strongest evidence that the minimum wage had a negative effect on employment in Puerto Rico comes from an aggregate time series analysis. The weakest evidence comes from cross-industry analyses. In general, however, I think one would have to consider the evidence surprisingly fragile...perhaps the conclusion that one should reach from the review of evidence is that the jury is still out on Puerto Rico's experience (p. 23)[25].

In the 14 years since then, statistical techniques for casual inference have come a long way, but the evidence from Puerto Rico still lies unexamined with a fresh set of statistical eyes. Even a prominent figure in research on the minimum wage like Arindrajit Dube only cites the Freeman and Krueger papers in a 2015 report on Puerto Rico[26], and David Neumark noted in 2018 that "surprisingly, to the best of my knowledge the evidence from Puerto Rico has not been revisited" (p. 9)[27].

Beyond Puerto Rico, Gregory & Zierahn (2020)[28] study another instance of a high relative minimum wage. In 2006, Germany's minimum wage for roofers increased, leading to a statutory minimum wage that was equal to or even exceeded the median wage in low-wage areas in eastern Germany. The authors conclude that the minimum wage caused the wages of low-skilled East Germans to increase by 5-6%, but also caused employment among that group to decline by 3.5%, suggesting an own-wage elasticity of -0.58 to -0.70.

2.3 Background

When the United States created its first national minimum wage through the Fair Labor Standards Act of 1938, Puerto Rico was exempted from the wage floor of $0.35 for fear that the resultant disemployment effects would "devastate the island's economy" (Castillo-Freeman & Freeman, 1992, p. 178)[24]. Instead, the law established committees to set separate minimum wages for various Puerto Rican industries and occupations. In the 1970s however, the industry-committee method of setting the island's minimum wage was gradually replaced until, by 1983, Puerto Rico was essentially covered by the federal minimum wage.

Table 2.1: Characteristics of Puerto Rican Workers

Group	Count	% Male	% Employed	Avg. Hrly Earnings	% Making $7.25/hr or less
Puerto Ricans (2006)	**34746**	**47.18**	**29.92**	**$11.34**	**46**
Ages 15-24	4900	51.78	23.53	$6.56	76
Accomodation and Food	822	50.49	N/A	$7.38	69
Puerto Ricans (2010)	**36032**	**47.15**	**29.26**	**$12.21**	**35**
Ages 15-24	4982	51.45	19.41	$8.36	58
Accomodation and Food	842	51.61	N/A	$8.77	55
All Americans (2006)	2969741	48.61	46.84	$21.49	12
Ages 15-24	380233	51.17	45.76	$10.19	30
Accomodation and Food	116917	41.92	N/A	$11.45	41
All Americans (2010)	3061692	48.63	44.87	$21.92	10
Ages 15-24	385912	51.37	40.10	$10.72	24
Accomodation and Food	112864	43.81	N/A	$11.98	32

The most recent increase in Puerto Rico's minimum wage occured with the Fair Minimum Wage Act of 2007, which was introduced by Representitive George Miller in January of 2007 and signed into law in May of the same year. The act increased the federal minimum wage of $5.15/hr first to $5.85/hr in July of 2007, then to $6.55/hr in July of 2008, and then finally to $7.25/hr one year later. Table 2.1 contains summary statistics for Puerto Rico and the United States in both 2006 prior to the phase-in of the $7.25/hr minimum wage and in 2010 after its completion.

It's possible that, although the statutory minimum wage for Puerto Rico is high, it's actual effect on workers' wages was mitigated by noncompliance on the part of employers. Perhaps workers shift into more informal work arrangements were the minimum wage is not in effect in reaction to the higher price floor. If this were the case, then Puerto Rico may actually be an inappropriate case study for examine the employment effect of the minimum wage. One way to test this possibility is to compare the distribution of hourly wages in Puerto Rico just before and just after the minimum wage increase. If the increase really did result in higher wages, then we should expect to see a dramatic decrease in the number of workers earning below the new minimum of \$7.25 and a corresponding increase in workers earning \$7.25. Using microdata from the 2000, 2006 and 2010 Puerto Rico Community Surveys (PRCS), I constructed the distribution of hourly wages on the island before the minimum wage increase in 2006 and after the completion of the phase-in in 2010. Comparing the wage distributions before and after the minimum wage increase, “spikes” are evident at the area of the minimum both before and after the increase, suggesting that a great number of Puerto Rican jobs were shifted into compliance with the new \$7.25 minimum wage (Figure 2.1). I also examine the difference between the 2000 and 2006 distributions in a placebo test, where the minimum wage was fixed at \$5.15 for both years. If the difference in the wage distributions between these two years is similarly dramatic, then we may be cautious to interpret the difference between the 2006 and 2010 distributions as a result of the minimum wage increase. The difference between the 2000 and 2006 distributions in Figure 2.2 are comparatively small, with both exhibiting a spike at the contemporary minimum of \$5.15. These results imply that the Fair Minimum Wage Act of 2007 was a major determinant of the actual wages that Puerto Ricans were paid. For another comparison, Figure 2.3 displays the change in the wage distribution in US states affected bound by the minimum wage increase. There is a clear decline in jobs below the new minimum with a corresponding increase in jobs paying the new minimum, but the effect is significantly less dramatic than the one seen in Puerto Rico.

2.4 Methods

The unique difficulty of addressing the impacts of any policy on Puerto Rico is the absence of an *a priori* sensible control group. When examining the impact of state minimum wage

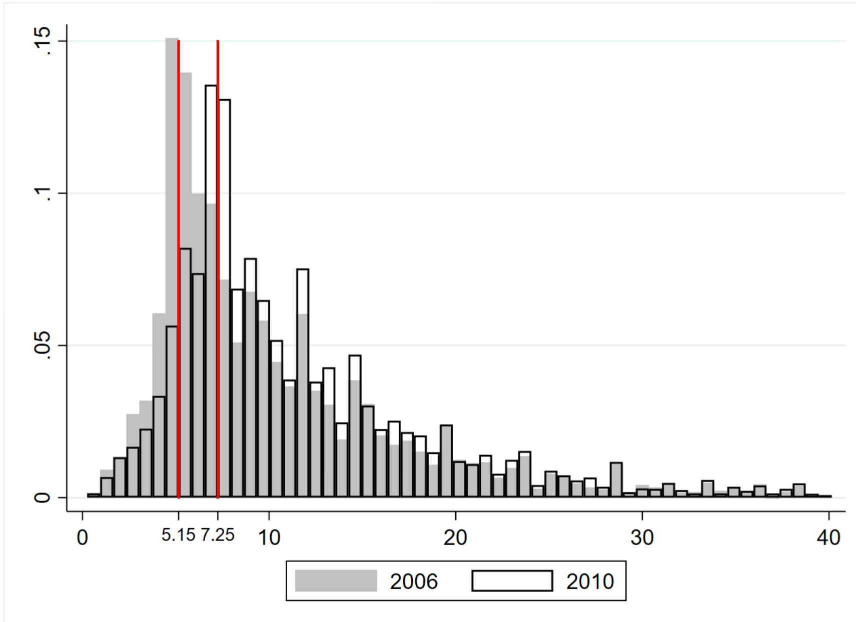

Figure 2.1: Distribution of Wages in Puerto Rico in 2006 and 2010

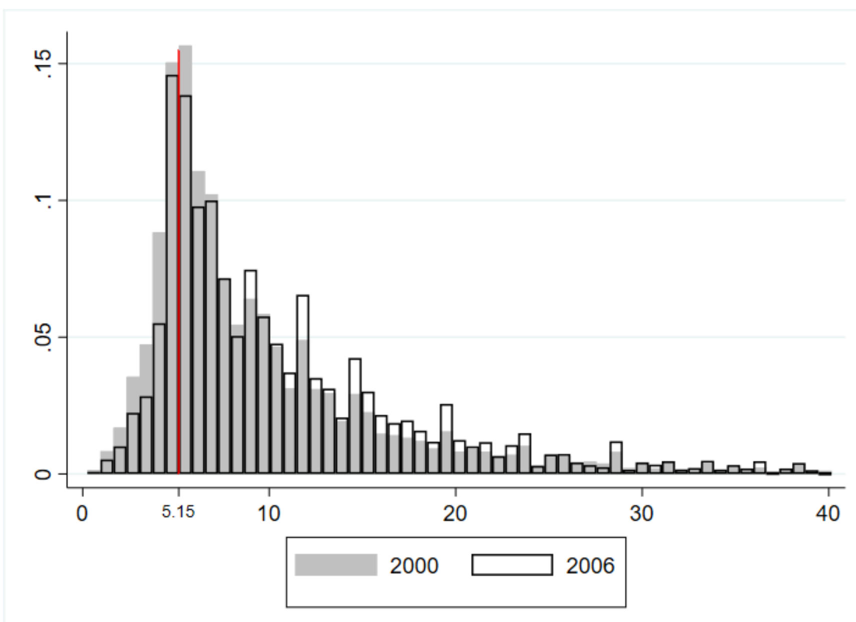

Figure 2.2: Distribution of Wages in Puerto Rico in 2000 and 2006

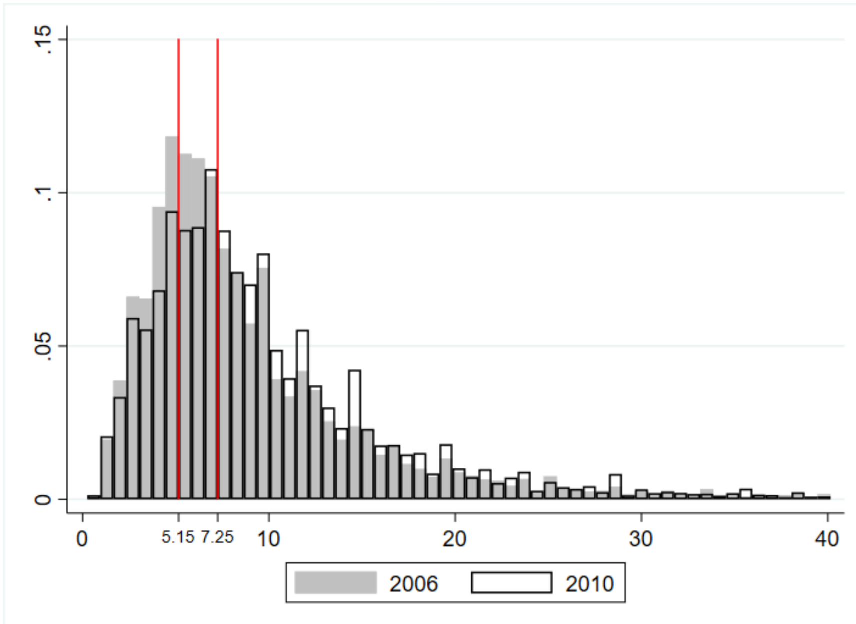

Figure 2.3: Distribution of Wages in U.S. States with Subfederal Minimum Wages in 2006 and 2010

increases on the continental United States, previous research most commonly utilized neighboring states which did not increase their minimum wage as a control group (Dube et al., 2010[29], Cengiz et al., 2017[21]). This strategy is inappropriate for Puerto Rico due to both the island's unique geographic position and relatively low level of economic development relative to the mainland United States. With this in mind, I chose to use a method of policy analysis unavailable to either Castillo-Freeman and Freeman or Krueger in the 1990s: the synthetic control methodology first introduced by Abadie et al. (2010)[22]. The intuition behind synthetic control methods are simple: by constructing a synthetic version of the treated united consisting of a weighted average of donor units which minimizes the pre-treatment root mean squared error in the dependent variable and other predictor variables, the post-treatment behavior of the synthetic unit can analyzed as the counterfactual of the treated united had the treatment not taken place. Thus the effect of the treatment on treated unit 1 in post treatment period t on dependent variable δ be written as:

$$\hat{\alpha}_{1t} = \delta_{1at} - \delta_{1st}$$

Where δ_{at} and δ_{st} are the levels of the dependent variable in the actual and synthetic treated unit at time t.

After constructing the treatment effects using the synthetic control for the treated units, statistical significance can be determined by running placebo tests. By estimating the same model on each untreated donor unit, while disallowing the treated unit to be used as a donor, one can generate a distribution of effect sizes for the placebo unit. If the size of the treatment effect for the treated unit is much larger than those generated for the untreated units, than it is unlikely that the estimated effect was the result of chance. If the distribution of placebo effects at time t is $\hat{\alpha}_{1t}^{PL} = \{\hat{\alpha}_{jt} : j \neq 1\}$, then the two-sided and one-sided p-values for period t are:

$$p - value_{2s} = Prob(|\hat{\alpha}_{1t}^{PL}| \geq |\hat{\alpha_{1t}}|)$$

$$p - value_{1s} = Prob(|\hat{\alpha} \geq \hat{\alpha_{1t}})$$

If some placebo units have poor matches than the p-values may be too conservative. Galiani and Quistorff (2016)[30] recommend two methods for adjusting p-values for the quality of pre-treatment fit. Firstly, donor units which exceed a certain pre-treatment root mean squared errors can be dropped from the distribution $\hat{\alpha}_{1t}^{PL}$ for the calculation of the p-values, or to divide all effects by the pre-treatment fit to generate "pre-treatment adjusted" p-values.

One final test for significance is the "placebo date test", where a model for the treated unit is estimated with the same parameters save for the treatment period. If the effects seen during the initial estimation are causally related to the treatment, then one should expect small and insignificant differences between the actual and synthetic unit following the placebo date.

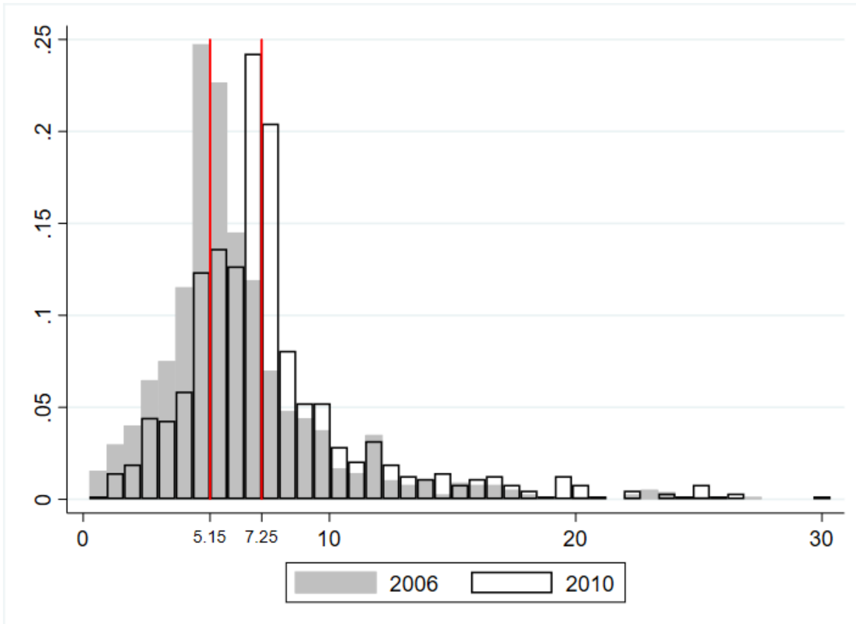

Figure 2.4: Distribution of Wages for Workers Aged 15 to 24 in Puerto Rico in 2006 and 2010

2.5 Results and Discussion

2.5.1 Teens/Young Adults

Full Donor Pool

Previous research on the minimum wage has used teenagers and other younger workers as a proxy for workers bound by the increase. In Puerto Rico, young workers between the ages of 15 and 24 were significantly more likely to be bound by the minimum wage increase, with 75% of workers aged 15 to 24 earning below $7.25 per hour in 2006 just before the passage of the Fair Minimum Wage Act of 2007, as seen in Figure 2.4.

Using the International Labor Organization's modeled estimate of the employment to population ratio for workers 15-24, as well as data from the World Bank on income per capita, GDP growth, and share of the population within the 15-24 year old range, I construct a synthetic control for Puerto Rico using 197 other countries as donors. The synthetic control algorithm, unsurprisingly, placed high weights on countries geographically close to Puerto

Rico (Suriname) or at similar levels of economic development (Gabon), but also, somewhat puzzlingly, placed a high weight on Norway (Tables 2.2 & 2.3). Puerto Rico experienced a substantial decline in teen and young adult employment relative to synthetic control, with employment in this group being, on average, 30.3% lower in Puerto Rico following the phase in of the minimum wage. In addition, Abadie's placebo test indicates high levels of significance, with results being significant at the 1% level following the completion of the minimum wage's phase in (Table B.1 and Figure 2.5). In addition, a placebo treatment date of 2000 for Puerto Rico yields treatment effects that are small and statistically insignificant; exactly what should be expected if the decline was related to the minimum wage increase (Table B.3). By dividing the estimated treatment effect on the log/employment population ratio, -0.303, by the percent increase in the minimum wage, 0.4, we can find the elasticity of teen/young adult employment to the minimum wage implied by these results to be -0.74

Table 2.2: Weights for Constructing Synthetic Puerto Rico

Country	Weight
Suriname	0.349
Norway	0.328
Macedonia	0.16
Gabon	0.16
Lesotho	0.003

A next step is to compare the treatment effects found for teens to those found for all workers aged 15 and older. If the decrease in employment relative to synthetic control for teens was due to the minimum wage increase, and not a local shock, we should expect smaller effects following the increase for all workers, for whom the minimum wage is less likely to bind. Since 35% of *all* Puerto Ricans earned $7.25/hr or less in 2010, compared to 58% of Puerto Ricans aged 15 to 24, they received a "dose" of the treatment that was 60% the size of the "dose" received by teens. With this in mind, we can combine synthetic

Table 2.3: Indicators in Puerto Rico vs. Synthetic Puerto Rico

Indicator	Puerto Rico	Synthetic Puerto Rico
GDP Growth	3.71%	2.18%
GDP per capita	$29,043	$25,731
Male Teen Population	9.0%	8.4%
Female Teen Population	8.2%	8.1%
Male Young Adult Population	8.25%	8.16%
Female Young Adult Population	7.73%	7.86%
$\ln(\frac{Employment}{Population})$ in 1991	3.03	3.08
$\ln(\frac{Employment}{Population})$ in 1995	3.18	3.21
$\ln(\frac{Employment}{Population})$ in 2000	3.33	3.25
$\ln(\frac{Employment}{Population})$ in 2005	3.14	3.17

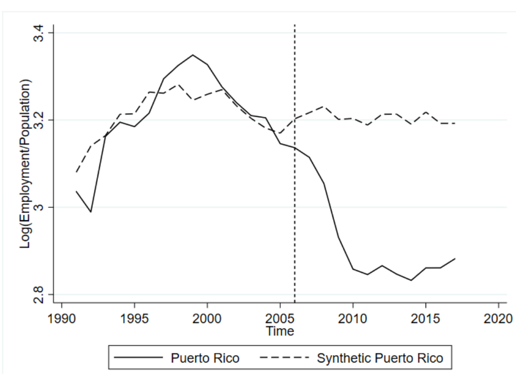

(a) Actual vs. Synthetic Control

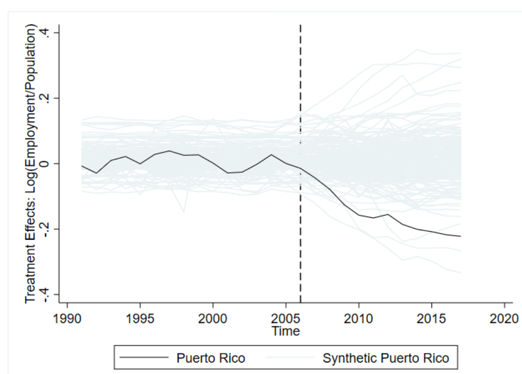

(b) Placebo Test

Figure 2.5: Treatment Effects and Placebo Test: Teen Employment (International Donors)

controls with difference-in-difference-in-differences (DDD) to generate a more accurate elasticity. Results showed that total employment in Puerto Rico was 11.1% lower than synthetic control following the minimum wage increase, with p-values ranging from 0.04 to 0.20, with treatment effects increasing in magnitude and significance as the minimum wage was phased in (Table B.4). To find the elasticity implied by this triple differences approach, we compute: $\varepsilon = \frac{(-0.303)-(-0.111)}{0.4-\frac{0.35}{0.58}(0.4)} = \frac{-0.192}{0.16} = -1.2$ The elasticity implied by the triple differences approach is -1.2, again substantially larger in magnitude than estimates from the mainland united states.

Limited Donor Pool

One alternative approach is to address concerns regarding the potential donor pool countries, given that Norway in particular seems an inappropriate control *ex ante*, by limiting the pool of donor countries to only those which are *a priori* sensible. In order to maintain a stock of placebo countries that is as large as possible, the donor pool was limited by dropping only inappropriate nations chosen by the synthetic control algorithm until the chosen donors for synthetic Puerto Rico had intuitive appeal. After dropping several western European countries, the synthetic control procedure placed high weights on the tropical island nations of Barbados and Comoros, with the remaining weight coming from Sri Lanka and the mainland United States (Tables B.4 & B.5). The limited donor pool reduced the size of the treatment effect, from an average of 30% to 16%, implying an elasticity of -0.4. Statistical significance varies, with evidence of a disemployment effect being strongest in the years following the completion of the phase-in of the minimum wage (p=0.04) (Table B.5). Finally, we can apply the triple differences approach to the limited donor pool by estimating treatment effects on total employment in Puerto Rico with a limited donor pool and subtracting. Total employment in Puerto Rico was 9.1% lower than the limited donor pool synthetic control following the phase-in of the minimum wage (Table B.6). Adjusting for the differential "bite" of the minimum wage across these two groups as above yields an

Table 2.4: Weights for Constructing Synthetic Puerto Rico (Limited Donor Pool)

Country	Weight
Barbados	0.48
Comoros	0.422
Sri Lanka	0.083
United States	0.015

Table 2.5: Indicators in Puerto Rico vs. Synthetic Puerto Rico (Limited Donor Pool)

Indicator	Puerto Rico	Synthetic Puerto Rico
GDP Growth	3.39	2.15
GDP per capita	29417.35	9061.07
Male Teen Population	8.98	9.41
Female Teen Population	8.12	8.96
Male Young Adult Population	8.23	8.68
Female Young Adult Population	7.71	8.31
$\ln(\frac{Employment}{Population})$ in 1991	3.36	3.37
$\ln(\frac{Employment}{Population})$ in 1995	3.39	3.41
$\ln(\frac{Employment}{Population})$ in 2000	3.49	3.43
$\ln(\frac{Employment}{Population})$ in 2005	3.33	3.37

elasticity of $\varepsilon = \frac{(-0.16)-(-0.091)}{0.4-\frac{0.35}{0.58}(0.4)} = \frac{-0.069}{0.16} = -0.43$.

GDP Matching

A threat to identification using the synthetic control method occurs when the treated unit experiences a unique shock at the same time as treatment. Since the phase-in of the new minimum wage coincided with the Great Recession of the late 2000s, it's possible that we are confusing the effects of the recession for the effects of the minimum wage increase. To test for this possibility, we can compare the path of GDP per capita in Puerto Rico to the synthetic Puerto Rico detailed in Table 2.2. When the two are compared in Figure 2.6, Puerto Rico's GDP grows faster than Synthetic Puerto Rico's during the pre-treatment period, while the two are mostly parallel during the post-treatment period. It's unclear from this examination alone whether or not the Great Recession is confounding the previous results.

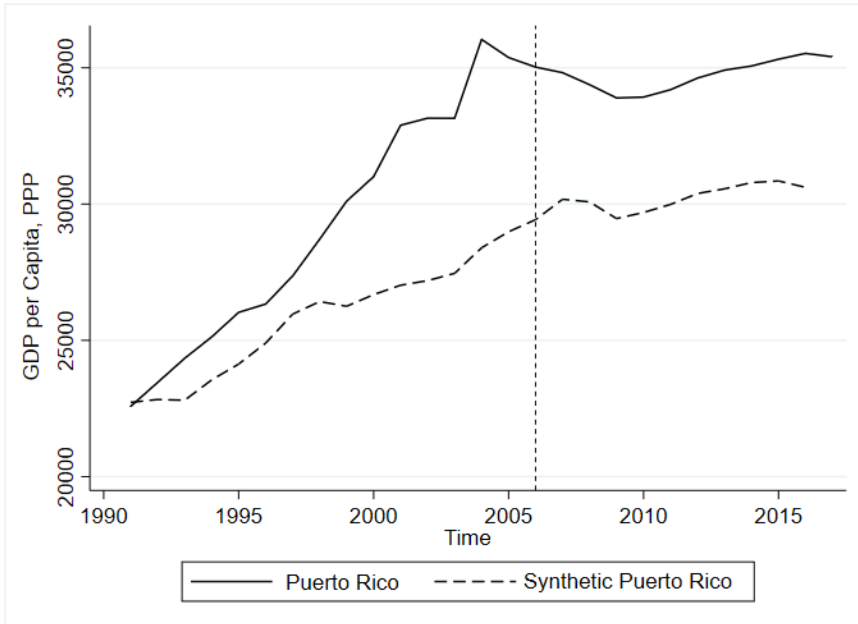

Figure 2.6: Path of GDP per Capita in Actual and Synthetic Puerto Rico

To try to account for the confounding effect of the Great Recession, we can change our strategy in creating synthetic controls. Instead of aiming to minimize the pre-treatment RMSE of employment, I create a synthetic Puerto Rico that minimizes the RMSE of GDP per capita throughout the entire sample period, pre- and post- treatment. This creates a counterfactual Puerto Rico that comes as close as possible to experiencing the same macroeconomic fluctuations as actual Puerto Rico (Table 2.6). Then, we can compare the path of employment in this new synthetic Puerto Rico to employment on the actual island.

As seen in Figure 2.7a and Table B.7, output in the "GDP matched" synthetic Puerto Rico closely follows that of the actual island. When comparing trends in employment in Figure 2.7b, teen/young adult employment in the "GDP matched" synthetic Puerto Rico is substantially higher than in actual Puerto Rico in *both* the pre-treatment and post-treatment periods. Thus, rather than calculating the treatment effect by taking the difference between synthetic and actual Puerto Rico in the post-treatment period, it's more appropriate in this case to compute a simple difference-in-differences estimator using the following linear

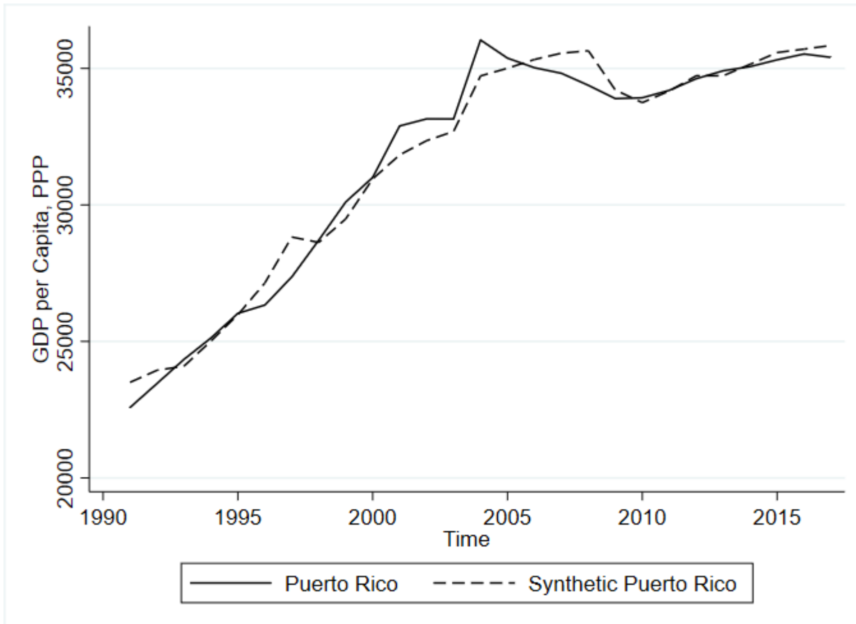

(a) Path of GDP

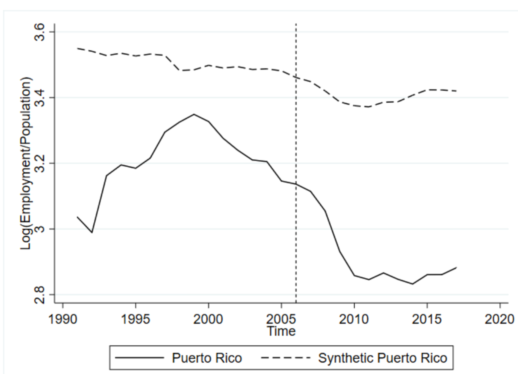

(b) Path of Teen/Young Adult Employment

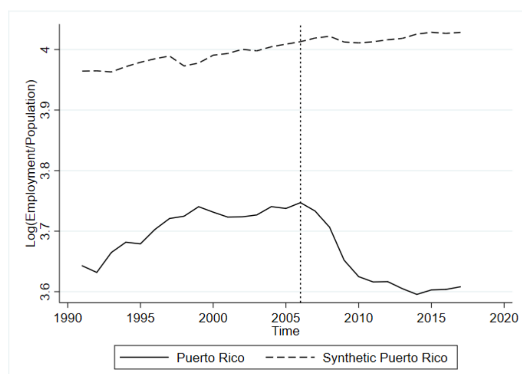

(c) Path of Total Employment

Figure 2.7: Comparisons Between Actual and Synthetic Puerto Rico (GDP Matching)

Table 2.6: Weights for Constructing Synthetic Puerto Rico (GDP Matching)

Country	Weight
United Arab Emirates	0.112
Equatorial Guinea	0.149
Ireland	0.046
Iraq	0.085
South Korea	0.312
Lebanon	0.190
Norway	0.106

regression model:

$$ln(\frac{employment}{population}_{it}) = \beta_0 + \beta_1(did_{it}) + \beta_2(post_t) + \beta_3(pri_i) + \varepsilon_i$$

where $pri_i = 1$ for the real Puerto Rico and 0 for synthetic Puerto Rico, $post_t = 1$ in years 2007 and later, and $did_{it} = pri_i * post_t$. The estimated treatment effect, β_1, was -0.196, with a standard error of 0.03. Thus, the implied elasticity of teen/young adult employment to the minimum wage is $\frac{-0.196}{0.4} = -0.49$. The "GDP matching" approach can also be used to estimate the effect on total employment, as seen in Figure 2.7c. Using the same difference-in-differences estimator from above yields an estimated treatment effect$\beta_1$for all adults of -0.108, with a standard error of 0.01. The implied elasticity of total employment to the minimum wage from this approach is $\frac{-0.108}{0.4} = -0.27$, and the implied "triple differences" elasticity is $\varepsilon = \frac{(-0.196)-(-0.108)}{0.4-\frac{0.35}{0.58}(0.4)} = \frac{-0.088}{0.16} = -0.55$.

2.5.2 Food Industry

Cross-Industry Comparisons

Recalling Krueger's finding that the weakest evidence for a disemployment effect of the minimum wage in Puerto Rico came from cross-industry comparisons, any exploration of the minimum wage's effect on Puerto Rico should utilize a similar technique. Additionally,

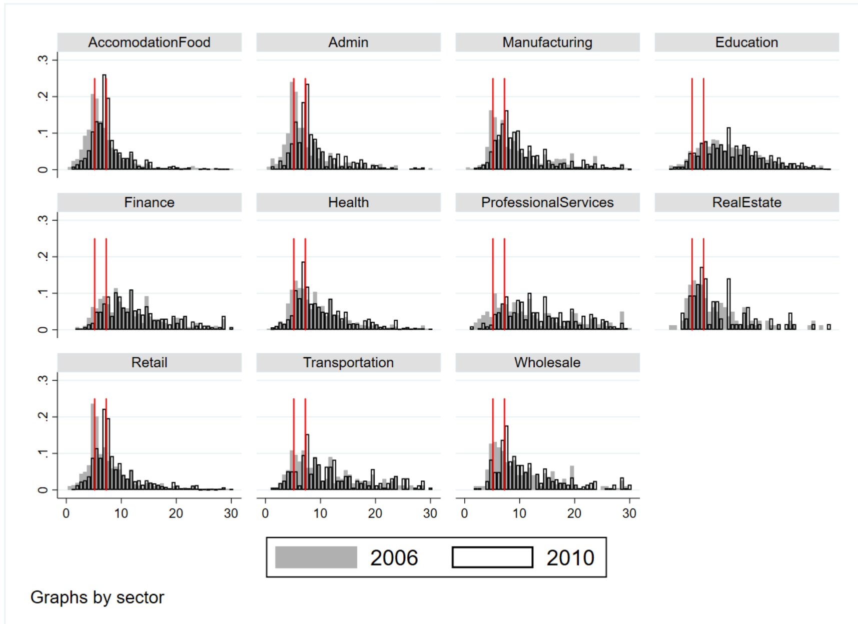

Figure 2.8: “Bite” of the Minimum Wage Across Puerto Rican Industries

since this approach only uses Puerto Rican industries as donors, there is less concern about Puerto Rican specific shocks contaminating the results. Figure 2.8 displays the "bite" of the minimum wage across Puerto Rican Industries. Using data from the Bureau of Labor Statistics' State and Area Employment Hours and Earnings (NAICS) program, my initial approach is to construct a synthetic control for the accommodation and food industry using all other island industries where the minimum wage is less likely to bind as donor.

The synthetic control algorithm constructed the synthetic accommodation and food industry using the health, retail, education, and professional services industries (Tables 2.7 & 2.8). Comparing the accommodation and food industry to its synthetic counterpart shows that total employment was, on average, 8.5% lower after that minimum wage was phased in. Unlike with results for teens, raw placebo p-values for the cross-industry synthetic controls generally failed to reach statistical significance ($0.07<p<0.31$). However, these results become significant or approach significance for all periods if the p-values are adjusted for the quality of the per-treatment fit (Table B.8, Figure 2.9). In addition, a placebo date

Table 2.7: Weights for Constructing Synthetic Accommodation and Food Industry (Puerto Rican Industry Donors)

Industry	Weight
Education	0.06
Health	0.587
Professional Services	0.187
Retail	0.22

Period	Accommodation & Food	Synthetic Accommodation & Food
Q1 1995	3.772761	3.814832
Q3 1997	3.972177	3.97396
Q1 2000	4.11741	4.117964
Q3 2002	4.149464	4.150361
Q1 2005	4.239887	4.227708

Table 2.8: Pre-Treatment Log Employment in Actual and Synthetic Accommodation and Food Industry (Puerto Rican Industry Donors)

test using Q4 2000 as the treatment date yielded treatment effects that were small and statistically insignificant (Table B.9).

To find the elasticity of accommodation and food industry employment to the minimum wage, we first need to find the coverage of the minimum wage in the constructed synthetic accommodation and food industry. By summing the products of each industry's share of the synthetic control and the share of workers earning \$7.25/hr or below in each industry, we can find that the synthetic accommodation and food industry had an effective coverage of 47%, compared to the actual accommodation and food industry's 69%. The elasticity of accommodation and food employment to the minimum wage implied by this approach is thus $\varepsilon = \frac{-0.085}{0.4-\frac{0.47}{0.69}(0.4)} = -0.66$.

2.5.3 Cross-City Comparisons

One additional strategy to identify the employment effect for bound industries is to construct a synthetic control for the accommodation and food industry in the San Juan metropolitan

(a) Actual vs. Synthetic Control

(b) Placebo Test (Donors with >2x PRI Pre Treatment RMSE Dropped)

Figure 2.9: Treatment Effects and Placebo Test: Accommodation and Food Industry (Island Industry Donors)

statistical area (MSA) using the same industry in MSAs on the mainland U.S. as donors. With this in mind, data on employment in the accommodation and food industry for San Juan and 99 other MSAs was obtained and combined with data from the 2000 census to construct the employment/population ratio for each MSA. The synthetic control procedure identified Visalia-Porterville, CA, Trenton, NJ, and Norwich-New London-Westerly, CT-RI as donors to construct San Jaun's synthetic accommodation and food industry (Tables 2.9 & 2.10). The employment/population ratio in San Juan's accommodation and food industry was found to be 9% lower than synthetic control on average following the phase in of the minimum wage. Fortunately, none of the three MSAs chosen as donors were bound by the minimum wage increase, so the implied elasticity is $\varepsilon = \frac{-0.09}{0.4} = -0.23$. The effects vary in their significance (0.02<p<0.37) depending on the post-treatment period, but this variance in significance may partially be the result of the fact that the data was not available with seasonal adjustments (Table B.10, Figure 2.10).

As an alternative, a synthetic Accommodation and Food industry using US MSAs was constructed using log employment, rather than the log employment/population ratio as the dependent variable of interest. In this case, the synthetic control procedure selected Miami FL, Tampa FL, Trenton NJ, and Tucson AZ as the donor cities comprising the synthetic San

Table 2.9: Weights for Constructing Synthetic Accommodation and Food Industry (USA Donors)

MSA	Weight
Visalia-Porterville, CA	0.665
Trenton, NJ	0.264
Norwich-New London-Westerly	0.071

Table 2.10: Pre-Treatment Log Employment/Population Ratio in Actual and Synthetic Accommodation and Food Industry (USA Donors)

Period	Accomodation and Food	Synthetic Accomodation and Food
Jan 2003	.663	.654
Jun 2003	.692	.707
Jan 2004	.732	.711
Jun 2004	.708	.735
Jan 2005	.749	.696
Jun 2005	.738	.735
Jan 2006	.765	.723
Jun 2006	.723	.750
Jan 2007	.710	.756
Jun 2007	.710	.798

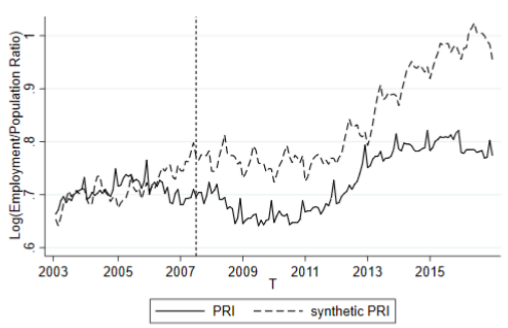

(a) Actual vs. Synthetic Control

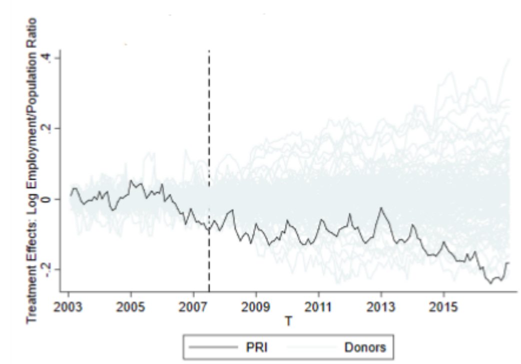

(b) Placebo Test (Donors with >2x PRI Pre Treatment RMSE Dropped)

Figure 2.10: Treatment Effects and Placebo Test: Accommodation and Food Industry (USA MSA Donors)

Table 2.11: Weights for Constructing Synthetic Accommodation and Food Industry (USA Donors)

MSA	Weight	A&F MW Coverage
Trenton, NJ	0.443	0%
Tampa-St. Petersburg, FL	0.106	37%
Miami-Ft. Lauderdale, FL	0.444	32%
Tucson, AZ	0.007	43%

Juan restaurant industry (Table 2.11). This specification also yielded a better pre-treatment fit than the previous synthetic control using the employment population ratio, without the concerning divergence between actual and synthetic employment observed in late 2006 prior to the wage increase that was previously observed. On average, employment in the San Juan restaurant industry was 4% lower than its synthetic counterpart, with effects also varying in significance depending on the post-treatment period ($0.01<p<0.93$) (Table B.8). Unlike the specification detailed in Table 2.9, some of the donor cities used to construct the synthetic San Juan restaurant industry were bound by the minimum wage increase. Using the values in Table 2.11, multiplying each donor city's coverage by its weight given by the synthetic control algorithm yields a synthetic minimum wage coverage of 18%, so the elasticity of employment to the minimum wage implied by this approach is $\varepsilon = \frac{-0.041}{0.4-\frac{0.18}{0.69}(0.4)} = -0.14$.

2.6 Discussion

Estimates of the elasticity of teen/young adult employment to the minimum wage in Puerto Rico, sumarized in Table 2.13, vary in magnitude and significance, but remain negative and at least modest in size. The average of all elasticities is -0.47, nearly equal equal to Castillo-Freeman's estimate of -0.41. Estimates of elasticity of accommodation and food employment to the minimum wage differ by method, with cross-industry comparisons yielding a higher elasticity than using US cities as donors. For comparison, Bellman & Wolfson (2019)[19] report the "consensus range" in the literature for the elasticity of teen employment to the minimum wage in the United States to be -0.13 to -0.07. Harasztosi & Lindner (2019)[31]

Table 2.12: Pre-Treatment Log Employment in Actual and Synthetic Accommodation and Food Industry (USA Donors)

Period	Accomodation and Food	Synthetic Accomodation and Food
Jan 2003	3.91	3.91
Jun 2003	3.94	3.94
Jan 2004	3.98	3.97
Jun 2004	3.96	3.98
Jan 2005	4.00	4.00
Jun 2005	4.01	3.99
Jan 2006	3.97	3.98
Jun 2006	3.97	3.98
Jan 2007	3.96	3.96
Jun 2007	3.96	3.98

report that 25% of teens in 2012 were directly affected by the minimum wage, so the "consensus range" of own-wage elasticities is -0.28 to -0.52 for the continental United States. Across all groups, the average own-wage elasticity of employment was -0.68, ranging from -0.22 to -1.65. The increase in the Federal Minimum wage from $5.15 to $7.25 appears to have depressed employment among *affected* workers in Puerto Rico to a greater degree than among *affected* workers in the continental United States.

There are multiple theoretical reasons why the *elasticity* of employment to the minimum wage may decrease at higher relative minimum wages. As Clemens (2021)[32] discusses, cuts in employment are only one way in which employers may respond to an increase in the minimum wage, but these alternative margins of adjustment "dry up" as the relative minimum wage grows higher. For example, while most of the literature on the employment effects of the minimum wage focuses on the "extensive" margin of employment, hiring and firing, there is emerging evidence that firms in the United States adjust along the "intensive" margin by cutting the number of hours each employee works. Jardim et al. (2017)[33] found that the minimum wage lead to a reduction of hours worked in Seattle, while also estimating a null effect on restaurant employment. Horton (2017)[34] also found that randomly imposed minimum wages lead to a large reduction in hours worked in an online labor market. An

Table 2.13: Estimated Own-Wage Elasticities of Employment

Model	Group	Method	Dependent Variable	MW Elasticity	MW Coverage	Own-Wage Elasticity
(1)	Ages 15-24	Synthetic Control	$\ln(\frac{employment}{population})$	-0.74	76%	-0.97
(2)	Ages 15-24	Triple Differences	$\ln(\frac{employment}{population})$	-1.26	76%	-1.65
(3)	Ages 15-24	Limited Donors	$\ln(\frac{employment}{population})$	-0.39	76%	-0.51
(4)	Ages 15-24	(2) & (3)	$\ln(\frac{employment}{population})$	-0.43	76%	-0.56
(5)	Ages 15-24	GDP Matching	$\ln(\frac{employment}{population})$	-0.49	76%	-0.65
(6)	Ages 15-24	(2) & (5)	$\ln(\frac{employment}{population})$	-0.55	76%	-0.72
(9)	All Workers	Synthetic Control	$\ln(\frac{employment}{population})$	-0.28	46%	-0.60
(10)	All Workers	Limited Donors	$\ln(\frac{employment}{population})$	-0.22	46%	-0.49
(11)	All Workers	GDP Matching	$\ln(\frac{employment}{population})$	-0.27	46%	-0.59
(13)	Restaurant Workers	Cross Industry	$\ln(employment)$	-0.66	69%	-0.95
(14)	Restaurant Workers	USA MSA Donors	$\ln(\frac{employment}{population})$	-0.23	65%	-0.35
(15)	Restaurant Workers	USA MSA Donors	$\ln(employment)$	-0.14	65%	-0.22

alternative response to cutting hours is to reduce non-cash “fringe benefits” afforded to workers such as health insurance or paid leave, a response documented by Clemens, Kahn, & Meer (2018)[35]. Since only 23.6% of Puerto Ricans receive employer-funded health insurance, compared to 49.6% of Americans in general (Census Bureau, 2019), employers in Puerto Rico were already less flexible in their ability to cut fringe benefits compared to employers in the United States. As the relative minimum wage climbs higher these alternative margins of adjustment become less capable of absorbing employers’ higher labor costs. Hours and fringe benefits can only go so low before employers are faced with the choice of either cutting back employment or shutting down entirely. Another mechanism by which the sensitivity of employment to the minimum wage increases at higher relative minimum wages has to do with the differential employment response across sectors of the economy. Research from Cengiz et al. (2019)[21] and Gopalan et al. (2020)[36] suggests that the employment response is significantly stronger in the tradable sector of the economy relative to the non-tradable sector. This heterogeneous response is likely due to higher product demand elasticity for the tradable sector limiting the ability of tradable firms to defray costs by raising output prices. In the mainland United States, the vast majority of workers bound by the minimum wage increase were concentrated in the non-tradable sector, where average wages are lower. In Puerto Rico, however, a quarter of manufacturing workers in 2010 earned \$7.25/hr or less (Table B.1), indicating that a greater share of affected workers in Puerto Rico belonged to the tradable sector where prices are less able to absorb the higher wage floor.

2.7 Conclusion

This paper contributes to the extensive literature on the employment effects of minimum wages by focusing on the 2007 increase in Puerto Rico's minimum wage, which led to a relative minimum wage for the island nation that was significantly higher than any found in the continental United States and thus affected the wages of a greater number of workers. Results indicate that employment in Puerto Rico fell relative to a data-constructed synthetic counterfactual following the phase-in of the higher minimum wage and employment for more affected subgroups like teens and restaurant workers fell more sharply proportional to their higher minimum wage coverage. Estimated elasticities of employment to the *minimum* wage for groups in Puerto Rico were higher than those found from studies of the mainland US, but elasticities of employment to the *own* wage were also larger than "consensus" estimates from the mainland. This may be due to the fact that firms are less able to adjust along non-employment margins such as hours or fringe benefits at higher relative minimum wages, and that higher relative minimum wages are more likely to bind in the tradable sector of the economy where output prices are less flexible. There still exist limitations to this particular strategy for estimating the effect of the minimum wage on employment in Puerto Rico. Limited availability of hourly wage data makes it difficult to assess the direct effect of the minimum wage on workers' incomes, and my analysis doesn't investigate alternative margins for adjustment such as output price increases, reductions in hours, or cuts to non-wage benefits.

The largest limitation of my analysis is my inability to completely rule out the effect of confounding shocks to Puerto Rico's labor market which coincided with the minimum wage increase. For example, between 1996 and 2006, Congress gradually phased out various tax incentives, most notably the possession tax credit under US Code section 936, conferred to Pueerto Rican companies, increasing the tax burden for many Puerto Rican firms[37]. While this phase-out occurred during the pre-treatment period, it's possible that its effects weren't fully felt until after the minimum wage increase in 2007. Similarly, as discussed in Part V, Section 1.3, cross country comparisons may fail to isolate the effect of the minimum wage increase from the effects of the Great Recession which hit Puerto Rico especially hard. However, several components of my analysis serve to control for Puerto Rican specific-shocks and still indicate a negative employment effect of the minimum wage. First, some of the comparisons in Part V, Section 2.1 exploit variation in the "bite" of the minimum wage across Puerto Rican industries, which would *all* be affected by Pureto Rican specific macroeconomic shocks. Secondly, in the results using international comparisons in Part V, Section 1, I find that the estimated treatment effect of the minimum wage was smaller for all workers than for teenage workers with higher levels of exposure to the minimum. The estimated employment elasticity from the "triple differences" approach was actually *larger* than the one obtained through international comparisons alone, the opposite of what would be expected if an island-specific shock was biasing the results. Finally, even when I create a synthetic Puerto Rico matching entirely on fluctuations in pre- and post- treatment output and growth, which assumes that the minimum wage increase had no effect on GDP levels or growth, I still find evidence of declines in employment.

Chapter 3: Confederate Monuments and Hate Crime

3.1 Introduction

Following the 2015 church shooting in Charleston, South Carolina, a national movement has urged policymakers to remove the over 700 standing monuments and other symbols of the Confederate States of America currently standing throughout the nation. As of 2019, more than 100 of these monuments had been removed (Birch & Brooks, 2019)[38]. In a June 2020 Press Conference announcing the take down of a statue of Robert E. Lee, Virginia Governor Ralph Northam summarized the intent behind the movement to removal symbols of the Confederacy from public spaces:

> "It's a time for Virginia, it's time for this country, to heal. And when there are symbols of divisiveness, such as these statues are, in order to heal that divisiveness, the statues need to come down" (State of Virginia, 2020)[39]

Economists often caution against judging policies by their intentions, rather than their results. Does the removal of Confederate monuments and other Confederate symbols help promote an agenda of unity or further fan the flames of divisiveness that the statues are said to represent? In this paper, I examine 54 monument removals that took place between 2009 and 2018 in 31 different locations. Using data from the FBI's annual supplement on hate crime, I implement a difference-in-differences research design using both OLS and negative binomial regression to estimate a positive and statistically significant increase in race-based hate crime following the removal of a Confederate monument in a census-designated place. Given the limited amount of data on covariates available at the level of census-designated places, I also construct a sample of 99 large U.S. cities which contain 18 of the 31 removals from the previous sample. This smaller sample allows me to control for a variety set of potential confounding variables, including shifts in real income and police employment. The city sample also allows me to account for a city's changing ideological landscape by including controls for Democrat's vote share in the most recent election, as well as Stephens-Davidowitz's (2014)[40] index of racial animosity constructed using Google searches. Examining heterogenieties in the treatment effect reveals that monuments removed in formally Confederate areas led to relatively larger spikes in hate crime than those removed from areas that were either Union states or not states at the time of the Civil War. The impact of monument removal is also dependent on the ordinal value of the monument being removed. Cities like Baton Rouge, which removed one monument while leaving five others standing, experience larger hate crime increases than cities like Boston, which removed the one and only monument in the city. [Interpretation] While my results can be used to justify leaving Confederate monuments standing, they also reinforce the link between Confederate monuments and racist ideology that motivates their removal in the first place.

3.2 Background and Previous Research

The first piece of federal hate crime legislation in the United States was Title I of the Civil Rights Act of 1968, which permitted federal prosecution of offenders who "willfully injures, intimidates or interferes with...any person because of his race, color, religion or national origin" (18 U.S.C. 245(b)(2)). The scope of federal hate crime laws were further expanded with the Violent Crime Control and Law Enforcement Act of 1994, which required increased penalties for crimes committed based on race, color, religion, national origin, ethnicity, or gender, and the Matthew Shepard and James Byrd, Jr. Hate Crimes Prevention Act of 2009, which added sexuality, gender identity, and disability status as protected categories (28 U.S.C. 994). The FBI began collecting data on "crimes that manifest evidence of prejudice based on race, gender or gender identity, religion, disability, sexual orientation, or ethnicity" with the Hate Crimes Statistics Act of 1990 (28 U.S.C. 534). In 2018, there were 7,120 hate crime incidents reported to the FBI, 4,047 (57%) of which were motivated by race, ethnicity, or ancestry (FBI, 2018).

Existing scholarship on the determinants of hate crime does not provide a clear-cut prediction as to the effects of Confederate monument removal on hate crime, with two main models of hateful behavior providing conflicting predictions. The first, the "group threat" model first developed by Blalock (1967)[41] posits that hate crimes occur as a reaction to a minority group threatening the privileged position occupied by a majority group in a society. Under a group threat framework, the removal of a Confederate monument represents a relative decline in the status of whites relative to nonwhites. In response, threatened white citizens increasingly victimize members of racial minorities in order to re-assert their claim to the privileged position in society that the monument represented. Thus, the group threat model predicts a *rise* in race-based hate crime following the removal of a Confederate monument in an area. The other primary model of hateful behavior, the "power-differential" model, makes predictions contrary to those of the "group threat" model. Under the power-differential model, hateful behavior is constrained by the ability of the minority group to resist such behavior. As the minority group rises in status relative to the majority group, majority group members become more hesitant to act on their bias. Thus, the power-differential model predicts a *fall* in race-based hate crime following the removal of a confederate monument, as the relative increase in the status of Black individuals makes it less likely that crimes which target them will be unnoticed, excused, or tolerated.

While empirical evidence exists supporting both models of hateful behavior, the literature in general has skewed towards vindicating the power-differential model. Most importantly, areas with larger and faster growing minority group populations typically experience *lower* levels of hate crime (Krueger & Pischke (1997)[42]u, Disha et al. (2014), Piatkowska et al. (2020)[43]). Valencia et al. (2019)[44] also provide evidence for the power-differential model by examining the effect of same-sex marriage legalization on the incidence of hate-crimes targeting LGBT individuals, concluding that hate crimes targeting LGBT individuals decline following the legalization of gay marriage. Nevertheless, there is still scholarship which connects a relative decline in status of a majority group to increased hate crime as the group threat model predicts. Mills (2019)[45] finds a positive relationship between anti-gay hate crime and the percentage of households in an area with same-sex unmarried couples, while Piatkowska et al. (2020)[43] conclude that “levels of Black out-group marriages with Whites are positively related to the Black hate crime victimization rate but not related to the incidence rate” (p. 105).

Empirical social science regarding Confederate monuments has thus far been limited, with Benjamin et al. (2020)[46] being one notable exception. The authors aim to examine the local government structures and demographic characteristics that best predict the removal of a monument in a given city, concluding that “the size of the Black population, the presence of a National Association for the Advancement of Colored People chapter, and the percentage of Democrats in a county in which a monument exists- as well as whether the monument exists in a state that constrains removal by legislative decree- best predict whether a Confederate monument will be taken down” (p. 237). In addition, they find that formally Confederate states are no more likely to see monument removals than non-confederate states (p. 241).

3.3 Data and Methods

For my initial investigation into the relationship between Confederate monument removal and hate crime, I began with data on Confederate Monuments’ locations, dates of erection, and dates of removal from the Southern Poverty Law Center’s report "Whose Heritage?". Next, using the FBI’s annual Hate Crime supplement to the Uniform Crime Reports, I constructed a panel data sample of the 4804 census-designated places that reported at least one race-based hate crime to law enforcement between 2004 and 2018. This sample contains 31 of the 37 places that removed monuments from 2014-2018, the remaining 6 had 0 reported hate crimes during the entire sample period, and thus had to be dropped due to perfect colinearity with fixed effects. Summary statistics for the sample of census-designated places are reported in Table 3.1.

While the sample of census-designated places captures as many removal events as is possible with a fixed effects research design, there are weaknesses to including so many

Table 3.1: Summary Statistics (Census-Designated Place Sample)

	(1) All Places	(2) Places with Population Data	(3) Treated Places
Total Incidents	0.59	0.69	8.57
(Count)	(3.09)	(3.52)	(17.38)
Population		43968	674403
		(189659)	(1442424)
Standing Monuments	0.03	0.04	2.1
(Count)	(0.41)	(0.48)	(0.07)
N	71955	49885	465

Standard Deviations in Parentheses

units. The first is the sheer rarity of hate crime outside of major metropolitan areas. Only about 5,000 of the 12,000 census designated places could be included in the sample, as the remaining places had zero reported incidents for the entire sample period. Additionally, about 2,000 of the 5,000 had only a single year where any incidents were reported, with only 46 places experiencing at least one hate crime for every single year of the sample. A second limitation is the lack of data on covariates available at the level of census-designated place. Changing demographic, economic, and political trends cannot be adequately be controlled for, potentially biasing the results.

To combat both of these issues, I also analyze a second, smaller panel of 98 major US cities for the same sample period. The advantage of the smaller sample of major metropolitan areas is the ability to include controls for formerly omitted variables. If a city experienced surging popularity of ideologies which emphasize historical injustices towards African Americans and the complicity of whites in modern systemic racism, then city officials would be more likely to choose to remove any standing Confederate monuments within city limits. Simultaneously, the "group threat" model of hate crime predicts that the relative increase in status of the minority group would result in higher levels of hate crime within those same cities. A naive statistical analysis would reveal a link between monument removals and hate crime when no causal relationship exists. Any attempt to examine the causal effect of monument removals on hate crime must account for this "backdoor path", but creating a quantitative measure of the popularity of newer anti racist ideology is difficult. The best measure that is widely available for cities in the sample is the share of the vote won by Democrats in the most recent election

Intra-city changes in the popularity of *racist* ideology could also confound my results. If racial animosity rises within a city for exogenous reasons, an increase in race-based hate crime will predictably follow. At the same time, policymakers may believe that removing any standing Confederate monuments can counteract the rise in hateful behavior, leading to another spurious correlation between monument removal and hate crime. While measuring racial animus is just as difficult as measuring the popularity of anti racist ideology, Stephens-Davidowitz (2014)[40] provide a novel way to measure the level of racial animosity within a city using search data from Google. By measuring the search rate for a racial epithet for African Americans within a city, Stephens-Davidowitz constructs a measure of racial animosity that correlates positively with opposition to interracial marriage, negatively with the percent of population with a bachelor's degree, and is associated with a reduction in votes for Barack Obama in 2008 relative to John Kerry in 2004. Using Google trends data, I reconstruct Stephens-Davidowitz's index for 78 of the 98 cities in my sample. Finally, I also include controls for each city's population, population density, fraction of the population that is African American, real income, and level of police employment. Table 3.2 contains summary statistics for the sample of cities.

Table 3.2: Summary Statistics (City Sample)

	(1) All Cities	(2) Treated	(3) Not Treated
Total Incidents	8.39	10.70	7.87
(Count)	(16.63)	(18.59)	(16.12)
Incident Rate	1.81	1.83	1.80
(Per 100,000 people)	(3.03)	(3.10)	(3.02)
Racial Animosity	20.41	21.53	20.07
(Searches)	(10.66)	(10.77)	(10.06)
Democrat Vote Share	56.10	58.55	55.54
(%)	(24.23)	(24.74)	(24.74)
Population	5.04	6.38	4.73
(Hundreds of Thousands)	(5.31)	(3.54)	(5.60)
Population Density	4.29	3.42	4.49
(Thousand/Sq. Mile)	(2.97)	(1.89)	(3.13)
African American Share	23.18	29.85	21.86
(%)	(18.40)	(18.99)	(17.93)
Real Income	34365.51	35631.82	34080.35
(2017 Dollars)	(8568.39)	(7241.61)	(8798.78)
Standing Monuments	1.07	2.91	0.66
(Count)	(2.49)	(3.19)	(2.09)
N	1470	270	1200

Standard Deviations in Parentheses

Table 3.3: Two-Way Fixed Effects Estimates (Place Sample)

	(1)	(2)	(3)
Cumulative Removals	0.298***	0.280***	0.311***
	(0.0432)	(0.0446)	(0.0576)
Population		0.000000135	-0.000000681***
		(8.24e-08)	(0.000000163)
Place Fixed Effects	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes
Place-Specific Time Trends	No	No	Yes
N	71949	49852	49852

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

3.4 Results

3.4.1 Main Results

Evidence from Census Designated Places

When investigating the effect of monument removal in census-designated places, I begin with a two-way fixed effect regression model of the following functional form:

$$\ln(hate_{it} + 1) = \beta_0 + \beta_1 R_{it} + \beta_2 Pop_{it} + \alpha_i + \lambda_t + T_{it}\delta + \varepsilon_{it}$$

Where $hate_{it}$ is the number of race-based hate crimes that occurred in place i in year t. R_{it} is the cumulative number of monument removals that have occurred in place i in year t, Pop_{it} is the population of place i in year t, α_i and λ_t are place and year fixed effects respectively, and T_{it} is a matrix of place-specific linear time trends. Alternatively, given the count nature of the incidents variable and the high level of dispersion, I also estimate a negative binomial regression with incidents as the dependent variable and the same predictors as the model above, using population as an exposure variable.

As seen in Tables 3.3 and 3.4, using both OLS and Negative Binomial Regression results

Table 3.4: Negative Binomial Regression Estimates (Place Sample)

	(1)	(2)
Cumulative Removals	0.302***	0.356***
	(0.0511)	(0.0725)
Population		-0.00000191***
		(8.48e-08)
Place Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes
N	49833	49833

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

suggest that the removal of a confederate monument is associated with a statistically significant increase in hate crime, even when place-specific linear time trends are added. Each monument removal increases the expected incidence of hate crime in a census-designated place by roughly 33%.

Evidence from Cities

Similar to my investigation of census-designated places, I estimate a two-way fixed effects research design of the following functional form:

$$\ln(hate_{it} + 1) = \beta_0 + \beta_1 R_{it} + \beta_2 DemShare_{it} + \beta_2 Animosity_{it} + \alpha_i + \lambda_t + T_{it}\delta + \varepsilon_{it}$$

Where $hate_{it}$ is the number of racially-biased hate crimes occurring in city i in year t, R_{it} is the cumulative number of monument removals that have occurred in city i in year t, α_i, $DemShare_{it}$ is the share of the vote won by democrats in the most recent presidential or senatorial election in city i in year t, $Animosity_{it}$ is Stephens-Davidowitz (2014)'s index of racial animosity, constructing using Google searches, $\lambda_t$are city and year fixed effects, respectively, and T_{it} is a matrix of control variables including population, population density, percent of the population which is African American, the level of police employment, and

Table 3.5: Two-Way Fixed Effects Estimates (City Sample)

	(1)	(2)	(3)	(4)	(5)
Cumulative Removals	0.409*	0.379*	0.494**	0.305	0.263
	(2.32)	(2.21)	(2.93)	(1.84)	(1.32)
Racial Animosity			-0.00662	-0.00948	-0.00890
			(-0.99)	(-1.67)	(-1.51)
Democrat Vote Share			-0.000793	-0.00317	-0.00191
			(-0.29)	(-1.13)	(-0.67)
City Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
City Controls	No	Yes	No	Yes	Yes
City-Specific Linear Time Trends	No	No	No	No	Yes
N	1470	1470	1155	1155	1155

t statistics in parentheses

Controls include Population, Density, African American Share, Real Income, and Police Employment

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

city-specific linear time trends, as well as a Negative Binomial Regression model of the same functional form with population also being used as the exposure variable.

As reported in Table 3.5 and 3.6, even when control variables are added, the results from cities are consistent with those found from investigating census-designated places. In both the OLS and negative binomial regression models, the coefficients are positive and economically significant in all specifications. Adding controls for either ideological changes or city demographics does not meaningfully effect the magnitude or significance of the estimated treatment effect, although the coefficient in the OLS model does lose statistical significance once *both* sets of controls are added. Adding city-specific linear time trends to both the OLS and negative binomial regression models also increases the standard error of the coefficient, but it remains positive and large even in these most restrictive specifications. The estimated treatment effects range from a 26% increase in the OLS model with city-specific time trends, to as high as a 40% increase.

Table 3.6: Negative Binomial Regression Estimates (City Sample)

	(1)	(2)	(3)	(4)
Cumulative Removals	0.331*	0.393*	0.349*	0.333
	(1.98)	(2.09)	(2.06)	(1.69)
Racial Animosity			-0.0101	-0.00958
			(-1.67)	(-1.55)
Democrat Vote Share			-0.00302	-0.00379
			(-0.76)	(-0.97)
City Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
City Controls	No	Yes	No	Yes
N	1470	1470	1155	1155

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

3.4.2 Subgroup Analysis

It's possible that the treatment effect of a monument removal may differ depending on characteristics of the area where the removal takes place. One potential source of heterogeneity is the historical affiliation of the area. Surprisingly, Confederate monuments were not only erected and removed in areas that were previously part of the Confederacy, with standing monuments being removed in places like Baltimore, Maryland that fought for the Union, as well as places like Tulsa, Oklahoma that were not states at the time of Civil War. To examine this heterogeneity, I split both the census-designated place and city samples into areas that did and did not secede during the Civil War and estimate both the OLS and negative binomial regression models for each subsample separately.

The effect of monument removal may also be heterogeneous by the margin on which the removal takes place. Some cities, like Boston, removed the one and only Confederate monument in their jurisdiction during the sample period, while others, like Baton Rouge, removed one monument while leaving five others standing. I generate a new variable that is equal to the ordinal value of the monument removed in treated years in treated cities and add it to both the OLS and Negative Binomial Specifications. For example, this variable

Table 3.7: Two-Way Fixed Effects Estimates, Confederate States (Place Sample)

	(1)	(2)	(3)
Cumulative Removals	0.254***	0.200***	0.427***
	(0.0509)	(0.0538)	(0.0697)
Population		0.000000924*	0.00000397***
		(0.000000379)	(0.00000112)
Place Fixed Effects	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes
Place-Specific Time Trends	No	No	Yes
N	17529	14574	14574

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

equals one in post-treated years for Boston and six in post-treated years for Baton Rouge. I then estimate the following model using both OLS and negative binomial regression:

Evidence from Census Designated Places

Of the 31 treated places in the place sample, 16 were in places that were part of the Confederacy during the Civil War. The results for previously Confederate places are given in Tables 3.7 and 3.8, while the results for non-Confederate places are given in Tables in C.1 and C.2. In the most restrictive specifications, the estimated treatment effect for previously Confederate places is about 21% larger than the estimated treatment effect for Union places.

Tables 3.9 and 3.10 contain the "margin of removal" analysis for the census designated place sample. The variable indicating the margin of removal is positive and statistically significant in all specifications, indicating that the magnitude of the increase in hate crime following a monument removal is highly dependent on how many monuments are present in a city to begin with. The model predicts that Boston's removal of a lone standing monument should increase hate crime incidence by 24% in the most restrictive specification, while Baton Rouge's removal of one of six monuments is predicted to increase hate crime by 44%, doubling the magnitude of the estimated treatment effect.

Table 3.8: Negative Binomial Regression Estimates, Confederate States (Place Sample)

	(1)	(2)
Cumulative Removals	0.283***	0.429***
	(0.0801)	(0.111)
Population		-0.00000265***
		(0.000000190)
Place Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes
N	14574	14574

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.9: Two Way Fixed Effects Estimates, Margin of Removal Effects (Place Sample)

	(1)	(2)	(3)
Cumulative Removals	0.0777	0.0736	0.177
	(0.57)	(0.54)	(1.56)
Margin of Removal	0.0710***	0.0700***	0.0455*
	(4.62)	(4.53)	(2.25)
Population		5.28e-08	-0.000000346**
		(0.42)	(-2.74)
City Fixed Effects	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes
Place-Specific Linear Time Trends	No	No	Yes
N	71949	49885	49885

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.10: Negative Binomial Regression Estimates, Margin of Removal Effects (Place Sample)

	(1)	(2)
Cumulative Removals	-0.0421	0.201*
	(-0.23)	(2.35)
Margin of Removal	0.0631***	0.0479***
	(4.73)	(3.47)
Population		-0.00000189***
		(-22.35)
City Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes
N	49833	49833

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Evidence from Cities

Similar to the place sample, 7 of the 18 monument removals in the city sample took place in states that either fought for the union or were not states during the time of the civil war (Maryland, Massachusetts, Missouri, California, Oklahoma, Kansas, and the District of Colombia). As above, I estimate the models from Tables 3.5 and 3.6 for cities in states the did and did not secede during the civil war separately.

Analyzing formally Confederate States separately in Tables 3.11 and 3.12 yields a treatment effect that is larger than the one obtained with the full sample, while the treatment effects using only Union states reported in Tables C.3 and C.4 is smaller and no longer statistically significant. The range of estimated treatment effects for Confederate states ranges from 35% to 52%, larger than the 26% to 40% range obtained when pooling all cities together.

Tables 3.13 and 3.14 report the margin of removal effects for the city sample. The results indicate that the treatment effect of monument removal is indeed heterogenous depending on the number of monuments that remain *standing* following a monument's removal. In both the OLS and Negative Binomial specifications, the "margin of removal" variable is positive

Table 3.11: Two-Way Fixed Effects Estimates, Previously Confederate States (City Sample)

	(1)	(2)	(3)	(4)	(5)
Cumulative Removals	0.496***	0.436**	0.525***	0.389*	0.353
	(4.21)	(3.13)	(5.18)	(2.56)	(1.44)
Racial Animosity			0.00456	-0.000155	0.00295
			(0.57)	(-0.02)	(0.32)
Democrat Vote Share			-0.00142	-0.00434	-0.000775
			(-0.35)	(-0.95)	(-0.24)
City Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
City Controls	No	Yes	No	Yes	Yes
City-Specific Linear Time Trends	No	No	No	No	Yes
N	510	510	450	450	450

t statistics in parentheses

Controls include Population, Density, African American Share, Real Income, and Police Employment

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.12: Negative Binomial Regression Estimates, Previously Confederate States (City Sample)

	(1)	(2)	(3)	(4)
Cumulative Removals	0.474**	0.400*	0.399*	0.354
	(3.16)	(1.99)	(2.45)	(1.62)
Racial Animosity			-0.00220	0.000468
			(-0.19)	(0.04)
Democrat Vote Share			-0.00562	-0.00628
			(-0.67)	(-0.80)
City Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
City Controls	No	Yes	No	Yes
N	510	510	450	450

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.13: Two Way Fixed Effects Estimates, Margin of Removal Effects (City Sample)

	(1)	(2)	(3)	(4)	(5)
Cumulative Removals	0.177	0.139	0.317	0.0804	0.0553
	(0.74)	(0.61)	(1.40)	(0.36)	(0.30)
Margin of Removal	0.0545*	0.0566**	0.0430*	0.0530*	0.0506*
	(2.55)	(2.69)	(2.10)	(2.43)	(2.22)
Racial Animosity			-0.00664	-0.00922	-0.00832
			(-0.99)	(-1.62)	(-1.41)
Democrat Vote Share			-0.000725	-0.00319	-0.00218
			(-0.27)	(-1.14)	(-0.76)
City Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
City Controls	No	Yes	No	Yes	Yes
City-Specific Linear Time Trends	No	No	No	No	Yes
N	1470	1470	1155	1155	1155

t statistics in parentheses

Controls include Population, Density, African American Share, Real Income, and Police Employment

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.14: Negative Binomial Regression Estimates, Margin of Removal Effects (City Sample)

	(1)	(2)	(3)	(4)
Margin of Removal	0.0752***	0.0800***	0.0740***	0.0744***
	(4.42)	(4.82)	(4.27)	(4.08)
Cumulative Removals	0.00399	0.0442	0.0255	0.00639
	(0.02)	(0.20)	(0.12)	(0.03)
Racial Animosity			-0.00935	-0.00895
			(-1.49)	(-1.40)
Democrat Vote Share			-0.00308	-0.00385
			(-0.78)	(-0.99)
City Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
City Controls	No	Yes	No	Yes
N	1470	1470	1155	1155

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

and statistically significant in all specifications, even when city-specific linear trends are added. Removing a lone monument is predicted to increase the incidence of race-based hate crime by only 5-7%, while removing one of 5 monuments is predicted to result in a 25-35% increase.

3.4.3 Robustness and Falsification Tests

As with all difference-in-differences estimators, a causal interpretation of my results depends on the assumption that the cities in the sample that did not remove their monuments serve as an appropriate counterfactual for those which did. If the observed increase in hate crime really is being driven by the removal of confederate monuments, then we should expect to see no difference in hate crime between treated and non-treated cities in the years preceding removal. While including area-specific time trends is one way to account for this bias, the parallel trends assumption can also be confirmed visually by creating a variable indicating the year that treated cities removed their monuments, and then regressing hate crime rates on a series of lags and leads of this variable. If the leads of the switching variable are positive, then it is likely that the parallel trends assumptions does not hold.

The results of this âevent studyâ analysis visualized in Figure 3.1 show no statistically significant difference between treated and untreated states in the years prior to removal, which strengthens the case for an underlying causal relationship between monument removal and increases in hate crime.

Another way to reinforce a causal interpretation of the above results is using a "placebo" regression of the rate of *all* violent crimes on monument removal, using the following functional form:

$$\ln(crimerate_{it}) = \beta_0 + \beta_1 Removals_{it} + \alpha_i + \lambda_t + T_{it}\delta + \varepsilon_{it}$$

where $crime_{it}$ is the violent crime rate in city i in year t, and the remaining variables are defined above. As an additional robustness test, the $crime_{it}$ variable can be replace with the rates of its component crimes: Homicide, Rape, Robbery, and Assault. If the removal

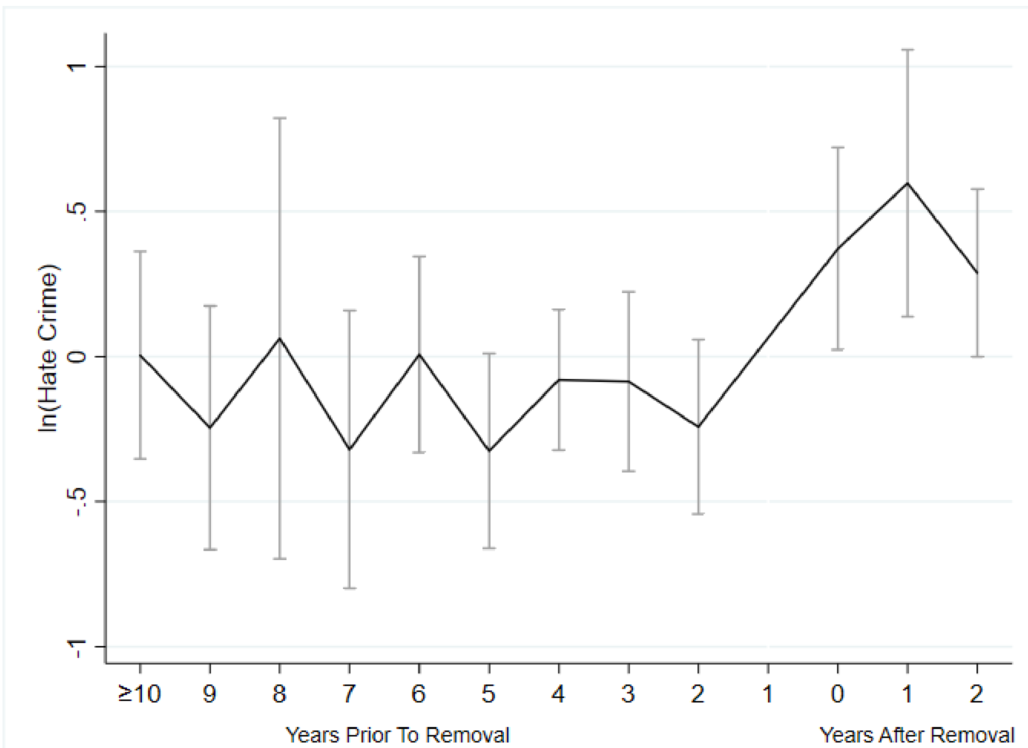

Figure 3.1: Event Study of Monument Removal Effects on Log Hate Crime

Table 3.15: The Effect of Monument Removal on Violent Crime Rates

	(1) All Violent Crime	(2) Homicide	(3) Rape	(4) Robbery	(5) Assault
Cumulative Removals	0.00824	-0.0298	0.00526	0.0410	-0.0317
	(0.36)	(-0.84)	(0.14)	(1.37)	(-1.20)
City Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
City Controls	Yes	Yes	Yes	Yes	Yes
N	1470	1470	1470	1470	1470

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

of Confederate monuments, rather than an omitted variable, is truly behind the relative increase in hate crime in treated states, then we would expect to find no effect of removal on violent crime as a whole. Conversely, if we *do* observe a relationship between monument removal and overall crime rates, we should be cautious in making a causal interpretation of any results.

Reassuringly, coefficients on the treatment variable reported in Table 3.15 are generally small and not statistically significant, which strengthens the case that the observed difference in hate crimes is indeed due to Confederate monument removal.

Finally, econometricians such as Cameron et al. (2008)[47] and Cameron Miller (2015)[48] note that cluster-robust standard errors have the potential to over-reject when the number of treated units is small, and recommend the use of cluster bootstrapping in order to lower the rejection rate. Since only 18 of the 98 cities ever experience a removal during the sample period, tables C.5 and C.6 in Appendix 3 replicate tables 3.5 and 3.11 using cluster bootstrapped standard errors. While the T-statistics using cluster bootstrapping are smaller than their counterparts, the coefficients on the treatment variable remains statistically significant in most specifications.

3.5 Limitations

The conclusions that can be drawn about the connection between Confederate monument removal and hate crime are partially limited by the depth and breadth of the data available. Since most of the take down incidents in the data set occurred after 2015, there are not enough post-treatment observations to draw conclusions about the dynamic effects of monument removal on hate crime. For example, it may be the case that the observed increase in hate crime "fades out" with time. Another limitation is that the FBI's hate crime statistics at the city level for the period 2004-2018 only disaggregate the total number of reported incidents by bias type *or* offense type, not both. It could be, for example, that the increase in racially-biased hate crimes is composed entirely of non-violent offenses like racist graffiti or vandalism. Additionally, city-level data does not contain a more detailed breakdown of bias incidents beyond broad categories like "Race-Based", "Religion-Based", "Gender Identity-Based", etc. for the entire sample period. For this reason, we cannot know whether the observed increase in hate crime is composed mostly of anti-Black hate crime or race-based hate crimes based on another bias.

3.6 Discussion and Conclusion

This paper exploits geographic and temporal variation in the removal of standing Confederate monuments to identify the effect that removing these monuments has had on the incidence of race-based hate crime between 2004 and 2018. Using both a large sample of census designated places, as well as a smaller sample of large cities with a richer set of covariates available, I estimate difference-in-differences models using both OLS and negative binomial regression and find that monument removal is associated with an increase in race-based hate crime in a city following a monument removal that is statistically significant in all but the most restrictive specifications. The result is often robust to the inclusion of area-specific linear time trends and leading coefficients in an event study analysis are statistically insignificant, which strengthens a causal interpretation. Additionally, monument removal is not found to increase violent crime overall, strengthening the observed link between monument removal and hate crime. The treatment effect is also heterogeneous along two margins. Firstly, areas that were part of the Confederacy during the civil war generally experienced larger hate crime increases than areas that were not. Secondly, the treatment effect differed depending on the number of monuments left standing following a removal. Area that removed one of many monuments had relatively larger spikes in hate crime compared to areas that removed the one and only standing monument. Most specifications report that each monument removal increases the incidence of race-based hate crime by roughly 30%, with this estimate being closer to 10% in formally Union states and closer to 40% in formally Confederate states.

The most straightforward interpretation of these results is that they vindicate the "group threat" model of hateful behavior. Removing a monument to Robert E. Lee can be viewed as a blow to the status of whites, men, Virginians, or even modern conservatives, while boosting the relative status of Blacks, northerners, or even modern progressives. Blalock (1967)'s model predicts that white conservatives will react to this by attempting to reinforce their privileged position in society through hateful behavior, which fits the results we observe. The fact that hate crimes are more sensitive to monument removal in formally Confederate areas strengthens this interpretation, as does my finding that hate crimes increase more following the removal of one of many monuments.

However, there are limitations to interpreting my results under the lens of the group threat model. Controlling for Stephens-Donowitz's (2014)[40] index of racial animosity does not meaningfully change the estimated treatment effect, which should not be the case if monument removals primarily increase hate crime by increasing racial animosity among dominant groups. Another issue is the fact that, in general, empirical research on hate crime has vindicated the "power differential" model over the group threat model. It could be that the group threat model is more valid in the case of *decreasing* ingroup status than it is in the case of *increasing* outgroup status, but more research examining exogenous changes in group status is necessary to verify this.

Another explanation for this phenomenon would be that groups which value the existence of the monuments engage in such behavior in order to make a credible threat of inflicting costs on society unless the remaining monuments are left standing. If this theory is true, that groups engage in increased levels of hate crime after a monument removal to make a credible threat of inflicting social costs should further monuments be removed, then we expect differing effects on hate crime depending on the number of monuments left standing in an area following a removal. If the sole monument in a city is removed, then hate groups have less to gain by retaliating; the damage is already done. However, if one of many monuments is removed, then groups have a much stronger incentive to inflict social costs: doing so sends a message to lawmakers and the voting public that further removal of monuments will result in more mayhem. The results presented in table X support this perspective: the increase in hate crime is larger when one monument is removed but others remain standing.

It is important to note that my results can be used normatively to justify both removing standing monuments as well as leaving them in place. While some of the debate regarding Confederate monuments has been about noticeable consequences of removal, the majority has been about what the monuments themselves represent. Anti-removal voices (see Nelson (2017) and Leigh (2019)) portray monuments to the Confederacy as historical artifacts first and foremost, with removing them being tantamount to forgetting a critical, albeit painful, part of our nation's history. Pro-removal perspectives (see SPLC (2019)) emphasize the link between Confederate monuments and the the ideological animus of the Confederacy, white supremacy. With the debate framed in these terms, my results are more consistent with the perspective of the *pro* removal advocates, as the link between monument removal and hate crime frames these monuments as intricately linked with racist ideology, rather than being benign historical artifacts. On the other hand, as the stated goal of removal is to reduce racial tensions, removal also appears counterproductive to a certain degree. Future policymakers will have to weigh the symbolic benefits of removal against the more tangible costs.

Appendix A: Supplementary Tables for Chapter 1

Table A.1: Time-Corrected Wald Estimates of Uber's Effect on the Unemployment Rate (Truncated)

Months Before/After Uber's Arrival	Estimated Effect	Bootstrapped Standard Error
-5	0.006	0.019
-4	-0.007	0.021
-3	0.041	0.018
-2	0.045	0.024
-1	-0.025	0.022
0	-0.030	0.021
1	-0.060	0.037
2	-0.059	0.041
3	-0.068	0.045
4	-0.070	0.046
5	-0.079	0.049
6	-0.103	0.048
7	-0.109	0.053
8	-0.109	0.053
9	-0.064	0.056
10	-0.058	0.042
11	-0.080	0.038
12	-0.103	0.043
13	-0.152	0.058
14	-0.197	0.063
15	-0.174	0.068

Appendix B: Supplementary Tables for Chapter 2

Table B.1: "Bite" of the Minimum Wage Across Groups

Group	% Earning 7.25/hr or less in 2006	% Earning 7.25/hr or less in 2010
All Puerto Ricans	46%	35%
Puerto Ricans Aged 15-24	76%	58%
Accommodation and Food Workers	69%	55%
Accommodation and Food Workers (San Juan)	65%	50%
Retail Workers	62%	47%
Education Workers	20%	17%
Health Workers	43%	37%
Manufacturing Workers	39%	25%
Professional Services Workers	37%	46%

Table B.2: Treatment Effects and Significance for Log Teen Employment/Population Ratio in Puerto Rico

Period	Estimated Effect	P-Value	Adjusted P-Value
2007	-0.065	.221	.662
2008	-0.082	.205	.610
2009	-0.137	.103	.558
2010	-0.238	.036	.426
2011	-0.350	.010	.338
2012	-0.341	.010	.359
2013	-0.349	.041	.354
2014	-0.343	.052	.374
2015	-0.339	.067	.4
2016	-0.314	.098	.436
2017	-0.300	.103	.477

Table B.3: Treatment Effects and Significance for Log Teen Employment/Population Ratio in Puerto Rico (Placebo Treatment Date of 2000)

Period	Estimated Effect	P-Value	Adjusted P-Value
2001	-0.031	.323	.810
2002	-0.046	.379	.795
2003	-0.047	.467	.851
2004	-0.061	.354	.800
2005	-0.081	.338	.784
2006	-0.152	.107	.677

Table B.4: Treatment Effects and Significance for Log Total Employment/Population in Puerto Rico

Period	Estimated Effect	P-Value	Adjusted P-Value
2007	-0.043	.312	.451
2008	-0.063	.179	.410
2009	-0.114	.046	.267
2010	-0.132	.041	.282
2011	-0.124	.046	.312
2012	-0.108	.097	.369
2013	-0.101	.138	.431
2014	-0.119	.102	.364
2015	-0.111	.144	.405
2016	-0.097	.201	.451
2017	-0.097	.200	.451

Table B.5: Treatment Effects and Significance for Log Teen Employment/Population Ratio in Puerto Rico (Limited Donor Pool)

Period	Estimated Effect	P-Value	Adjusted P-Value
2007	-0.058	.244	.694
2008	-0.101	.178	.644
2009	-0.171	.111	.550
2010	-0.217	.078	.527
2011	-0.263	.044	.439
2012	-0.169	.138	.605
2013	-0.190	.122	.594
2014	-0.153	.222	.617
2015	-0.101	.389	.783
2016	-0.117	.316	.750
2017	-0.048	.656	.894

Table B.6: Treatment Effects and Significance for Log Total Employment/Population Ratio in Puerto Rico (Limited Donor Pool)

Period	Estimated Effect	P-Value	Adjusted P-Value
2007	-0.037	.294	.567
2008	-0.059	.239	.478
2009	-0.097	.133	.378
2010	-0.113	.106	.367
2011	-0.106	.128	.439
2012	-0.084	.183	.500
2013	-0.075	.250	.561
2014	-0.092	.206	.494
2015	-0.087	.217	.522
2016	-0.083	.239	.511
2017	-0.089	.228	.522

Table B.7: Indicators in Puerto Rico vs. Synthetic Puerto Rico (GDP Matching)

Indicator	Puerto Rico	Synthetic Puerto Rico
GDP per capita in 1991	22579.09	23501.23
GDP per capita in 1995	26027.23	25976.47
GDP per capita in 2000	31005.18	30950.21
GDP per capita in 2003	33147.79	32690.81
GDP per capita in 2005	35375.3	35005.24
GDP per capita in 2007	34820.77	35557.17
GDP per capita in 2009	33892.85	34206.46
GDP per capita in 2010	33924.06	33746.06
GDP per capita in 2011	34195.65	34173.33
GDP per capita in 2013	34913.63	34731.05
GDP per capita in 2015	35314.43	35581.21
GDP per capita in 2017	35403.01	35840.44

Table B.8: Treatment Effects and Significance for Log Employment in Puerto Rican Accommodation and Food Industry (Puerto Rican Industry Donors)

Period	Estimated Effect	P-Value (One Sided)	Adjusted P-Value
Q4 2007	-0.0502	.214	0
Q1 2008	-0.0673	.071	0
Q2 2008	-0.0590	.214	.071
Q3 2008	-0.0764	.071	0
Q4 2008	-0.0887	0	0
Q1 2009	-0.1049	.071	0
Q2 2009	-0.1044	.142	0
Q3 2009	-0.1081	.214	0
Q4 2009	-0.0998	.285	.071
Q1 2010	-0.0906	.214	.142
Q2 2010	-0.0829	.285	.142
Q3 2010	-0.0916	.214	.071
Q4 2010	-0.0879	.214	.071
Q1 2011	-0.0775	.214	.071
Q2 2011	-0.0901	.285	.071
Q3 2011	-0.0972	.214	.071
Q4 2011	-0.0806	.357	.071
Q1 2012	-0.0868	.214	.071
Q2 2012	-0.0851	.285	.071
Q3 2012	-0.0751	.285	.071

Table B.9: Treatment Effects and Significance for Log Employment in Puerto Rican Accommodation and Food Industry (Placebo Treatment Date of Q4 2000)

Period	Estimated Effect	P-Values (One Sided)	Adjusted P-Values
Q1 2001	-0.0111	.285	.642
Q2 2001	0.0079	.642	.714
Q3 2001	-0.0103	.214	.642
Q4 2001	-0.0183	.214	.571
Q1 2002	-0.0491	.214	.142
Q2 2002	-0.0567	.357	.214
Q3 2003	-0.0352	.357	.428
Q4 2003	-0.0437	.428	.357
Q1 2004	-0.0356	.428	.428
Q2 2004	-0.0439	.428	.428
Q3 2004	-0.0417	.357	.571
Q4 2004	-0.0524	.428	.357
Q1 2005	-0.0304	.357	.571
Q2 2005	-0.0450	.285	.428
Q3 2005	-0.0649	.428	.285
Q4 2005	-0.0413	.357	.500
Q1 2006	-0.0337	.357	.571
Q2 2006	-0.0218	.357	.714
Q3 2006	-0.0117	.357	.857
Q4 2006	-0.0064	.5	.857
Q1 2007	-0.0120	.357	.785

Table B.10: Treatment Effects and Significance for Log Employment/Population Ratio in Puerto Rican Accommodation and Food Industry (USA Donors) (Truncated to Quarterly for Space)

Period	Estimated Effect	P-Value (One-Sided)	Adjusted P-Value
Jul-07	-.0754	0.049	0.0693
Oct-07	-.0892	0.049	0.049
Jan-08	-.0417	0.2376	0.5445
Apr-08	-.0830	0.0693	0.2277
Jul-08	-.1050	0.039	0.2079
Oct-08	-.1256	0.029	0.1485
Jan-09	-.0863	0.1089	0.3960
Apr-09	-.1110	0.0990	0.3168
Jul-09	-.1180	0.1089	0.2277
Oct-09	-.0911	0.1386	0.3069
Jan-10	-.0760	0.1782	0.4158
Apr-10	-.1021	0.1485	0.3663
Jul-10	-.1276	0.0792	0.2277
Oct-10	-.1220	0.0891	0.3267
Jan-11	-.0579	0.3267	0.6336
Apr-11	-.0918	0.1782	0.4851
Jul-11	-.1058	0.1584	0.4158
Oct-11	-.0780	0.2277	0.5841
Jan-12	-.0749	0.2475	0.6336
Apr-12	-.098	0.1683	0.5049
Jul-12	-.1172	0.1287	0.3861
Oct-12	-.0727	0.2772	0.6336
Jan-13	-.0416	0.5544	0.7425
Apr-13	-.0919	0.1881	0.5049
Jul-13	-.1155	0.1386	0.3762
Oct-13	-.1157	0.1584	0.4059
Jan-14	-.0829	0.2574	0.4653
Apr-14	-.1382	0.1386	0.3564
Jul-14	-.1579	0.0891	0.3168
Oct-14	-.1520	0.1188	0.3564

Table B.11: Treatment Effects and Significance for Log Employment in Puerto Rican Accommodation and Food Industry (USA Donors) (Truncated to Quarterly for Space)

Period	Estimated Effect	P-Value (One-Sided)	Adjusted P-Value
Jul-07	-0.008	0.67	0.72
Oct-07	-0.026	0.29	0.31
Jan-08	-0.013	0.60	0.65
Apr-08	-0.063	0.08	0.04
Jul-08	-0.041	0.22	0.23
Oct-08	-0.033	0.41	0.43
Jan-09	-0.061	0.22	0.24
Apr-09	-0.039	0.33	0.36
Jul-09	-0.029	0.47	0.53
Oct-09	-0.011	0.73	0.76
Jan-10	-0.036	0.46	0.48
Apr-10	-0.046	0.37	0.41
Jul-10	-0.054	0.31	0.34
Oct-10	-0.035	0.5	0.54
Jan-11	-0.051	0.39	0.41
Apr-11	-0.040	0.54	0.58
Jul-11	-0.026	0.69	0.75
Oct-11	-0.005	0.93	0.94
Jan-12	-0.046	0.51	0.54
Apr-12	-0.042	0.50	0.54
Jul-12	-0.026	0.68	0.71
Oct-12	0.016	0.83	0.87
Jan-13	-0.027	0.73	0.79
Apr-13	-0.031	0.61	0.63
Jul-13	-0.03	0.62	0.65
Oct-13	-0.008	0.88	0.91
Jan-14	-0.051	0.50	0.5
Apr-14	-0.061	0.39	0.38
Jul-14	-0.054	0.43	0.48
Oct-14	-0.072	0.38	0.36

Appendix C: Supplementary Tables for Chapter 3

Table C.1: Two-Way Fixed Effects Estimates, Union States (Place Sample)

	(1)	(2)	(3)
Cumulative Removals	0.391***	0.373***	0.107
	(0.0733)	(0.0764)	(0.0930)
Population		0.000000105	-0.000000779***
		(8.66e-08)	(0.000000168)
Place Fixed Effects	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes
Place-Specific Time Trends	No	No	Yes
N	54420	35278	35278

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.2: Negative Binomial Regression Estimates, Union States (Place Sample)

	(1)	(2)
Cumulative Removals	0.345***	0.353**
	(0.0666)	(0.108)
Population		-0.00000173***
		(8.74e-08)
Place Fixed Effects	Yes	Yes
Year Fixed Effects	Yes	Yes
N	35259	35259

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.3: Two-Way Fixed Effects Estimates, Union States (City Sample)

	(1)	(2)	(3)	(4)	(5)
Cumulative Removals	0.340	0.294	0.472	0.206	-0.0462
	(0.51)	(0.51)	(0.74)	(0.39)	(-0.15)
Racial Animosity			-0.0117	-0.0128	-0.0135
			(-1.32)	(-1.97)	(-1.83)
Democrat Vote Share			-0.0000362	-0.00228	-0.00221
			(-0.01)	(-0.71)	(-0.49)
City Fixed Effects	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes
City Controls	No	Yes	No	Yes	Yes
City-Specific Linear Time Trends	No	No	No	No	Yes
N	960	960	705	705	705

t statistics in parentheses

Controls include Population, Density, African American Share, Real Income, and Police Employment

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.4: Negative Binomial Regression Estimates, Union States (City Sample)

	(1)	(2)	(3)	(4)
Cumulative Removals	0.122	0.183	0.190	0.140
	(0.17)	(0.26)	(0.26)	(0.20)
Racial Animosity			-0.0123	-0.0124
			(-1.73)	(-1.70)
Democrat Vote Share			-0.00203	-0.00346
			(-0.51)	(-0.92)
City Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
City Controls	No	Yes	No	Yes
N	960	960	705	705

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.5: Effect of Monument Removal on Log Racial Hate Crime Incidents (Cluster Bootstrap)

	(1)	(2)	(3)	(4)
Cumulative Removals	0.409*	0.379*	0.389*	0.305
	(0.184)	(0.181)	(0.184)	(0.178)
Racial Animosity			-0.00984	-0.00948
			(0.00534)	(0.00554)
Democrat Vote Share			-0.00195	-0.00317
			(0.00259)	(0.00273)
City Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
City Controls	No	Yes	No	Yes
N	1470	1470	1470	1470

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.6: Effect of Monument Removal on Log Racial Hate Crime Incidents in Confederate States (Cluster Bootstrap)

	(1)	(2)	(3)	(4)
Cumulative Removals	0.496**	0.436*	0.436*	0.389
	(0.171)	(0.187)	(0.182)	(0.211)
Racial Animosity			-0.00271	-0.000155
			(0.00962)	(0.0100)
Democrat Vote Share			-0.00327	-0.00434
			(0.00482)	(0.00512)
City Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
City Controls				
N	1470	1470	1470	1470

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Bibliography

[1] A. J. Ravenelle, *Hustle and Gig.* University of California Press, 2019.

[2] M. C. Munger, *Tomorrow 3.0.* Cambridge University Press, mar 2018.

[3] P. Oyer, "The gig economy," *IZA World of Labor*, 2020.

[4] C. de Chaisemartin, X. D'Haultfoeuille, and Y. Guyonvarch, "Stata module to estiamte sharp difference-in-difference designs with multiple groups and periods," *Statistical Software Comparents S458643, Boston College of Economics*, 2019.

[5] B. Callaway and P. H. Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, vol. 225, no. 2, pp. 200–230, dec 2021.

[6] A. Abadie, A. Diamond, and J. Hainmueller, "Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program," *Journal of the American Statistical Association*, vol. 105, no. 490, pp. 493–505, jun 2010.

[7] J. D. Hall, C. Palsson, and J. Price, "Is uber a substitute or complement for public transit?" *Journal of Urban Economics*, vol. 108, pp. 36–50, nov 2018.

[8] J. M. Barrios, Y. V. Hochberg, and H. Yi, "Launching with a parachute: The gig economy and new business formation," *Journal of Financial Economics*, vol. 144, no. 1, pp. 22–43, apr 2022.

[9] A. Tsotsis, "Uber opens up platform to non-limo vehicles with uber x, service will be 35% less expensive," *TechCrunch*, Jul. 2012.

[10] J. V. Hall and A. B. Krueger, "An analysis of the labor market for uberâs driver-partners in the united states," *ILR Review*, vol. 71, no. 3, pp. 705–732, 2018.

[11] C. Brandtner, A. Lunn, and C. Young, "Spatial mismatch and youth unemployment in US cities: public transportation as a labor market institution," *Socio-Economic Review*, vol. 17, no. 2, pp. 357–379, mar 2017.

[12] D. Hernandez, M. Hansz, and R. Massobrio, "Job accessibility through public transport and unemployment in latin america: The case of montevideo (uruguay)," *Journal of Transport Geography*, vol. 85, p. 102742, may 2020.

[13] M. Bertrand, E. Duflo, and S. Mullainathan, "How much should we trust differences-in-differences estimates?" *The Quarterly Journal of Economics*, vol. 119, no. 1, pp. 249–275, feb 2004.

[14] C. de Chaisemartin and X. DâHaultfÅuille, "Two-way fixed effectsestimators with heterogeneous treatment effects," *American Economic Review*, vol. 110, no. 9, p. 2964â2996, 2020.

[15] S. Galiani and B. Quistorff, "The synth_runner package: Utilities to automate synthetic control estimation using synth," *The Stata Journal*, vol. 17, no. 4, pp. 834–849, 2017.

[16] E. Cavallo, S. Galiani, I. Noy, and J. Pantano, "Catastrophic natural disasters and economic growth," *Review of Economics and Statistics*, vol. 95, no. 5, pp. 1549–1561, dec 2013.

[17] D. Card and A. B. Krueger, "Minimum wages and employment: A case study of the fast food industry in new jersey and pennsylvania," Tech. Rep., 1993.

[18] D. Neumark and W. Wascher, "The effect of new jersey's minimum wage increase on fast-food employment: A re-evaluation using payroll records," *NBER Working Paper*, no. w5224, 1995.

[19] P. Wolfson and D. Belman, "15 years of research on us employment and the minimum wage," *Labour*, vol. 33, no. 4, pp. 488–506, 2019.

[20] D. Neumark and P. Shirley, "Myth or measurement: What does the new minimum wage research say about minimum wages and job loss in the united states?" National Bureau of Economic Research, Tech. Rep., 2021.

[21] D. Cengiz, A. Dube, A. Lindner, and B. Zipperer, "The effect of minimum wages on low-wage jobs," *The Quarterly Journal of Economics*, vol. 134, no. 3, pp. 1405–1454, 2019.

[22] A. Abadie, A. Diamond, and J. Hainmueller, "Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program," *Journal of the American statistical Association*, vol. 105, no. 490, pp. 493–505, 2010.

[23] C. E. Santiago, "Closing the gap: The employment and unemployment effects of minimum wage policy in puerto rico," *Journal of Development Economics*, vol. 23, no. 2, pp. 293–311, 1986.

[24] A. C. Freeman and R. B. Freeman, "Minimum wages in puerto rico: Textbook case of a wage floor?" *NBER working paper*, no. w3759, 1991.

[25] A. B. Krueger, "The effect of the minimum wage when it really bites: A reexamination of the evidence from puerto rico," National Bureau of Economic Research, Tech. Rep., 1994.

[26] A. Dube and B. Zipperer, "Puerto rico's predicaments: Is its minimum wage the culprit?" Washington Center for Equitable Growth, Tech. Rep., 2015.

[27] D. Neumark, "The econometrics and economics of the employment effects of minimum wages: Getting from known unknowns to known knowns," sep 2018.

[28] T. Gregory and U. Zierahn, "When the minimum wage really bites hard: Impact on top earners and skill supply," *IZA Discussion Paper Series*, 2020.

[29] A. Dube, T. W. Lester, and M. Reich, "Minimum wage effects across state borders: Estimates using contiguous counties," *The review of economics and statistics*, vol. 92, no. 4, pp. 945–964, 2010.

[30] S. Galiani and B. Quistorff, "The synth runner package: Utilities to automate synthetic control estimation using synth," *Stata Journal*, vol. 17, no. 4, pp. 834–849, 2017. [Online]. Available: https://EconPapers.repec.org/RePEc:tsj:stataj:v:17:y:2017:i:4:p:834-849

[31] P. Harasztosi and A. Lindner, "Who pays for the minimum wage?" *American Economic Review*, vol. 109, no. 8, pp. 2693–2727, aug 2019.

[32] J. Clemens, "How do firms respond to minimum wage increases? understanding the relevance of non-employment margins," *Journal of Economic Perspectives*, vol. 35, no. 1, pp. 51–72, feb 2021.

[33] E. Jardim, M. Long, R. Plotnick, E. van Inwegen, J. Vigdor, and H. Wething, "Minimum wage increases, wages, and low-wage employment: Evidence from seattle," jun 2017.

[34] J. J. Horton, "Price floors and employer preferences: Evidence from a minimum wage experiment," *SSRN Electronic Journal*, 2017.

[35] J. Clemens, L. Kahn, and J. Meer, "The minimum wage, fringe benefits, and worker welfare," may 2018.

[36] R. Gopalan, B. H. Hamilton, A. Kalda, and D. Sovich, "State minimum wages, employment, and wage spillovers: Evidence from administrative payroll data," *Journal of Labor Economics*, pp. 000–000, may 2021.

[37] "An overview of the special tax rules related to puerto rico and an analysis of the tax and economic policy implications ofrecent legislative options," Joint Committee on Taxation, Tech. Rep. (JCX-24-06), Jun. 2006.

[38] H. Beirch and L. Brooks, "Whose heritage? public symbols of the confederacy," Southern Poverty Law Center, Tech. Rep., 2019.

[39] Virginia, News Conference, Jun. 2020.

[40] S. Stephens-Davidowitz, "The cost of racial animus on a black candidate: Evidence using google search data," *Journal of Public Economics*, vol. 118, pp. 26–40, 2014.

[41] H. M. Blalock, *Toward a theory of minority-group relations.* John Wiley Sons, 1967.

[42] A. B. Krueger and J.-S. Pischke, "A statistical analysis of crime against foreigners in unified germany," *Journal of Human Resources*, pp. 182–209, 1997.

[43] S. J. Piatkowska, S. F. Messner, and A. Hövermann, "Black out-group marriages and hate crime rates: A cross-sectional analysis of us metropolitan areas," *Journal of research in crime and delinquency*, vol. 57, no. 1, pp. 105–135, 2020.

[44] Z. Valencia, B. Williams, and R. Pettis, "Pride and prejudice: Same-sex marriage legalization announcements and lgbt hate-crimes," *Available at SSRN 3362835*, 2019.

[45] C. E. Mills, "Gay visibility and disorganized and strained communities: A community-level analysis of anti-gay hate crime in new york city," *Journal of interpersonal violence*, vol. 36, no. 17-18, pp. 8070–8091, 2021.

[46] A. Benjamin, R. Block, J. Clemons, C. Laird, and J. Wamble, "Set in stone? predicting confederate monument removal," *PS: Political Science & Politics*, vol. 53, no. 2, pp. 237–242, jan 2020.

[47] A. C. Cameron, J. B. Gelbach, and D. L. Miller, "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, vol. 90, no. 3, pp. 414–427, 2008.

[48] A. C. Cameron and D. L. Miller, "A practitioners guide to cluster-robust inference," *Journal of human resources*, vol. 50, no. 2, pp. 317–372, 2015.

Robert Tucker Omberg

804-833-7327 | romberg@gmu.edu | tuckeromberg.com |

Education

George Mason University — Fairfax, VA
Doctor of Philosophy in Economics — *Expected May 2022*

- Mercatus Center PhD Fellow
- Bradley Fellow

George Mason University — Fairfax, VA
Master of Arts in Economics — *December 2020*

George Mason University — Fairfax, VA
Bachelor of Science in Economics — *May 2017*

- Honors College

Experience

Assistant Professor of Economics — August 2022 – Present
Jacksonville University — *Jacksonville, FL*

Graduate Lecturer — November 2021 – Present
George Mason University — *Fairfax, VA*

Research Affiliate — August 2020 – May 2021
New York University — *New York, NY*

Publications

Estimating the Elasticity of Labor Supply to a Firm: Evidence from a Field Experiment
Applied Economics Letters, vol.22, https://doi.org/10.1080/13504851.2022.2030030

Puerto Rico's Minimum Wage: Revisiting a Price Floor with "Bite"
IZA Journal of Labor Policy,Institute of Labor Economics, vol.11(1), https://doi.org/10.2478/izajolp-2021-0009

Alternatives for Paying Efficiency Wages: Why no PEOPLEFAX?
Journal of Private Enterprise, The Association of Private Enterprise Education, vol.35(Summer 20),pages 77-88

Is It Possible to Prepare for a Pandemic?
Forthcoming in Oxford Review of Economic Policy (w. Alex Tabarrok)

Courses Taught

Principles of Macroeconomics | *Jacksonville University* — F2022

Economics of Globalization | *Jacksonville University* — F2022

Introductory Econometrics | *George Mason University* — F2020-F2021

Economic Problems and Public Policy | *George Mason University* — F2019-S2022