

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Three Empirical Analyses Within Applied Microeconomics

Permalink

<https://escholarship.org/uc/item/576857tx>

Author

Ruiz Junco, Pablo

Publication Date

2019

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Three Empirical Analyses Within Applied Microeconomics

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Economics

by

Pablo Ruiz Junco

Committee in charge:

Professor Julie Cullen, Chair
Professor Julian Betts, Co-Chair
Professor Gordon Dahl
Professor James Fowler
Professor Sally Sadoff

2019

Copyright

Pablo Ruiz Junco, 2019

All rights reserved.

The dissertation of Pablo Ruiz Junco is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Co-Chair

Chair

University of California San Diego

2019

DEDICATION

A mis padres, José Luis e Isabel.

EPIGRAPH

Actually, it's polite to arrive early, and smart.

Only really good friends show up early.

Ergo de facto, show up early, become a really good friend.

—Michael Scott,

“Cocktails,” *The Office* (U.S.), season 3, episode 18, NBC.

TABLE OF CONTENTS

Signature Page	iii
Dedication	iv
Epigraph	v
Table of Contents	vi
List of Figures	ix
List of Tables	x
Acknowledgements	xii
Vita	xiv
Abstract of the Dissertation	xv
Chapter 1 Social Interactions With Close Neighbors: Do They Matter for Educational Outcomes?	1
1.1 Introduction	1
1.2 Data	8
1.2.1 Address Data and Coverage of the Neighbor Network	11
1.2.2 Neighborhoods and Micro-Neighborhoods	13
1.3 Empirical Strategy	14
1.3.1 Sorting Across and Within Neighborhoods	14
1.3.2 Detecting Social Interactions Using Cross-Sectional Data	17
1.3.3 Detecting Social Interactions Using Longitudinal Data	20
1.3.4 Additional Outcomes	25
1.4 Results	27
1.4.1 Detecting Social Interactions Using Cross-Sectional Data	27
1.4.2 Detecting Social Interactions Using Longitudinal Data	31
1.4.3 Additional Outcomes	40
1.4.4 Mechanisms and Limitations	42
1.5 Conclusion	43
1.6 Acknowledgement	45
1.7 Appendix	46
1.7.1 Additional Figures	46
1.7.2 Additional Tables	49
1.7.3 Clustering of Students Based on Educational Outcomes	52

Chapter 2	Extreme Weather and the Politics of Climate Change: a Study of Campaign Contributions and Elections	54
2.1	Introduction	54
2.2	Data	60
2.2.1	Database on Ideology, Money in Politics, and Elections	60
2.2.2	ActBlue	61
2.2.3	League of Conservation Voters Scorecard	63
2.2.4	Weather Shocks	65
2.2.5	Natural Disasters	67
2.3	Empirical Framework	67
2.3.1	Short-Run Weather Impacts	68
2.3.2	Medium-Run Natural Disaster Impacts	70
2.3.3	Medium-Run Weather Impacts	72
2.4	Results	73
2.4.1	Short-Run Weather Impacts	75
2.4.2	Medium-Run Natural Disaster Impacts	83
2.4.3	Medium-Run Weather Impacts	91
2.4.4	Mechanisms and Limitations	92
2.5	Conclusion	94
2.6	Acknowledgement	97
2.7	Appendix	98
2.7.1	Additional Figures	98
2.7.2	Additional Tables	99
2.7.3	Additional Specifications in the Short-Run Analysis	107
Chapter 3	Short and Long Term Effects of Extended Unemployment Insurance for First Time Claimants	109
3.1	Introduction	109
3.2	Institutional Framework	115
3.2.1	The Spanish Unemployment Insurance System	115
3.2.2	The Spanish Economy	117
3.3	Data	119
3.4	Empirical Strategy	122
3.4.1	Regression Discontinuity	122
3.4.2	Validity of the Regression Discontinuity Design	123
3.5	Results	128
3.5.1	Short Term Effects	128
3.5.2	Long Term Effects	132
3.5.3	Robustness	135
3.6	Discussion	143
3.7	Conclusion	145
3.8	Acknowledgement	147
3.9	Appendix	148

3.9.1	Additional Figures	148
3.9.2	Additional Tables	156
3.9.3	Prediction and Characterization of Workers at Spikes	165
	Bibliography	168

LIST OF FIGURES

Figure 1.1:	Scatterplots of children in the census versus the school district	46
Figure 1.2:	Distribution of the number of close neighbors	47
Figure 1.3:	Distribution of the number of same-grade close neighbors	47
Figure 1.4:	Example of a block group and a micro-neighborhood	48
Figure 1.5:	Elbow curve for cluster selection	52
Figure 1.6:	Visual representation of clustering results	53
Figure 2.1:	LCV score distribution by party affiliation	65
Figure 2.2:	Variations of estimates across quarters in election cycle	82
Figure 2.3:	Example of an ActBlue donation	98
Figure 2.4:	Distribution of individual ActBlue donation amounts	98
Figure 3.1:	The potential UI duration function	117
Figure 3.2:	Histogram of days worked over the past six years before an unemployment spell	126
Figure 3.3:	The short term effect of extended UI duration on days spent on UI	137
Figure 3.4:	The short term effect of extended UI duration on days in nonemployment . .	139
Figure 3.5:	Estimates of short and long term effects on days on UI for both high and low labor force attachment samples	141
Figure 3.6:	Estimates of short and long term effects on days in nonemployment for both high and low labor force attachment samples	141
Figure 3.7:	Unemployment in Spain 1990-2015	148
Figure 3.8:	Full histogram	148
Figure 3.9:	Truncated histogram	149
Figure 3.10:	Graphical results of McCrary test at different cutoffs	149
Figure 3.11:	Monte Carlo analysis of McCrary test results	150
Figure 3.12:	Estimates of short term effects on days on UI at different cutoffs	150
Figure 3.13:	Estimates of short term effects on days on UI for both high and low labor force attachment samples	151
Figure 3.14:	Estimates of short term effects on days spent in nonemployment at different cutoffs	151
Figure 3.15:	Estimates of short term effects on days spent in nonemployment for both high and low labor force attachment samples	152
Figure 3.16:	The long term effect of extended UI duration on days spent on UI	153
Figure 3.17:	The long term effect of extended UI duration on days in nonemployment . .	154
Figure 3.18:	Estimates of short and long term effects on days on UI at different cutoffs .	155
Figure 3.19:	Estimates of short and long term effects on days in nonemployment at differ- ent cutoffs	155
Figure 3.20:	ROC curve for spike worker detection	167

LIST OF TABLES

Table 1.1:	Full sample summary statistics	10
Table 1.2:	Examining sorting on observable characteristics	16
Table 1.3:	Cross-sectional sample summary statistics	19
Table 1.4:	Breakdown of relevant peer groups	20
Table 1.5:	Peer group acronyms	21
Table 1.6:	Size of relevant peer groups	22
Table 1.7:	Longitudinal sample summary statistics	25
Table 1.8:	Cross-sectional social interaction effects on educational outcomes	30
Table 1.9:	Social interaction effects on English grades	34
Table 1.10:	Social interaction effects on behavioral grades	36
Table 1.11:	Social interaction effects on English grades by school attended	37
Table 1.12:	Social interaction effects on behavioral grades by school attended	38
Table 1.13:	Social interaction effects on the school attended by students	41
Table 1.14:	Comparison of single-family and other block groups	49
Table 1.15:	English outcomes for math and behavioral sample	49
Table 1.16:	Predicting current geographical proximity with past school attendance	50
Table 1.17:	Predicting pair-level outcome differences using parental education similarity	50
Table 1.18:	School attended effects by income level	51
Table 1.19:	Results of k -means clustering of student outcomes ($k = 5$)	53
Table 2.1:	Summary statistics	74
Table 2.2:	Actblue donation responses to short-run temperature shocks	76
Table 2.3:	Heterogeneous effects by incumbent characteristics	80
Table 2.4:	The effects of natural disasters on amount raised (\$1,000)	86
Table 2.5:	The effects of natural disasters on the number of donors	88
Table 2.6:	The effects of natural disasters on elections	90
Table 2.7:	FEMA disaster declarations 1990-2012	99
Table 2.8:	Predicting total Democratic donations using ActBlue donations	100
Table 2.9:	Positive and negative temperature shocks on ActBlue contributions	101
Table 2.10:	The effect of extreme temperature events on ActBlue contributions	102
Table 2.11:	Heterogeneous effects across quarters in election cycle	103
Table 2.12:	The effects of extreme temperature on amount raised (\$1,000)	104
Table 2.13:	The effects of extreme temperature on the number of donors	105
Table 2.14:	The effects of extreme temperature on elections	106
Table 3.1:	Summary statistics	121
Table 3.2:	The short term effect of extended UI duration on days spent on UI	138
Table 3.3:	The short term effect of extended UI duration on days in nonemployment	140
Table 3.4:	Short and long term effects of extended UI duration for high and low labor force attachment samples	142
Table 3.5:	Occupational category groups	156

Table 3.6:	Results of McCrary test at different cutoffs	157
Table 3.7:	Short term effects for both high and low labor force attachment samples . .	157
Table 3.8:	The long term effect of extended UI duration on days spent on UI	158
Table 3.9:	The long term effect of extended UI duration on days in nonemployment . .	159
Table 3.10:	Robustness of short term effects on days spent on UI	160
Table 3.11:	Robustness of short term effects on days spent in nonemployment	161
Table 3.12:	Robustness of long term effects on days spent on UI	162
Table 3.13:	Robustness of long term effects on days spent in nonemployment	163
Table 3.14:	The effect of extended UI duration on transition to a permanent contract . .	164
Table 3.15:	Characterizing workers at spikes in terms of predetermined variables	165
Table 3.16:	Characterizing workers at spikes in terms of outcome variables	166
Table 3.17:	Evaluating random forest classifier performance at different threshold values	166

ACKNOWLEDGEMENTS

I would like to thank my committee chairs, Julie Cullen and Julian Betts, for their constant guidance and encouragement. Their feedback and insights have been invaluable and they have taught me a great deal about conducting academic research. Furthermore, I am especially grateful to them for always being supportive of both my academic and my professional goals.

I am also grateful to my other committee members, Gordon Dahl, James Fowler, and Sally Sadoff. I would like to thank them for their time and valuable insights. Gordon Dahl has also taught me much about causal inference and has contributed to my continued interest in the subject.

I would also like to acknowledge my colleagues and professors at UCSD. Yanjun Liao has been a fantastic co-author and working with her has made research more exciting and enjoyable. Andrew Zau has helped me immensely with collecting and understanding the data necessary for chapter 1. Other people from UCSD that deserve special mention include, but are not limited to, Roger Gordon, Mark Jacobsen, Jeffrey Clemens, Kate Antonovics, Eli Berman, Prashant Bharadwaj, Karen Bachofer, Diego Vera Cossio, Jason Bigenho, Seung-Keun Martinez, Eul Noh, Grant Johnson, Kristen Duke, Vinayak Alladi, Kilian Heilmann, Claudio Labanca, Sieuwerd Gaastra, Miles Berg, Paul Feldman, Yann Pannasié, Julián Martínez Iriarte, Alex Kellogg, Jonathan Leganza, and Bruno López-Videla.

Outside of UCSD, there are several people whom I would like to acknowledge. First, I would like to thank Andrew Chamberlain from Glassdoor and Yohsuke Miyamoto from Google. Over two summers, they helped me gain confidence in my abilities and showed me how the skills I acquired during graduate school can be applied outside of academia. Second, I would like to thank my sister Natalia. With her extensive academic experience, she has guided me through this whole process, and our frequent conversations helped me regain motivation when I needed it. Finally, I want to thank my dear Emilie, who kept me grounded from the start and always gave me the strength to keep on going.

Chapter 1 is currently being prepared for submission for publication of the material. Ruiz Junco, Pablo; “Social Interactions With Close Neighbors: Do They Matter for Educational Outcomes?” The dissertation author was the sole investigator and author of this material.

Chapter 2 is currently being prepared for submission for publication of the material. Liao, Yanjun; Ruiz Junco, Pablo; “Extreme Weather and the Politics of Climate Change: a Study of Campaign Contributions and Elections.” The dissertation author was a principal co-investigator and co-author of this material.

Chapter 3 is currently being prepared for submission for publication of the material. Ruiz Junco, Pablo; “Short and Long Term Effects of Extended Unemployment Insurance for First Time Claimants.” The dissertation author was the sole investigator and author of this material.

VITA

- 2012 *Licenciatura* in Economics, Universidad Carlos III de Madrid
- 2015 M. A. in Economics, University of California San Diego
- 2017 C. Phil. in Economics, University of California San Diego
- 2019 Ph. D. in Economics, University of California San Diego

ABSTRACT OF THE DISSERTATION

Three Empirical Analyses Within Applied Microeconomics

by

Pablo Ruiz Junco

Doctor of Philosophy in Economics

University of California San Diego, 2019

Professor Julie Cullen, Chair
Professor Julian Betts, Co-Chair

This dissertation is composed of three chapters on distinct topics within applied microeconomics. Each chapter uses econometric methods to analyze data from administrative sources, aiming to establish causal relationships between variables of interest.

In chapter 1, I study whether social interactions with residential neighbors in very close geographical proximity influence students' educational outcomes. In a cross-sectional analysis, I show that close neighbors have more similar levels of English achievement than pairs of students living in the same neighborhood but farther away. In a longitudinal analysis, I show that in the case of both English and behavior, close neighbors' outcomes are positively associated with own

outcomes. I also show that pairs of close neighbors are more likely to attend the same school, compared to those living in the same neighborhood but farther away.

In chapter 2, we study how extreme weather and natural disasters affect political outcomes such as campaign contributions and elections. We suggest that these weather events increase the salience of climate change and also lead some individuals to update their beliefs about it. In a short-run analysis, we find that the number of contributions to the Democratic Party increases in response to weekly average temperature, and the effect is stronger in counties with more anti-environment incumbent politicians. In a medium-run analysis, we find suggestive evidence that, following a natural disaster, the level of competition of elections rises, as measured by the probability of a challenger entering the race, the number of donors, and total funds raised.

In chapter 3, I study the effects of increased unemployment insurance benefit duration on individuals' short and long term labor market outcomes, using a regression discontinuity design. In the short term, I show that for a subset of workers, receiving 60 additional days of UI benefits induces them to spend 21 more days on UI and 14 more days in nonemployment. In the long term, a similar sample of workers spends 7 additional days on UI as a result of extended benefits, compared to the short term. Furthermore, I find suggestive evidence that the effect on days spent in nonemployment in the long term is lower than in the short term, which is consistent with increased post-unemployment job stability.

Chapter 1

Social Interactions With Close Neighbors: Do They Matter for Educational Outcomes?

1.1 Introduction

Social interactions among students are considered to play a role in shaping educational outcomes. An important dimension of these interactions is where they occur; for example, students may interact at school or in the residential neighborhood. As a result, both school and neighborhood contexts have received significant attention from academics and policymakers. However, less is known about the degree to which different sets of students within these settings influence individual outcomes.

When it comes to social interactions among students in the same neighborhood, a distinction can be made between neighbors and close neighbors. While social interactions may occur with both groups of students, it is possible that the increased geographical proximity in the case of the latter intensifies these interactions, giving rise to *micro-neighborhoods*. If this is the case,

we may expect close neighbors to have an influence on students' educational outcomes beyond that of regular neighbors or other neighborhood factors. Further, it is also worth considering the role that schools might play in the influence of close neighbors on educational outcomes. On one hand, attending the same school may act as a catalyst for the development of strong social connections with close neighbors, or provide a channel through which knowledge and behaviors can be transmitted. In this case, we would expect close neighbors to matter only if they are in the same school. On the other hand, close geographical proximity may foster social interactions with all neighbors, regardless of the school attended. In this case, we would expect all close neighbors to matter.

In this paper, I use administrative data from the San Diego Unified School District (SDUSD) to study whether social interactions with close neighbors influence students' educational outcomes. I primarily focus on three outcomes: English test scores, math test scores, and classroom behavioral grades. I use two distinct approaches – both geared toward the detection of social interaction effects – which exploit different sources of variation. Firstly, I use a cross-sectional approach based on Bayer, Ross, Topa, *Journal of Applied Econometrics*, December (2008), comparing the similarity in outcomes of pairs of students living in the same neighborhood and close by to that of pairs living in the same neighborhood but farther away. I address the issue of sorting within neighborhoods, extend the approach to account for other factors that may vary at the micro-geographical level, and explore differential effects by whether the students in the pair attend the same school. Secondly, I use a longitudinal approach that employs student-address-school fixed effects, isolating variation within students who are residentially immobile and attend the same school over a given time period. Using this variation, I estimate the relationship between close neighbors' educational outcomes and own educational outcomes, while accounting for the influence of neighbors from the same school, neighbors from different schools, students in the same school, and general shocks at the neighborhood and micro-neighborhood levels. I then extend this approach to study both social interactions with close neighbors in the same school

and with neighbors in different schools. Following these analyses, I return to the cross-sectional approach to study how these interactions may influence additional outcomes, such as the decision of which school to attend. Finally, I discuss alternative mechanisms and explanations for these results.

In the cross-sectional analysis, I start by examining the extent of sorting within neighborhoods. I find evidence of sorting within neighborhoods based on race and parental education in the main sample, but no evidence of sorting in a subsample of neighborhoods comprised of single family homes. In the subsample, I find evidence that social interactions with close neighbors matter in the case of English achievement, with outcomes being 3% of a standard deviation closer for pairs of students in the same grade who live close by, compared to pairs in the same grade and neighborhood but living farther away. In a validation exercise involving students in different grades, I find that the effect is about half of that in the main regression, suggesting that the observed effect is picking up on social interactions. I argue that the size of the effect is economically non-negligible by comparing it to the influence of certain shared covariates on outcome proximity. In the case of math and behavioral outcomes, I do not find evidence that social interactions matter, although the resulting level of statistical significance does not allow me to rule out small effects. For all three outcomes, the results are uninformative of whether the effect of proximity varies by whether the students attend the same school.

In the longitudinal analysis, I find suggestive evidence that close neighbors matter for behavioral outcomes. Specifically, I find that a one standard deviation increase in same-grade close neighbors' behavioral outcomes is associated with a 0.021 standard deviation increase in own behavioral outcomes. I do not find evidence that close neighbors matter for English achievement in the main sample, although I do find suggestive evidence of this in a subsample, which includes students after they have been at an address for more than one year. Specifically, I find that a one standard deviation increase in same-grade close neighbors' English outcomes is associated with a 0.018 standard deviation increase in own English outcomes. This sample

restriction is meant to address attenuation bias, which may arise if students find it more difficult to make strong social connections with close neighbors during the first year at an address. Further, I examine whether these associations exist among students in the same school, students in different schools, or both. In the case of English outcomes, I lack statistical power to determine whether students in the same or different schools matter. However, the estimated association is three times larger in the case of students in the same grade and school, compared to that of students in the same grade but not in the same school. This relative magnitude is consistent with students in the same school being more influential than other students in the neighborhood when it comes to English outcomes. In terms of behavioral outcomes, I again lack statistical power, although the estimates are closer to being statistically significant than for English outcomes. However, in this case the estimated coefficients for students in the same grade and school and those of students in the same grade but not the same school are quite similar in magnitude. These results are consistent with close neighbors in both the same and in different schools mattering for behavioral outcomes.

In my analysis of additional outcomes, I study how social interactions with close neighbors influence outcomes related to which school students choose to attend. I find that students living in close proximity are 1.5% more likely to attend the same school, compared to students living in the same neighborhood but farther away. I also find that students living in close proximity are 1.5% more likely to make the same decision to attend or not attend their home school. Further, I examine heterogeneity by income level and show that the estimated effects are stronger in the case of wealthier neighborhoods. Through statistical tests, I rule out that these effects are driven by reverse causality, which may arise if students who have attended the same school in the past move in next door to each other.

While the cross-sectional and longitudinal analyses focus on similar outcomes, I argue that they capture distinct dimensions of social interactions, and therefore should be interpreted separately. On one hand, the cross-sectional analysis focuses on outcome *similarity*, i.e. whether

student outcomes are closer to each other as a result of close geographical proximity. On the other hand, the longitudinal analysis focuses how student outcomes respond to *changes* in neighbors' outcomes, i.e. whether the outcomes of students and their close neighbors co-move over time, independently of how close they are. Further, the cross-sectional analysis may capture the cumulative effect of exposure to close neighbors over past years, whereas the longitudinal analysis focuses on contemporaneous changes in the outcomes of neighbors living in close proximity during each given year.

This paper is related to various literatures. In particular, it is positioned at the intersection of the literatures on neighborhood effects and peer effects in education. The literature on peer effects in education is wide and has been reviewed in Sacerdote (2011). Among the papers discussed by Sacerdote (2011), a commonly used method to study peer effects is the linear-in-means model, which I employ in the longitudinal analysis section of this paper.¹ These studies typically focus on classroom peer effects and find that peer educational achievement has a positive effect on a student's own achievement (Whitmore, 2005; Hanushek, Kain, Markman, & Rivkin, 2003; Hoxby, 2000; Angrist & Lang, 2004; Ammermueller & Pischke, 2009). There is also an extensive literature focusing on neighborhood effects. Durlauf (2004) and Graham (2018) review theory, estimation, and previous work on neighborhood effects, with the latter emphasizing neighborhood segregation. Earlier studies compare students' educational and behavioral outcomes to neighborhood characteristics, arguing that there are significant neighborhood effects (Brooks-Gunn, Duncan, Klebanov, & Sealander, 1993; Ainsworth, 2002; Case & Katz, 1991; Turley, 2003). Other approaches include Goux and Maurin (2007), who study French neighborhoods and use variation in the proportion of neighbors born later in the year. They focus on educational advancement and conclude that there are strong neighborhood effects. Agrawala, Altonji, and Mansfield (2018) use a control function approach to account for sorting into neighborhoods and

¹For a review of the estimation of social interaction models, see Manski (1993), Moffitt (2000), and Blume, Brock, Durlauf, and Ioannides (2011). Other methods to study peer effects in education include the excess variance approach outlined in Graham (2008).

study effects on long term outcomes. They find that experiencing a more favorable school and location combination has a positive effect on high school graduation and college enrollment. Furthermore, a separate set of papers study neighborhood exposure effects by exploiting variation in the age of children when families move (Aaronson, 1998; Chetty & Hendren, 2016a, 2016b). For example, Chetty and Hendren (2016a) study long term outcomes such as employment, college attendance, and teenage births, and find that these improve linearly in proportion to the amount of time children spend growing up in a better neighborhood. Finally, several studies have made use of housing programs and social experiments which aim to assign families to improved neighborhoods (Oreopoulos, 2003; Rosenbaum, 1995; Galster, Santiago, Stack, & Cutsinger, 2016; J. Kling, Liebman, & Katz, 2007; Jeffrey R Kling, Liebman, Katz, & Sanbonmatsu, 2004; Ludwig et al., 2013; Jeffrey R. Kling, Ludwig, & Katz, 2005; Sanbonmatsu, Kling, Duncan, & Brooks-Gunn, 2006). The evidence resulting from these studies is mixed, or at least program dependent. For example, the studies focusing on the Moving to Opportunity (MTO) program conclude that while there was no impact on program participants in terms of labor market outcomes and educational achievement, there is evidence of effects on mental and physical health (J. Kling et al., 2007; Jeffrey R Kling et al., 2004; Ludwig et al., 2013; Jeffrey R. Kling et al., 2005; Sanbonmatsu et al., 2006). On the other hand, Rosenbaum (1995) and Galster et al. (2016) suggest that there are effects on outcomes such as employment and educational performance.

An important distinction between previous work on neighborhood effects and this study is the size of the geographical area of interest. Most studies focus on the neighborhood as a whole and define it as either the ZIP code or the census tract, with some studies using census block groups.² In this study I focus on neighbors in very close geographical proximity, i.e. close neighbors. I define close neighbors as students living in the same block group, within a 0.05 mile radius of each other, and on the same street. Further, I study the impact of close neighbors on educational outcomes *beyond* that of the wider neighborhood, by controlling for block group level

²One exception is Goux and Maurin (2007), who study the impact of neighbors on educational advancement. They focus on French *aires*, which are smaller than U.S. census block groups.

factors. This strategy is based on Bayer et al. (2008), who study the impact of social interactions on job referrals by examining whether individuals living in the same block group are more likely to work in the same place if they also live in the same individual census block. Other studies employing this strategy or utilizing variation at the micro-geographical level have examined the effect of fast food restaurants on obesity (Currie, DellaVigna, Moretti, & Pathania, 2010), how proximity to a sex offender's residence affects property values (Linden & Rockoff, 2008), and neighbor spillovers in bankruptcy decisions (Scholnick, 2013), housing aspirations (Galiani, Gertler, & Undurraga, 2018), and automobile purchases (Grinblatt, Keloharju, & Ikäheimo, 2008).

On top of studying the impact of close neighbors, this study aims to separately examine the role of neighbors in the same school and neighbors in different schools in shaping educational outcomes. This has proven difficult in the past due to data limitations and the often overlapping nature of schools and neighborhoods, although a few recent studies have attempted to shed light on the issue.³ Gibbons, Silva, and Weinhardt (2013) use administrative data from England to study the role of neighbors in shaping educational outcomes. Using a changes-in-changes specification, they exploit variation in neighborhood composition experienced by residentially immobile students, and are able to study effects from neighbors in the same school and in different schools separately. They conclude that there are no peer effects from neighbors, regardless of whether they attend the same school or not, although they do find some effects on behavioral outcomes. While their approach is similar to the longitudinal strategy in this paper, there are some important differences. First, my approach is distinct in that it studies the impact of close neighbors beyond that of regular neighbors, whereas they focus on changes in the neighborhood as a whole. Second, their definition of neighborhood is the Output Area (OA) from the British

³The role of schools and neighborhoods in shaping student outcomes has also received attention in the sociology literature on neighborhoods and education, which is reviewed in Nieuwenhuis and Hooimeijer (2016). Some of these studies use multilevel models to explain variation in student outcomes and find that once schools are accounted for the variance explained by neighborhoods drops significantly (Sykes & Musterd, 2011; Brännström, 2008). Kauppinen (2008) suggests that neighborhood effects operate through the school context. Owens (2010) shows that a student's relative neighborhood socioeconomic status within a school is predictive of educational attainment.

census, which is larger and less precise than my definition of micro-neighborhood.⁴ Third, the behavioral outcomes they consider are self-reported, whereas those studied in this paper are determined by teachers. Another related study is Del Bello, Patacchini, and Zenou (2015). They use cross-sectional social network data including friendship nominations for a sample of schools in the U.S. to study the impact of peers at school and peers in the neighborhood on educational achievement. They show that, when simultaneously modeling peer influences and the friendship link formation process, the effect of peers in the neighborhood on own achievement becomes insignificant, suggesting that geographical proximity may not play a major role in facilitating peer effects. However, there are two aspects of their study that are worth noting. First, their estimated effect of peers in the neighborhood, while insignificant, is non-negligible. If we consider this as well as their relatively low sample size ($N = 837$), we may not rule out that peers in the neighborhood matter.⁵ Second, their network data only include students in the same school, so they are not able to study peer effects from neighbors in different schools.

The remainder of this article is organized as follows. Section 2 describes the data structure and sources. Section 3 outlines the cross-sectional and longitudinal empirical strategies used in this paper. In section 4 I present the results of the analysis, interpret them, and discuss alternative mechanisms. I conclude in section 5.

1.2 Data

This paper utilizes student-level administrative data from the San Diego Unified School District (SDUSD) for the 2003-2013 time period. With enrollment close to 120,000, this school district is the second largest in California and among the largest in the United States. Not only

⁴British OAs have 125 households on average and approximately five students in the same grade, which is approximately twice as much as the number of same-grade students in a micro-neighborhood in my sample. Further, I am able to observe specific student addresses which allows me to restrict micro-neighborhoods to students living on the same street, where social interactions should be more likely.

⁵Without accounting for link formation, the size of their estimated effect of peers in the neighborhood is 0.049 and significant. When they account for link formation, the size drops to 0.033 and becomes insignificant.

does SDUSD serve a large number of students, it also serves students from a wide variety of ethnic and socioeconomic backgrounds, as reflected in table 1.1. In this table we can see that around 45% of the students in the district are Hispanic, 25% are White, 17% are Asian, and 13% are Black. Further, 79% of parents have at least a high school degree, and 33% have at least a college degree. SDUSD runs multiple school choice and open enrollment programs, so only 63% of students are *residents*, i.e. they attend their neighborhood school. As a result of its size and diversity, SDUSD is an excellent setting for educational research in general. Additionally, the non-overlapping nature of neighborhoods and schools attended described above makes SDUSD an excellent setting for this specific research project.⁶

Throughout this article, I primarily focus on three educational outcomes. The first two outcomes are test scores for English and math on the California Standards Test (CST), a component of the Standardized Testing and Reporting (STAR) Program.⁷ This test was offered through the 2012-2013 academic year for grade levels 2-11, which constitutes my main sample.⁸ The third outcome I study is a secondary behavioral GPA. This outcome is available starting in secondary school and is an average of the classroom behavior grade for each course the student took in a year.⁹ I standardize these variables at the grade by year level, as reflected in table 1.1.

In order to better understand these three outcomes and the relationships between them I perform an exploratory data analysis. Using a clustering algorithm, I assign students to groups (clusters) such that the students in each group have similar values of the outcome variables, with each group characterized by its center in Euclidean space.¹⁰ Using this method, I find five distinct groups of students, described in table 1.19 and in figure 1.6. This analysis suggests that students

⁶For more information regarding SDUSD, visit www.sandiegounified.org.

⁷For more information, visit <https://www.cde.ca.gov/re/pr/star.asp>.

⁸While the math CST is offered for grades 2-11 and scores are available in the data, after grade seven students enter separate tracks for this exam. Therefore, I exclude these grades from the analysis when studying math outcomes.

⁹Teachers assign classroom behavior grades based on the following five point scale: “0=Unsatisfactory, 1=Needs improvement, 2=Satisfactory, 3=Good, and 4=Excellent.” As a result, it is important to note that since behavioral grades are assigned by teachers, they may be dependent on classroom and school cultures. This is a limitation of the behavioral outcome compared to the two academic outcomes, since the latter come from a single objective test taken by all students in the school district.

¹⁰The details of this procedure are in the appendix.

typically fall into groups with similar values of English and math, while behavioral grades may take on different values, regardless of the value of the other two outcomes. In table 1.19 we can see that the resulting groups all have simultaneously high, medium, or low English and math values, whereas the value of behavioral grades varies. This is evidenced by the positive correlation between English and math in panel (a) of figure 1.6 and the weaker relationship between English and behavioral in panel (b).¹¹ This exercise is suggestive that while English and math scores may be driven by some of the same factors, such as test-taking or general academic ability, there are other factors at play in determining behavioral grades. Importantly, some of these factors may be more transmittable than others, so the influence of social interactions on these variables may be different.

Table 1.1: Full sample summary statistics

	N	Mean	Std. Dev.	Min.	Max.
Year	938,546	2,008.102	3.089	2,003	2,013
Grade level	938,546	6.243	2.785	2	11
Parents H.S.	895,325	0.791	0.406	0	1
Parents college	895,325	0.332	0.471	0	1
Hispanic	938,546	0.447	0.497	0	1
White	938,546	0.248	0.432	0	1
Asian	938,546	0.166	0.372	0	1
Black	938,546	0.130	0.336	0	1
Charter	938,546	0.100	0.300	0	1
Resident	937,743	0.631	0.482	0	1
English	880,482	0.00000	1.000	-3.935	5.362
Math	874,111	0.00001	1.000	-3.450	7.698
Behavior	472,878	0.00003	1.000	-4.552	2.588

Notes: the above table includes summary statistics for the full sample at the student-year level.

¹¹The relationship between math and behavioral looks similar to panel (b).

1.2.1 Address Data and Coverage of the Neighbor Network

An important building block of this study is the geographical information that underlies it. In order to determine student location, I use residential addresses reported to the school district by the students and their parents. I geocode the addresses and use this information to calculate distances between addresses as well as assign students to census tracts and block groups. For certain supplementary analyses, I incorporate information from the 2010 U.S. Census contained in the Neighborhood Change Database (NCDB) (Geolytics Inc., 2012). These data provide useful socioeconomic information at the census tract level.

It is important for the validity of the results in this paper that the district data capture the majority of students in each neighborhood.¹² If this is not the case, then coverage of neighbors will be incomplete and the neighbor network will not be adequately captured. Further, it is also crucial that the address data are as accurate as possible and that they are not strategically misreported. This may occur if students from outside the district have an incentive to report a different address in order to gain access to SDUSD schools. For example, students may report the address of a relative instead of their own. In order to assess the extent of these problems, I calculate the number of children in each census tract according to the district data, and compare these values to the official census ones.

I start by studying the relationship between the number of children in a census tract from both sources graphically in figure 1.1. If the district data adequately cover the neighbor network and there is no strategic misreporting, we would expect the number of children in the district to slightly underestimate the number of children according to the census. This is because some children will not be enrolled in the district if they attend private schools, charter schools not affiliated with the district, or schools in a different district. If strategic misreporting of addresses is a pervasive problem, we would expect the district populations to slightly overestimate the

¹²Students may also attend private schools, charter schools not affiliated with the district, or schools in a different district.

census populations.

Subfigure (a) shows the relationship for all census tracts present in the district data. It is clear from this figure that there are two distinct types of tract in the data: the first type of tract has reasonable amounts of children in both the district data and the census data, and the populations align quite well at first sight. The second type of tract has very few students in the district data, regardless of the number of children in the census data.¹³ This figure reveals that tracts of the second type are populated mostly by non-district students and therefore the district data will not adequately capture the potential social network in these cases.¹⁴ As a result, in order to ensure accuracy of the neighbor information, in this article I focus only on tracts of the first type.¹⁵ Note that this sample selection is not related to the strategic misreporting concern, but it ensures that the district data capture the majority of students in each neighborhood, which is also important for the validity of the research design.

In subfigure (b) of figure 1.1, I show the relationship after performing the above sample selection. It is clear that in this sample the data align very well, with most points falling close to the 45-degree line. This means that the number of students in each tract according to the district data and the number of students in each tract according to the census are quite close. In fact, using the district population as a predictor for the census population results in an R^2 of about 0.8, which confirms the strength of the fit. Further, as expected, the data points generally fall slightly above the 45-degree line, indicating that the district populations include most students in the tracts except for a small share of non-district students, and that an excess of district students due to strategic misreporting is unlikely. Overall, these figures show that the address data are reliable and for the selected tracts, the neighbor network will be adequately captured.

¹³These are the ones bunched to the left of the graph.

¹⁴These tracts have so few students in the district because the majority of them attend private schools, charter schools not affiliated with the district, or schools in a different district altogether.

¹⁵In order to do this, I restrict the sample to tracts with 100 or more students in the district. This represents the great majority of the data.

1.2.2 Neighborhoods and Micro-Neighborhoods

Another important element in this setting is to determine how neighborhoods and close neighbor relationships are defined. In this article, I use census block groups as the definition of neighborhood. Therefore, I consider students to be neighbors if they reside in the same census block group.¹⁶ According to the U.S. Census Bureau, census block groups typically contain between 600 and 3,000 people and usually cover a contiguous area.¹⁷ Block groups are subdivisions of census tracts and are themselves comprised of individual census blocks. In the San Diego Unified School District, there are 138 students per block group on average.¹⁸

When it comes to close neighbor relationships, I consider a pair of students to be close neighbors if they live in the same block group, within a 0.05 mile radius of each other, and on the same street.¹⁹ Throughout this article, I refer to close neighbors as living in the same micro-neighborhood. Micro-neighborhoods are intended to capture smaller areas within a neighborhood where social interactions are more likely as a result of close geographical proximity. This idea is captured by Lee and Campbell (1999), who define micro-neighborhoods as 10 adjacent housing units (five on either side of the street), and report that 31% of these neighbors are judged close or very close by respondents.²⁰

In order to further characterize micro-neighborhoods, I plot the distribution of the number of close neighbors from all grade levels in figure 1.2. As we can see in this figure, the average number of close neighbors is 12.7. In terms of same-grade close neighbors, the distribution is in figure 1.3. In this case the average number of same-grade close neighbors is 1.3. If we

¹⁶Students in the same household are not considered neighbors.

¹⁷As explained in Bayer et al. (2008), local participants in the Census Bureau's Participant Statistical Areas Program have a major role in the delineation of census block groups. These participants are encouraged by the Census Bureau to consider features that encompass similar community patterns when drawing these lines (U.S. Census Bureau, 1997). For more information on census geographical definitions, visit www.census.gov/geo/reference/.

¹⁸Using 2010 as the year of reference and students in grades 2-11.

¹⁹With this approach, I avoid categorizing students as close neighbors if their homes face away from each other despite being on the same block. Even though such pairs of students are close geographically, social interactions may be less likely if they do not run into each other as often.

²⁰While at a larger scale, Bailey, Cao, Kuchler, Stroebel, and Wong (2018) use Facebook data to study social connectedness at the U.S. county level, finding that geographical proximity is an important determinant.

consider the sample of students with at least one close neighbor, then the average number of close neighbors from all grades is 14. Similarly, in the sample of students with at least one same-grade close neighbor, the average number of same-grade close neighbors is 2.5. Finally, in figure 1.4 I provide a visual example of a neighborhood and a micro-neighborhood.

1.3 Empirical Strategy

In this section, I go over the empirical framework and outline two distinct approaches for studying how social interactions with close neighbors may influence educational outcomes. While both approaches utilize the same data source, the first approach uses variation across pairs of students and the second approach uses variation within individual students over time.

1.3.1 Sorting Across and Within Neighborhoods

The identification strategy used in Bayer et al. (2008) provides the starting point for my cross-sectional study. This identification strategy revolves around the idea that while families choose which neighborhood to live in, they may not be able to target a small geographical area within a neighborhood; i.e., while there is sorting *across* neighborhoods, there might not be sorting *within* neighborhoods. This may be due to few housing units being available at a given point in time for small areas within the neighborhood. Further, targeting specific areas may also be difficult if, before moving in, it is hard to identify close neighbor characteristics beyond those of the wider neighborhood.

If there is no sorting within neighborhoods, then a correlation in outcomes among close neighbors beyond that expected by living in the same neighborhood could be interpreted as the result of social interactions fostered by geographical proximity, barring any additional confounders. Therefore, it will be important to determine whether this assumption holds in this setting.

One way to test for the validity of the above assumption is to test whether there is sorting

on observable characteristics within a neighborhood. In order to do this, I include all pairs of students living in the same neighborhood (block group), and estimate the following model:²¹

$$D_{ij} = \rho Close_{ij} + \delta_b + \varepsilon_{ij} \quad (1.1)$$

where D_{ij} is an indicator variable equal to one if students i and j in the pair share a demographic characteristic, and zero otherwise. Further, $Close_{ij}$ is an indicator for whether the pair of students are close neighbors, δ_b represents a block group fixed effect that captures the probability that students living in the same neighborhood share demographic characteristics, and ε_{ij} is an error term. The coefficient of interest is ρ , which captures the differential likelihood of a pair sharing a demographic characteristic if they live close by compared to farther away in the neighborhood.²²

The results of this test are below in table 1.2. The first column reports the estimate of ρ for different demographic characteristics, but without the block group fixed effects. This gives us an idea of the extent of sorting across neighborhoods. The second column includes block group fixed effects to estimate the model above. In the case of race, we can see that conditional on a pair of students living in the same neighborhood, they are almost 7% more likely to be of the same race if they live in close proximity, indicating that there is in fact sorting within neighborhoods based on race. In the case of parental education, we can see that a pair is about 2% more likely to share a level of parental education if they live in close proximity, which also points to sorting. Finally, there does not appear to be sorting in the case of gender or grade level.

The existence of sorting based on race and parental education within neighborhoods is problematic for the identification strategy, as these characteristics correlate to a number of outcomes of interest. One possible explanation for this result is that not all neighborhoods are homogeneous in terms of the housing stock. For example, if part of the neighborhood is comprised

²¹I focus on the year 2010 for the entire cross-sectional analysis, as it is a census year and ensures that the census information has the highest possible accuracy.

²²I cluster standard errors at the block group level for this and subsequent cross-sectional analyses.

Table 1.2: Examining sorting on observable characteristics

	(1) No Controls	(2) Block Group F.E.	(3) Single-family B.G.
Race	0.584*** (0.017)	0.068*** (0.017)	0.0004 (0.005)
Parental Education	0.327*** (0.005)	0.021*** (0.004)	0.004 (0.005)
Grade Level	0.102*** (0.0005)	0.001 (0.0005)	0.001 (0.002)
Gender	0.500*** (0.001)	-0.0001 (0.001)	0.001 (0.003)
Observations	17,051,046	17,051,046	2,379,338

Notes: the above table includes results of pair-level O.L.S. regressions of whether the pair of students shares a demographic characteristic (outcome) on whether the students in the pair are close neighbors. The second column adds block group fixed effects and the third column restricts the sample to block groups comprised of single-family homes, while also including block group fixed effects. Pairs included in the estimation sample are all pairs residing in the same census block group in 2010. Statistical significance levels are indicated by * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

of single family homes and there is a block with an apartment tower, then it is possible that families will sort into the tower versus the single family homes based on the above characteristics. In order to explore whether this could be behind the observed sorting results, I restrict the sample to block groups where all students live in single family homes. Not only will the housing stock be homogeneous in this case, but the density of available housing units will also be lower, making sorting even more difficult.

I report results for this subset of block groups in column (3) above. We can see that there is no within neighborhood sorting on any of these characteristics, as all coefficients in this column are very close to zero and precisely estimated. As a result of this exercise, for the remainder of the cross-sectional analysis I choose to focus on this subset of block groups, where I can be confident that sorting on demographic characteristics is not driving the results. I do this while keeping in

mind two caveats. Firstly, it may still be possible that there is sorting on other characteristics, such as unobservables that are not captured by the variables in table 1.2.²³ This caveat motivates the research design I propose later on which involves a validation exercise to account for factors such as these. Secondly, note that the selected sample may not be representative of the school district as a whole, as shown in table 1.14 in the appendix. In particular, this sample of block groups is wealthier, less likely to have minority residents, and more likely to have college educated parents compared to the rest of the school district neighborhoods.

1.3.2 Detecting Social Interactions Using Cross-Sectional Data

Building on the above approach, in this section I propose a method for detecting social interaction effects in educational outcomes. I posit that social interactions occur primarily among students in the same grade, and include all pairs of students living in the same block group and in the same grade to estimate the following model:

$$|\Delta Y_{ij}| = \rho \text{Close}_{ij} + \beta X_{ij} + \delta_b + \delta_g + \varepsilon_{ij} \quad (1.2)$$

where $|\Delta Y_{ij}|$ is the absolute difference in an educational outcome between students i and j in the pair. Further, Close_{ij} is an indicator for whether the pair of students are close neighbors, δ_b represents a block group fixed effect, δ_g represents a grade level fixed effect, and ε_{ij} is an error term. The block group fixed effects capture the proximity in educational outcomes of students in each block group and the grade fixed effects capture the proximity in educational outcomes among students in each grade level. Finally X_{ij} is a vector of control variables that includes indicator variables for the students being of the same race, gender, or parental education level, as well as indicators for whether they attend the same school, whether they both attend their home

²³Note that if sorting is the result of families targeting micro-neighborhoods with similar families, then sorting on unobservables should be less likely than sorting on observables, given that unobservable characteristics are not directly visible to the agents making the decisions.

school or they both attend a school other than their home school, and whether they both attend a charter school or both attend a regular school.²⁴

The coefficient of interest is ρ , which captures the differential proximity in educational outcomes of students living close by, compared to students living in the same neighborhood but not close by. If we would like to interpret ρ as resulting from social interactions, it is important to rule out that it is driven by other factors that may vary at the micro-geographical level and are correlated with proximity in outcomes. For example, it may still be possible that there is sorting on characteristics beyond those included in table 1.2, such as unobservables. In order to examine whether the above coefficient may be driven by these other factors, I conduct a validation exercise by estimating the above model but only including pairs of students who are in different grade levels. If these factors are influencing proximity in outcomes, they will do so for all grades, so both within-grade and cross-grade estimates will be driven by these to the same extent, and may both be non-zero. Further, since social interactions should be stronger among students in the same grade, we expect the within-grade estimate of ρ to be greater in absolute value than the cross-grade estimate, with both possibly being non-zero if micro-geographical factors are present. Finally, note that the resulting model is similar to the one above except that now the grade level fixed effect δ_g varies with each pair of grades and captures baseline proximity in outcomes for students in each combination of grade levels.²⁵

Moreover, using this setup it is also possible to explore whether the observed effects are driven through students in the same school, students in different schools, or both. To accomplish this, I again take all pairs of students in the same grade and neighborhood and estimate the following model:

$$|\Delta Y_{ij}| = \rho_1 Close_{ij} + \rho_2 Close_{ij} \times School_{ij} + \beta X_{ij} + \delta_b + \delta_g + \varepsilon_{ij} \quad (1.3)$$

²⁴Standard errors are clustered at the block group level.

²⁵For example, all pairs of students in which one of the students is in grade 6 and the other is in grade 8 will share a fixed effect value, as will all pairs of students in which one student is in grade 6 and the other is in grade 9.

where all variables are defined as above and the difference is that now I interact the *Close* variable with *School*, which is an indicator for whether the students in the pair attend the same school. Note that *School* was already included in the vector of control variables X_{ij} . Therefore, I interpret ρ_2 as the differential impact of geographical proximity on outcome proximity for students in the same school, compared to students in different schools.

Summary statistics for the cross-sectional analysis are in table 1.3. With the exception of the outcome variables, all other variables are binary and are equal to one if the students in the pair share a characteristic. For example, 50% of pairs in the sample are of the same gender, and 43% of pairs are of the same race. In terms of the outcome variables, as explained above they are the absolute difference in the standardized educational outcomes between the students in the pair. As a result, a noticeable feature is that they cannot take on values below zero, as reflected in table 1.3.

Table 1.3: Cross-sectional sample summary statistics

	N	Mean	Std. Dev.	Min.	Max.
Grade level	2,586,650	0.103	0.304	0	1
Close	2,586,650	0.030	0.172	0	1
Gender	2,586,650	0.501	0.500	0	1
Race	2,586,650	0.434	0.496	0	1
Parental education	2,379,338	0.312	0.463	0	1
English	2,263,336	0.996	0.779	0.000	6.724
Math	2,234,542	1.065	0.828	0.000	7.221
Behavior	760,444	0.842	0.758	0.000	5.012
School	2,586,650	0.204	0.403	0	1
Home	2,496,304	0.325	0.468	0	1
Resident	2,585,060	0.693	0.461	0	1
Charter	2,586,650	0.895	0.306	0	1

Notes: the above table includes summary statistics for the cross-sectional sample. These are pair level data for the year 2010.

1.3.3 Detecting Social Interactions Using Longitudinal Data

In this section, I go over the empirical strategy for detecting social interaction effects using variation within students over time. An advantage of this approach is that it is possible to control for both observed and unobserved time-invariant student characteristics by employing fixed effects, as outlined in the models below. However, prior to going over the model specifications, it is important to understand the different peer groups that will be relevant in this setting. Table 1.4 includes a breakdown of these groups for each student of reference based on joint geographical, grade, and school attendance characteristics.

Table 1.4: Breakdown of relevant peer groups

Block Group (B)	Close (C)	Grade (G)	School (S)	Description
1	1	1	1	Close neighbors, same grade, same school
1	1	1	0	Close neighbors, same grade, not same school
1	0/1	1	1	Neighbors, same grade, same school
1	0/1	1	0	Neighbors, same grade, not same school
0/1	0/1	1	1	Same grade, same school

The first two rows of table 1.4 include students who live in the same neighborhood (block group B), live in close proximity (C), and are in the same grade (G), regardless of school attended.²⁶ We refer to this group of students as the BCG group. When focusing only on the first row of the table, we see that it includes students who live in the same neighborhood, live in close proximity, are in the same grade, and attend the same school (S). We refer to this group as the $BCGS$ group. Similarly, the second row of the table comprises the $BCGNS$ group, as they do not attend the same school as the student of reference.²⁷ The third row includes students who live in the same neighborhood and are in the same grade and school, denoted as the BGS

²⁶Close neighbors are defined in the same way as in the cross-sectional section.

²⁷Note that $BCG = BCGS \cup BCGNS$.

group. The fourth row includes students in the same neighborhood, in the same grade, but in a different school, comprising the *BGNS* group. Finally, the last row includes all students in the same grade and school, the *SG* group. For reference, table 1.5 includes the above acronyms and their descriptions. Throughout this article, I follow the same naming convention.

Table 1.5: Peer group acronyms

Acronym	Description
BCG	Same neighborhood, close proximity, same grade
BCGS	Same neighborhood, close proximity, same grade, same school
BCGNS	Same neighborhood, close proximity, same grade, different school
BGS	Same neighborhood, same grade, same school
BGNS	Same neighborhood, same grade, different school
SG	Same grade, same school

In addition to outlining the relevant peer groups, it is also important to understand the relative sizes of these groups. This information is useful for better understanding the setting as well as for interpreting the results later on. I report group size statistics in table 1.6 below. We can see in the first three rows that the number of same-grade close neighbors is in the range of 2-3 students. The fifth and sixth rows indicate that the number of same-grade neighbors in the block group is on the order of 10-20 students. Finally, there are around 360 students in the same school and grade in this sample, on average.

Keeping the above information in mind, I next outline two linear-in-means regression models for detecting social interaction effects among close neighbors. Both models postulate that the outcomes of students in the same school and grade can have an effect on student outcomes and that same-grade peers in the neighborhood may also influence outcomes, but this separate effect may be different depending on whether the students attend the same school or not. However, the models take different stances on how neighbors in close geographical proximity influence student outcomes.

The first model attempts to understand whether social interactions with close neighbors matter for educational outcomes, independently of whether students attend the same school.

Table 1.6: Size of relevant peer groups

Group	Close neighbors		Close neighbors by school	
	Mean	Median	Mean	Median
BCG	2.6	2	-	-
BCGS	-	-	2.2	1
BCGNS	-	-	2.6	2
BGS	11.2	8	12.9	11
BGNS	16.5	14	18.3	16
SG	360.9	375	357.9	370

Notes: the two columns to the left include mean and median peer group sizes for students with same-grade close neighbors, regardless of the school they attend. The two columns on the right correspond to the sample of students with both same school and different school same-grade close neighbors. Students in both samples also have populated BGS, BGNS, and SG peer groups.

This model assumes that same-grade peers living in close geographical proximity may have an additional effect compared to other neighborhood peers, but that this effect is the same regardless of the school attended. The motivation for this feature is that while it is likely that attending the same school increases the probability of social interactions, at the micro-geographical level these interactions may be equivalent among students in different schools, given that close geographical proximity may foster social interactions as well. The regression specification for this model is:

$$Y_{iabst} = \beta_1 \bar{Y}_{iabgt}^{BCG} + \beta_2 \bar{Y}_{ibgst}^{BGS} + \beta_3 \bar{Y}_{ibgst}^{BGNS} + \beta_4 \bar{Y}_{igst}^{SG} + \beta_5 B_{bgt} + \beta_6 C_{abgt} + \delta_{ias} + \epsilon_{iabgst} \quad (1.4)$$

where Y_{iabst} is an educational outcome for student i with address a , block group b , grade level g , school s , and during year t . Further, \bar{Y}_{iabgt}^{BCG} is the average achievement of students who are close neighbors and are in the same grade, \bar{Y}_{ibgst}^{BGS} is the average achievement of same-grade students in the same neighborhood who attend the same school, \bar{Y}_{ibgst}^{BGNS} is the average achievement of same-grade students in the same neighborhood who attend a different school, and \bar{Y}_{igst}^{SG} is the average achievement of students in the same school and grade. Further, $B_{bgt} \equiv \bar{Y}_{bgt}^{BNG}$ and $C_{abgt} \equiv \bar{Y}_{abgt}^{BCNG}$ are control variables which flexibly capture general trends or shocks at the neighborhood (B) and

micro-neighborhood (C) level by utilizing the outcomes of students in different grades. Moreover, the variable δ_{ias} represents a student-address-school (SAS) fixed effect which isolates variation over time for a student residing in a given address while attending a given school. Finally, ϵ_{iabgst} is an error term.²⁸

The coefficient of interest is β_1 , which tests for the existence of social interactions with close neighbors. The identifying assumption for β_1 is that there do not exist any factors that affect changes in the outcomes of all students in a micro-neighborhood and in the same grade over time, beyond the influence of neighbors from the same school, neighbors from different schools, students in the same school, and general shocks at the neighborhood and micro-neighborhood levels. It is important to note that β_1 does not estimate a peer effect *per se*, since the above model is subject to the reflection problem, i.e. the student of reference may influence the outcomes of other students in their peer group, as well as the other way around. However, if peer effects did not exist, then β_1 would be equal to zero. As a result, I mainly focus on the sign of β_1 , testing for the existence of social interactions through a significance test of β_1 .

The second model goes beyond the first by attempting to disentangle whether social interactions with close neighbors matter for students in the same school, students in different schools, or both. This model postulates that same-grade peers in close geographical proximity can have differential effects depending on whether they attend the same school or not, just as I previously allowed to be the case at the neighborhood level. This allows for the possibility that despite being in close geographical proximity, social interactions with students in different schools are weaker than those with students in the same school. The regression specification for this model is:

$$\begin{aligned}
 Y_{iabgst} = & \beta_1 \bar{Y}_{iabgst}^{BCGS} + \beta_2 \bar{Y}_{iabgst}^{BCGNS} + \beta_3 \bar{Y}_{ibgst}^{BGS} + \beta_4 \bar{Y}_{ibgst}^{BGNS} \\
 & + \beta_5 \bar{Y}_{igst}^{SG} + \beta_6 B_{bgt} + \beta_7 C_{abgt} + \delta_{ias} + \epsilon_{iabgst}
 \end{aligned}
 \tag{1.5}$$

where all variables are defined as above except that now I separate \bar{Y}_{iabgt}^{BCG} into \bar{Y}_{iabgt}^{BCGS} and \bar{Y}_{iabgt}^{BCGNS} ,

²⁸Standard errors are clustered at the student and block group level.

depending on whether the close neighbors are in the same school or not. As explained above, this model allows for the possibility that social interactions with close neighbors are differential depending on the school attended.²⁹

The coefficients of interest are β_1 and β_2 , which capture whether there are social interactions with close neighbors in the same school or in different schools, respectively. I primarily test for these interactions through separate significance tests of β_1 and β_2 . The identifying assumption for β_1 is that there do not exist any factors that affect changes in the outcomes of all students in a micro-neighborhood, in the same grade, and in the same school over time, beyond those accounted for by the rest of the variables in the model. Likewise, identifying β_2 requires that there do not exist any factors that affect changes in the outcomes of all students in a micro-neighborhood and in the same grade over time, after controlling for all other variables.

Another issue that arises when it comes to interpreting these coefficients is that of measurement error. If the outcomes in this study are noisy proxies for underlying abilities and behaviors, then random shocks to these variables may lead to measurement error and subsequently to attenuation bias, especially in the case of smaller peer groups. If we look at the peer group sizes in table 1.6, we can see that some of the groups comprised of close neighbors are quite small, with 2-3 students on average. Therefore, we would expect some degree of attenuation bias for these coefficients, but we would expect this to be less of a problem as the groups become larger, as in the case of the school-grade variables (SG). As a result, it also becomes difficult to compare coefficients from peer groups with noticeably different sizes, so my primary focus will be to conduct significance tests of individual coefficients, as described above. Furthermore, it is important to note that attenuation bias would shift coefficients towards zero, so detecting social interactions through the proposed significance tests becomes more difficult, but is still meaningful.

Summary statistics for the longitudinal analysis are in table 1.7. The dependent variables

²⁹Standard errors are clustered at the student and block group level.

of interest are English and behavior. The independent variables are named according to their corresponding outcome and the peer group used to calculate them. For example, “English BCG” refers to the average English achievement of students who live in the same neighborhood, are in close proximity, and are in the same grade.

Table 1.7: Longitudinal sample summary statistics

	N	Mean	Std. Dev.	Min.	Max.
Year	170,800	2,008.269	2.945	2,003	2,013
Grade level	170,800	8.739	1.337	7	11
English	157,381	-0.208	0.958	-3.577	4.458
English BCG	164,915	-0.196	0.819	-3.492	4.391
English BGS	168,184	-0.198	0.641	-3.577	4.681
English BGNS	169,271	-0.146	0.546	-3.310	4.052
English SG	170,756	-0.138	0.429	-2.066	1.092
English C	155,949	-0.224	0.645	-3.577	3.997
English B	170,796	-0.165	0.418	-1.901	1.879
English BCGS	96,655	-0.170	0.899	-3.492	4.400
English BCGNS	110,093	-0.286	0.797	-3.474	3.851
Behavior	161,341	-0.127	1.022	-4.508	2.340
Behavior BCG	163,077	-0.131	0.845	-4.226	2.588
Behavior BGS	163,507	-0.117	0.641	-4.434	1.889
Behavior BGNS	164,579	-0.143	0.504	-3.542	2.588
Behavior SG	168,625	-0.102	0.413	-3.224	2.409
Behavior C	155,190	-0.155	0.642	-4.434	2.409
Behavior B	170,796	-0.113	0.357	-1.596	1.293
Behavior BCGS	97,658	-0.073	0.907	-4.214	2.340
Behavior BCGNS	105,780	-0.241	0.872	-4.508	2.588

Notes: the above table includes summary statistics for the longitudinal sample at the student-year level.

1.3.4 Additional Outcomes

While English, math, and behavioral outcomes are of first-order concern, there are other important educational outcomes and decisions that may be influenced by social interactions with close neighbors. In this section, I go over a methodology to study the influence of geographical proximity on outcomes related to which school students choose to attend. For example, I ask

whether living in close proximity leads students to attend the same school as their neighbors. The empirical strategy below utilizes variation across pairs of students and is therefore closely related to the cross-sectional methodology previously outlined for the main outcomes.

When it comes to making a decision about which school to attend, it is likely that parents have significant input in this process. They may also receive information from other parents nearby, regardless of the specific grade of their children. Therefore, I no longer focus on within- and cross-grade comparisons for these outcomes. Instead, I focus on pairs of students who live in the same neighborhood and are assigned to the same home school, regardless of their grade level.³⁰ I estimate the following model:

$$S_{ij} = \rho Close_{ij} + \beta X_{ij} + \delta_b + \delta_g + \varepsilon_{ij} \quad (1.6)$$

where S_{ij} is an indicator variable related to the school or type of school attended (e.g. whether the students in the pair attend the same school) and X_{ij} is a vector of control variables that includes indicator variables for the students being of the same race, gender, or parental education level. The remaining variables are defined as in the previous section. The coefficient of interest is ρ , which captures the differential likelihood that the pair of students matches on a school attendance decision if they live close by, compared to students living in the same neighborhood but not close by.³¹

The main threat to identification in this setting is the possibility that pairs of students who attend the same school (or have attended the same school in the past) move in right next to each other, driving a correlation between *Close* and school attendance decisions. Firstly, this is unlikely because while families may move in to the same neighborhoods as their friends, it may be difficult to locate an available unit or home in such a small area as the one captured by *Close*.

³⁰Restricting the sample to pairs of students with the same home school implicitly keeps students who are in similar grade levels, although not necessarily the same grade level.

³¹Standard errors are clustered at the block group level.

Secondly, I test for this possibility explicitly by selecting the sample of pairs where at least one of the students in the pair is new to the neighborhood in that year and estimating the following model:

$$Close_{ij} = \lambda SchoolPast_{ij} + \beta X_{ij} + \delta_b + \delta_g + \varepsilon_{ij} \quad (1.7)$$

where variables and fixed effects are defined as above and the variable $SchoolPast_{ij}$ is an indicator variable for whether the students in the pair have attended the same school in any of the previous three years. The coefficient of interest is λ , which captures whether having attended the same school in the past is predictive of living close by in the present, conditional on having moved into the same neighborhood.

1.4 Results

In this section, I go over the results of detecting social interaction effects with close neighbors. I present and interpret results for the cross-sectional study, the longitudinal study, and the study focusing on additional outcomes. I then go over alternative mechanisms that could explain the results.

1.4.1 Detecting Social Interactions Using Cross-Sectional Data

I estimate models (1.2) and (1.3) for English, math, and behavioral outcomes. However, it is important to note that these outcomes are not available for the exact same grades, so whenever possible I report the results for an outcome for a set of grades that matches the other two.³² I report the estimation results in table 1.8 below, with each panel corresponding to an outcome.

³²English scores are available for all grades in the sample. Behavioral scores become available starting in secondary school. Math scores are available for all grades as well, but after grade seven students enter separate tracks for the math CST exam so these are excluded from the analysis. Further, in the case of behavioral grades there is an additional sample selection, which is that students in charter schools are removed. This sample selection is largely implicit as the great majority of charter schools do not collect secondary behavioral grades, so they are missing from the data.

In the case of English outcomes, we can see in column (1) that students who live close by have English outcomes which are 3% of a standard deviation closer than those of students who live farther away, on average. In column (2), I conduct a validation exercise by estimating the same model but using pairs of students in different grades. We can see that the estimate is also negative and significant, indicating that either there are also social interactions with other grades or that there are factors varying at the micro-geographical level that influence proximity in outcomes for all grades, or both. However, the point estimate is about half the size than for students in the same grade, suggesting that the stronger observed effect in column (1) is due to social interactions among students in the same grade.³³ In order to attempt to interpret the magnitude of this coefficient, I compare it with the differential proximity in outcomes for students with the same parental education level, versus students with different parental education levels, using data from the entire school district.³⁴ The results of this exercise are in table 1.17. We can see that pairs of students with the same level of parental education are predicted to have outcomes that are 12% of a standard deviation closer than pairs of students with different levels of parental education. If we take the 3% estimate at face value, it would be 25% of the predicted difference using parental education. However, this is likely to be an upper bound for the effect since, as shown in column (2), it is possible that there are other factors that influence proximity in outcomes. If we conservatively attributed the entire effect in column (2) to confounding factors and none to interactions with other grades, then the additional correlation in column (1) would still be on the order of 1.5% of a standard deviation, or 12.5% of the predicted difference in outcomes using parental education. While this is certainly a back-of-the-envelope calculation, it hints that even though the effects are small in size, they are economically non-negligible.

In column (3) I explore whether the differential proximity is driven by students in the same

³³There are no baseline differences in columns (1)-(3), so the estimates can be easily compared from one column to another.

³⁴It is preferable to use data from the entire school district for this exercise since the selected block groups are already more homogeneous in terms of covariates and neighborhood characteristics. On top of utilizing data from the entire school district, this exercise also includes pairs of students who are in the same grade but reside in different neighborhoods.

school or in different schools, by interacting *Close* with an indicator for whether the students in the pair attend the same school. Given the significance level on the resulting estimate, we cannot rule out that the effect of proximity is the same for pairs of students in both the same and different schools. However, the implied confidence interval on this estimate is wide and includes both positive and negative values which are economically significant, both compared to the estimate on *Close* and to the predicted differences using parental education. Therefore, we do not gain any valuable insights from this particular regression about whether the effects are driven through students in the same or in different schools.

In terms of math scores, the results are in panel B of table 1.8. In column (1), we see that the point estimate on *Close* is positive. However, it is imprecisely estimated and insignificant, with its confidence interval including non-negligible positive and negative results, so this regression is not very informative (see table 1.17). In column (2), we use pairs in different grades as part of a validation exercise, which increases the sample size and in turn the precision of our estimates. The estimate in this case is close to zero and is more precisely estimated, indicating that we do not find any effect of proximity in this sample. Column (3) does not add any insight in the case of math achievement either, as the estimate on the interaction term is centered around zero but with a wide confidence interval. In general, these results are inconclusive, since the results in column (2) do not rule out that social interactions with close neighbors in the same grade matter for math achievement. Even though we are not confident in the null estimate in column (1), math ability is likely harder to transmit than English or behavioral traits, which could explain these results.

When it comes to behavioral outcomes, the effect of proximity in column (1) seems to be centered around zero, although the standard error allows for non-negligible effects in either direction.³⁵ In the validation exercise, the estimate is negative and significant, indicating that either there exist social interactions with other grades or that there are factors varying at the

³⁵I take the results in table 1.17 for this outcome with a grain of salt. While we do expect parental education to be predictive of English and math achievement *a priori*, this is not necessarily the case for behavioral marks. Therefore, it is possible that the estimate in table 1.17 is not a good benchmark.

Table 1.8: Cross-sectional social interaction effects on educational outcomes

	(1) Same Grade	(2) Different Grades	(3) Same Grade
<i>Panel A: English Outcomes</i>			
Close	-0.030*** (0.011)	-0.017** (0.008)	-0.039** (0.017)
Close×School			0.015 (0.027)
Block Group F.E.	Yes	Yes	Yes
Grade Level F.E.	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Observations	221,472	1,941,910	221,472
R ²	0.014	0.011	0.014
<i>Panel B: Math Outcomes</i>			
Close	0.013 (0.020)	-0.002 (0.011)	0.015 (0.027)
Close×School			-0.003 (0.039)
Block Group F.E.	Yes	Yes	Yes
Grade Level F.E.	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Observations	136,382	672,858	136,382
R ²	0.015	0.013	0.015
<i>Panel C: Behavioral Outcomes</i>			
Close	0.004 (0.018)	-0.018* (0.009)	0.004 (0.040)
Close×School			-0.00002 (0.054)
Block Group F.E.	Yes	Yes	Yes
Grade Level F.E.	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Observations	120,620	587,216	120,620
R ²	0.078	0.076	0.078

Notes: Statistical significance levels are indicated by *p<0.1; **p<0.05; ***p<0.01.

micro-geographical level that influence proximity in outcomes for all grades, or both.³⁶ However, it is worth noting that the significance could be explained by the larger sample size for this column. Further, the confidence interval for the estimate in column (1) includes negative effects of similar magnitude, so the results are not necessarily contradictory, but rather inconclusive. In column (3) we do not learn much either, as the estimate on the interaction term is centered around zero but with a wide confidence interval.

In table 1.15 in the appendix I report the within-grade estimates for English outcomes but instead use data from the grade levels available for math and behavioral grades, in order to understand whether the observed results are due to differences in the outcome or differences in the grade levels studied. In both samples, the estimates are only slightly weaker than in the full sample but they are no longer significant, which is expected given that the sample size has been approximately halved. This seems to indicate that the observed differences in results are due to differences in how the outcome is influenced by social interactions, instead of differences in the grade levels studied.

1.4.2 Detecting Social Interactions Using Longitudinal Data

In this section, I go over the results of detecting social interaction effects with close neighbors using variation over time within students. I estimate models (1.4) and (1.5) while focusing on secondary school English and behavioral outcomes.³⁷ The results of estimating these models, while varying which independent variables are included, are in columns (1)-(4) of tables 1.9-1.12. In order to address a possible source of attenuation bias, I also re-estimate these models for a subsample including students residing at addresses where they have lived for more than one

³⁶A possible explanation for this is that behavior may be more transmittable across grades than academic outcomes such as English or math. For example, children may observe and copy behaviors from students in other grades, especially those they admire, even though they may not interact with them as often as with peers in the same grade.

³⁷Math outcomes are excluded from this analysis because students enter separate tracks after grade seven. Further, since behavioral outcomes are only available starting in secondary school, and secondary school starts by grade seven in the majority of cases, the final analysis sample includes grades 7-11.

year. To facilitate comparison with results from the main sample, I include the results from this analysis in column (5) of the aforementioned tables, even though I address them towards the end of this section.

The results of estimating model (1.4) to detect social interaction effects for English outcomes are below in table 1.9. In the first column, we can see that there is a positive and significant association between close neighbors' English achievement and own English achievement. The association becomes weaker in column (2) when adding the achievement of students in the wider neighborhood, in both the same and in different schools. In column (3) the achievement of all students in the same school is added, making the coefficient on close neighbors' English insignificant. In column (4) I add controls for neighborhood and micro-neighborhood factors, and they do not influence the results. After having added all relevant controls, the coefficient on BCG is positive but not significant, so based on this analysis we cannot say that social interactions with close neighbors matter for English achievement.

Even though the coefficient on BCG is not significant using conventional tests, I argue that in this context it is also possible to use a one-sided significance test with $\beta_1 > 0$ as the alternative, albeit as supplementary evidence. The reason for this is that there is a strong prior for the sign of the coefficient to be positive; if we look at the review article by Sacerdote (2011), we see that 12 out of the 13 linear-in-means models featured show zero or positive peer effects in student outcomes. Among these 12 studies, Betts and Zau (2004) in fact use the same data source as this article. If we look at the bottom of table 1.9, column (4), we can see that the coefficient on BCG is significant at the 10% level. This result, combined with the possibility of measurement error driving coefficients towards zero, hints that social interactions with close neighbors may be driving English achievement, even though the evidence is weak.

In terms of the other variables in the model, we see that the association is stronger for neighbors in the same school (BGS) than for neighbors in different schools (BGNS). While this may hint towards same-school peers in the neighborhood being more influential, it is also

possible that the coefficient on BGS may capture the influence of shared schooling inputs. For example, students in the same neighborhood and school may have similar baseline achievement levels which influence achievement growth relative to the school's average growth rate. When it comes to the coefficient on SG, we see there is a strong correlation between average school English achievement and own achievement: an increase in average school English achievement of one standard deviation is associated with a 0.6 standard deviation increase in own English achievement. This coefficient stands out as it is noticeably larger than the rest of the coefficients in the table. I discuss three possible explanations for this. First, under the linear-in-means model, note that changes in the mean achievement of larger groups require changes in the achievement of a larger number of peers, on average. As a result, we might expect mean achievement of the SG group, often comprised of hundreds of students, to be more influential than the mean achievement of close neighbors, which usually includes 2-3 students.³⁸ Second, it is possible that this difference has to do with measurement error. As can be seen in table 1.6, groups other than SG may be susceptible to measurement error since their group sizes are smaller. If this is the case, we would expect some degree of attenuation bias in the rest of coefficients but none in the coefficient on SG. Third, while peer effects are likely to be part of the story, it is important to note that this variable may also proxy for school-level factors. For example, it may capture time-varying shocks that affect all students in that school or school effectiveness which drives achievement growth for all students.

Next, I examine the results of estimating model (1.4) for behavioral outcomes in table 1.10. The pattern is similar to that for English outcomes. We can see in column (1) that there is a positive association between close neighbors' behavioral outcomes and own outcomes, which is statistically significant. When we add more variables to the model, we can see that while the coefficient on BCG becomes smaller, it remains significant. In column (4), when all relevant variables are added, we find that a one standard deviation increase in close neighbors' behavioral

³⁸A related point is that the standard deviation of the SG variable is lower than that of most other variables in the model, since they have been calculated using smaller group sizes. This can be seen in table 1.7.

Table 1.9: Social interaction effects on English grades

	(1)	(2)	(3)	(4)	(5)
English BCG	0.023** (0.009)	0.016* (0.010)	0.013 (0.009)	0.013 (0.009)	0.018* (0.010)
English BGS		0.061*** (0.014)	0.024* (0.014)	0.024* (0.014)	0.026* (0.015)
English BGNS		0.016 (0.018)	0.019 (0.017)	0.018 (0.017)	0.024 (0.019)
English SG			0.605*** (0.053)	0.605*** (0.053)	0.612*** (0.065)
English B				0.030 (0.043)	0.031 (0.050)
English C				0.003 (0.010)	0.005 (0.011)
SAS F.E.	Yes	Yes	Yes	Yes	Yes
Settled Subsample	No	No	No	No	Yes
P-val. $H_1 : \beta_1 > 0$	0.006***	0.048**	0.088*	0.088*	0.045**
Observations	136,154	136,154	136,154	136,154	97,841
R ²	0.945	0.946	0.946	0.946	0.944
Adjusted R ²	0.828	0.828	0.830	0.830	0.830

Notes: This table reports the results of estimating model (1.4) for English achievement. Standard errors are clustered at the student and block group level. Statistical significance levels are indicated by *p<0.1; **p<0.05; ***p<0.01.

outcomes is associated with a 0.02 standard deviation increase in own behavioral outcomes, which is significant at the 5% level using both conventional and one-sided significance tests. This result provides evidence that social interactions with close neighbors matter for behavioral outcomes. If we compare this result to the one for English outcomes, it seems like social interactions with close neighbors are more relevant in the case of behavior. A possible explanation for this is that behavior is more transmittable than English, which relies on other inputs such as parental inputs or general test taking ability, which may be harder to learn from peers.

When it comes to the remaining variables in the model, the only coefficient that we can statistically distinguish from zero is the one on SG, the behavior of students in the same grade and school. In this case, an increase in average school behavior of one standard deviation is associated with a 0.7 standard deviation increase in own behavior, which is a strong relationship. As was the case when interpreting the coefficients for English outcomes, both group size and measurement error may play a role in its relative size. Moreover, this coefficient might also capture time-varying school level shocks, school level changes in behavioral grade assignment policies, and teacher composition changes that affect how behavioral grades are assigned.

Up until this point I have studied social interactions with close neighbors regardless of the school attended by the neighbors. However, it is possible that close neighbors may have a differential effect on outcomes depending on if they attend the same school, as social interactions may be stronger in this case. In order to explore this question, I restrict the sample to students who have neighbors attending both the same school and different schools. Using this sample, I estimate model (1.5) for English and behavioral outcomes.

I present the results of estimating model (1.5) for English achievement in table 1.11. In the first column it appears that the English achievement level of close neighbors in the same school is positively and significantly associated with own English achievement, but that this is not the case for close neighbors attending different schools. As more variables are added, the coefficient on the BCGS variable becomes weaker and insignificant, while the coefficient on BCGNS is unchanged. In column (4), the point estimate is 0.027 for BCGS and 0.011 for BCGNS. Even though these point estimates are of similar order than those in the first row of table 1.9, the estimation sample for model (1.5) is smaller, leading to a loss of statistical power. Further, even though neither coefficient is significant, it is worth noting that the coefficient on BCGS is over twice the size of the BCGNS coefficient, which is small and close to zero. Keeping in mind the potential problem of measurement error, this may hint that same school peers are more influential when it comes to English outcomes. A possible explanation for this would be

Table 1.10: Social interaction effects on behavioral grades

	(1)	(2)	(3)	(4)	(5)
Behavior BCG	0.041*** (0.010)	0.028*** (0.010)	0.021** (0.009)	0.021** (0.009)	0.024** (0.010)
Behavior BGS		0.105*** (0.019)	0.012 (0.016)	0.012 (0.016)	0.012 (0.017)
Behavior BGNS		0.007 (0.017)	0.015 (0.017)	0.015 (0.017)	0.006 (0.018)
Behavior SG			0.692*** (0.062)	0.692*** (0.061)	0.682*** (0.060)
Behavior B				0.055 (0.056)	0.057 (0.060)
Behavior C				0.015 (0.012)	0.019 (0.012)
SAS F.E.	Yes	Yes	Yes	Yes	Yes
Settled Subsample	No	No	No	No	Yes
P-val. $H_1 : \beta_1 > 0$	0.000***	0.002***	0.013**	0.014**	0.009***
Observations	137,269	137,269	137,269	137,269	97,393
R ²	0.925	0.925	0.928	0.928	0.926
Adjusted R ²	0.758	0.760	0.767	0.767	0.772

Notes: This table reports the results of estimating model (1.4) for behavioral outcomes. Standard errors are clustered at the student and block group level. Statistical significance levels are indicated by *p<0.1; **p<0.05; ***p<0.01.

that social interactions with students at school happen in a setting where the transmission of academic knowledge is possible, whereas interactions that happen only in the neighborhood are not as conducive to these types of exchanges.

In terms of behavioral grades, the results of estimating model (1.5) are in table 1.12. If we focus on column (4), we can see that the coefficients on BCGS and BCGNS are both positive, with the point estimate being 0.033 for BCGS and 0.034 for BCGNS. However, neither coefficient is statistically significant using conventional tests. If we consider the one-sided test instead, then the

Table 1.11: Social interaction effects on English grades by school attended

	(1)	(2)	(3)	(4)	(5)
English BCGS	0.045** (0.023)	0.030 (0.025)	0.027 (0.025)	0.027 (0.025)	0.029 (0.028)
English BCGNS	0.011 (0.022)	0.010 (0.021)	0.011 (0.021)	0.011 (0.021)	0.012 (0.024)
English BGS		0.050 (0.038)	-0.006 (0.037)	-0.005 (0.037)	-0.002 (0.044)
English BGNS		0.006 (0.042)	0.008 (0.042)	0.008 (0.042)	0.018 (0.050)
English SG			0.635*** (0.125)	0.634*** (0.125)	0.613*** (0.160)
English B				0.068 (0.104)	0.069 (0.125)
English C				0.003 (0.034)	0.017 (0.041)
SAS F.E.	Yes	Yes	Yes	Yes	Yes
Settled Subsample	No	No	No	No	Yes
P-val. $H_1 : \beta_1 > 0$	0.024**	0.113	0.138	0.14	0.151
P-val. $H_1 : \beta_2 > 0$	0.308	0.324	0.302	0.306	0.315
Observations	37,506	37,506	37,506	37,506	25,596
R ²	0.951	0.951	0.952	0.952	0.948
Adjusted R ²	0.808	0.808	0.811	0.811	0.806

Notes: This table reports the results of estimating model (1.5) for English achievement. Standard errors are clustered at the student and block group level. Statistical significance levels are indicated by * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

coefficient on BCGNS becomes significant at the 10% level. While these results are suggestive at best they seem to indicate that, when it comes to behavior, students in different schools may be just as influential as students in the same school. This would align with the idea that behavior can be transmitted through interactions in the neighborhood whereas English or test taking skills may

be harder to transmit outside of the school environment.

Table 1.12: Social interaction effects on behavioral grades by school attended

	(1)	(2)	(3)	(4)	(5)
Behavior BCGS	0.072*** (0.026)	0.039 (0.028)	0.034 (0.028)	0.033 (0.028)	0.040 (0.030)
Behavior BCGNS	0.033 (0.021)	0.032 (0.023)	0.034 (0.022)	0.034 (0.022)	0.036 (0.023)
Behavior BGS		0.111* (0.062)	-0.009 (0.057)	-0.009 (0.057)	-0.024 (0.053)
Behavior BGNS		0.005 (0.053)	0.003 (0.051)	0.004 (0.051)	-0.018 (0.056)
Behavior SG			0.632*** (0.134)	0.629*** (0.134)	0.647*** (0.141)
Behavior B				0.133 (0.119)	0.119 (0.127)
Behavior C				0.016 (0.040)	0.034 (0.044)
SAS F.E.	Yes	Yes	Yes	Yes	Yes
Settled Subsample	No	No	No	No	Yes
P-val. $H_1 : \beta_1 > 0$	0.003***	0.08*	0.115	0.119	0.09*
P-val. $H_1 : \beta_2 > 0$	0.056*	0.084*	0.06*	0.062*	0.061*
Observations	39,149	39,149	39,149	39,149	26,481
R ²	0.936	0.936	0.937	0.937	0.934
Adjusted R ²	0.742	0.743	0.749	0.749	0.751

Notes: This table reports the results of estimating model (1.5) for behavioral outcomes. Standard errors are clustered at the student and block group level. Statistical significance levels are indicated by *p<0.1; **p<0.05; ***p<0.01.

While interpreting the results in this section, I have discussed measurement error as source of attenuation bias. However, I argue that there is potentially another source of attenuation bias in this setting. When a student arrives at an address, it is possible that during the first year it is

more difficult to make strong social connections with close neighbors. If this is the case, it would result in neighbors' outcomes mattering less for own outcomes for that year, thus leading to noise and biasing estimates towards zero.³⁹ In order to address this problem, I re-estimate models (1.4) and (1.5) on a sample of students who have been at an address for more than one year. That is, I exclude data from the first year a student is at an address. I report the results for each outcome in column (5) of tables 1.9-1.12, while including all relevant control variables.

In the case of English outcomes, we can see in column (5) of table 1.9 that the coefficient on BCG is now stronger and statistically significant at the 10% level. This supports the notion that the lack of social connections during the first year at an address may lead to attenuation bias. Further, it shows that social interactions with close neighbors matter for English outcomes, at least in this subsample. In table 1.11 I report the results of estimating model (1.5) for this subsample. Focusing on column (5), we see the results are quite similar to those of the main sample; they are slightly larger but neither is statistically significant. However, the pattern that the estimate on BCGS is over twice as large as the estimate on BCGNS persists.

In terms of behavioral outcomes, the results are similar to those using the main sample. In column (5) of table 1.10, the estimate on BCG is only slightly stronger than in the main sample, so the implications are unchanged. In table 1.12, the estimates are again slightly stronger but not yet significant using two-sided significance tests. However, in this case both are significant at the 10% level using a one-sided test. Considering that measurement error may be biasing coefficients towards zero, these results are suggestive that social interactions with close neighbors influence behavioral outcomes both when it comes to students in the same school and students in different schools. However, the evidence is certainly weak given the level of statistical significance of the estimates.

³⁹Turley (2003) suggest that neighborhood effects are stronger for students who have lived in a neighborhood for a longer period of time.

1.4.3 Additional Outcomes

In this section, I go over the results of estimating model (1.6) for three outcomes related to the school attended by students. The first outcome is an indicator for whether students attend the same school. The second outcome is an indicator for whether the students in the pair either both attend their home school or both attend a school other than their home school (I refer to a student as a *resident* if they attend their home school). The third is an indicator for whether the students in the pair either both attend a charter school or both attend a regular school. The estimation results are in table 1.13.

The first column of table 1.13 studies how geographical proximity influences whether students with the same home school attend the same school. We can see that students living in close proximity are 1.5% more likely to attend the same school, compared to students living in the same neighborhood but farther away. With 64% of students in the same attendance zone actually attending the same school in the sample, this effect represents 2.4% of the baseline, which is economically meaningful. Similarly, in column (2) we see that students living in close proximity are also 1.5% more likely to make the same decision to attend or not attend their home school, compared to students living in the same neighborhood but farther away. In this case 71% of students in the same attendance zone make the same residency decision, so the effect is about 2.1% of the baseline, which is also meaningful. In the case of the decision to attend or not attend a charter school, we cannot say whether students appear to be more likely to make the same decision if they live in close proximity, as the estimate in column (3) is not significant and is close to zero.

As discussed above, a possible threat to identification is that students who have attended the same school in the past may move in next door to each other, driving a correlation between geographical proximity and school attended, if that makes them more likely to attend the same school in the present. I explicitly test for this by estimating model (5) and report the results in table 1.16 in the appendix. In column (1), I include all pairs of students where there is a newcomer

Table 1.13: Social interaction effects on the school attended by students

	(1)	(2)	(3)
	School	Resident	Charter
Close	0.015** (0.007)	0.015** (0.007)	0.005 (0.004)
Block Group F.E.	Yes	Yes	Yes
Grade Level F.E.	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Observations	737,980	737,980	737,980
R ²	0.155	0.062	0.091

Notes: Statistical significance levels are indicated by *p<0.1; **p<0.05; ***p<0.01.

to the neighborhood, regardless of grade level (but ensuring that in the present they are assigned to the same home school). In column (2), I include only pairs in the same grade level, as the results may be different for this sample. As we can see, in both columns the estimates are very close to zero and precisely estimated. This indicates that having attended the same school in the past is not predictive of living in close geographical proximity later on, conditional on having moved to the same neighborhood. These results lend greater validity to the hypothesis that the estimates observed in table 1.13 are the result of social interactions among close neighbors.

I next explore whether there is any heterogeneity in these effects by income level. This exercise is motivated by Schneider, Teske, and Marschall (2002), who suggest that families with higher socioeconomic status rely more on word of mouth when picking schools, compared to families with lower socioeconomic status, who may rely more on publicly available test scores. In order to do this, I estimate model (1.6) separately for high income and low income neighborhoods and report the results in table 1.18.⁴⁰ According to the estimates in this table, it does seem like social interactions are more influential when choosing schools for families with higher socioeconomic status. For example, we can see that students living in close proximity are 1.8%

⁴⁰I use tract level income data from the 2010 U.S. Census (Geolytics Inc., 2012). I then categorize block groups in tracts with above median income as high income, and the rest as low income.

more likely to attend the same school in high income neighborhoods, whereas this effect is 1.3% in low income neighborhoods and is not significant. When it comes to the decision to attend the assigned home school, the effect is 2.3% in high income neighborhoods compared to a statistically insignificant estimate of 0.8% in the case of low income neighborhoods. There is no difference in the case of charter school attendance.

1.4.4 Mechanisms and Limitations

Throughout this article, I have suggested that the observed correlations in outcomes are due to social interactions among close neighbors. In this section, I address alternative mechanisms that could be driving these results.

In the case of English outcomes, in the longitudinal study I suggested that students in the same neighborhood may share schooling inputs (such as baseline achievement levels) which drive outcome growth relative to the school average. One concern is that students in the same micro-neighborhood are even more likely to share schooling inputs than are students in the same neighborhood, which would explain joint outcome growth at that level compared to the neighborhood level. While this is a possibility, I argue that it is unlikely to influence the results since I control for micro-neighborhood level factors through the outcomes of close neighbors in different grades (the *C* variable). Even though this variable includes students from other grades, these students should share the same micro-neighborhood level schooling inputs and experience similar differential growth rates.

In terms of behavioral outcomes, one concern is that students at the micro-neighborhood level are more likely to take the same classes, compared to students living in the same neighborhood but not in the same micro-neighborhood. If this is the case, then changes over time in the classes taken could lead to changes in behavioral outcomes driven by teacher composition, since teachers may follow idiosyncratic standards for assigning behavioral grades. While this mechanism is not the focus of this study, it may still suggest that social interactions with close

neighbors influence important educational decisions, such as what classes to take, which can have long-term economic impacts. Another possible mechanism in this case is that teachers assign similar behavioral grades to students who are friends, so while geographical proximity may drive behavioral grades it may not be due to any real-life changes in behavior but instead due to friendship links. However, if either of these mechanisms were the driving force behind the results, the coefficient on BCGS in table 1.12 would be significantly stronger than the coefficient on BCGNS, since these are not plausible mechanisms unless the students attend the same school. As we can see in table 1.12, the coefficients are quite similar in magnitude, so these explanations are unlikely.

When it comes to the the school attended variables in the additional outcomes section, it is possible that close neighbors are more likely to attend the same school if hyper-local amenities influence the desirability of some schools over others. For example, if the students in a micro-neighborhood are very close to a bus line for a school, then they may be more likely to attend the same school compared to students farther away from the bus line. However, this seems unlikely since the observed effects would have to be explained by differences in marginal distances to these amenities, which are quite small given the magnitude of block groups. Finally, another reason why close neighbors may be more likely to attend the same school is the convenience this offers as far as joint transport. Even so, this decision may also be fostered by social interactions.

1.5 Conclusion

In this paper, I study how social interactions with close neighbors influence students' educational outcomes. Using two distinct identification strategies – one cross-sectional and one longitudinal – I find evidence that social interactions with close neighbors matter both in the case of English and behavioral outcomes. I also find evidence that they influence students' decision of which school to attend. Further, my results are consistent with social interactions with close

neighbors in different schools mattering as much as those with students in the same school for behavioral outcomes, but not for English outcomes. However, I cannot statistically rule out that these are null results. Finally, I do not find enough evidence to determine whether these interactions matter for math outcomes.

The results in this paper point to the importance of close neighbors in determining educational outcomes. They also suggest that micro-neighborhoods are to neighborhoods what classrooms are to schools: a setting deserving of special attention where students can have an additional influence on each other. More generally, these results provide further evidence that social interactions among students shape educational outcomes as well as choices. In terms of policy implications, this article highlights an additional avenue that policymakers should take into account when considering spillover effects of educational interventions. For example, even if some interventions are restricted to certain groups of students within a school, they may also have an effect on the close neighbors of the affected students. Further, while more evidence is needed on the influence of close neighbors in different schools, policymakers should consider close neighbor relationships as a channel through which the effects of school-specific policies can affect students in other schools, especially in settings with school choice.

Finally, more research is needed to fully understand the role of close neighbors in shaping educational outcomes. In particular, the use of social network data that include information on students in different schools is promising. If it were possible to gather these data, it would permit researchers to study friendship link formation among close neighbors who do not attend the same school, thus shedding light on possible mechanisms. Further, more evidence is needed on the magnitude of the effects. One possible solution is to study policy spillovers on close neighbors, as described above, provided that these policies have large enough first order effects on the affected students.

1.6 Acknowledgement

Chapter 1 is currently being prepared for submission for publication of the material. Ruiz Junco, Pablo; “Social Interactions With Close Neighbors: Do They Matter for Educational Outcomes?” The dissertation author was the sole investigator and author of this material.

1.7 Appendix

1.7.1 Additional Figures

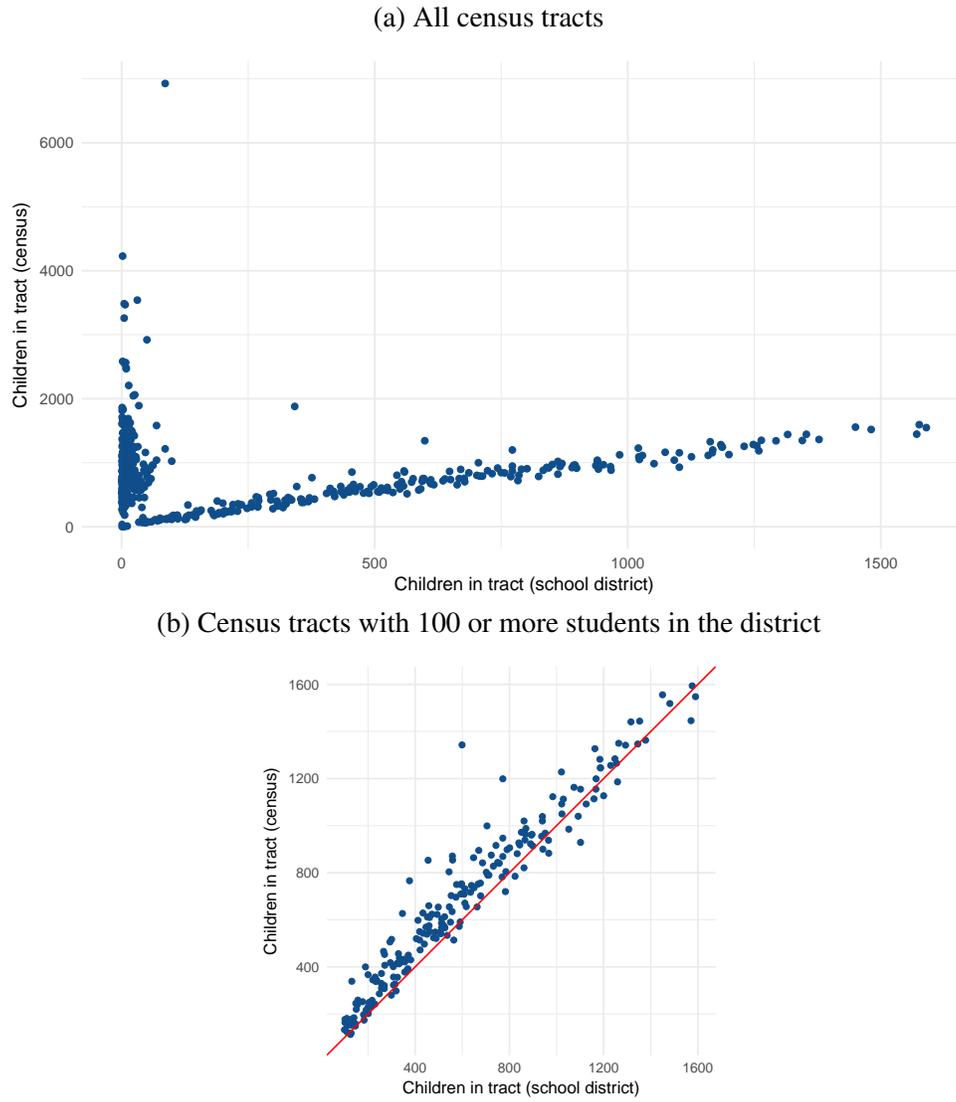


Figure 1.1: Scatterplots of children in the census versus the school district

Notes: the above figures plot the amount of children aged 5-18 in a census tract (y-axis) compared to the amount of enrolled SDUSD students in grades K-12 in each census tract (x-axis). The data are from the year 2010. Subfigure (b) also includes a red line denoting the 45-degree line.

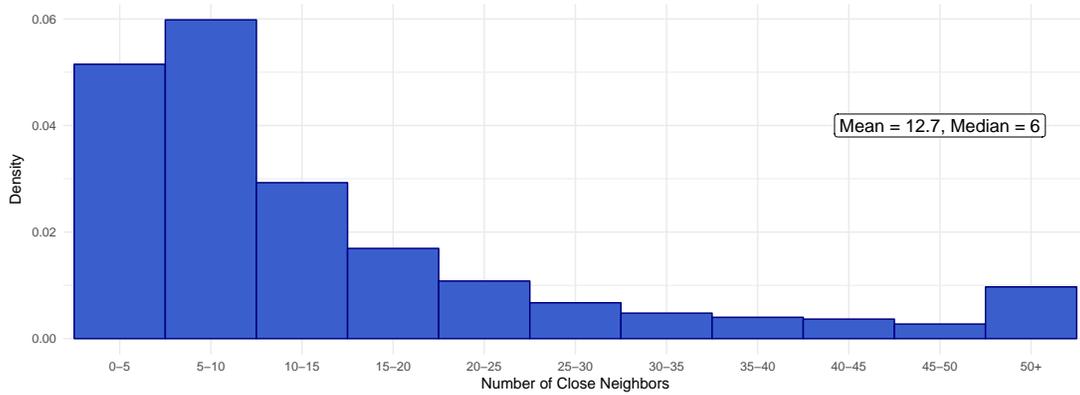


Figure 1.2: Distribution of the number of close neighbors

Notes: the above figure plots the density of the number of close neighbors for each SDUSD student in 2010. The figure includes neighbors of all grade levels in the main sample (2-11). If we restrict the sample to students with at least one close neighbor, the average number of close neighbors is 14.

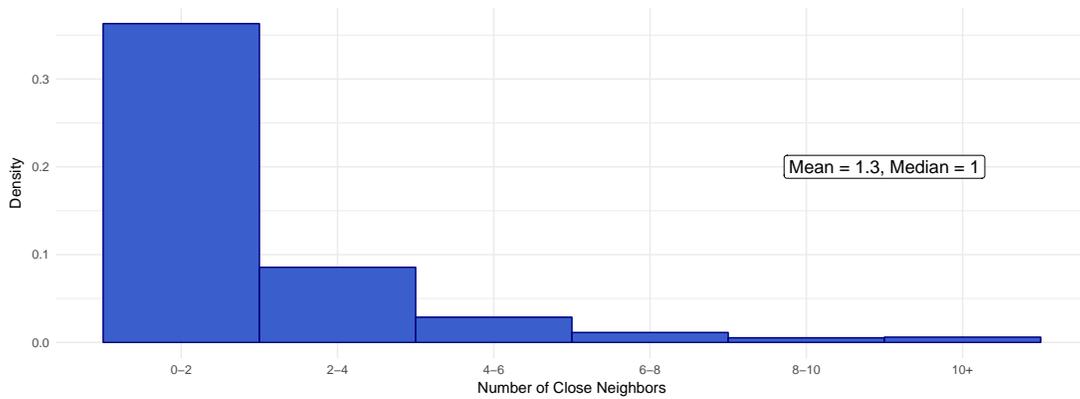


Figure 1.3: Distribution of the number of same-grade close neighbors

Notes: the above figure plots the density of the number of same-grade close neighbors for each SDUSD student in 2010. The figure includes students in grades 2-11. If we restrict the sample to students with at least one same-grade close neighbor, the average number of same-grade close neighbors is 2.5



Figure 1.4: Example of a block group and a micro-neighborhood

1.7.2 Additional Tables

Table 1.14: Comparison of single-family and other block groups

	Mixed B.G.	Single-family B.G.
Parents H.S.	0.83	0.96
Parents college	0.34	0.57
White	0.28	0.58
Hispanic	0.44	0.20
Black	0.12	0.06
Asian	0.15	0.15
Fraction B.G.	0.78	0.22
Fraction students	0.83	0.17
Median income	\$67,212	\$98,681

Table 1.15: English outcomes for math and behavioral sample

	Math Grades Sample (1)	Behavioral Grades Sample (2)
Close	-0.024 (0.016)	-0.022 (0.018)
Block Group F.E.	Yes	Yes
Grade Level F.E.	Yes	Yes
Controls	Yes	Yes
Observations	137,110	132,314
R ²	0.019	0.020

Notes: Statistical significance levels are indicated by *p<0.1; **p<0.05; ***p<0.01.

Table 1.16: Predicting current geographical proximity with past school attendance

	(1) All Grades	(2) Same Grade
School Past	-0.002 (0.003)	0.0001 (0.004)
Block Group F.E.	Yes	Yes
Grade Level F.E.	Yes	Yes
Controls	Yes	Yes
Observations	49,518	15,286
R ²	0.017	0.032

Notes: Statistical significance levels are indicated by *p<0.1; **p<0.05; ***p<0.01.

Table 1.17: Predicting pair-level outcome differences using parental education similarity

	(1) English	(2) Math	(3) Behavioral
Parental Education	-0.123*** (0.004)	-0.098*** (0.005)	-0.066*** (0.004)
Observations	12,454,002	7,295,563	6,384,043
R ²	0.003	0.002	0.001

Notes: statistical significance levels are indicated by *p<0.1; **p<0.05; ***p<0.01.

Table 1.18: School attended effects by income level

	<i>High income</i>			<i>Low income</i>		
	School (1)	Resident (2)	Charter (3)	School (4)	Resident (5)	Charter (6)
Close	0.018** (0.008)	0.023*** (0.007)	0.005 (0.005)	0.013 (0.011)	0.008 (0.010)	0.005 (0.006)
Block Group F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Grade Level F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	441,958	441,958	441,958	296,022	296,022	296,022
R ²	0.178	0.086	0.086	0.117	0.026	0.102

Notes: Statistical significance levels are indicated by *p<0.1; **p<0.05; ***p<0.01.

1.7.3 Clustering of Students Based on Educational Outcomes

I use a k -means clustering algorithm to organize students into similar groups according to their educational outcomes. For details on the k -means algorithm, see Hastie, Tibshirani, and Friedman (2009). While k -means is unsupervised, it requires setting the number of clusters (k) manually. In order to pick an adequate number of clusters, I follow the elbow method outlined in figure 1.5 below. Notice from the figure that as the number of clusters grows, the within-cluster sum-of-squares decreases, indicating that the students inside each cluster are more similar to each other. The goal of the elbow method is then to choose the number of clusters such that adding more clusters does not significantly reduce the within-cluster sum-of-squares, thus balancing within-cluster similarity and insight. To accomplish this, I choose five as the number of clusters and report the results in table 1.19 and provide a visual representation of the groups in figure 1.6.

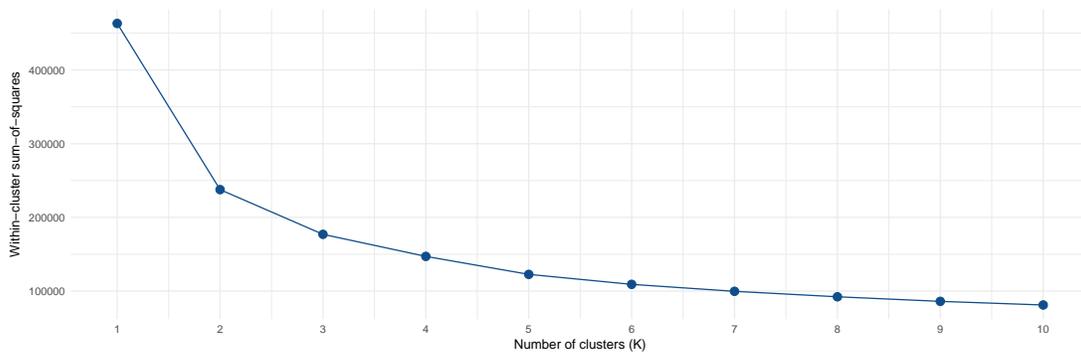


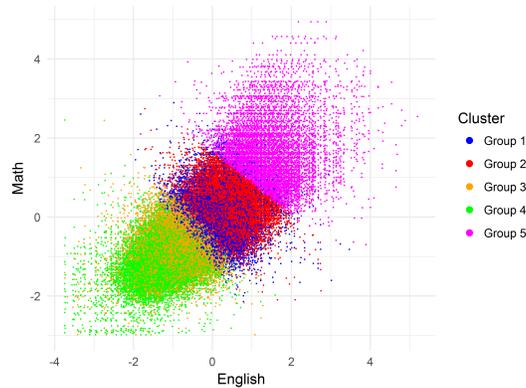
Figure 1.5: Elbow curve for cluster selection

Table 1.19: Results of k -means clustering of student outcomes ($k = 5$)

Group	Description ($H/M/L$: high/medium/low)	Fraction	<i>Group centroid:</i>		
			English	Math	Behavior
1	English M , math M , behavior M	0.17	0.15	0.04	-0.79
2	English M , math M , behavior H	0.29	0.41	0.30	0.70
3	English L , math L , behavior M	0.22	-0.81	-0.77	0.19
4	English L , math L , behavior L	0.15	-1.19	-1.09	-1.50
5	English H , math H , behavior H	0.17	1.40	1.57	0.85

Notes: this table reports cluster centroids and cluster sizes resulting from running k -means on the sample of students having complete and valid English, math, and behavioral outcomes, for all years in the sample pooled together. $N = 153,551$.

(a) Math vs. English dimension



(b) Behavior vs. English dimension

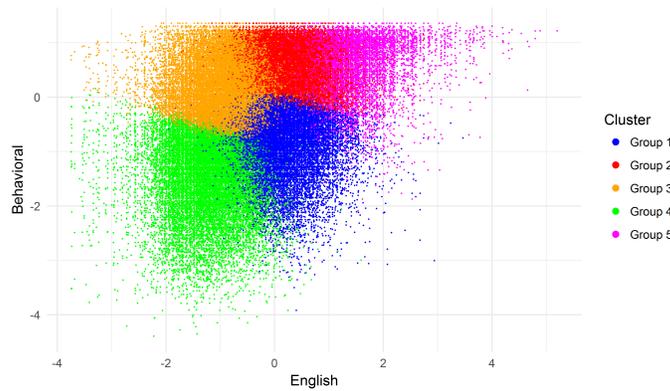


Figure 1.6: Visual representation of clustering results

Chapter 2

Extreme Weather and the Politics of Climate Change: a Study of Campaign Contributions and Elections

2.1 Introduction

Public opinion on key issues is thought to play a crucial role in shaping policies and elections in a democracy. Therefore, it is important to understand what factors contribute to the formation of these opinions as well as their political ramifications. One such key issue is climate change, which has received significant policy attention for decades.

In the United States, both the public and legislators remain divided when it comes to climate change, which has prompted several studies of Americans' attitudes towards it (Akerlof, Maibach, Fitzgerald, Ceden, & Neuman, 2013; Howe, Mildemberger, Marlon, & Leiserowitz, 2015; Myers, Maibach, Roser-Renouf, Akerlof, & Leiserowitz, 2013; Spence, Poortinga, Butler, & Pidgeon, 2011; Zaval, Keenan, Johnson, & Weber, 2014). These studies document great heterogeneity in beliefs across the country and provide important insights into the factors that

shape such beliefs. In particular, as extreme weather events and natural disasters have become widely associated with climate change (IPCC, 2013), one recurring finding is that exposure to these events leads individuals to change their reported perceptions of the issue.

Even though weather anomalies have been shown to influence stated beliefs regarding climate change, less is known about how they may impact costly, real-world actions. Importantly, people may misreport their true preferences due to social or strategic considerations, so surveys of stated beliefs can be misleading. Furthermore, it is unclear whether these extreme weather events can, through changes in beliefs, have real-world political consequences. In particular, it remains an open question whether politicians will be held accountable for their positions on environmental issues as beliefs regarding climate change vary.¹

In this paper, we present direct evidence of campaign finance and electoral responses to extreme weather events. We assemble a comprehensive dataset of extreme weather shocks, natural disasters, and U.S. House of Representative elections. These data allow us to examine various response margins, from campaign contributions to the competitiveness of elections and their outcomes. Moreover, we collect information on the environmental voting records of members of Congress to assess where they stand on the *anti-environment* to *pro-environment* spectrum.² Importantly, we test for differential effects of weather and disaster shocks depending on the environmental stance of incumbent politicians. This approach helps us shed light on whether environmental ideology is a driver of campaign contributions and political support for candidates in elections. Our results document a margin of political behavior in this context that is, to the best of our knowledge, novel in the literature. They also shed light on the mechanisms through which public opinion may shape climate change adaptation and mitigation policies.

Our study follows the literature closely in terms of choosing regression frameworks and constructing measures of weather shocks. In terms of their empirical strategies, previous studies

¹One reason changes in beliefs might not lead to political consequences is that climate change may not always be a top policy priority (Davis & Wurth, 2003; Guber, 2001).

²These terms are used for concise communication with the reader and do not necessarily represent the views of the authors on these issues or the politicians involved.

can be classified into three categories. The first set of studies focuses on short-run weather shocks, over a period of a month or less (Egan & Mullin, 2012; Hamilton & Stampone, 2013; Joireman, Truelove, & Duell, 2010; Li, Johnson, & Zaval, 2011; Zaval et al., 2014). The second set of studies uses medium-run temperature shocks, over a period of a month to a year (Deryugina, 2013). The third set of studies focuses on medium-run natural disaster shocks, also over a period of a month to a year (Lang & Ryder, 2016; Sisco, Bosetti, & Weber, 2017; Spence et al., 2011). In order to relate to the existing literature, in our study we focus on both short-run and medium-run temperature variations as well as medium-run natural disaster shocks.

In the short-run analysis, we examine how weekly temperature shocks affect contributions to Democratic candidates through ActBlue, an online fundraising platform, over the 2006-2012 period. The identification relies on two features. First, temperature shocks are measured by deviations of weekly mean temperature from the historical average in the same month and location, which eliminates most cross-sectional variation and seasonality that may be correlated with unobserved confounding factors. Second, we control for a rich set of fixed effects including county, week-in-sample, and state-by-election cycle. The results reveal an extensive-margin response: a 1°F increase in weekly average temperature has a contemporaneous effect of a 1.2% increase in the contribution rate, and a cumulative monthly effect of 2.7%. In our main analysis, we do not detect any intensive-margin effect. When looking across incumbent characteristics such as party membership and environmental attitudes, we find stronger responses to temperature shocks among constituents with more anti-environment incumbents. Together, these results suggest that following a temperature shock, Democratic candidates are rewarded for their pro-environment stance when the incumbents in the contributor's district are more anti-environment. This effect is mainly driven by an increase in the number of contributions.

In the medium-run analysis, we start by exploring how natural disasters in an election cycle interact with an incumbent's stance on environmental issues to influence both campaign finance and electoral outcomes. We examine the universe of individual and political action

committee (PAC) contributions to candidates in U.S. House of Representatives races during election cycles 1990-2012. Importantly, this analysis includes both Republican and Democratic candidates. We use two distinct regression specifications. The first is our main specification and exploits variation in the incumbent's stance on environmental issues regardless of party affiliation, whereas the second utilizes within-party variation. Using the main specification, we find that after a natural disaster, total fundraising and the number of donors in an election cycle is higher if the incumbent has a more anti-environment stance, and there is some suggestive evidence that the effect is stronger for donations to challengers than for incumbents. Further, we find that after a disaster, the more anti-environment the incumbent is the higher the chance of a challenger entering the race, leading to a slightly lower re-election probability for the incumbent. When we use within-party variation in incumbents' environmental stances, we find that estimates generally point in the same direction as before, but they are now imprecisely estimated and statistically indistinguishable from zero. The one exception is the probability of an incumbent entering the race as a response to a natural disaster, which is still higher the more anti-environment the incumbent is and remains statistically significant. In addition to natural disasters, we also focus on medium-run temperature shocks by using a similar methodology but treating hot and cold cycles as separate events. While we do not have enough power to rule out null results, the magnitude and direction of the effects of hot weather events are similar to that of natural disasters. In the case of cold weather events, we find that the effects are typically opposite in sign, which points to the possibility that in this context people react differently to hot and cold weather abnormalities. These results complement the short-run analysis since we are able to utilize the universe of contributions to House candidates, both Republican and Democrat, online and offline.

We explore a number of possible explanations for our results and argue that the most plausible explanation is the *environmental preference mechanism*. We suggest that there are two margins at play within this mechanism. First, extreme weather events prompt people to update their beliefs about climate change and become more politically active as a result. Second, extreme

weather events make environmental issues more salient, serving as a call to action for people with static beliefs who are politically active. Furthermore, we note that in a competitive environment such as the electoral system, people on both sides of the climate change debate may be galvanized by extreme weather events, either independently or as a response to the other side's actions.

This paper contributes to several research areas. Firstly, it is one of the few existing studies that adopt a revealed preference approach to studying the effects of weather shocks on people's beliefs about climate change. Previous studies have focused on outcomes such as Google searches (Herrnstadt & Muehlegger, 2014; Lang & Ryder, 2016) and Twitter posts (Sisco et al., 2017; Moore, Obradovich, Lehner, & Baylis, 2019). Li et al. (2011) show that respondents in a survey donated more money to an environmental charity if they thought that day was warmer than usual, although this donation came from the fee they were awarded for completing the study. Importantly, the outcomes in this study are more costly and spontaneous, thereby providing more meaningful reflections of people's beliefs. Moreover, we also directly examine the political processes through which public opinion on climate change might shape policies, which has been a goal of this literature.

Secondly, our results contribute to the current understanding of the motivations for political giving. Our results show that the responses to short-run temperature shocks in terms of ActBlue contributions occur mostly on the extensive margin. This is consistent with the mainstream view that individuals make campaign contributions for ideological reasons (Barber, 2016; Bonica, 2014; Ensley, 2009; Francia, Green, Herrnson, Powell, & Wilcox, 2003) and that they derive direct utility from contributing to their candidate of choice, as if they were consuming an ideologically-motivated consumption good (Ansolabehere, de Figueiredo, & Snyder, 2003). Our findings are also consistent with the idea that online fundraising platforms like ActBlue have enabled such "political consumption" by significantly lowering transaction costs (Karpf, 2013). Our medium-run results provide insights into PAC contributions, whose motivation has not been unanimously agreed upon in the literature. While a prevalent theory is that PAC contributions

have a *quid pro quo* nature, recent studies reveal that ideological considerations are also at play (Barber, 2016; Bonica, 2013, 2014, 2013; Snyder, 1990). Our evidence that PAC contributions also respond to natural disaster shocks lends further support to the ideological mechanism.

Thirdly, our results shed light on whether politicians are held accountable by constituents for their policy positions. On one hand, there is evidence to suggest that voters make seemingly irrational decisions, since same-day weather conditions or financial windfalls from lotteries have been shown to affect voting outcomes (Gomez, Hansford, & Krause, 2007; Bagues & Esteve-Volart, 2016; Meier, Schmid, & Stutzer, 2016; Jachimowicz, Menges, & Galinsky, 2017). On the other hand, there is also evidence that incumbents are held partially accountable for their roles in disaster preparedness and post-disaster relief. For example, Arceneaux and Stein (2006) show that voters punish the incumbent mayor following a flood if they believed the city was responsible for flood preparation. Gasper and Reeves (2011) show that electorates may punish presidents and governors for severe weather damage, but that only the president is punished if a request by the government for federal assistance is rejected. Healy and Malhotra (2009) show that voters reward the incumbent presidential party for delivering disaster relief spending. Our analysis complements these previous studies by exploring legislative elections to the U.S. House of Representatives. We also go beyond the direct impact of natural disasters in our context, by focusing on differential political consequences for incumbents with varying environmental stances. Further, our results shed light on those in Herrnstadt and Muehlegger (2014), who show that congresspersons are more likely to vote in favor of environmental legislation following natural disasters in their state. While there are multiple possible channels for their results, a higher probability of being challenged could certainly put pressure on incumbents to change voting behavior.

Finally, our findings also have important policy implications. Even though environmental issues are usually not front-and-center in U.S. elections, we demonstrate that the electorate is responsive to the salience of these issues. However, we caution that these responses may not be

rational, since people may process shocks with psychological biases.³ Further, these responses may also reflect a suboptimal allocation of attention. Nevertheless, our findings suggest that scientific outreach or other approaches that raise issue salience have the potential to induce substantial changes in political behavior.

The remainder of this article is organized as follows. Section 2.2 describes the data sources we draw upon for our analysis, while section 2.3 describes our empirical strategy. In section 2.4 we report and discuss the results. We conclude in section 2.5.

2.2 Data

2.2.1 Database on Ideology, Money in Politics, and Elections

The political data we use come from the Database on Ideology, Money in Politics, and Elections (DIME) (Bonica, 2016). This database includes over 100 million campaign contributions made by individuals and organizations to candidates in local, state, and federal elections from 1979 to 2016. The main source of information is administrative records from the Federal Election Commission (FEC). In addition to campaign finance data, the database includes characteristics of the candidates receiving contributions, as well as information on election outcomes.⁴

For our study of the impact of short-run weather shocks on political campaign contributions, we use a subsample of the individual contributions data from DIME. The reason is that while individual contributions have dates assigned to them, these dates do not always match the contribution date. Instead, they may indicate the date the campaign or candidate filed these contributions. Since we are interested in people's response to short-run, time-varying weather shocks, it is essential for us to make use of accurate date information. In order to circumvent this

³For example, Gallagher (2014) examines flood insurance take-up following flood events and finds a pattern indicative of availability bias or other forms of Bayesian learning with incomplete information.

⁴For a detailed description of the database and data sources, please visit <https://data.stanford.edu/dime>.

problem, we focus on contributions made through the online fundraising platform ActBlue, since the reported date matches the date of the contribution in this sample. We assess the implications of using ActBlue data in the following section.

For our study of the political consequences of natural disasters and medium-run weather shocks, we use the recipients file of the DIME database. This file contains information at the election cycle-by-candidate level and includes the total amount of funds raised by candidates from different sources, the seat sought, and the result of the election.

2.2.2 ActBlue

In our short-term study, we focus on campaign contributions made through ActBlue, which is an online fundraising platform for Democratic candidates. The site was founded in 2004 and its popularity rose quickly thereafter. In our sample, ActBlue accounts for 4.3% of contributions and 0.8% of the total amount contributed to Democratic candidates.⁵

The main advantage of using ActBlue data is that the dates on ActBlue records are accurate, as they are electronically recorded at the time the contribution is made. Offline contributions, on the other hand, are at risk of being inaccurate in a non-random fashion, with the associated date often corresponding to the campaign's filing date. Naturally, relying on accurate date information is crucial for estimating responses to short-run weather variations. Another advantage of using ActBlue is that the contributions made on the site are typically small in quantity (see figure 2.4). For our purposes, these donations are very relevant, since they correspond to more spontaneous, lower stakes contribution decisions that may be affected by short-run weather variations.⁶

However, there are two concerns with using only ActBlue data for our short term study. The first concern is that the lack of an established Republican equivalent of ActBlue leaves us with

⁵Conversely, the total amount of contributions to Democrats is about 24 times the number of ActBlue contributions, and the total amount contributed to Democrats is 122 times the amount contributed through ActBlue. We keep these numbers in mind when assessing the magnitude of our coefficients later on.

⁶For an example of how contributions are made to Democratic candidates through ActBlue, see figure 2.3.

only donations to Democrats.⁷ This feature of our data does not allow us to see how donations to Republicans would respond, which is a shortcoming of our strategy. However, in the following sections we propose alternative methodologies to address this concern. The second concern is that it is unclear whether using ActBlue data will yield results that are representative of how contributions to Democrats as a whole respond to weather shocks. For example, it is possible that Internet contributors may be different in fundamental ways from the rest of contributors who still make up the majority. This issue has been investigated by Karpf (2013) and Wilcox (2008). Karpf (2013) suggests that the Internet brings about an increase in small donors by lowering transaction costs. They also suggest that this change has facilitated the flow of campaign funds towards more polarizing candidates. Meanwhile, Wilcox (2008) finds that Internet donors are much younger than other donors, but that those giving small amounts to Democrats online are actually similarly likely to consider themselves “ideologically extreme” as larger donors are. However, these findings are taken from surveys conducted in the year 2000, while our main time period is 2006-2012, during which Internet use was more prevalent among the general population. For our purposes, even though Internet contributors may not be a mirror image of the general contributing population, focusing on these contributions allows us to hone in on lower cost, spontaneous decisions that may be affected by weather variations.

In order to further alleviate these representativeness concerns, we would ideally correlate changes over time in ActBlue donations to changes in non-ActBlue donations, given that we exploit time-varying weather shocks in our analysis. However, there are two difficulties associated with doing this. First, as stated above, the date information for the non-ActBlue data is unreliable. Second, ActBlue was founded in 2004 and has become more popular since then, meaning that the trend of donations made through ActBlue will likely differ from the trend of overall Democratic donations. However, even though exploiting the time dimension may be difficult, we can explore whether ActBlue data do a good job of explaining the cross-section of total

⁷Rightroots, Big Red Tent, and Slatecard are examples, but their popularity has been far lower than ActBlue’s.

donations to Democrats. In order to do this, we regress total donation amounts and counts at the state-by-election cycle level on ActBlue donations and counts. If the cross-section of ActBlue donations is representative of the total Democratic cross-section, it should have high explanatory power. Additionally, to account for the fact that ActBlue becomes more popular over time and may represent a larger portion of total donations, we let our coefficients vary by election cycle in alternative regressions.

The results of these regressions are in table 2.8. The first two columns refer to the total amount contributed and the next two refer to the number of contributions. As can be seen in column (1), simply including the amount donated through ActBlue is a strong predictor of total donations, leading to an R^2 of 0.74. When we allow the effect to vary by election cycle, as in column (2), the explanatory is even higher, with an R^2 of 0.86. When we consider counts of donations instead of amounts donated, the fit is slightly better, with an R^2 of 0.83 and 0.88 in columns (3) and (4), respectively. Finally, an interesting feature of table 2.8 is the time-varying estimates in columns (2) and (4). The estimates for earlier years tend to be larger than in later years, revealing that over time the portion of ActBlue donations in total Democratic donations is rising.⁸

2.2.3 League of Conservation Voters Scorecard

In order to capture the stance of incumbent politicians on environmental issues, we use the League of Conservation Voters (LCV) scorecard (also known as the *National Environmental Scorecard*). The LCV scorecard assigns percentage scores to U.S. congresspersons based on their voting records regarding environmental legislation introduced during a particular year.⁹ According to the terminology used by the LCV, if a politician aligns with the LCV opinion on a vote, it is marked as a *pro-environment* action; conversely, if the politician does not align with

⁸It is worth pointing out that this trend stabilizes during the 2012 election cycle.

⁹The legislation included in the scorecard arises from a consensus among leading environmental and conservation organizations in the U.S.

the LCV on a vote, it is marked as an *anti-environment* action (League of Conservation Voters, 2007). For conciseness, in this paper we will follow this terminology and refer to politicians who frequently align with the LCV as pro-environment and to those who don't as anti-environment.¹⁰

More specifically, LCV scores range from zero to one with pro- and anti-environment voting records on either side of the spectrum. In this paper, we subtract the original scores from one so that a score of zero indicates that the politician has disagreed with the LCV on 0% of the votes selected (pro-environment); conversely, a score of one indicates that the politician has disagreed with the LCV on 100% of the votes selected (anti-environment).¹¹

If politicians tend to vote along party lines when it comes to environmental issues, then we would expect to see a divide in the LCV scores of Democrats versus Republicans. This is certainly the case, as can be seen in figure 2.1, where we plot the LCV score of U.S. congresspersons. As shown in the figure, most Democrats fall into the 0-0.25 range, meaning that they disagree with the LCV on less than 25% of the relevant votes. On the other hand, most Republicans fall in the 0.75-1 range, meaning that they disagree with the LCV more than 75% of the time. However, judging from the remaining mass in the 0.25-0.75 region, there is still substantial within-party variation in environmental voting records.

Additionally, the LCV score is an important indicator of whether the politician is a climate change denier. To show this, we obtain information on which congresspersons in the 112th caucus are climate change deniers from the site ThinkProgress.org.¹² Linking this information with LCV score data, we show that the probability of being a climate change denier is 51% for politicians with LCV scores above 0.5. Conversely, the probability of being a climate change denier for politicians with LCV scores below 0.5 is zero.

¹⁰Disclaimer: these terms are used to facilitate communication with the reader and do not necessarily represent the views of the authors on these issues or the politicians involved.

¹¹For more information about the LCV scorecard, please visit the LCV website at <http://scorecard.lcv.org/>.

¹²See the article "The Climate Zombie Caucus Of The 112th Congress" at <https://thinkprogress.org/the-climate-zombie-caucus-of-the-112th-congress-2ee9c4f9e46/>.

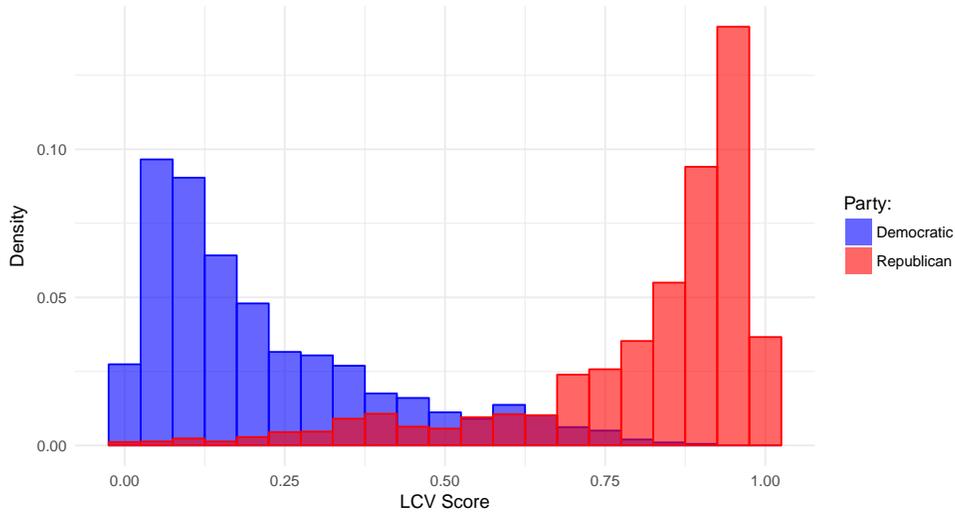


Figure 2.1: LCV score distribution by party affiliation

2.2.4 Weather Shocks

We obtain historical weather data from the Global Historical Climatology Network Daily (GHCN-D) database. This database contains daily observations of maximum temperature and precipitation from more than 8,000 weather stations throughout the United States during 1960-2014. Using this information, we construct measures of county-level weather.¹³

We construct two measures of daily temperature shocks, which we later aggregate to the appropriate time frames for our analyses. The first measure is the daily deviation in maximum temperature from the historical climate normal in each county and month:

$$TmaxDev_{cmd} = Tmax_{cmd} - \overline{Tmax}_{cm}$$

where c is county, m is month, and d is day-in-sample. $Tmax_{cmd}$ is the contemporaneous daily maximum temperature in county c . \overline{Tmax}_{cm} is the long-run average of maximum temperature for this county in the same month, calculated over the 30 preceding years. The second measure is a pair of indicators for whether the maximum daily temperature is abnormally high or low,

¹³If there is more than one weather station present in a given county, we take the average over all weather stations.

compared to historical temperature distributions:

$$TmaxLow_{5,cmd} = 1(Tmax_{cmd} \leq Tmax_{5,cm})$$

$$TmaxHigh_{5,cmd} = 1(Tmax_{cmd} \geq Tmax_{95,cm})$$

where $Tmax_{5,cm}$ is the 5th percentile of the distribution of maximum temperatures in the same county and month over the 30 preceding years, and $Tmax_{95,cm}$ is the corresponding 95th percentile. As a result, $TmaxLow_{5,cmd}$ is an indicator for whether the contemporaneous temperature is lower than the 5th percentile of the historical distribution, whereas $TmaxHigh_{95,cmd}$ indicates whether it is higher than the 95th percentile of that distribution.

For our short-run analysis at the county-week level, we aggregate these daily measures by week, in order to obtain our independent variables of interest. We construct $TmaxDev_{cw}$, which is the average of $TmaxDev_{cmd}$ over the week, and is our primary temperature shock measure. We also construct $TmaxHigh_{cw}$ and $TmaxLow_{cw}$, which are the sums over the week of $TmaxHigh_{5,cmd}$ and $TmaxLow_{5,cmd}$, respectively. Importantly, these measures capture different aspects of weather shocks. Further, we also construct similar measures of precipitation deviations which we use as controls in our regressions.

For our medium-run analysis, we first calculate the number of hot days, defined as those above the 90th percentile of the historical distribution, experienced by the average person in each congressional district and election cycle.¹⁴ We then rank district-cycle observations by this variable and assign *hot* status to those cycles in the top quartile. Similarly, we assign *cold* status to a district-cycle if it is in the top quartile ranked by number of cold days, defined as those below the 10th percentile of the historical distribution.

¹⁴Our procedure makes use of the MABLE/Geocorr crosswalks developed by Missouri Census Data Center (2017), which partitions the population in a congressional district into its overlapping counties using Census data.

2.2.5 Natural Disasters

We obtain official disaster declaration data from the Federal Emergency Management Agency (FEMA) between 1990 and 2012. There are a total of 2,206 climate-related disasters, with the largest categories being severe storms, hurricanes, floods, fires and snow events (see table 2.7 for a detailed breakdown of disaster types). Importantly, these official records contain the period of the incident and the specific counties affected. Most declarations are not statewide.

Because we analyze the impact of natural disasters at the congressional district level, we need to assign disaster status to these for each election cycle. In order to do this, we first calculate the fraction of the population in a district living in counties hit by disasters,¹⁵ and then assign disaster status to a district if this fraction exceeds 50%. It should be noted that this might not be the exact threshold at which natural disasters become salient politically, potentially leading to measurement error. However, in our data the majority of district-cycle observations have a fraction of the population affected of either zero or one, so adjustments to the threshold would not have substantial impact on our results.

2.3 Empirical Framework

Existing literature suggests that people might update their beliefs about climate change as a result of contemporaneous weather shocks, or even after a prolonged period of unusual temperature. Similarly, natural disasters may also lead to belief updating, given the salience of these events. Given that there are multiple possible relevant time frames and types of weather event, we examine the impacts of weather shocks both in the short-run and in the medium-run, as well as the medium-run impacts of natural disasters. In this section, we outline our empirical strategy for doing so.

¹⁵This procedure also uses MABLE/Geocorr crosswalks (Missouri Census Data Center, 2017).

2.3.1 Short-Run Weather Impacts

We first analyze the impact of weekly weather shocks on contributions to Democrats through ActBlue. Since Democratic candidates tend to be more pro-environment than non-Democratic candidates, we expect these donations to be responsive to weather shocks through the environmental preference mechanism we have proposed.

The estimating equations we use to study the impact of weather shocks on campaign contributions take the following general form:

$$Y_{cw} = \gamma' Weather_{cw} + \delta_w + \delta_c + \delta_{se} + \epsilon_{cw} \quad (2.1)$$

where c is county, w is week-in-sample, s is state, and e is election cycle. Y_{cw} is the outcome of interest, which can be either (1) the contribution rate (the number of contributions per million people), or (2) the average amount per contribution. These outcomes are calculated based on contributions made through the ActBlue platform. Furthermore, $Weather_{cw}$ is a vector of weather variables, which includes measures of temperature conditions as the key regressors and those of precipitation conditions as controls. Our coefficients of interest are the ones in γ' corresponding to the temperature variables.

We use three different specifications of $Weather_{cw}$ in our main analysis. The first specification includes only weekly contemporaneous deviations of temperature and precipitation from long-run climate normals, as they are expected to be the most relevant. It takes the following form:

$$Weather_{cw} = [TmaxDev_{cw}, PrcpDev_{cw}]^T.$$

One thing to note about this specification is that it does not allow for weather from past weeks to influence contribution decisions. However, it may be possible that the decision to contribute is

delayed with respect to exposure to weather shocks. To allow for this, our second specification is:

$$Weather_{cw} = [TmaxDev_{cw}, \dots, TmaxDev_{c,w-4}, PrcpDev_{cw}, \dots, PrcpDev_{c,w-4}]^T$$

where we have added four lags of both temperature and precipitation deviations. This way, we capture the influence of weather shocks occurring during the past month. Furthermore, another weakness of our first specification is that it may not adequately capture current weather for contributions occurring earlier on in the week, since our measures are weekly averages. In order to address this, we propose the following third specification:

$$Weather_{cw} = [\overline{TmaxDev}_{c,w}^{2W}, \overline{PrcpDev}_{c,w}^{2W}]^T$$

where $\overline{TmaxDev}_{c,w}^{2W} = \frac{1}{2}(TmaxDev_{cw} + TmaxDev_{c,w-1})$ in the case of temperature, and similarly $\overline{PrcpDev}_{c,w}^{2W} = \frac{1}{2}(PrcpDev_{cw} + PrcpDev_{c,w-1})$ in the case of precipitation. Therefore, we average deviations in temperature and precipitation over the current and previous weeks.

In all of our specifications we also control for week-in-sample (δ_w), county (δ_c), and state-by-election cycle (δ_{se}) fixed effects. The county fixed effects absorb time-invariant factors in each county such as general political preferences and contribution behavior. The week-in-sample fixed effects control for confounding national events and the exponential growth of the platform itself. Finally, the state-by-cycle fixed effects account for slower-moving changes across states, such as whether the current president is politically aligned with the state, or new policies adopted by the State. We cluster standard errors at the county level.

Aside from the main specifications outlined above, we also consider four extensions. First, we estimate the effects of positive and negative temperature deviations separately, to shed light on whether people respond to these differently. Second, we estimate effects separately for each quarter of the election cycle, to study how these vary with the progression of campaigns. Third, we consider alternative measures of our independent variables by focusing on counts of extreme

temperature events instead of average temperature. Fourth, we extend our main regressions to allow for heterogeneous effects depending on the environmental stance of the incumbents in the contributor's place of residence, to rule out unobservable confounding factors that may drive all contributions across time and location, and not only those that are environmentally motivated. More specifically, we examine whether counties where the majority of the population lives in districts represented by anti-environment incumbents exhibit stronger responses to weather shocks.¹⁶ For more details on these additional specifications, see the appendix.

2.3.2 Medium-Run Natural Disaster Impacts

Aside from studying weather shocks, we are interested in how fundraising and elections are affected by natural disasters in the medium run. Specifically, we study how this relationship varies depending on the environmental stance of the incumbent politician. In order to explore this issue, we focus on races for the U.S. House of Representatives during election cycles 1990-2012. We study campaign finance outcomes, such as the total funds raised and the funds raised by the challenger or incumbent separately. Further, we consider electoral outcomes such as the probability of the race being competitive and the probability of the incumbent being re-elected.

One concern that our approach will have to address is that natural disasters may have significant effects on campaign contributions and other political outcomes through channels unrelated to environmental preferences and beliefs. For example, Stevens (2001) documents that following the September 11 terror attacks, individuals substituted away from campaign contributions and towards charitable giving. We expect this to be relevant in the case of natural disasters, since they often entail tragic consequences and loss of property.

In order to address the above concern, our research design consists in comparing the outcomes of congressional districts experiencing natural disasters whose incumbent politicians have

¹⁶This is what we would expect as long as the Democratic candidates receiving contributions on ActBlue are more pro-environment on average.

an anti-environment voting record to the outcomes of other districts experiencing natural disasters but whose incumbents exhibit pro-environment voting records. By studying differential impacts by the environmental stance of incumbents, we hope to isolate the environmental preference mechanism. From an econometric standpoint, we run regressions of the following form:

$$Y_{de} = \beta_1 Disaster_{de} + \beta_2 LCV_{de} + \beta_3 Disaster_{de} \times LCV_{de} + \delta_d + \delta_e + \varepsilon_{de} \quad (2.2)$$

where Y_{de} is an outcome for a race in congressional district d during election cycle e ; $Disaster_{de}$ is an indicator variable for whether over half of the population in a congressional district lives in counties affected by a natural disaster during that cycle; LCV_{de} is the LCV score of the incumbent;¹⁷ and δ_d and δ_e are fixed effects for congressional district and election cycle, respectively. We cluster standard errors at the state level.

Our coefficient of interest is β_3 . We interpret this coefficient as the difference in the outcome of a congressional district affected by a natural disaster whose incumbent congressperson has the most anti-environment voting record ($LCV = 1$), and the outcome of a similar, disaster-struck congressional district whose incumbent congressperson has the most pro-environment voting record possible ($LCV = 0$). Given that a one unit difference in the LCV score is a very large difference, we suggest scaling our estimates by the standard deviation of the LCV score in order to interpret them properly. Since the standard deviation of the LCV score is 0.2, we interpret our coefficients by dividing them by five.¹⁸

While the LCV score captures precisely the political dimension we care about, it is important to consider how it relates to party affiliation. As we can see in figure 2.1, even though there is certainly within-party variation in the LCV score, the two are closely related. This raises a limitation of model (2.2): following a natural disaster, if people react differently to incumbents

¹⁷In order to incorporate all available information at the time of the race, we average the LCV score of politicians for that election cycle and all past election cycles, using this measure throughout in our regressions.

¹⁸The standard deviation we use is that of the LCV score after controlling for the politician's party, which is 0.2. Without controlling for the politician's party the standard deviation is 0.32, which is slightly larger.

from different parties for non-environmental reasons, then the above coefficient of interest would be picking up on these factors as well. To address this issue, we propose an extension of the above model. The proposed model is:

$$\begin{aligned}
 Y_{de} = & \beta_1 Disaster_{de} + \beta_2 LCV_{de} + \beta_3 Disaster_{de} \times LCV_{de} \\
 & + \beta_4 R_{de} + \beta_5 Disaster_{de} \times R_{de} + \delta_d + \delta_e + \varepsilon_{de}
 \end{aligned}
 \tag{2.3}$$

where all variables are defined as in model (2.2) and we have now added R_{de} , an indicator variable for whether the incumbent is a Republican, as well as an interaction of this variable with $Disaster_{de}$. Our coefficient of interest is still β_3 , which is now identified using variation in the LCV score within the incumbent's political party.

It is important to note that model (2.3) comes with both advantages and disadvantages with respect to model (2.2). The main advantage, as previously discussed, is that it addresses the concern that people may respond to disasters differently depending on the incumbent party for reasons unrelated to the environment. The disadvantage is that the model does not make use of of valuable cross-party variation in environmental stances, which is perhaps the most visible and available to people when making decisions. Therefore, we consider these models to be complementary and keep these features in mind when interpreting results.

2.3.3 Medium-Run Weather Impacts

In the same way that campaign contributions and elections may be affected by natural disasters, they may also respond to shocks to medium-run temperature. The main difference with respect to natural disasters is that in this case there can be two distinct shocks: a hot weather shock and a cold weather shock. Since people may respond to these shocks differently, they enter our regressions separately. Using the same notation as in the natural disasters section, we study

the impact of medium-run weather as follows:

$$\begin{aligned}
 Y_{de} = & \beta_1 Hot_{de} + \beta_2 Cold_{de} + \beta_3 LCV_{de} + \beta_4 Hot_{de} \times LCV_{de} \\
 & + \beta_5 Cold_{de} \times LCV_{de} + \delta_d + \delta_e + \varepsilon_{de}
 \end{aligned}
 \tag{2.4}$$

where all variables are as previously defined, and Hot_{de} and $Cold_{de}$ are indicators for whether the election cycle was particularly hot or cold for a given district in an election cycle, constructed as described in section 2. Aside from including these variables, we also add their interaction with the LCV score, following our earlier methodology.¹⁹

Our coefficients of interest are β_4 and β_5 . As above, we interpret these coefficients as the difference in the outcome of a congressional district undergoing an unusually hot (cold) cycle, whose incumbent congressperson has the most anti-environment voting record ($LCV = 1$), and the outcome of a similar district whose incumbent congressperson has the most pro-environment voting record possible ($LCV = 0$). Again, to ensure a reasonable interpretation of estimates we scale them by dividing by five. Finally, it is straightforward to extend this methodology to make use of within-party variation in the LCV score following a disaster, by adding R_{de} and its interaction with LCV as in model (2.3) above.

2.4 Results

In this section, we present our results in three parts: (1) short-run weather impacts on ActBlue contributions, (2) medium-run natural disaster impacts on campaign contributions and election outcomes, and (3) medium-run weather impacts on the same set of outcomes as in (2).

¹⁹Standard errors are clustered at the state level.

Table 2.1: Summary statistics

Variable	N	Mean	Std. Dev.	Min.	Max.
<i>ActBlue, 2006-2012 (county-week)</i>					
Amount (\$)	935,201	151.72	2061.05	0	583663.8
Count	935,201	2.42	23.10	0	5315
Average amount	935,201	13.19	125.67	0	32500
Count (per 1M pop)	935,201	15.40	135.84	0	38848.92
Population	935,201	110665.3	337187.7	403	9974868
Mean LCV	830,316	0.6720	0.3219	0	0.9800
Republican incumbent	830,316	0.6220	0.4849	0	1
Anti-env incumbent	830,316	0.6945	0.4606	0	1
<i>Short-run Weather, 2006-2012 (county-week)</i>					
Tmax dev. (F)	935,201	0.4453	6.621	-37.60	37.51
Tmax positive dev. (F)	935,201	2.789	4.036	0	37.51
Tmax negative dev. (F)	935,201	2.344	3.805	0	37.60
Tmax low (< 5th pct)	935,201	0.3185	0.7855	0	7
Tmax high (> 95th pct)	935,201	0.4719	1.0652	0	7
Prcp dev. (1/10mm)	935,201	0.0609	13.5375	-49.91	468.54
<i>Natural Disasters, 1990-2012 (congressional district-cycle)</i>					
Election cycle	4,478	2001.03	7.01	1990	2012
Disaster indicator	4,478	0.58	0.49	0	1
Num. donors (C)	4,095	232.18	1074.94	0	27122
Receipts (\$1,000) (C)	4,095	326.72	703.71	0	9825.57
Num. donors (I)	4,095	536.83	1151.29	0	43718
Receipts (\$1,000) (I)	4,095	1076.15	1108.93	6.77	25894.72
Receipts PACs (\$1,000) (C)	4,095	42.57	101.39	0	1122.45
Receipts PACs (\$1,000) (I)	4,095	466.07	387.56	0	4082.6
Receipts Ind. (\$1,000) (C)	4,095	194.61	436.17	0	6248.91
Receipts Ind. (\$1,000) (I)	4,095	545.11	745.35	0	21836.54
Competitive election	4,478	0.74	0.44	0	1
Unopposed election	4,478	0.18	0.38	0	1
Open race election	4,478	0.09	0.28	0	1
LCV score	4,478	0.5	0.34	0	1
Republican incumbent	4,478	0.47	0.5	0	1
Incumbent wins	4,095	0.95	0.22	0	1
Hot indicator	4,478	0.25	0.43	0	1
Cold indicator	4,478	0.25	0.43	0	1

2.4.1 Short-Run Weather Impacts

In the short-run analysis, we investigate how ActBlue contributions are affected by temperature shocks in current and previous weeks. We examine two main outcomes. The first outcome is the contribution rate, defined as the number of contributions per million people in a county. This variable captures extensive-margin responses, i.e. whether temperature shocks motivate more or fewer contributions. The second outcome is the average amount per contribution, calculated as the total amount contributed divided by the number of contributions for each county-week. Absent any extensive-margin responses, this outcome measures intensive-margin responses, i.e. whether temperature shocks motivate larger or smaller donations from regular contributors. However, if extensive-margin responses are present, this outcome captures both the intensive-margin responses and potential changes in the composition of contributors.

In our sample period of 2006-2012, each county receives on average 2.4 donations per week or around 15.4 per million people, as can be seen in table 2.1. The average amount of these donations is \$13.2, demonstrating that ActBlue contributions are usually small. Table 2.1 also provides summary statistics on weather variables. The mean temperature deviation is 0.45°F, since positive deviations are larger than negative ones on average. This pattern is also captured by the extreme temperature bins, as the number of extremely hot days exceeds the number of extremely cold ones.²⁰ This general warming trend is a common observation in the literature.

The estimates from equation (2.1) are reported in table 2.2. In columns (1)-(3), we examine responses in the contribution rate. In column (1), the main weather variable is temperature in the current week. As detailed in Section 2.2.4, this variable is constructed as the deviation from the long-run temperature normal. Precipitation deviation is also included in the model as a control. The estimate on temperature is positive and significant at the 1% level. This suggests that more positive or less negative deviations are associated with a higher number of contributions.

²⁰Extremely hot days are those above the 95th percentile of the historical distribution, while extremely cold days are those below the 5th percentile of that distribution.

Table 2.2: Actblue donation responses to short-run temperature shocks

Dep. Var.	(1) Count /1M pop	(2) Count /1M pop	(3) Count /1M pop	(4) Avg. amount	(5) Avg. amount	(6) Avg. amount
$TmaxDev_{c,w}$	0.186*** (0.057)	0.134*** (0.045)		0.016 (0.047)	0.013 (0.046)	
$TmaxDev_{c,w-1}$		0.103*** (0.035)			-0.043 (0.034)	
$TmaxDev_{c,w-2}$		0.054*** (0.017)			0.055 (0.041)	
$TmaxDev_{c,w-3}$		0.072*** (0.025)			0.035 (0.035)	
$TmaxDev_{c,w-4}$		0.055** (0.023)			-0.035 (0.032)	
$\overline{TmaxDev}_{c,w}^{2W}$			0.287*** (0.083)			-0.007 (0.054)
R^2	0.209	0.204	0.209	0.054	0.054	0.054
N	944,172	935,201	941,672	944,172	935,201	941,672
Mean D.V.	15.45	15.40	15.42	13.13	13.19	13.15
County F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Week F.E.	Yes	Yes	Yes	Yes	Yes	Yes
State-cycle F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Notes: point estimates from equation (2.1) are shown. The dependent variable in columns (1)-(3) is the number of ActBlue contributions per 1 million population. The dependent variable in columns (4)-(6) is the average amount per ActBlue contribution. The sample consists of ActBlue contributions by week and county. Standard errors (in parenthesis) are clustered by county. Statistical significance: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

In particular, a 1°F increase in weekly temperature is associated with about 0.19 additional contributions per million people, or an approximately 1.2% increase in the contribution rate relative to the mean.²¹

²¹ $\hat{\gamma}_0$ is the coefficient on $TmaxDev_{c,w}$. The contemporary effect as a percentage of the mean of the dependent variable is $\hat{\gamma}_0 / MeanD.V. = 0.186 / 15.40 \approx 1.2\%$.

Column (2) adds four lags of temperature deviations. This specification allows us to examine whether temperature deviations from previous weeks might affect contemporaneous contributions. The estimated dynamics have two remarkable features. First, the estimates are across-the-board positive and significant. This pattern is inconsistent with a harvesting mechanism, where a temperature shock simply shifts the timing of contributions but not the overall amount. Instead, these effects could represent a net increase in contributions. Second, the impact of a temperature shock appears to decay over time, as the estimates are lower for temperature deviations that took place longer ago. The contemporaneous effect is smaller than in column (1), but the cumulative effect is larger: a 1°F increase in weekly temperature is associated with a cumulative effect of 0.42 additional contributions per million population, or a 2.7% increase in the contribution rate relative to the mean.²²

The comparison between columns (1) and (2) suggests that omitting temperature lags might have led us to overestimate the contemporaneous effect and underestimate the overall effects of a temperature shock. Furthermore, column (2) shows that the current and previous week temperature deviations have the largest impact. In column (3), we take the average of deviations in these two weeks and use it as our main independent variable. As expected, the estimate is again positive and significant at the 1% level. Below, we will use this as our specification of reference and estimate variants of it in later extensions of our analysis.

Next, we re-estimate our models using the average contribution amount as the outcome variable. The results are in columns (4)-(6) of table 2.2. In this case, all estimates are small and statistically insignificant, with no recognizable pattern. This could mean that temperature shocks do not induce intensive-margin responses or changes in contributor composition that are strong enough to be statistically detectable. It is also possible that these changes go in opposite directions and cancel each other out.

A feature of the specification used above is the implicit assumption that reducing a

²² $(\hat{\gamma}_0 + \dots + \hat{\gamma}_4)/MeanD.V. = (0.134 + 0.103 + 0.054 + 0.072 + 0.055)/15.40 \approx 2.7\%$.

negative shock by one degree and increasing a positive one by the same amount has the same quantitative effect. However, positive and negative shocks might not be viewed symmetrically by individuals when it comes to belief formation. In table 2.9, we report separate estimates of the effects of positive and negative shocks. This is implemented by replacing each deviation regressor above with a pair of variables that separately capture the absolute values of its positive and negative components.²³ The results suggest that the observed effects are mainly driven by variation in positive deviations.

So far, our temperature measures are all in the form of deviations from the long-run norm. An alternative is to use measures of extreme temperature events because they might be more salient. As detailed in Section 2.2.4, we also construct variables that count the number of extremely hot or cold days in a week. Table 2.10 reports estimates based on these measures. Again, we find effects on the contribution rate but not the average amount. Specifically, one more extreme-hot day in a week is associated with a contemporaneous increase of 0.35 contributions per million people, or 2.3% of the mean contribution rate. Further, the cumulative monthly increase in this case is 7%.²⁴ For one more extreme-cold day, the contemporaneous and cumulative effects are decreases of 6.6% and 15.8%, respectively. These results share a number of qualitative similarities with the ones in table 2.9. They both show that heat shocks have extensive-margin impacts, but not intensive-margin ones. They both find these effects lasting up to a month and decreasing over time. On the other hand, the two sets of results differ on the effects of cold shocks. The average deviation specification suggests no effect, while the cumulative extreme weather approach suggests significant and negative effects. This suggests that a cold spell might not be as salient as one additional extremely cold day. It might also be related to the political discourse surrounding extremely cold weather events. In the past, some politicians have used extremely cold weather to argue against climate change, which could plausibly change public opinion in the

²³For more details on the specification, see the appendix.

²⁴The contemporaneous effect of an additional hot day is $\hat{\gamma}_0 / \text{MeanD.V.} = 0.353 / 15.40 \approx 2.3\%$ of the mean of the dependent variable.

opposite direction (Pierre-Louis, 2017).

Overall, our results so far are consistent with two potential mechanisms. First, temperature shocks prompt people to update their beliefs about climate change and become more politically aligned with Democratic candidates, who are typically pro-environment. Second, temperature shocks make environmental issues more salient in the election, reminding pro-environment donors to express their preferences through contributions. As these two mechanisms are similar, we do not attempt to disentangle them and refer to them collectively as the environmental preference mechanism. However, there are still a few alternative explanations for our results that are unrelated to environmental reasons. For example, weather is known to change voting behavior through psychological channels (Meier et al., 2016; Jachimowicz et al., 2017). It might also affect time use or the expediency of online versus other contribution channels.²⁵

To address these concerns and shed light on the mechanisms behind our results, we examine heterogeneous effects based on incumbent characteristics.²⁶ To enhance statistical power, we estimate effects by using a single temperature variable, the mean deviation in the current and past week.²⁷ Furthermore, we also include a measure that characterizes the population-weighted environmental stance of the incumbent congresspersons in a county, and its interaction with the temperature variable. We are interested in the coefficient associated with the interaction term, which shows how the effects of weather shocks vary according to incumbent characteristics.²⁸ Three characteristics are examined in separate regressions: (1) population-weighted mean LCV score (mean = 0.672); (2) whether over half of the population has a Republican incumbent (mean = 62.2%); and (3) whether over half of the population has an incumbent with an LCV score above 0.5 (mean = 69.5%). In this analysis, we restrict our sample to competitive races. The

²⁵Section 2.4.4 presents a detailed discussion of the alternative mechanisms.

²⁶We do not observe which candidates receive the contribution in the ActBlue data, only the place of residence of the donor. This limits our investigation to incumbent characteristics. While many contributions are directed to candidates outside of the district of residence of the contributor, we think this is a meaningful margin of giving behavior to study. For example, environmentally motivated donors may look to other congressional district races if the district they reside in is a very safe seat held by an anti-environment politician.

²⁷Baseline estimates are reported in table 2.2, columns (3) and (6).

²⁸For more details on the specification, see the appendix.

Table 2.3: Heterogeneous effects by incumbent characteristics

Incumbent Char.	LCV		Republican		1(LCV>0.5)	
	(1) Count /1M pop	(2) Avg. amount	(3) Count /1M pop	(4) Avg. amount	(5) Count /1M pop	(6) Avg. amount
$\overline{TmaxDev}_{c,w}^{2W}$	0.193** (0.098)	-0.115 (0.118)	0.228** (0.103)	-0.102 (0.075)	0.241** (0.110)	-0.110 (0.088)
<i>IncChar</i>	10.299** (4.658)	2.016 (1.318)	5.336** (2.674)	0.764 (0.854)	7.043** (3.129)	1.538 (0.954)
$\overline{TmaxDev}_{c,w}^{2W} \times IncChar$	0.152** (0.068)	0.128 (0.110)	0.107*** (0.037)	0.117** (0.056)	0.079** (0.037)	0.118* (0.067)
<i>N</i>	830,316	830,316	830,316	830,316	830,316	830,316
<i>R</i> ²	0.2070	0.0550	0.2070	0.0551	0.2070	0.0551
Mean D.V. (All/R/Anti)	12.29	11.09	11.12	8.16	10.33	7.48
Mean D.V. (All/D/Pro)	12.29	11.09	14.22	15.93	16.75	19.30
County F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Week F.E.	Yes	Yes	Yes	Yes	Yes	Yes
State-cycle F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Notes: point estimates from equation (2.7) are shown. The dependent variables in columns (1), (3), and (5) is the number of ActBlue contributions per 1 million population. The dependent variable in columns (2), (4), and (6) is the average amount per ActBlue contribution. The temperature shock measure is the average temperature deviation in the current and past week. The *IncumbChar* variable is the mean LCV in columns (1)-(2), whether the county is dominated by Republican incumbents in columns (3)-(4), and whether the county is dominated by incumbents with anti-environment voting records in columns (5)-(6). Standard errors (in parenthesis) are clustered by county. Statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

summary statistics suggest that the counties in our sample tend to have incumbents who are unfavorable to environmental protection and more likely to be Republicans. This could be due to people supporting a Democratic challenger being more active in online contributions than those supporting a Democratic incumbent.

The results are reported in table 2.3. In terms of the baseline effects, as before, we find a positive and significant effect of temperature deviations on the contribution rate, but not on the average contribution amount. Importantly, by examining the interaction term we can now see that the effect on contributions is larger when the incumbents have a more unfavorable view

of environmental protection. This is true for all three measures of incumbents' environmental positions. The scale is also economically important. According to column (1), when the mean LCV score increases by one standard deviation, the scale of the positive effect goes up by 15.8% of the baseline effect, corresponding to that of a county represented fully by pro-environment incumbents.²⁹ Further, the effect of temperature shocks in Republican-dominated counties is 47% larger than in Democrat-dominated ones, according to column (2). The corresponding difference in column (3) is 32.7%.³⁰ These results suggest that following a temperature shock, Democratic candidates are rewarded for their pro-environment stance when the incumbents in the contributor's district are more anti-environment.³¹

Next, we explore how effects vary depending on the progression of campaigns. We use a similar specification, interacting the two-week average deviation measure with a set of eight indicators for quarters in the election cycle. This allows us to obtain a separate estimate for each quarter. The results are plotted in figure 2.2 and reported in table 2.11.

For the contribution rate, the estimates are positive and significant for quarters 2-5 and 8. As expected, the effect is especially pronounced in the last quarter, as the election date draws near and campaigning efforts ramp up. The second largest effect is in quarter 3, or June to August in the year prior to the election. This might be driven by presidential candidates declaring their candidacy around that time. When we examine average amount as an outcome, the pattern is flipped. The effect is large and negative in quarter 8, followed by quarter 3, and fluctuates around 0 for the rest. This pattern is most likely driven by a change in the composition of the contributors – heat shocks draw in more small-amount contributions in this case.³² Given this negative impact

²⁹LCV incremental effect: $SD(\overline{LCV}) \times \hat{\beta}_3 / \hat{\beta}_1 = 0.2 \times 0.152 / 0.193 \approx 15.8\%$.

³⁰Republican incremental effect: $\hat{\beta}_3 / \hat{\beta}_1 = 0.107 / 0.228 \approx 47\%$. Anti-environment incremental effect: $\hat{\beta}_3 / \hat{\beta}_1 = 0.079 / 0.241 \approx 32.7\%$.

³¹It is also possible to examine the coefficients on the incumbent characteristics, which are all positive and significant when looking at contribution rates. Since we have controlled for county fixed effects, this parameter identifies the increase in the contribution rate corresponding to an increase in anti-environment incumbents within a county. Together, these results suggest that people make compensatory contributions when politicians ideologically different from them are elected in their district.

³²When interpreting the pattern, one caveat is that the average-amount variable is constructed by dividing total

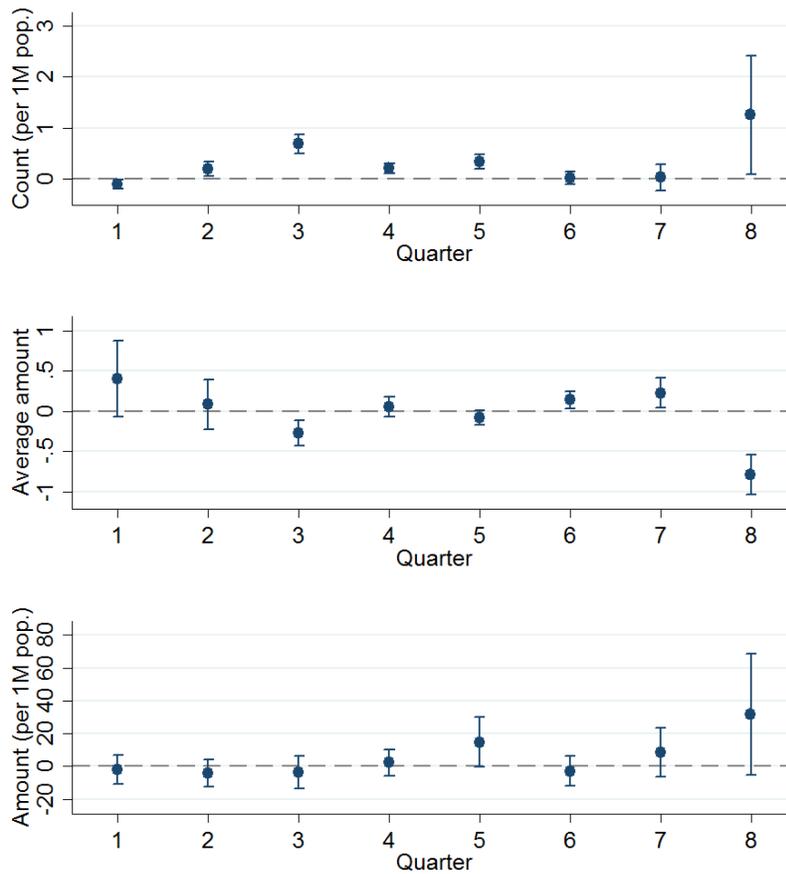


Figure 2.2: Variations of estimates across quarters in election cycle

Notes: Point estimates from equation (2.8) and 95 percent confidence intervals are shown. The outcome variables, as displayed next to the y-axis, are based on ActBlue records. Standard errors are clustered by county. All regressions control for county, week-in-sample, and state-by-cycle fixed effects.

on the average contribution amount, we also examine the contribution amount per million people, which is a combined outcome of the two margins. As the extensive and intensive margins often change in opposite directions, the overall effect is mostly small except for quarters 5 and 8, where the increase in the extensive margin appears to dominate, leading to increased contribution

contribution amount by the number of contributions. This procedure might transform any measurement error in the denominator into non-classical ones in the outcome. The estimates might thus be biased, the confidence interval invalid, especially when the mean of the denominator is close to zero. A remedy of this issue is beyond the scope of this paper. In quarters 3 and 8, however, the number of contributions is sufficiently large that we are more confident in the estimates.

amounts in those quarters.

Lastly, we perform a back-of-the-envelope calculation to infer the effects of temperature shocks on total Democratic contributions, instead of only ActBlue contributions. This calculation allows us to gauge the actual magnitudes of our previous estimates. In our sample, the total number of Democratic contributions is 24 times that of ActBlue contributions, and the total amount is 122 times larger. Using these numbers and our estimates in table 2.2, we find that the contemporaneous effect of a 1°F increase in weekly mean temperature corresponds to a total increase of 3.2 contributions or \$215.6 per million people per week.³³ The corresponding cumulative effects are 10 contributions and \$672.6. It should be noted that this calculation relies on the assumption that total Democratic donations and ActBlue donations react similarly to weather shocks. In reality, we may expect ActBlue donations to react more strongly given the small and spontaneous nature of these contributions, meaning that these calculations are likely to represent upper bounds on the actual effects.

2.4.2 Medium-Run Natural Disaster Impacts

In the previous section, we show that heat shocks are associated with more contributions to Democratic candidates through ActBlue. As they are often more pro-environment, it is consistent with the mechanism that people update their beliefs about climate change and express their preference for pro-environmental policies through contributions. However, with only contribution records for the Democrats, it is difficult to fully disentangle this environmental preference channel from other factors that drive contributions in general.

In this section, we study the impact of natural disasters on campaign finance and elections. This analysis complements the previous results, as we explicitly account for politicians' environmental attitudes and include contributions to both Democratic and Republican candidates.

³³ Δ number of contribution = $\hat{\gamma}_0 \times ratio(Dem/ActB) = 0.134 \times 24 = 3.22$. Δ total amount = Δ number of contributions \times average amount $\times ratio(Dem/ActB) = 0.134 \times 13.19 \times 122 = 215.63$.

As suggested by previous studies, natural disasters can draw public attention to climate change (Lang & Ryder, 2016; Sisco et al., 2017) and they also bring about political ramifications for the incumbents (Arceneaux & Stein, 2006; Gasper & Reeves, 2011; Healy & Malhotra, 2009). Building on this literature, we hypothesize that anti-environment leaning incumbents will be held accountable for their environmental stance when a natural disaster strikes, leading to increased support for challengers. However, we do not necessarily expect support for these incumbents to remain unchanged, since they may intensify their fundraising efforts as a response to the increased support for challengers. In other words, we are agnostic about what happens to incumbent support.

In order to study the impact of natural disasters on campaign finance and elections, we examine how the impact of a natural disaster varies depending on the environmental voting record of the incumbent politician (LCV score), while focusing on competitive races. We run regressions that follow the specifications in equations (2.2) and (2.3). When interpreting the magnitude of our coefficients, we divide them by five so that they correspond to the effect of a one-standard-deviation difference in the LCV score of a candidate.³⁴ Summary statistics for the sample used to estimate these models can be found in the bottom panel of table 2.1. This sample includes congressional races during the 1990-2012 election cycles. Out of the races in our data, 74% are competitive while in 18% of them the incumbent runs unopposed. We can also see from the table that both the number of donors and the total amount of funds raised are higher for incumbents than for challengers.

We start by studying the effects of natural disasters on the amount of funds raised during the election in table 2.4. Panel A contains the results using our main specification, described in equation (2.2). In column (1) we can see that total fundraising following a natural disaster is

³⁴By doing this, we make sure we are taking into account the variation in the LCV variable. Since the standard deviation of the LCV score (after controlling for the politician's party) is 0.2, we interpret our coefficients by dividing them by five. This allows us to interpret results as those corresponding to a one-standard-deviation difference in the LCV score of a candidate, instead of a one-unit difference, which would correspond to the entire range of the variable.

higher in those districts with more anti-environment incumbents. More precisely, a one-standard-deviation increase in the LCV score of the incumbent translates to a \$78,000 increase in total fundraising during a cycle (5.1% D.V. mean), when a natural disaster strikes. We can further break down the fundraising effect by distinguishing between the types of donor. In columns (2) and (3) we can see that funds from PACs go up by about \$20,000 (3.8% D.V. mean) and funds from individuals go up by about \$38,000 (4.6% D.V. mean).

In order to assess whether there is increased support for challengers, incumbents, or both, we turn to columns (4)-(6) in table 2.4. In column (4), we find that a one-standard-deviation increase in the LCV score of the incumbent translates to a \$39,000 increase in fundraising by challengers during a cycle, when a natural disaster strikes. This result substantiates the hypothesis of increased support for challengers in races with anti-environment incumbents. In column (5), we find that a one-standard-deviation increase in the LCV score of the incumbent also translates to a \$39,000 increase in fundraising by incumbents following a natural disaster. These results are consistent with the hypothesis that incumbents may perceive the strengthened support for the challengers and react by increasing fundraising efforts themselves. Furthermore, even though fundraising increases for both the incumbent and the challengers, we can assess the relative strength of these effects by studying the share of funds raised by the challenger. The results of this exercise are in column (6). According to these estimates, the share of funds going to challengers is higher when the incumbent is anti-environment, although the estimate is not statistically significant. However, it is worth pointing out that the magnitude of the effect is economically non-negligible: a one-standard-deviation increase in the LCV score is associated with a 0.6 p.p. increase in the share of funds going to challengers, over a baseline of 19.6 p.p., when a natural disaster strikes. Further, another way to assess whether the increase in support is greater for challengers or incumbents is by comparing the effects in columns (4) and (5) to their respective dependent variable means. In this case, the effect for challengers appears to be stronger, as the baseline level of donations to incumbents is higher and the estimated effects are similar. While

the evidence is weak, it is consistent with the hypothesis that the increase in support may be stronger for challenges than for incumbents.

Table 2.4: The effects of natural disasters on amount raised (\$1,000)

Dep. Var.	(1) Total	(2) Total (PAC)	(3) Total (Ind.)	(4) Challenger	(5) Incumbent	(6) Share (C)
<i>Panel A: Main Specification</i>						
Disaster	-316.4*** (70.39)	-78.15*** (24.16)	-164.8*** (55.17)	-184.7*** (46.02)	-131.7** (54.61)	-0.0328*** (0.0122)
LCV	-253.2 (305.8)	-87.08* (48.97)	-39.90 (259.3)	-226.7** (107.3)	-26.53 (221.0)	-0.0453* (0.0249)
Disaster × LCV	389.7*** (125.6)	101.8*** (35.78)	191.8** (88.96)	194.1** (93.66)	195.5** (73.00)	0.0299 (0.0245)
R^2	0.413	0.533	0.388	0.262	0.457	0.287
<i>Panel B: Within-Party Specification</i>						
Disaster	-300.2*** (72.39)	-79.88*** (24.44)	-152.3*** (55.54)	-168.3*** (48.70)	-131.8** (56.14)	-0.0247** (0.0112)
LCV	228.1 (406.4)	33.41 (86.51)	200.2 (332.0)	-11.27 (176.0)	239.4 (295.7)	0.0371 (0.0385)
Disaster × LCV	255.7 (227.9)	122.2* (61.79)	84.72 (154.2)	49.67 (156.7)	206.1* (115.5)	-0.0422 (0.0272)
R^2	0.414	0.534	0.388	0.263	0.458	0.290
Observations	3299	3299	3299	3299	3299	3299
Mean D.V.	1535.8	535.7	823.8	405.6	1130.2	0.196
Cycle F.E.	Yes	Yes	Yes	Yes	Yes	Yes
C.D. F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Notes: point estimates from equations (2.2) and (2.3) are shown. The dependent variable in columns (1)-(5) is the amount of money raised in an election cycle from different sources in a given district, expressed in thousands of dollars. The dependent variable in column (6) is the share of total funds raised by the challengers. The sample includes competitive elections. Standard errors are clustered at the State level. Statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel B of table 2.4 includes the results of our within-party specification described in equation (2.3). As previously mentioned, this specification is not necessarily superior to our main specification, since within-party differences in environmental stances may be harder for people to

perceive. With the exception of the share of funds raised by challengers, all the estimates on the interaction term are of the same sign as those in Panel A, albeit with varying relative magnitudes and levels of statistical significance. For example, the effects on giving from PACs and giving to the incumbent remain significant and of similar size as those in Panel A, whereas the effects on giving to challengers and by individuals are reduced considerably and are not significant. Therefore, under this specification, we do not find evidence that giving to challengers is higher following a disaster if the incumbent has a more anti-environment voting record.

Having examined natural disaster impacts on fundraising, it is natural to also look at the number of donors to challengers and incumbents. We do so in table 2.5, and start by examining the results for our main specification in Panel A. The coefficients in columns (1)-(3) are all positive and significant, indicating that, following a natural disaster, there is a higher number of donors from all sources if the incumbent has a more anti-environment stance. For example, in column (1) we can see that a one-standard-deviation increase in the LCV score of the incumbent translates to 97 additional donors when there is a natural disaster. When we break it down by giving to incumbent versus challengers, the effect sizes are about 42 and 55 additional donors, respectively. As in the case of fundraising, we have seen that the number of donors increases both in the case of challengers and incumbents, and we would again like to explore the relative strength of these effects by looking at the share of donors contributing to challengers. In column (4), we find a similar effect, indicating that the share of donors to challengers is slightly higher, although it is not significant. However, we again note that the effect for challengers is larger when comparing it to its average, suggesting that the effect on giving to challengers might be relatively stronger.

In panel B, we report estimates from our within-party specification. In this case, none of these estimates is statistically significant. The estimated effect on giving to challengers is now negative, although the confidence interval is very wide, including economically significant effects in either direction. Therefore, the results in this panel are uninformative of whether within-party

Table 2.5: The effects of natural disasters on the number of donors

Dep. Var.	(1) Total	(2) Challenger	(3) Incumbent	(4) Share (C)
<i>Panel A: Main Specification</i>				
Disaster	-335.1*** (118.5)	-118.2** (50.78)	-216.8** (105.0)	-0.0301** (0.0139)
LCV	-515.5 (492.9)	47.11 (153.4)	-562.6 (392.1)	0.00104 (0.0279)
Disaster × LCV	484.9** (201.1)	211.4* (111.1)	273.5* (155.5)	0.0334 (0.0264)
R^2	0.297	0.216	0.320	0.283
<i>Panel B: Within-Party Specification</i>				
Disaster	-305.3** (122.0)	-88.64 (63.64)	-216.7** (106.2)	-0.0207 (0.0130)
LCV	-167.3 (611.3)	184.4 (236.5)	-351.7 (507.9)	0.0978** (0.0407)
Disaster × LCV	220.9 (306.4)	-58.37 (252.9)	279.2 (172.6)	-0.0507 (0.0333)
R^2	0.297	0.217	0.321	0.286
Observations	3299	3299	3299	3299
Mean D.V.	862.6	288.2	574.4	0.208
Cycle F.E.	Yes	Yes	Yes	Yes
C.D. F.E.	Yes	Yes	Yes	Yes

Notes: point estimates from equations (2.2) and (2.3) are shown. The dependent variable in columns (1)-(3) is the number of donors in an election cycle from different sources in a given district. The dependent variable in column (4) is the share of total donors corresponding to the challengers. The sample includes competitive elections. Standard errors are clustered at the State level. Statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

variation in the LCV score influences the number of donors following a natural disaster.

Next, we take the analysis a step further by exploring how natural disasters affect elections. There are several reasons for this. First, the campaign finance consequences of natural disasters, as shown above, may in turn affect electoral outcomes. Second, natural disasters may affect

the composition of elections by influencing the choice of prospective challengers to enter the race. Third, natural disasters may prompt issue voting by environmentally conscious constituents and directly influence the results of the election. Therefore, we focus on the following four outcomes: whether the election is competitive (i.e. there is a challenger), whether the incumbent runs unopposed, whether there is an open seat election (i.e. the incumbent does not run for re-election), and whether the incumbent is re-elected.

The results of this analysis are in table 2.6. In panel A we present results from our main specification, with the first three columns focusing on the election type and the fourth column focusing on the election outcomes. When it comes to the election type, we can see that if the incumbent in a disaster-struck district has a higher LCV score, they are more likely to face a challenger in the race. For a one-standard-deviation increase in the LCV score, the probability of the race being competitive following a disaster is almost 3% higher. As a result, the effect on the incumbent running unopposed is negative and similar in magnitude. The probability of an open seat election taking place is unaffected. Importantly, these results are in line with our hypothesis of increased support for challengers, although they relate to the presence of challengers (extensive margin) and not the strength of the support they receive (intensive margin). Potential challengers may join the race simply because of the increased funds they are able to raise, or because they recognize an opportunity to run on a pro-environment platform given the incumbent's record. As for the election outcome, we examine the impacts on the incumbent's re-election probability in column (4). The estimates show that the probability of the incumbent winning the election is lower following natural disasters if the incumbent has an anti-environment voting record. Specifically, for a one-standard-deviation difference in the LCV score this effect is about 1.3%. Note that there are at least two possible mechanisms behind this effect: first, given that we observe an increased probability of the election being competitive, this can mechanically lead to a reduced probability of an incumbent victory, as the incumbent is now more likely to face a challenger. Second, the effect may be driven by stronger opposition to the incumbent as a result of their voting record,

independently of the presence of a challenger.

Table 2.6: The effects of natural disasters on elections

Dep. Var.	(1) Competitive	(2) Unopposed	(3) Open Seat	(4) Incumbent Re-Election
<i>Panel A: Main Specification</i>				
Disaster	-0.0352 (0.0246)	0.0391 (0.0264)	-0.00390 (0.0158)	0.0470** (0.0189)
LCV	-0.103** (0.0493)	0.0484 (0.0512)	0.0547 (0.0410)	0.104** (0.0436)
Disaster × LCV	0.141*** (0.0411)	-0.117** (0.0444)	-0.0239 (0.0232)	-0.0629** (0.0254)
R^2	0.169	0.232	0.103	0.171
<i>Panel B: Within-Party Specification</i>				
Disaster	-0.0358 (0.0241)	0.0353 (0.0265)	0.000543 (0.0176)	0.0420** (0.0181)
LCV	-0.0853 (0.0786)	-0.0402 (0.0602)	0.125** (0.0614)	0.0243 (0.0477)
Disaster × LCV	0.147** (0.0640)	-0.0842 (0.0686)	-0.0632 (0.0531)	-0.0166 (0.0370)
R^2	0.169	0.233	0.104	0.173
Observations	4478	4478	4478	4095
Mean D.V.	0.737	0.178	0.0855	0.948
Cycle F.E.	Yes	Yes	Yes	Yes
C.D. F.E.	Yes	Yes	Yes	Yes

Notes: point estimates from equations (2.2) and (2.3) are shown. The dependent variable in columns (1)-(3) is the probability of a congressional race being of a certain type (competitive, unopposed, open seat). The dependent variable in column (4) is the probability that the incumbent is re-elected. Columns (1)-(3) include all elections and column (4) excludes open seat elections. Standard errors are clustered at the State level. Statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In panel B we examine the effects of natural disasters on elections using our within-party specification. The results look similar to those in panel A. Importantly, we still find evidence that if the incumbent in a disaster-struck district has a higher LCV score, they are more likely to face a

challenger in the race. This result, which is robust and significant across specifications, is in line with our hypothesis of increased support for challengers. However, under this specification we no longer find evidence that the probability of the incumbent being re-elected is reduced, although the estimate is still negative.³⁵

2.4.3 Medium-Run Weather Impacts

In the previous section we studied how natural disasters affected campaign finance and elections in the medium run. However, it is possible for other weather events to have a similar effect in the medium run, such as temperature shocks. In this section, we focus on how these shocks influence our political outcomes. Importantly, we treat temperature shocks as separate from natural disasters for two reasons: first, temperature shocks can be either hot weather shocks or cold weather shocks, each with their own possible ramifications. While people may interpret extremely hot weather as evidence of climate change, this may not be the case for extremely cold weather, which some may interpret as evidence against climate change, as we argued in the section on short-run impacts. Second, temperature shocks may be less salient than natural disasters, in part because the latter often results in property damage.

We start by looking at the effect of extreme temperature on the amount raised in an election in table 2.12. As explained in previous sections, we now study the effect of two events at once: whether the election cycle was abnormally hot or whether it was abnormally cold. In line with our previous analysis, we interpret the interaction terms to see how the effects of weather events vary with the LCV score. We focus on specifications of the type in equation (2.4). We can see that for abnormally hot cycles, the estimated coefficients on the interaction

³⁵Note how these findings related to election composition may affect the interpretation of our earlier results on amount raised and the number of donors. Since we focus on competitive elections, and the interaction of *LCV* and *Disaster* influences whether an election is competitive in the first place, our earlier results may capture a selection effect as well as a direct effect. The selection effect is likely to induce weaker challengers – those who would have not entered the race in absence of a natural disaster – to enter the race. This would likely bias our results towards lower fundraising, especially in the case of challengers, which could explain the absence of an effect on the share of funds going to challengers in table 2.4.

term are positive, as in the case of natural disasters, but they are insignificant and imprecisely estimated. The effects for abnormally cold cycles are also typically positive (with the exception of donations to PACs) and imprecisely estimated. In general, these results are not informative of the fundraising consequences of abnormal weather in the medium run. Next, we look at effects on the number of donors in table 2.13. The estimated effects are typically positive for abnormally hot election cycles and typically negative (with the exception of donors to challengers) for abnormally cold cycles. However, as before, these are imprecisely estimated. As in the natural disasters section, we also explore the effects of extreme temperatures on elections in table 2.14. For hot cycles, the effects point in the same direction as natural disasters, with races more likely to be competitive and a lower probability of incumbent re-election. However, these effects are smaller and not statistically significant. In the case of cold cycles, these relationships are largely reversed, although the estimates are still not statistically significant.

Given the level of statistical significance of our estimates, we cannot rule out that these are null results. A possible explanation for this is the salience of temperature events compared to natural disasters, as mentioned above. However, it is worth noting that in some of these results suggest the effects of hot and cold cycles are opposite in sign, with the effects of hot cycles looking more like those of natural disasters. While suggestive at best, these results point to the possibility that in this context people react differently to hot and cold weather abnormalities.

2.4.4 Mechanisms and Limitations

Throughout this paper, we have proposed the environmental preference mechanism – which contains two separate margins – as the driver of our results. First, we suggest that natural disasters and weather shocks lead people to update their beliefs about climate change, which in turn enables them to act on these beliefs politically. Second, we suggest that extreme weather events simply increase the salience of climate change and other environmental issues, which in turn incentivizes those who have already made up their mind on these issues to take action. We do

not attempt to distinguish between these two mechanisms, and acknowledge that our observations may stem from either or both of these. However, there may still be other possible explanations for these results. In this section, we address limitations and alternative mechanisms, discussing them in the context of our findings.

One mechanism that is especially relevant to our results regarding weather shocks is time use. If short- or medium-run weather shocks affect time use, which in turn may affect giving behavior, then this could be a driving force behind our results. This is especially important in the case of online giving, since if weather leads people to spend more time indoors then this could expose them to more opportunities for online giving. Importantly, if time spent indoors is driving our results, then results should be similar for hot weather shocks and cold weather shocks since it has been shown that both types of shock can lead to more time spent indoors (Graff Zivin and Neidell, 2014). However, in both short-run and medium-run analyses, we find that hot and cold events generally have different, sometimes opposite effects on campaign contributions. In addition, in the medium-run analysis we include contributions from all sources, and the patterns are generally similar even though the effects are often not statistically significant. Therefore, time spent indoors is not likely to account for our results.

In the case of natural disasters, an important alternative mechanism is that of factors that are correlated with the LCV score but unrelated to incumbents' stances on environmental issues. For example, if pro-environment candidates are also more willing to pass disaster relief packages for those affected, this may explain the increase in funds and support for these candidates following disasters. While this is a possibility, we argue that these factors are not likely to be driving our results for two reasons. First, the observed effects of unusually hot weather in the short and medium-run are in the same direction as those of natural disasters, and there is no obvious policy position regarding hot weather other than a politician's stance on environmental issues. This is especially true in the case of short-run weather variations. Second, to alleviate this concern we study the effects of natural disasters using a within-party specification, since policy

positions on several issues are determined along party lines. We find that estimates generally point in the same direction as in our main specification, despite party affiliation being perhaps the most visible source of information on politicians' environmental positions.

Finally, it is possible that there are other psychological explanations for our results. For example, Meier et al. (2016) explore the link between rainy weather, risk aversion, and voting for status quo candidates. This link between short term weather and emotions could be a confounder to the extent that emotions affect individuals' incentives to make political campaign contributions. However, we observe similar patterns in the medium term, as well as stronger short-term effects in counties with more anti-environment incumbents, so our results cannot be entirely driven by the emotional consequences of short-run weather.

2.5 Conclusion

In this paper, we study the impacts of extreme weather events on campaign contributions and electoral outcomes in the United States. As these events are often considered signs of climate change, our analyses place particular emphasis on testing for differential constituent responses based on the incumbent politician's views on environmental issues. In a short-run analysis, we find that weekly temperature shocks lead to a higher number of online donations to Democratic candidates, especially in counties with a greater share of anti-environment incumbents. In a medium-run analysis, we find that when a natural disaster strikes, challengers are more likely to enter congressional races the more anti-environment the incumbent is. We also find suggestive evidence that natural disasters lead to increased overall funding in congressional races where the incumbent is more anti-environment.

Throughout the paper, we have argued that our results are driven by the environmental preference mechanism. We have suggested that this mechanism is a combination of two margins. First, extreme weather may lead people to update their beliefs in climate change, which in turn

enables them to act on these beliefs politically. Second, extreme weather events may simply increase the salience of climate change, which in turn incentivizes those who have already made up their mind on these issues to take action.

Irrespective of the margin at play, the results in this paper suggest that politicians' policy positions on environmental issues are taken into account by people when responding to extreme weather events, and that these responses may have political consequences. Further, these findings suggest additional mechanisms for results in previous studies showing that congresspersons are more likely to vote in favor of environmental legislation following natural disasters in their state (Herrnstadt & Muehlegger, 2014). Put together, these behaviors from constituents, candidates, and legislators are consistent with a representative democracy at work.

Even though our results have important implications as they stand, both margins of the environmental preference mechanism have their own ramifications. If the belief-updating margin is at play, one concern is the extent to which such updating is rational; as is often pointed out in the literature, people tend to process shocks with psychological and ideological biases (Tversky & Kahneman, 1974). If this is the case, our findings caution that such biases might be expressed through the political system with unclear welfare consequences. On the other hand, if the salience margin is at play, one might instead be concerned about the electorate having limited attention. If this is the case, election outcomes might not represent well-defined collective preferences, since the salience of different policy issues may vary due to stochastic events. Given these ramifications, it is important to distinguish which margin of the environmental preference mechanism is predominantly at play. Therefore, we believe this is an important question for future work.

In addition to distinguishing between the two margins of the environmental preference mechanism, the results presented in this paper pose a series of additional questions and possible extensions. Firstly, a question raised by this work is whether the politicians themselves react to the salience of climate change by adjusting their narratives when it comes to speeches and

soliciting contributions. Secondly, an important player which is missing from our analysis are environmental advocacy groups. Future research should focus on the role these groups play in disseminating information and forming opinions following extreme weather events. Finally, it is an open question whether the behavior observed here generalizes to other policy areas in which the event of interest has a stochastic component, like terrorism or gun violence.

2.6 Acknowledgement

Chapter 2 is currently being prepared for submission for publication of the material. Liao, Yanjun; Ruiz Junco, Pablo; “Extreme Weather and the Politics of Climate Change: a Study of Campaign Contributions and Elections.” The dissertation author was a principal co-investigator and co-author of this material.

2.7 Appendix

2.7.1 Additional Figures



Figure 2.3: Example of an ActBlue donation

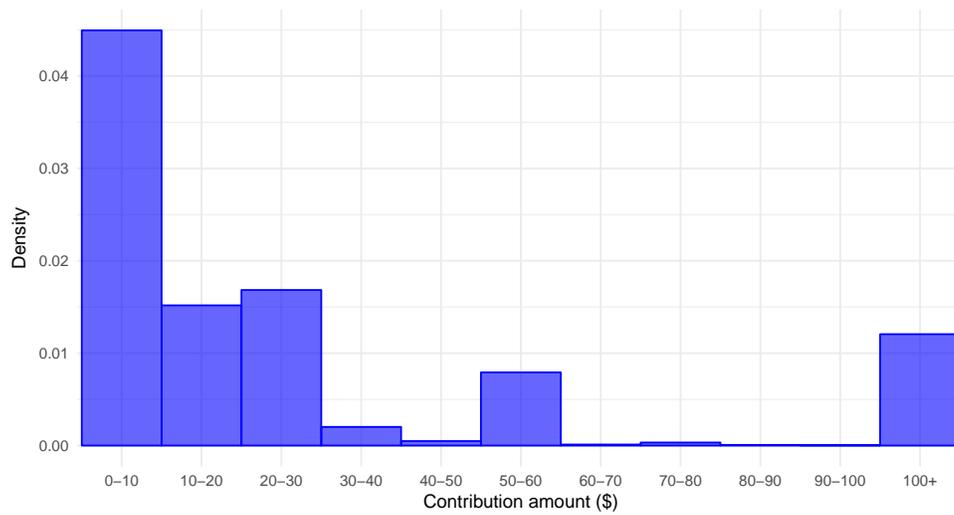


Figure 2.4: Distribution of individual ActBlue donation amounts

2.7.2 Additional Tables

Table 2.7: FEMA disaster declarations 1990-2012

Type	# Disasters	# County-Year Obs.
Coastal Storm	18	402
Drought	5	178
Fire	785	2,535
Flood	178	3,198
Freezing	6	127
Hurricane	233	7,680
Severe Ice Storm	44	1,566
Severe Storm(s)	738	13,187
Snow	126	2,745
Tornado	41	480
Tsunami	3	9
Typhoon	29	69
Total	2,206	32,176

Table 2.8: Predicting total Democratic donations using ActBlue donations

Dep. Var.	Amount		Number	
	(1)	(2)	(3)	(4)
ActBlue	85.33*** (5.63)		14.67*** (1.30)	
ActBlue × 2006		209.93*** (24.98)		32.58*** (8.14)
ActBlue × 2008		99.09*** (6.14)		21.95*** (2.97)
ActBlue × 2010		57.51*** (6.32)		12.15*** (2.44)
ActBlue × 2012		111.23*** (6.55)		14.63*** (1.00)
Observations	200	200	200	200
R ²	0.74	0.86	0.83	0.88

Notes: the above table includes point estimates from various O.L.S. regressions of the amount and number of donations to Democrats from all sources, to the amount and donations from ActBlue sources. All regressions include an intercept term. Standard errors are clustered at the state level. Statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.9: Positive and negative temperature shocks on ActBlue contributions

Dep. Var.	(1) Count/1M pop	(2) Avg. amount
$TmaxDev_{c,w}^+$	0.274*** (0.098)	-0.070 (0.052)
$TmaxDev_{c,w-1}^+$	0.110*** (0.029)	-0.075 (0.059)
$TmaxDev_{c,w-2}^+$	0.165*** (0.041)	0.072 (0.059)
$TmaxDev_{c,w-3}^+$	0.173*** (0.037)	-0.096** (0.042)
$TmaxDev_{c,w-4}^+$	0.124*** (0.033)	-0.035 (0.039)
$TmaxDev_{c,w}^-$	0.015 (0.031)	-0.096 (0.059)
$TmaxDev_{c,w-1}^-$	-0.095 (0.074)	0.003 (0.039)
$TmaxDev_{c,w-2}^-$	0.070* (0.041)	-0.041 (0.051)
$TmaxDev_{c,w-3}^-$	0.036 (0.032)	-0.175** (0.075)
$TmaxDev_{c,w-4}^-$	0.029 (0.031)	0.031 (0.060)
N	935,201	935,201
R^2	0.204	0.054
Mean D.V.	15.40	13.19
County F.E.	Yes	Yes
Week F.E.	Yes	Yes
State-cycle F.E.	Yes	Yes

Notes: point estimates from equation (2.5) are shown. The dependent variable in column (1) is the number of ActBlue contributions per 1 million population. The dependent variable in column (2) is the average amount per ActBlue contribution. The sample consists of ActBlue contributions by week and county. Standard errors (in parenthesis) are clustered by county. Statistical significance: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 2.10: The effect of extreme temperature events on ActBlue contributions

Dep. Var.	(1) Count/1M pop	(2) Avg. amount
$TmaxHigh_{c,w}$	0.353* (0.183)	-0.043 (0.138)
$TmaxHigh_{c,w-1}$	0.211** (0.092)	0.058 (0.187)
$TmaxHigh_{c,w-2}$	0.224*** (0.082)	0.215 (0.211)
$TmaxHigh_{c,w-3}$	0.277** (0.135)	-0.058 (0.126)
$TmaxHigh_{c,w-4}$	0.010 (0.159)	-0.032 (0.125)
$TmaxLow_{c,w}$	-1.020*** (0.368)	-0.318 (0.229)
$TmaxLow_{c,w-1}$	-0.641* (0.336)	0.176 (0.183)
$TmaxLow_{c,w-2}$	-0.083 (0.188)	-0.063 (0.211)
$TmaxLow_{c,w-3}$	-0.315* (0.187)	-0.140 (0.220)
$TmaxLow_{c,w-4}$	-0.372** (0.170)	0.442 (0.368)
N	936,954	936,954
R^2	0.203	0.054
Mean D.V.	15.40	13.19
County F.E.	Yes	Yes
Week F.E.	Yes	Yes
State-cycle F.E.	Yes	Yes

Notes: point estimates from equation (2.6) are shown. The dependent variable in column (1) is the number of ActBlue contributions per 1 million population. The dependent variable in column (2) is the average amount per ActBlue contribution. The sample consists of ActBlue contributions by week and county. Standard errors (in parenthesis) are clustered by county. Statistical significance: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 2.11: Heterogeneous effects across quarters in election cycle

Dep. Var.	(1) Count/1M pop	(2) Avg. amount	(3) Amount/1M pop
$\overline{TmaxDev}_{c,w}^{2W} \times$			
Q1 (Dec-Feb)	-0.114** (0.049)	0.401* (0.241)	-1.901 (4.431)
Q2 (Mar-May)	0.183** (0.072)	0.086 (0.158)	-4.271 (4.233)
Q3 (Jun-Aug)	0.681*** (0.093)	-0.273*** (0.078)	-3.407 (5.062)
Q4 (Sep-Nov)	0.200*** (0.051)	0.054 (0.062)	2.314 (4.076)
Q5 (Dec-Feb)	0.324*** (0.073)	-0.077* (0.046)	14.826* (7.697)
Q6 (Mar-May)	0.010 (0.063)	0.142** (0.055)	-2.968 (4.634)
Q7 (Jun-Aug)	0.023 (0.131)	0.223** (0.095)	8.640 (7.540)
Q8 (Sep-Nov)	1.256** (0.596)	-0.795*** (0.126)	31.664* (18.835)
<i>N</i>	941,672	941,672	941,672
<i>R</i> ²	0.209	0.054	0.073
Mean D.V.	15.40	13.19	639.65
County F.E.	Yes	Yes	Yes
Week F.E.	Yes	Yes	Yes
State-cycle F.E.	Yes	Yes	Yes

Notes: point estimates from equation (2.8) are shown, which allows the estimates to differ by quarter-in-cycle. The dependent variable in column (1) is the number of ActBlue contributions per 1 million population. The dependent variable in column (2) is the average amount per ActBlue contribution. The dependent variable in column (3) is the ActBlue contribution amount per 1 million population. The sample consists of ActBlue contributions by week and county. Standard errors (in parenthesis) are clustered by county. Statistical significance: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 2.12: The effects of extreme temperature on amount raised (\$1,000)

Dep. Var.	(1) Total	(2) Total (PAC)	(3) Total (Ind.)	(4) Challenger	(5) Incumbent	(6) Share (C)
Hot	-106.4 (137.8)	-6.998 (29.85)	-56.97 (92.21)	-28.24 (78.99)	-78.15 (83.71)	0.00160 (0.0187)
LCV	-75.00 (257.1)	-21.50 (40.37)	37.19 (210.4)	-121.5 (97.15)	46.53 (192.8)	-0.0269 (0.0232)
Hot × LCV	123.5 (231.0)	0.871 (52.98)	65.81 (157.6)	12.90 (108.9)	110.6 (153.7)	0.00538 (0.0277)
Cold	-43.86 (117.2)	0.333 (22.41)	-18.95 (91.79)	-28.16 (78.13)	-15.70 (61.89)	-0.00420 (0.0136)
Cold × LCV	92.25 (154.3)	-23.10 (34.22)	85.53 (129.0)	22.17 (93.32)	70.08 (91.45)	-0.0122 (0.0178)
Observations	3299	3299	3299	3299	3299	3299
R^2	0.411	0.531	0.386	0.259	0.456	0.285
Mean D.V.	1535.8	535.7	823.8	405.6	1130.2	0.196
Cycle F.E.	Yes	Yes	Yes	Yes	Yes	Yes
C.D. F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Notes: point estimates from equation (2.4) are shown. The dependent variable in columns (1)-(5) is the amount of money raised in an election cycle from different sources in a given district, expressed in thousands of dollars. The dependent variable in column (6) is the share of total funds raised by the challengers. The sample includes competitive elections. Standard errors are clustered at the State level. Statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.13: The effects of extreme temperature on the number of donors

Dep. Var.	(1) Total	(2) Challenger	(3) Incumbent	(4) Share (C)
Hot	-204.2 (125.7)	-170.8** (76.48)	-33.37 (94.47)	-0.0114 (0.0219)
LCV	-342.7 (371.5)	116.7 (171.5)	-459.4 (285.3)	0.0225 (0.0250)
Hot × LCV	476.0 (311.7)	230.6 (197.1)	245.4 (188.0)	0.00832 (0.0369)
Cold	28.87 (242.4)	-33.83 (91.73)	62.70 (179.7)	0.00152 (0.0171)
Cold × LCV	-16.09 (351.2)	7.949 (137.7)	-24.03 (254.6)	-0.0145 (0.0344)
Observations	3299	3299	3299	3299
R^2	0.296	0.217	0.320	0.282
Mean D.V.	862.6	288.2	574.4	0.208
Cycle F.E.	Yes	Yes	Yes	Yes
C.D. F.E.	Yes	Yes	Yes	Yes

Notes: point estimates from equation (2.4) are shown. The dependent variable in columns (1)-(3) is the number of donors in an election cycle from different sources in a given district. The dependent variable in column (4) is the share of total donors corresponding to the challengers. The sample includes competitive elections. Standard errors are clustered at the State level. Statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.14: The effects of extreme temperature on elections

	(1) Competitive	(2) Unopposed	(3) Open Seat	(4) Incumbent victory
Hot	-0.0574* (0.0289)	0.0516* (0.0294)	0.00582 (0.0181)	0.00733 (0.0138)
LCV	-0.0144 (0.0467)	-0.0220 (0.0492)	0.0364 (0.0310)	0.0732* (0.0385)
Hot × LCV	0.0568 (0.0411)	-0.0608 (0.0374)	0.00397 (0.0309)	-0.0197 (0.0278)
Cold	0.0245 (0.0292)	-0.00792 (0.0260)	-0.0166 (0.0167)	0.0124 (0.0145)
Cold × LCV	-0.0722 (0.0545)	0.0622 (0.0448)	0.0101 (0.0314)	-0.00588 (0.0208)
Observations	4478	4478	4478	4095
R^2	0.166	0.231	0.103	0.169
Mean D.V.	0.737	0.178	0.0855	0.948
Cycle F.E.	Yes	Yes	Yes	Yes
C.D. F.E.	Yes	Yes	Yes	Yes

Notes: point estimates from equation (2.4) are shown. The dependent variable in columns (1)-(3) is the probability of a congressional race being of a certain type (competitive, unopposed, open seat). The dependent variable in column (4) is the probability that the incumbent is re-elected. Columns (1)-(3) include all elections and column (4) excludes open seat elections. Standard errors are clustered at the State level. Statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.7.3 Additional Specifications in the Short-Run Analysis

As shown in Section 2.3.1, our main estimating equation takes the form

$$Y_{cw} = \gamma' Weather_{cw} + \delta_w + \delta_c + \delta_{se} + \varepsilon_{cw}.$$

Specifications for table 2.2:

– Columns (1) and (4): $Weather_{cw} = [TmaxDev_{cw}, PrcpDev_{cw}]^T$.

– Columns (2) and (5):

$$Weather_{cw} = [TmaxDev_{cw}, \dots, TmaxDev_{c,w-4}, PrcpDev_{cw}, \dots, PrcpDev_{c,w-4}]^T$$

– Columns (3) and (6): $Weather_{cw} = [\overline{TmaxDev}_{c,w}^{2W}, \overline{PrcpDev}_{c,w}^{2W}]^T$ where

$$\overline{TmaxDev}_{c,w}^{2W} = \frac{1}{2}(TmaxDev_{cw} + TmaxDev_{c,w-1})$$

and

$$\overline{PrcpDev}_{c,w}^{2W} = \frac{1}{2}(PrcpDev_{cw} + PrcpDev_{c,w-1}).$$

Specification for table 2.9:

$$Weather_{cw} = [TmaxDev_{cw}^+, \dots, TmaxDev_{c,w-4}^+, TmaxDev_{cw}^-, \dots, TmaxDev_{c,w-4}^-, PrcpDev_{cw}, \dots, PrcpDev_{c,w-4}]^T \quad (2.5)$$

where $TmaxDev^+ = TmaxDev \times (TmaxDev > 0)$ and $TmaxDev^- = |TmaxDev| \times (TmaxDev < 0)$. This specification allows us to estimate the effects of positive and negative deviations separately.

Specification for table 2.10:

$$\begin{aligned} Weather_{cw} = [TmaxHigh_{cw}, \dots, TmaxHigh_{c,w-4}, TmaxLow_{cw}, \dots, TmaxLow_{c,w-4}, \\ PrcpDev_{cw}, \dots, PrcpDev_{c,w-4}]^T, \end{aligned} \quad (2.6)$$

where $TmaxHigh_{cw}$ is the total number of days in week w when the maximum temperature exceeds 95th percentile of the historical distribution in the month, and $TmaxLow_{cw}$ counts days with temperature below the 5th percentile.

Specification for table 2.3:

$$\begin{aligned} Y_{cw} = \beta_1 \overline{TmaxDev}_{c,w}^{2W} + \beta_2 IncChar + \beta_3 \overline{TmaxDev}_{c,w}^{2W} \times IncChar + \\ \gamma \overline{PrcpDev}_{c,w}^{2W} + \delta_w + \delta_c + \delta_{se} + \epsilon_{cw}, \end{aligned} \quad (2.7)$$

where $\overline{TmaxDev}_{c,w}^{2W}$ and $\overline{PrcpDev}_{c,w}^{2W}$ are defined as above. $IncChar$ is a measure that characterizes the environmental stance of the incumbents associated with a county. Our coefficient of interest is β_3 , which shows how the effects of weather shocks vary according to the incumbent characteristic.

Specification for figure 2.2 and table 2.11:

$$Y_{cw} = \sum_{t=1}^8 \beta_t \overline{TmaxDev}_{c,w}^{2W} \times Q_t + \gamma \overline{PrcpDev}_{c,w}^{2W} + \delta_w + \delta_c + \delta_{se} + \epsilon_{cw}, \quad (2.8)$$

where $\overline{TmaxDev}_{c,w}^{2W}$ and $\overline{PrcpDev}_{c,w}^{2W}$ are defined as above. Q_t is a set of eight indicators for quarters in the election cycle. This specification allows us to obtain a separate estimate for each quarter-in-cycle.

Chapter 3

Short and Long Term Effects of Extended Unemployment Insurance for First Time Claimants

3.1 Introduction

The effect of unemployment insurance (UI) on individuals' labor market outcomes has long been a widely debated topic in both policy and research circles. While the debate is not a new one, it has resurfaced over the past decade in light of the Great Recession and the emergence of new data sources. Furthermore, a growing body of research has started to examine the long term effects of public programs, noting that these effects – and as a result their policy implications – may be different than those in the short term.

In this paper, I use administrative data from Spain to study the effects of increased unemployment insurance (UI) benefit duration on individuals' short term and long term labor market outcomes. I make use of the shape of the function determining UI benefit duration to implement a regression discontinuity design at nine different cutoffs. This function relates the

number of days worked over the past six years before the unemployment spell to UI benefit duration, and jumps discontinuously by 60 days at each of the cutoffs. Throughout this article, I focus on a sample of young workers who are claiming UI benefits for the first time and follow them up to five years after the start of their unemployment spell. For conciseness, I divide these workers into two groups. The first group is comprised of workers with low labor force attachment who are located at the earlier cutoffs, since they have been working for fewer days over the past six years at the start of their unemployment spell. The second group is comprised of workers with high labor force attachment, located at the later cutoffs, since they have been working longer. These two groups also differ in terms of treatment intensity, since those who have been working longer have a higher baseline UI benefit duration, and the treatment is always 60 days. In a short term analysis, I show that for workers with low levels of labor force attachment, receiving an additional 60 days of UI benefits induces them to spend about 21 more days on UI and 14 more days in nonemployment. In a long term analysis, I show that these workers spend 7 additional days on UI as a result of extended benefits, compared to the short term. Furthermore, I find suggestive evidence that the effect on days spent in nonemployment in the long term is in fact lower than in the short term, which is consistent with increased post-unemployment job stability. However, I interpret this result with caution due to robustness concerns. Further, I do not find evidence that this effect is driven by increased probability of transitioning into a permanent contract.

The existing literature focusing on UI benefit duration typically studies how the probability of exiting unemployment varies with the time remaining until benefit exhaustion, how extended duration affects time spent receiving benefits, and how extended duration affects the time until an individual finds a new job.¹ In a seminal paper, Meyer (1990) uses U.S. data to document that the probability of leaving unemployment sharply increases close to benefit exhaustion. This feature of the hazard function is also documented by Katz and Meyer (1990), who suggest that a one week

¹While UI benefit duration has been more widely studied, there is also an important literature on UI benefit amount. For examples, see Røed and Zhang (2003) and Sovago (2015).

increase in benefits increases unemployment duration up to 0.2 weeks. Card and Levine (2000) study a program that offered 13 weeks of extended benefits in New Jersey and calculate that the average recipient would have collected benefits for an extra week. Furthermore, there have been a number of papers examining the effects of extended UI duration in European countries as well. Card, Chetty, and Weber (2007) use Austrian data to emphasize the distinction between using exits from registered unemployment versus re-employment to measure spikes, documenting much smaller spikes using the latter definition. Lalive (2007) also uses Austrian data and finds that large UI duration extensions increase unemployment duration and duration until a new job is taken. Boone and van Ours (2012) use Slovenian data and suggest that workers exploit UI benefits for subsidized leisure (moral hazard), although Chetty (2008) suggests that 60% of the observed increase in unemployment duration as a result of UI benefits can be explained by a liquidity effect and not moral hazard. Johannes F Schmieder, von Wachter, and Bender (2012) estimate the nonemployment effects of UI extensions in Germany throughout the business cycle, and argue that these effects are not more severe during recessions and in fact may be slightly weaker during them. Caliendo, Tatsiramos, and Uhlendorff (2013) also study the German UI system and show that unemployment duration is higher for workers receiving 18 months of UI compared to those receiving 12 months. Finally, Gonzalez-Rozada and Ruffo (2014) and Gonzalez-Rozada and Ruffo (2016) study the UI system in Argentina, an economy with a high degree of informality, and show that extended UI duration increases the length of unemployment spells but that increased benefits do not.

There is also a literature studying unemployment and UI in Spain, which is the context of this study. Bover, Arellano, and Bentolila (2002) suggest that the receipt of unemployment benefits reduces the hazard rate of leaving unemployment and that this rate can be up to twice as large for workers without benefits. Furthermore, J Ignacio García-Pérez and Muñoz-Bullón (2011) emphasize the role of individual heterogeneity in estimating exit rates and suggest that those receiving unemployment benefits exit unemployment more slowly. Arranz and Muro Romero

(2004) compare UI to another public program, unemployment assistance (UA) which assists certain workers who are not eligible for UI and does not depend on work history. They find that while the hazard rate for UI rises dramatically at benefit exhaustion, this is not the case for UA. Alba-Ramírez, Arranz, and Muñoz-Bullón (2007) distinguishes between workers who are recalled to their old job and those who find new ones. They find that the job finding hazard rate increases with unemployment duration but that the recall hazard rate drops to zero, suggesting that individuals intensify search when they realize they will not be recalled by their former employers. On the other hand, Y. Rebollo-Sanz (2012) studies job turnover and distinguishes between workers who are laid off (eligible for UI benefits) and those who voluntarily leave their jobs (not eligible for UI benefits). They suggest that the layoff hazard rate increases as workers qualify for UI but remains stable for quits, concluding that UI benefits seem to favor job turnover. In the first part of this paper, I contribute to this literature by studying short term effects of UI extensions for a sample of particularly young workers in Spain, while achieving identification by using a regression discontinuity design.

A key takeaway from the above literatures is that extended UI duration does seem to lead to longer periods of unemployment. In more recent years, research has gone beyond measuring the effects of UI extensions on unemployment duration and has attempted to study the period following unemployment. The main question this literature seeks to answer is whether longer periods of unemployment can lead to better job matches, which would be beneficial for workers and firms. However, there is no clear consensus yet. For example, José Ignacio García-Pérez and Rebollo-Sanz (2005) study European countries and find that job mobility through unemployment benefits has negative returns in all of the countries they examine. Tatsiramos (2009) also focuses on Europe and suggests that the effect of benefits on employment stability is stronger in countries with more developed and generous UI systems. In the case of Slovenia, Van Ours and Vodopivec (2008) find no detectable effects on post-UI wages, transitions to permanent contracts, or duration of the post-UI job. In Argentina, Gonzalez-Rozada and Ruffo (2014, 2016)

suggest that extended UI duration has no effect on re-employment wages, but raising the benefit level does. In Germany, Caliendo et al. (2013) show lower wages and lower job stability for those finding jobs close to benefit expiration, compared to workers who receive extended benefit duration. In the Netherlands, Sovago (2015) studies the effect of benefit amount on labor market outcomes, finding that while increased benefits do affect unemployment duration, there is no evidence that they affect re-employment wages or the number of working hours. Finally, Lalive (2007) studies the Austrian UI system and does not find any effects on wages.

While the literature in the above paragraph is suggestive that the *short term* effect of extended UI duration on employment may differ from the *long term* effect, the empirical evidence on long term outcomes is scarcer. These two effects may differ for the reason stated above: if longer unemployment duration leads to better match quality for workers, then employment stability for them may rise and lead to an increase in employment in the long term. On the other hand, there are reasons to believe that extended periods of unemployment may be harmful in the long term. One reason this could be the case is that human capital may depreciate during unemployment, leading to worsened outcomes in the future. Another reason long term outcomes may be worsened is that firms may prefer workers who have been unemployed for the least amount of time, putting those with longer unemployment duration at a disadvantage (Blanchard & Diamond, 1994). Further, Lemieux and MacLeod (2000) point out that first-time exposure to the UI system may increase lifetime use of UI, which they describe as “supply side hysteresis.” In terms of more recent studies, there are only a few that take a long term perspective to study the effects of extended UI duration on employment outcomes. One such study is Johannes F. Schmieder, von Wachter, and Bender (2012), which I follow closely in terms of methodology. This paper studies the German UI system and takes advantage of a sharp increase of six months in benefit duration for workers aged 42 or older to implement a regression discontinuity design. They find that while workers with extended benefit duration seem to use more days of UI benefits in the long term, the amount of days they spend nonemployed is actually *lower*. In Spain, a paper

that takes a more long term perspective is Y. F. Rebollo-Sanz and García-Pérez (2015). This paper uses a multivariate mixed proportional hazard model to study how the receipt of benefits affects subsequent job stability. They show that benefits do encourage job stability for workers on temporary contracts and part of this effect can be explained by an increased probability of transitioning into a permanent contract.

In the second part of this paper, I contribute to the literature on long term effects of UI. I focus on estimating the long term effects of extended UI duration on the amount of days spent on UI and the amount of days spent in nonemployment over the five years following an unemployment spell. In particular, I study the case of workers who claim UI benefits for the first time, which is a particularly young sample of workers. I believe this is interesting for a number of reasons. First, this contrasts with other papers in the literature which focus on older workers, such as those studying Austria or Germany. Second, workers are thought to accumulate human capital and occupation-specific capital faster earlier in their careers, so we may expect the effects of UI to be stronger for younger workers who experience unemployment spells. Third, the long term effects of extended UI duration may be stronger the first time a worker is exposed to the UI system (Lemieux & MacLeod, 2000), which is the case for the workers in my sample. Furthermore, this study also contributes to this literature in the Spanish context by using a regression discontinuity design to achieve identification. Finally, in this study there are nine different discontinuities at which the treatment intensity and the types of workers affected vary, allowing me to study effects for different populations.

The remainder of the paper proceeds as follows. In section 3.2 I introduce the institutional framework, in section 3.3 I describe the data, and in section 3.4 I go over the empirical methodology in detail. In section 3.5 I present the results, and then proceed to discuss these in section 3.6. Finally, I conclude in section 3.7.

3.2 Institutional Framework

3.2.1 The Spanish Unemployment Insurance System

The Spanish Unemployment Insurance (UI) system is the main form of social insurance in the country protecting workers against negative labor shocks. The UI system is administered by the Spanish Social Security Administration (SSSA) and financed through a payroll tax. Whenever a firm hires a worker, a worker becomes *affiliated* to the SSSA through their relationship with the firm. For the duration of the contract, a payroll tax is applied to the worker's income and deducted from their paycheck, which is then deposited at the SSSA to fund the nationwide program.

A worker may become eligible for UI benefits for a variety of reasons, including termination with cause, termination without cause, if the worker is laid off, or if the worker is on a temporary contract which expires and is not renewed by the employer. However, a worker may not voluntarily quit their job and become eligible for UI benefits. After a worker becomes unemployed, they must sign up to receive UI benefits at the regional office of the SSSA. Once the worker has been deemed eligible, the UI award amount will be proportional to the worker's income during the last six months. The specific amount will depend on the "regulation base," which corresponds to a portion of the worker's income, starting at a minimum amount for all workers and maxing out for the highest incomes. The regulation base will also loosely depend on the worker's contribution group, an occupational classification based on the level of skills and education required for a job, used internally by the SSSA.² Specifically, a worker will receive a UI award of 70% of their regulation base up to six months, and 50% of the base thereafter, until UI benefits are exhausted.

Once a worker is receiving UI benefits, they must follow a set of rules to remain eligible. Failure to follow these rules can result in sanctions, such as losing days of UI eligibility. Specifically, if a worker is charged with a minor offense, they will lose a full month of UI benefits, and

²Contribution groups will be discussed in further detail in the next section.

will be at risk of losing all remaining days of benefits after four minor offenses. Workers can also be charged with serious offenses, which will lead to a loss of 3 months of UI benefits, and loss of all benefits after three serious offenses. Minor offenses include failure to periodically verify unemployed status at the regional employment office or failure to make oneself physically present at the office when required, unless the worker has a valid excuse. Serious offenses include refusal to accept an employment offer that is deemed to be adequate given the worker's characteristics or refusal to participate in employment or retraining programs.

Duration of UI benefits will depend on a worker's employment history. In particular, duration will depend on the number of days a worker has been affiliated to the SSSA during the six years preceding the start of the worker's unemployment spell. Days of affiliation for full-time workers are calculated by adding up the days between the start of a work contract and the end of that contract, while avoiding double-counting for those working more than one job.³ Benefit duration will be longer for those with longer affiliation periods, since these workers have been contributing to SSSA for a longer period of time. In particular, duration will be determined by a step function, depicted in figure 3.1.⁴

According to figure 3.1, workers that become unemployed and have been affiliated to the SSSA for less than 360 days in the past six years are not eligible for UI. However, if a worker has been affiliated to the SSSA for 360 days, then the worker becomes eligible for 120 days of UI benefits. This is also the case for workers that were affiliated to the SSSA for more than 360 days but less than 540 days, at which point the worker would become eligible for 180 days of UI. The step function continues to follow this pattern, increasing UI duration by 60 days for every additional 180 days of affiliation to SSSA over the past 6 years. Finally, if a worker has been affiliated to the SSSA for 2160 days or longer, they are eligible for the maximum award duration, 720 days.

³In this study, I will focus on full-time workers, so this is the relevant calculation. However, the rules for calculating days of affiliation for part-time workers are different than those for full-time workers.

⁴Throughout the paper, I will use the terms "affiliation" and "labor force attachment" interchangeably.

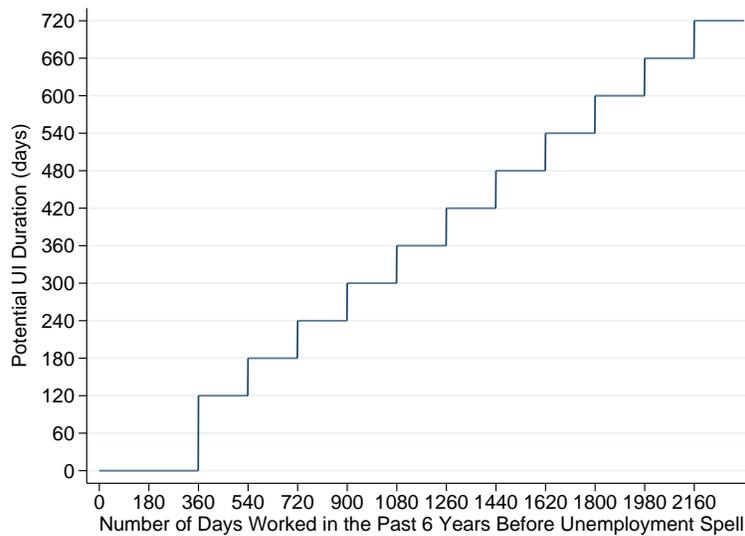


Figure 3.1: The potential UI duration function

As explained in detail later on, I make use of this feature of the Spanish UI system to implement a regression discontinuity design, comparing the outcomes of workers close to the cutoffs in figure 3.1, but on opposite sides. I focus on the *interior* cutoffs, namely those at 540, 720, 900, 1080, 1260, 1440, 1620, 1800, and 1980 days; a total of nine cutoffs. Importantly, this setup allows me to estimate each discontinuity separately, or in bundles, to learn about parameters of interest for workers with different levels of labor force attachment, and for whom the size of the treatment is different in proportion to the baseline duration of benefits.⁵

3.2.2 The Spanish Economy

The Spanish economy is characterized by a persistently high level of unemployment compared to other European Union (EU) and Organisation for Economic Co-operation and Development (OECD) countries. In figure 3.7 I plot the evolution of the unemployment rate for the period 1990-2015.⁶ During this period, the lowest unemployment rate was 8% in 2007, while

⁵Note that while the jump is always 60 days, the baseline duration of UI is increasing for workers with higher levels of labor force attachment.

⁶In this study, I focus on unemployment spells starting during 1994 or later, up to early 2015.

the highest unemployment rate was 27% in 2013. The average was 17%. Importantly, the period of interest is representative of the business cycle, since it contains two recessions and a long expansion.

There are a number of macroeconomic trends that drive the unemployment rate over the period of reference. Prior to this period, Spain had joined the EU in 1986, furthering economic integration between Spain and other European countries. In the early 1990s, the global recession affecting other developed countries reached Spain, driving up the unemployment rate. This recession was followed by a strong expansion which includes the introduction of the European common currency, the Euro, between 1999 and 2002. During this period, unemployment came down. In 2008, the unemployment rate started rising again as the country was strongly affected by the global economic recession. In Spain, the recession was characterized by the collapse of housing prices and the construction sector, which had performed strongly throughout the previous expansion. The financial crisis of 2008 was followed by a European sovereign debt crisis, and as a result the unemployment rate continued to rise until it peaked in 2013, falling thereafter.

Perhaps the most distinctive characteristic of the Spanish economy is the *duality* of its labor market. This duality arises from the great divide in job protections between permanent and temporary contracts. Permanent contracts have significant mandatory severance payments and are characterized by high levels of job protection, whereas temporary contracts are characterized by low levels of protection. During the early 1980s, permanent contracts comprised around 90% of all contracts. However, due to a series of regulatory changes, temporary contracts comprised around 35% of all contracts by the early 1990s (Dolado, García-Serrano, & Jimeno, 2002). Therefore, most hiring during that period was in the form of temporary contracts. In fact, Miller and Llull (2016) report that between 1991 and 2012, 84% of employment contracts signed were temporary ones. Bentolila, Cahuc, Dolado, Le Barbanchon, et al. (2010) argue that this gap in dismissal costs is an explanation for the higher unemployment rate observed in Spain recently compared to France, which does not have such a gap in dismissal costs between workers in temporary

and permanent contracts. Finally, as a result of this duality, I later consider the probability of transitioning to a permanent contract as an additional outcome which may be affected by UI benefit extensions.

3.3 Data

I use data from the 2014 edition of the *Muestra Continua de Vidas Laborales*, henceforth referred to as MCVL, a longitudinal sample of individuals assembled from Spanish Social Security records (Seguridad Social, 2014).⁷ Individuals are selected based on a 4% random sample of individuals affiliated to the Spanish Social Security Administration (SSSA) in a given year, which typically amounts to around 1.1 million people. Affiliated individuals include the employed, the self-employed, unemployed individuals receiving UI benefits, retirees receiving retirement benefits, and other non-employed individuals receiving benefits or pensions from the SSSA (e.g. those on disability insurance).

Complete work histories and unemployment spells can be observed for all individuals in the sample, ranging back to the mid-1960s. However, in the case of most pensions and benefits, information on these is only available starting in the 1990s. The data also include employer and employee identifiers, employer location, industry code, type of contract, a broad measure of occupation, and the portion of income that is subject to payroll tax. Starting in 2004, there is some tax return information available for the working, non self-employed individuals in the sample. Aside from these employment variables, the data include a range of demographic variables such as age, gender, country of birth, province of birth, and a noisy measure of educational attainment.

The first edition of the MCVL is from 2004, the year when the first sample is drawn. All individuals included in the 2004 edition remain in subsequent samples as long as they are still affiliated to the SSSA during the year of reference. To ensure that the sample remains

⁷I use a version of this database that includes certain tax information.

representative over time, it is updated each year to account for new entrants into the labor force and attrition in the original sample (e.g. due to emigration or death). For the purposes of this study, I focus on the 2014 edition, which allows me to study work histories up until that year.⁸

Since the goal of this paper is to study UI extensions, I start by selecting the sample of workers in the MCVL who at one point claim UI benefits. Further, I select individuals who claim UI benefits for the first time during 1994 or later, the year the UI duration policy described in figure 3.1 came into effect. Finally, since the rules for accruing days of affiliation with the SSSA are distinct for full-time workers and part-time workers, I restrict the sample to full-time workers.⁹ These restrictions result in a sample of 44,754 individuals.

Summary statistics for these individuals can be found in table 3.1. The average worker in the sample has been working for 2.8 years before claiming UI benefits, spends 170 days on UI, and 296 days in nonemployment before returning to work. Since these are workers claiming UI benefits for the first time, it is a young sample with an average age of 29. Further, there are more men than women in the sample (35% female), and about 20% are foreign-born workers. UI spells begin between 1994 and 2014, with more spells occurring in the second half of the period (the mean is 2006), likely due to the financial crisis and subsequent recession starting in 2008. Finally, 65% of workers were on a temporary contract before claiming UI benefits.

Table 3.1 also contains information on the “worker group.” This categorization is used by the SSSA to categorize workers into groups depending on their level of skill and the difficulty of the tasks carried out on the job. These groups are used by the SSSA to determine what part of a worker’s income is subject to payroll tax. There are a total of 11 groups in my sample, depicted in table 3.5. I translate the original Spanish names into English and consolidate them into five broad categories, largely following Miller and Llull (2016): skilled workers, staff, technical workers, laborers, and underage workers. The distribution of workers across these categories is described in table 3.1. In my sample, 9% of workers fall into the skilled category, 24% fall into the staff

⁸It is possible to request access to newer editions.

⁹The information contained in the MCVL only allows construction of days of affiliation for full-time workers.

Table 3.1: Summary statistics

Variable	N	Mean	Std. Dev.	Min.	Max.
Attachment	44754	1034.271	587.532	360	2159
Days on UI	44754	170.571	148.059	1	670
Days in Nonemployment	44754	296.201	340.563	2	2000
Year of UI	44754	2006.093	3.831	1994	2014
Temporary Contract	44754	0.680	0.467	0	1
Occup. Skilled	44754	0.089	0.284	0	1
Occup. Staff	44754	0.241	0.427	0	1
Occup. Technical	44754	0.4	0.49	0	1
Occup. Laborer	44754	0.262	0.44	0	1
Occup Underage	44754	0.009	0.094	0	1
Num. Work Spells	44754	4.594	6.559	1	323
Num. Employers	44754	2.73	2.165	1	38
Female	44754	0.346	0.476	0	1
Foreign	44754	0.208	0.406	0	1
Age	44754	29.756	8.429	18	70

category, 40% in the technical category, 26% are laborers, and 1% are underage.

3.4 Empirical Strategy

3.4.1 Regression Discontinuity

I study the effects of extending the duration of UI benefits on days spent on unemployment insurance and on days spent in nonemployment using a Regression Discontinuity design (RD). This approach makes use of the discontinuities generated by the step function in figure 3.1 by comparing the outcomes of individuals close to each of the cutoffs but on opposite sides. If the individuals on either side of the cutoffs are comparable, this approach makes it possible to estimate a causal effect of extended UI duration on the outcomes of interest.¹⁰

For each of the nine discontinuities, I estimate models of the following form:

$$y_{ia} = \alpha + \tau D_{iac} + f_{a \geq c}(a - c) + f_{a < c}(a - c) + \delta X_i + \varepsilon_{ia} \quad (3.1)$$

where y_{ia} is an outcome of interest for individual i with labor force attachment a , i.e. days of affiliation to the SSSA over the past six years. $D_{iac} \equiv 1\{a \geq c\}$ is an indicator for whether the individual's labor force attachment is to the right of the cutoff, and $f_{a \geq c}(a - c)$ and $f_{a < c}(a - c)$ are unknown functions of the worker's labor force attachment relative to the cutoff, allowing for a different functional relationship on either side of the cutoff. Further, X_i is a vector containing individual-specific control variables, such as gender, year of birth, province in which the worker became affiliated to the SSSA, type of affiliation relationship to the SSSA, occupational group, foreign status, year in which the unemployment spell started, age at which the unemployment spell occurred, and whether the worker was on a permanent or temporary contract. Finally, ε_{ia} is an individual-specific error term.

The coefficient of interest in the above model is τ . If the RD identification assumptions are met, τ is the causal effect of extending UI duration by 60 days on an outcome y , for workers

¹⁰I provide a detailed discussion of the assumptions necessary for the validity of this design later on.

near the cutoff c .

Even though I estimate different variations of the above model to assess robustness, the main model will be a linear one. This model assumes separate linear functional relationships on either side of each cutoff and takes the following form:

$$y_{ia} = \alpha + \tau D_{iac} + \beta_1(a - c) + \beta_2 D_{iac} \times (a - c) + \delta X_i + \varepsilon_{ia} \quad (3.2)$$

where the terms are defined as above and τ is still the coefficient of interest.

Even though the models above include individual covariates, I will also estimate models separately without covariates, to assess whether the inclusion of covariates significantly changes the estimate of τ . Following Johannes F. Schmieider et al. (2012), I cluster standard errors at the day (attachment) level. Finally, models will be estimated using triangular weights in order to prioritize observations that are closest to the cutoff.

3.4.2 Validity of the Regression Discontinuity Design

A crucial step in any RD design is to determine whether it is valid, i.e. whether the RD identification assumptions hold. In this section, I discuss various identification assumptions in the context of the Spanish UI system and conduct tests to determine whether they hold.

An important assumption of RD is that the individuals on either side of each cutoff are comparable or similar. If, for some reason, workers to one side of the cutoff have different characteristics than those on the other side, the RD estimate may pick up on the influence of these characteristics and fail to adequately capture the causal effect of interest. One way I will address this issue is by the inclusion of covariates in the models, as reflected in equations (3.1) and (3.2).¹¹ While this will ensure that discontinuities in the included observable covariates are not driving the ultimate results, it may still be possible for other characteristics not captured by

¹¹I also estimate models with and without covariates. The robustness of estimates to adding covariates is a way to determine whether characteristics varying at the cutoffs is problematic for the design or not.

these covariates to vary at the cutoff. For example, since UI duration is determined based on work histories, if certain *types* of workers with distinct work histories happen to be concentrated near cutoffs but only on one side, then this could be challenging for the RD design.

To determine whether this is an issue, I start by plotting the distribution of individuals in the sample by labor force attachment in figure 3.8. The distribution has a number of density points throughout, with the tallest one located around 365 days. I argue that these density points, or *spikes*, are a result of the characteristics of the labor market and indeed reflect the concentration of certain types of workers with distinct work histories at different points of the attachment distribution. For example, the spike at 365 can be explained by the fact that many workers who become eligible for UI have been on a one-year contract which has expired. Likewise, the rest of the spikes are located at the 1.5-year mark, the 2-year mark, the 2.5-year mark, and so on.¹² In order to provide a closer look at the distribution, I plot only workers to the right of the 365-day density point in figure 3.9. This figure reveals the underlying jumpiness of the distribution and further emphasizes the density points just discussed.¹³

Having established that there is a concentration of workers with distinct work histories at different points of the attachment distribution, we would like to determine whether these workers at spikes are of a different *type* than the rest.¹⁴ Workers may be considered of a different type if they are fundamentally different in terms of predetermined variables and outcomes than the surrounding workers. To determine this, I perform three different exercises, which are detailed

¹²This can also be deduced from the distance of the spikes relative to the cutoff as days of affiliation increase. To see this, note that each two cutoffs amounts to 360 days on the horizontal axis, while a one year contract will last for 365 days, or five more days. A two year contract will last for 730 days, or ten more than 720, which is where the cutoff is located. This pattern drives the spikes progressively to the right of the cutoff, which is visible in the figure.

¹³It is important to point out that the density points do not necessarily take up exactly one day on the distribution, which is apparent from figure 3.9. This arises because some contracts lasting for discrete years may occur during a leap year, or because contracts that do not last a discrete number of years may have different effective lengths depending on the start and end months. While this does not necessarily mean that the worker will get paid for an extra day, it does translate into extra days of affiliation to the SSSA, which is the running variable of interest. For example, if a one-month contract starts halfway through a month with 31 days lasting until the same day next month, then a worker on this contract will have been affiliated to the SSSA for one extra day compared to the same contract starting during a month with 30 days only. As a result, I consider the days surrounding the main density points to be part of the spikes as well.

¹⁴I consider a worker to be at a spike if they have a work history corresponding to discrete years or half-years.

in the appendix. First, in table 3.15 I regress worker characteristics on an indicator for whether that worker is at a spike. I show that workers at spikes are on average 10% more likely to be female, 10% more likely to be on a temporary contract before claiming UI benefits, have on average two fewer work spells in their work history, almost one fewer employer, and are 40% more likely to have been continuously employed before claiming UI benefits. Second, in table 3.16 I regress worker outcomes on the spike indicator. I show that while there does not seem to be a difference between these workers and the rest in terms of days spent on UI, workers at spikes do seem to take 34 days longer, on average, to find jobs after an unemployment spell. Third, I employ a random forest classifier to predict whether a worker is at a spike, using worker characteristics.¹⁵ At certain classification threshold values, I show that around 60% of workers at spikes are correctly identified as such by the classifier, while around 95% of workers not at spikes are correctly identified as such.

Given the results of the above exercises, I conclude that the workers located at these spikes are indeed of a different type than the rest of the population, and that these workers can be reasonably well identified based on observable characteristics. For these reasons, I move forward with the analysis by selecting only workers not located at discrete year and half-year regions of the attachment distribution. By doing this, I avoid pooling both types of workers together in the analysis, which could lead to biased estimates.¹⁶

Another important assumption of RD is that there is no perfect manipulation in this setting. In other words, it should not be possible for workers to effectively decide which side of the cutoff they are on. There are a number of reasons why I expect manipulation to be difficult in this setting. First, in order to become eligible for UI benefits, the firm must make the decision to terminate or lay off the worker, or at least refuse to renew their temporary contract. This means that a worker cannot directly *choose* exactly how many days to work, at least in the neighborhood of a cutoff

¹⁵This method flexibly models interactions between predictor variables and makes use of those that are most predictive, doing so in a recursive manner. For more details, see Breiman (2001).

¹⁶This small additional sample selection brings the sample size down from 44,754 to 40,947 observations.

value. Second, even though a worker may request a contract extension to reach a different benefit duration region, the firm would be unlikely to accept because they would have to continue paying the worker’s salary and Social Security payroll contributions.¹⁷ Third, the workers in the sample are all first-time claimants, so they may not yet be savvy about how benefit duration is determined. Even if they are aware that working 360 or more days makes them eligible for UI benefits, they might not know the specific steps that determine benefit duration beyond the initial eligibility cutoff, at least until they actually claim benefits.

In order to determine whether manipulation is an issue, one approach is to examine the density of workers near the cutoffs of interest and look for any discontinuities or irregularities. Given that higher UI duration amounts are desirable, if manipulation is a problem we would expect increased density to the right of each cutoff. To examine whether this is the case, I plot the final distribution of individuals by labor force attachment, after having carried out the sample selection described earlier in this section. The resulting density plot can be found in figure 3.2, with vertical dotted lines denoting each cutoff point.

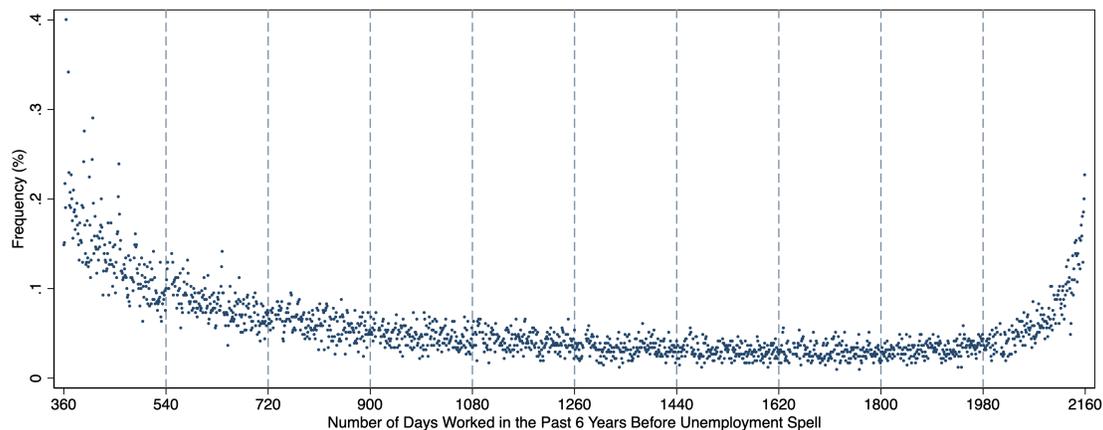


Figure 3.2: Histogram of days worked over the past six years before an unemployment spell

¹⁷It is worth noting that Y. Rebollo-Sanz (2012) estimate that the layoff hazard does increase when a worker becomes eligible for UI benefits, at 360 days of affiliation. This cutoff is not included in my analysis given that there are no workers to the left of 360 receiving benefits, since they are not eligible to begin with. They also depict the layoff hazard up to 16 months after the worker becomes eligible for UI and it is quite stable in that entire region, indicating that firms do not consider the particular steps of the potential UI duration function when making layoff decisions.

When inspecting figure 3.2, there does not appear to be any evidence of manipulation, since the density of the distribution seems to be stable around the cutoff points. However, in order to test this assumption more formally, I follow the literature and perform a series of McCrary tests (McCrary, 2008), which estimate whether there is a significant discontinuity in the density at the cutoff points. The results of these tests are in figure 3.10 and table 3.6. Figure 3.10 provides a graphical depiction of the distribution near each cutoff. With the exception of perhaps the 540-day cutoff, there does not seem to be any evidence of manipulation. In fact, some cutoffs appear to have lower density to the right of the cutoff, even though manipulation would imply the opposite. The numerical estimates of the difference in density at each cutoff provided by the McCrary tests can be found in table 3.6. As reflected in the graphs, there is no systematic pattern indicating manipulation at the cutoffs. The one exception is the 540-day cutoff; at this cutoff, the density jump is significant at the 10% level and approaches stronger levels of statistical significance. However, it is worth noting that this result may be due to the inherent jumpiness of the distribution. If the distribution is too irregular, the McCrary test may be inappropriate, since a core assumption of this test is that the underlying distribution can be described by a continuously differentiable function. Therefore, to determine whether the estimated discontinuities are rare given the nature of the distribution, I perform a Monte Carlo analysis by randomly picking points on the distribution and performing McCrary tests at these points. The distribution of the estimated discontinuities is depicted in figure 3.11. If we compare these results with those in table 3.6, it does not appear that the values in the table are rare, regardless of the sign. In fact, the distribution appears to be centered around 0.1, which is close to the estimated discontinuity at the 540-day cutoff. Therefore, considering the appearance of the histogram, and the McCrary test results in conjunction with the Monte Carlo exercise, I conclude that manipulation is not likely to be a problem in this setting.

Another threat to the validity of the RD design that is worth discussing is that of selection. Selection may be an issue if the increased duration of UI benefits at the cutoff induces a different

set of workers to apply for UI benefits. If this is the case, then workers on opposite sides of the cutoff may not be comparable. As discussed above, this would be problematic for the RD assumptions. However, I argue that this is unlikely to be a concern in this setting for various reasons. First, notice that all workers in the sample are eligible for at least 120 days of UI, and in many cases they are eligible for significantly longer periods. If the increased UI duration is leading different kinds of workers to apply for benefits, it would have to be a set of workers that would not have applied to UI benefits if the offer was for 120 days, but would apply if the offer is for 180 days. Given that 120 days is already a considerable amount of time, it seems unlikely that the marginal increase from 120 to 180 would change people's minds about whether to apply for benefits. Further, this argument is even stronger in the case of higher baseline duration amounts for later cutoffs. Second, note that the density tests described above are also informative of the selection issue. In the same way that manipulation would imply increased density to the right of each cutoff, extended UI duration would have the same effect, since the offer is now strictly more attractive. Therefore, the above density tests also embed and address the selection issue.¹⁸

3.5 Results

3.5.1 Short Term Effects

In this section I study how receiving 60 additional days of UI affects two outcomes: days spent on UI and days spent in nonemployment. Days spent on UI are the total number of days spent receiving UI benefits before an unemployed worker starts a new job. Days in nonemployment is defined as the duration in days between the start of the unemployment spell

¹⁸Note that by adding workers who do not apply for UI to the analysis, it would be possible to carry out two separate tests, one for manipulation and one for selection. These tests, like the proposed RD specifications, would seek to estimate discontinuities at the cutoffs in a similar fashion. The outcomes of interest would instead be the probability of a worker losing their job (for the manipulation test) and the probability of a worker applying for benefits (for the selection test).

and the date a worker starts working at a new job, regardless of the number of days spent on UI.

In order to do this, I estimate regression discontinuity models of the type depicted in equation (3.2). It is important to point out that I do not expect the effect to be the same at each of the cutoffs. The main reason for this is that there are two important factors that vary across cutoffs and along the attachment distribution. The first of these factors is the labor force attachment of the workers at each of the cutoffs. Workers at the earlier cutoffs will have lower labor force attachment than workers at the later cutoffs, since they will have been employed fewer days in the past six years prior to the unemployment spell. If workers with different levels of labor force attachment have different responses to extended UI benefit duration, the estimate at each of the cutoffs is likely to be different as a result. *A priori*, the difference in the treatment effect between high and low labor force attachment workers, for a fixed UI benefit duration, is ambiguous. On one hand, higher attachment workers might have more savings than low attachment workers, which allows them to better sustain consumption at pre-unemployment levels.¹⁹ As a result, a UI extension may reduce search effort relatively more for high attachment workers, leading to a stronger treatment effect for them. On the other hand, if the labor market is tighter for high attachment workers, perhaps because they possess skills in high demand, then the effect of extended UI duration might be weaker for them. The second of these factors is the baseline level of UI benefit duration at each cutoff. To the left of the first cutoff, workers are eligible for 120 days of UI, whereas to the right they become eligible for 180 days. This translates into an increase of 60 days over a baseline of 120, or an increase of 50%. In the case of the last cutoff, workers to the right of the cutoff receive 60 more days of UI over a baseline of 600, or a 10% increase. Since the baseline duration is increasing at later cutoffs, and the percentage increase in UI duration is lower as a result, I would expect a weaker response at later cutoffs compared to earlier ones, holding the type of worker constant. As a result of these two factors, the observed magnitude of treatment effects at each of the cutoffs will be a combination of the varying marginal increase in

¹⁹Recall that UI benefits do not replace 100% of income.

UI duration and differences in labor force attachment. While keeping both factors in mind, for exposition I refer to the effects at the earlier cutoffs as the effects on low labor force attachment workers, and I refer to the effects at the later cutoffs as the effects on high labor force attachment workers.

I start by presenting graphical results of the effects of receiving 60 more days of UI on total days spent on UI in figure 3.3. There are nine graphs, one for each discontinuity, with separate weighted linear trends on each side of the cutoff. There is a visible positive jump in days spent on UI at the first four cutoffs, whereas the effect is still positive but less visible at the later cutoffs. In order to properly assess the magnitude of these effects, I present them in regression format in table 3.2, starting with the linear unweighted model in the first column. In the case of the 540-day cutoff, a 60-day extension of UI benefits increases the number of days spent on UI by 23 days, or about a third of the extension, and the difference is strongly statistically significant. The magnitude and significance level is similar at the 720-, 900-, and 1080-day cutoffs. The point estimates at later cutoffs, though still positive across the board, are in general not statistically significant and lower in magnitude than the effects at the earlier cutoffs. In the second column I add control variables to the regression, which allows me to assess whether any covariates jump discontinuously at the cutoff, a possible threat to identification. For the first four cutoffs, the pattern is again similar, and the discontinuity estimates barely change, indicating that covariates are smooth at the cutoff and that the effect is due to the UI extension. For the last five cutoffs, the estimate does not change much either, with the exception of the 1800-day cutoff, for which the estimate drops significantly. However, it is important to note that the sample size is significantly smaller at later cutoffs, which could explain some of the variability in the point estimates after adding covariates. Finally, in column three I add triangular weights in order to prioritize observations closer to the cutoff, which are the most relevant for identification purposes. This is my preferred specification. Adding weights does not significantly change the estimates at earlier cutoffs but does seem to change the point estimate at some of the later cutoffs, although

these remain insignificant. I plot the estimates of the third column of table 3.2 in figure 3.12, which reveals a decreasing overall trend, indicating that the effect seems to be lower at later cutoffs.

In order to simplify the analysis and draw more statistically precise conclusions, I pool some discontinuities together. In particular, I pool the first four discontinuities together and take these to be representative of low labor force attachment workers. Similarly, I pool the last five discontinuities together and take them to be representative of high labor force attachment workers. The regression results are in the first two columns of table 3.7 and a graphical depiction of the point estimates is in figure 3.13. According to table 3.7, an additional 60 days of UI leads low labor force attachment workers to stay on UI for an additional 21 days, whereas the same increase leads high labor force attachment workers to stay on UI for an additional 8 days, although the estimate is only marginally significant. The estimates depicted in figure 3.13 tell the same story. Finally, it should be noted that these differences may not be due to the characteristics of the workers alone, since the baseline level of UI is also higher for high labor force attachment workers, as discussed above.

I now move on to analyze the effect of receiving an additional 60 days of UI on the duration of a worker's nonemployment spell. As above, I present graphs for each discontinuity with separate weighted linear trends on each side of the cutoff, this time in figure 3.4. At the 540-day cutoff, there is a visible increase in the days spent in nonemployment as a result of increased UI duration. Even though this effect is not visible at the 720-day cutoff, there seems to be a discernible positive jump at the 900-, 1080-, and 1260-day cutoffs.²⁰ For the remaining cutoffs, there exists no visible jump, with the exception of the cutoff at 1800 days. I present the corresponding regression results in table 3.3, where the first column contains the results of a linear unweighted model with no controls, the second column adds controls, and the third column includes triangular weights. As before, adding controls does not seem to change the

²⁰The jump at the 900-day cutoff may be a result of the choice of linear trend.

coefficients significantly, indicating that covariates are smooth across the cutoffs. On the other hand, adding triangular weights does seem to influence the estimates in this case. For this reason, I interpret the results in column three of table 3.3 only. The estimate at the 540-day cutoff suggests that receiving 60 additional days of UI extends nonemployment duration by 23 days, and the estimate is statistically significant. At the 720-day cutoff, there appears to be no effect once the observations around the cutoff are prioritized. For the remaining cutoffs, the effect is not statistically different from zero, although for the cutoffs at 900, 1080, 1260, 1800, and 1980 days, the estimates are economically meaningful when compared to the estimate at the first cutoff. For exposition, these point estimates are plotted in figure 3.14.

As above, I next pool the the early discontinuities and the later discontinuities to draw conclusions about workers with low and high labor force attachment. The regression results are in the third and fourth columns of table 3.7 and a graphical depiction of the point estimates is in figure 3.15. The third column of table 3.7 reports the estimate for the low labor force attachment sample, located at the earlier cutoffs. This seems to indicate that there is a significant increase of 14 days of nonemployment as a result of receiving an additional 60 days of UI. For the high labor force attachment sample at the later cutoffs, the effect is similar in magnitude, although it is not statistically distinguishable from zero. Figure 3.15 tells the same story.

3.5.2 Long Term Effects

In this section, I move on to analyze the long term effects of extended UI duration on my two outcomes of interest. I focus on the same outcomes as in the previous section, but instead measure them over the next five years following the beginning of the unemployment spell. More specifically, I add up total days spent on UI following the start of the unemployment spell up to five years to recalculate the *days on UI* variable. Similarly, I add up the total number of days an individual is not working over the next five years to recalculate the *days in nonemployment* variable. I also compare the estimates in this section to their short term counterparts and examine

whether the short term effects are different from the long term effects. When making these comparisons, the short term effects I present will be slightly different from those in the previous section, since I can only include individuals in the sample for whom I observe at least five years of data following the start of their unemployment spell.²¹

I start by graphically examining the long term effect of extended UI duration on days spent on UI in figure 3.16. The graphs are visibly similar to those in figure 3.3, which examined short term effects.²² As in the previous section, the effect is more visible and consistent at earlier cutoffs, with some of the later cutoffs showing no effect. I include the regression results for these discontinuity estimates in the third column of table 3.8, as well as the short term estimate in this sample in the second column, for comparison. For the first cutoff, the long term estimate is 35 days whereas the short term estimate is 22, and this difference of 13 days is statistically significant, as reflected in the first column. This means that for the workers around that cutoff, extending UI by 60 days leads to an increase in use of UI by 21 days in the short term, and 13 more days up until five years from the start of the unemployment spell. For the next three discontinuities, the long term estimate is higher than the short term estimate in two out of three cases, although the difference is not statistically significant. For the remainder of the cutoffs the point estimates are smaller and in general not statistically significant, with the exception of the point estimate at 1620 days, which is large, statistically significant, and much larger than the short term estimate. I present each of these point estimates graphically in figure 3.18, with short term point estimates in red and long term estimates and their confidence intervals in blue. Again, there is a clear decreasing pattern, with long term estimates usually higher than short term ones. The only exception to the trend is the cutoff at 1620, which I interpret with caution given its anomalous nature.

²¹The dataset contains work spells starting during April 2015 or earlier, so in order to study long term effects over a five year period, I must restrict the sample to those unemployment spells starting during April 2010 or earlier. Requesting more recent versions of the data would therefore help increase statistical precision, especially in this section.

²²Note that that the values on the vertical axes are now higher, since I am now counting days on UI over a five year period.

To make interpretation easier and increase statistical power, I again pool earlier and later cutoffs together into high and low labor force attachment groups. I present these results graphically in figure 3.5. In both cases, there is a positive and significant long term effect of extended UI duration, and this effect is larger than the short term estimate in both cases. Furthermore, the long term effect at the earlier cutoffs is higher than the effect for the later cutoffs. In the case of the high attachment sample, however, I note that the long term estimate may be driven by the anomalous estimate at the 1620-day cutoff. The regression version of these estimates can be found in the first two columns of both panels in table 3.4. These results indicate that for the low labor force attachment workers, extending UI by 60 days leads to an increase in use of UI by 21 days in the short term, and 7 more days up until five years from the start of the unemployment spell. For the high labor force attachment workers, the difference between the long term estimate and the short term estimate is about 12 days.

I now move on to analyzing the long term effects of extended UI duration on days spent in nonemployment over a five year period. Each discontinuity is depicted graphically in figure 3.17. In this case there appears to be no clear pattern, with some discontinuities showing positive jumps, others showing negative jumps, and the remainder showing no jump at all. However, none of these jumps is particularly visible and it seems like there is no real effect, indicating that the non zero jumps could be due to noise or model choice. The regression analogue of these graphs is presented in table 3.9, along with short term estimates. The observations from the figure are confirmed: there is no clear pattern. Importantly, none of the long term estimates are statistically significant, which seems to hint that there is no long term effect on nonemployment duration. I also report the short term effects for this sample in the second column of table 3.9. I report short term point estimates, long term point estimates, and long term confidence intervals for each discontinuity in figure 3.19. For the first six discontinuities, the short term estimate is above the long term estimate, whereas for the last three it is below.

Pooled regression results for high and low labor force attachment samples are reported

in the third and fourth columns of both panels in table 3.4. The graphical equivalent of these estimates is in figure 3.6. For the later discontinuities with workers with high labor force attachment, both the short term and long term estimates seem to be very close to zero, indicating that at these discontinuities there is no differential long term effect. For the earlier discontinuities capturing the low labor force attachment sample, the long term effect is again very close to zero, while the short term effect remains positive and significant, as in the previous section.

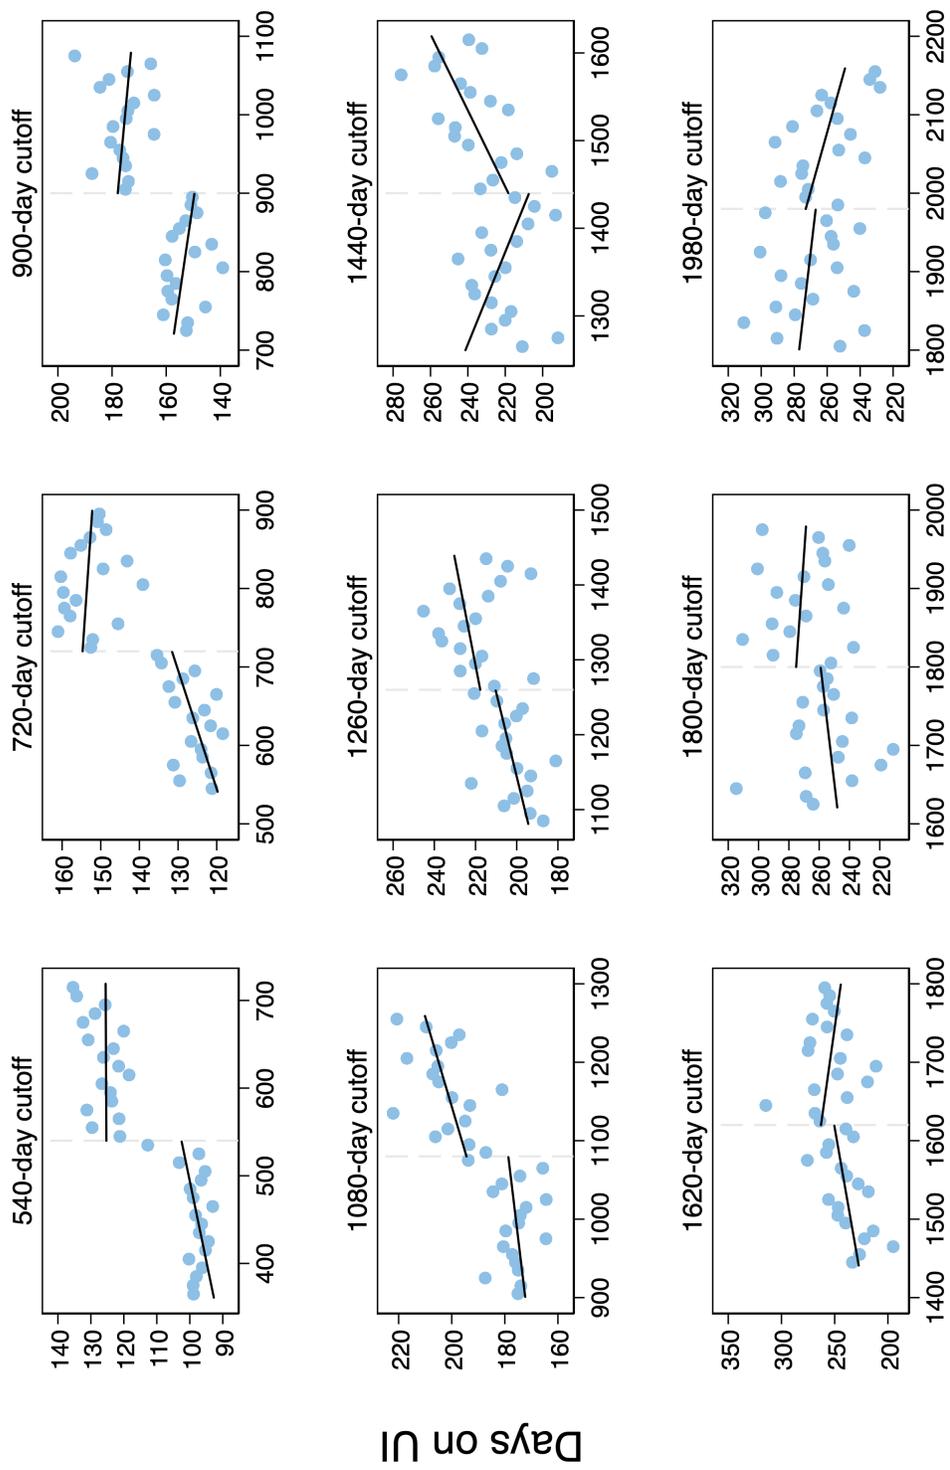
3.5.3 Robustness

In this section I assess the robustness of the above results to various different specifications. First, I implement a “donut” RD design by dropping observations immediately around the cutoffs and re-estimating the linear models with the weights from the previous sections. Second, I relax the linearity assumption and estimate RD models with separate quadratic trends at either side of the cutoffs. Third, I estimate the models via local linear regression, selecting bandwidths optimally according to Imbens and Kalyanaraman (2012).

The robustness of the short term effects on days spent on UI are in table 3.10. In the first column, I present the result from the previous section, using linear trends; in the second column, I present the donut RD; in the third column, I present the results of an RD with quadratic trends; finally, in the fourth column, I present the results of the local linear regression. For the first three cutoffs, the results are robust to all three alternative specifications. For the fourth cutoff, results are weaker for the quadratic and local linear models, indicating that results at this cutoff may be driven by the choice of linear specification. The same is true of the fifth cutoff, at which the sign of the estimate changes depending on the model of choice. For the remaining cutoffs, the choice of model does not in general significantly change the results, although the sample size is now smaller which will likely cause results to be more sensitive model choice. In table 3.11 I examine the robustness of the short term effects on days spent in nonemployment, with the same column structure described above. The results for the first cutoff at 540 vary somewhat in magnitude with

the choice of model but are consistently large, positive, and significant in every column. This is also true of the third cutoff, which has consistently large estimates across the board with two of them being significant. For the rest of the cutoffs, estimates appear to be sensitive to model choice, so I interpret them with caution.

In table 3.12 I assess the robustness of the long term effects of increased UI duration on days spent on UI. The first two cutoffs are quite robust to alternative specifications, while the following four seem to be sensitive to relaxing the linearity assumption. Out of the last three cutoffs, two seem stable across specifications while one is also sensitive to the linearity assumption. Finally, in table 3.13 I examine the robustness of the results for long term effects on days spent in nonemployment. The results in this table are in the majority of cases not significant and do change substantially depending on the choice of model. This is in line with the visual examination of figure 3.17, which does not show any clear jumps in any of the graphs.



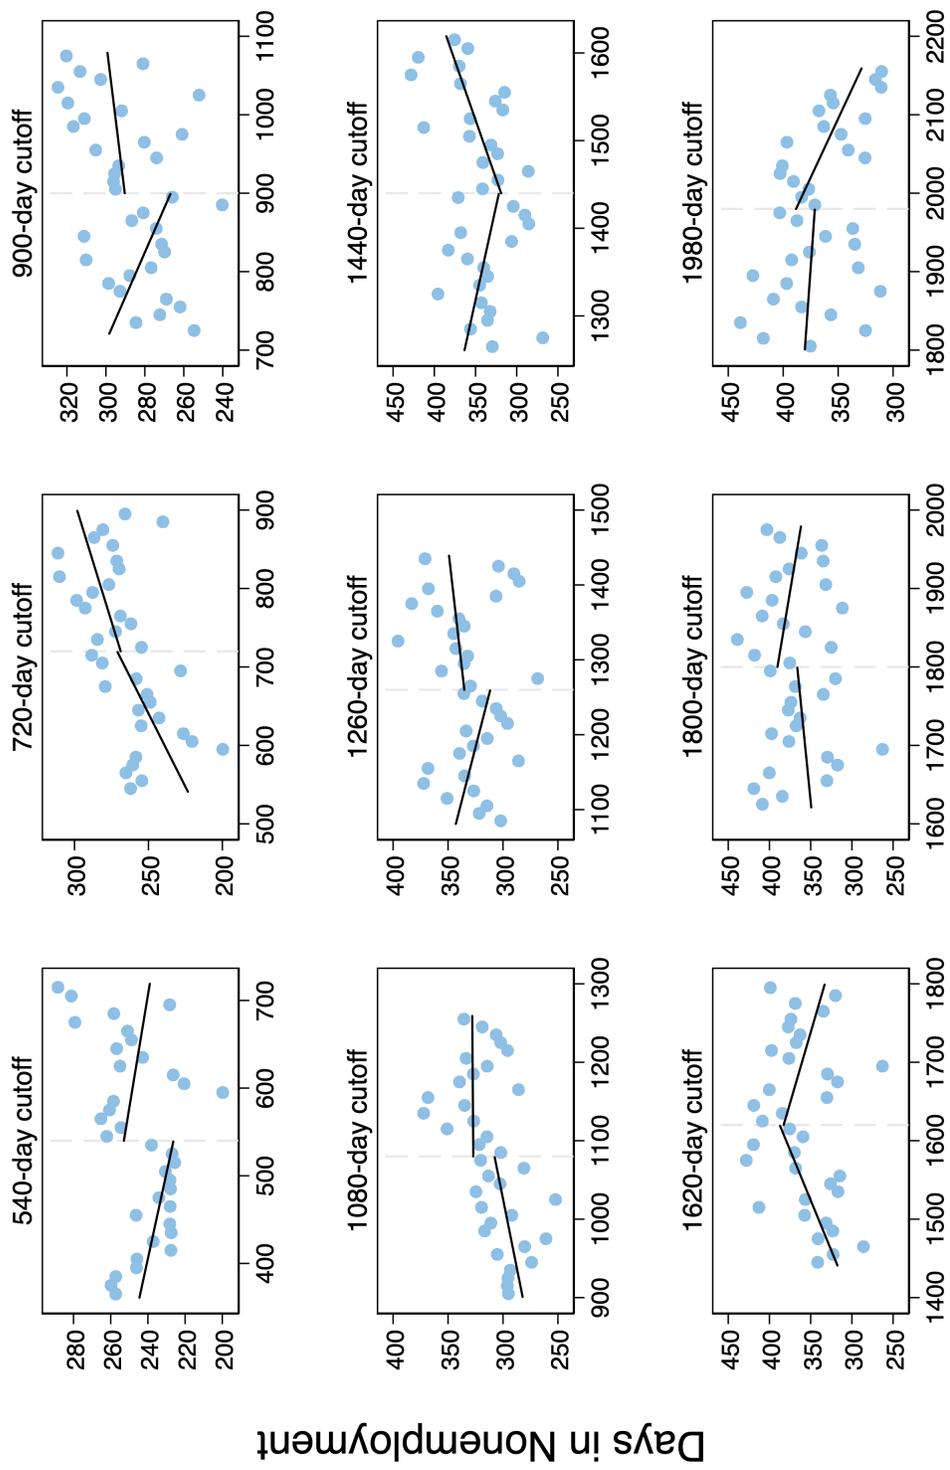
Attachment

Figure 3.3: The short term effect of extended UI duration on days spent on UI

Table 3.2: The short term effect of extended UI duration on days spent on UI

	(1)	(2)	(3)
D_{540}	23.26*** (2.186)	22.77*** (2.140)	22.52*** (2.281)
Obs. (R^2)	15947 (0.046)	15947 (0.100)	15947 (0.107)
D_{720}	25.98*** (3.079)	24.78*** (3.128)	22.09*** (3.245)
Obs. (R^2)	10291 (0.025)	10291 (0.097)	10291 (0.103)
D_{900}	25.81*** (4.601)	24.47*** (4.644)	27.46*** (4.944)
Obs. (R^2)	7730 (0.012)	7730 (0.091)	7730 (0.096)
D_{1080}	17.90*** (5.966)	17.72*** (5.903)	16.12** (6.434)
Obs. (R^2)	6363 (0.011)	6363 (0.102)	6363 (0.109)
D_{1260}	15.33* (8.572)	12.43 (8.219)	4.363 (8.761)
Obs. (R^2)	5426 (0.004)	5426 (0.098)	5426 (0.116)
D_{1440}	6.930 (10.03)	5.004 (9.698)	12.25 (10.03)
Obs. (R^2)	4622 (0.004)	4622 (0.111)	4622 (0.118)
D_{1620}	4.185 (11.64)	3.141 (11.16)	13.44 (11.64)
Obs. (R^2)	4338 (0.004)	4338 (0.139)	4338 (0.148)
D_{1800}	21.52* (12.41)	13.72 (12.12)	9.316 (13.53)
Obs. (R^2)	4499 (0.001)	4499 (0.138)	4499 (0.147)
D_{1980}	14.56 (10.85)	9.112 (10.84)	-0.785 (11.77)
Obs. (R^2)	7372 (0.00374)	7372 (0.105)	7372 (0.121)
Controls (weights)	No (No)	Yes (No)	Yes (Yes)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.



Attachment

Figure 3.4: The short term effect of extended UI duration on days in nonemployment

Table 3.3: The short term effect of extended UI duration on days in nonemployment

	(1)	(2)	(3)
D_{540}	22.14** (10.18)	18.28* (10.20)	23.43** (10.68)
Obs. (R^2)	15947 (0.001)	15947 (0.081)	15947 (0.086)
D_{720}	18.62 (12.47)	14.45 (12.20)	-5.841 (13.37)
Obs. (R^2)	10291 (0.002)	10291 (0.097)	10291 (0.104)
D_{900}	9.661 (13.39)	5.200 (13.43)	15.25 (13.83)
Obs. (R^2)	7730 (0.001)	7730 (0.092)	7730 (0.088)
D_{1080}	28.63* (16.17)	24.44 (15.58)	18.00 (16.80)
Obs. (R^2)	6363 (0.002)	6363 (0.101)	6363 (0.120)
D_{1260}	28.09 (18.15)	16.75 (17.99)	13.53 (19.98)
Obs. (R^2)	5426 (0.000)	5426 (0.101)	5426 (0.118)
D_{1440}	-12.05 (21.30)	-12.74 (20.31)	2.656 (21.34)
Obs. (R^2)	4622 (0.002)	4622 (0.114)	4622 (0.117)
D_{1620}	-14.20 (23.65)	-15.86 (22.76)	-1.951 (24.91)
Obs. (R^2)	4338 (0.002)	4338 (0.128)	4338 (0.138)
D_{1800}	30.01 (23.62)	21.03 (23.63)	14.83 (26.51)
Obs. (R^2)	4499 (0.001)	4499 (0.118)	4499 (0.127)
D_{1980}	29.03 (19.37)	30.25 (19.20)	17.70 (20.99)
Obs. (R^2)	7372 (0.00390)	7372 (0.0894)	7372 (0.101)
Controls (weights)	No (No)	Yes (No)	Yes (Yes)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.

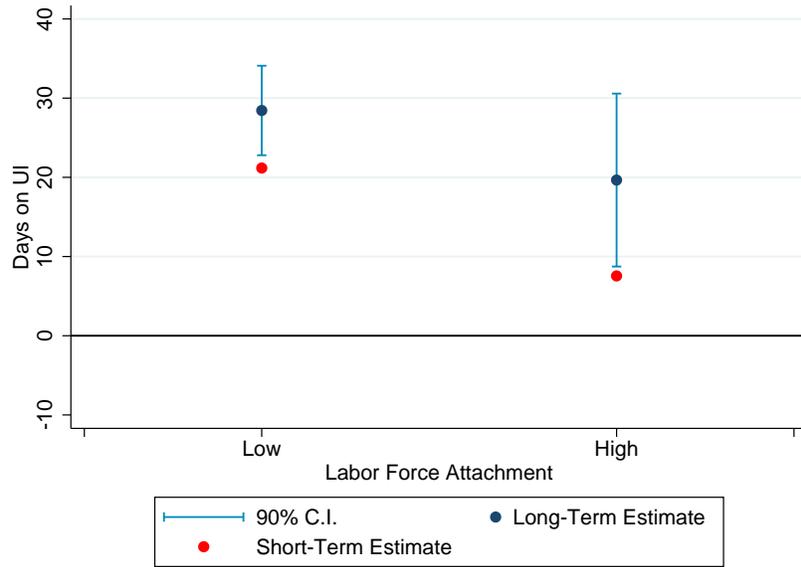


Figure 3.5: Estimates of short and long term effects on days on UI for both high and low labor force attachment samples

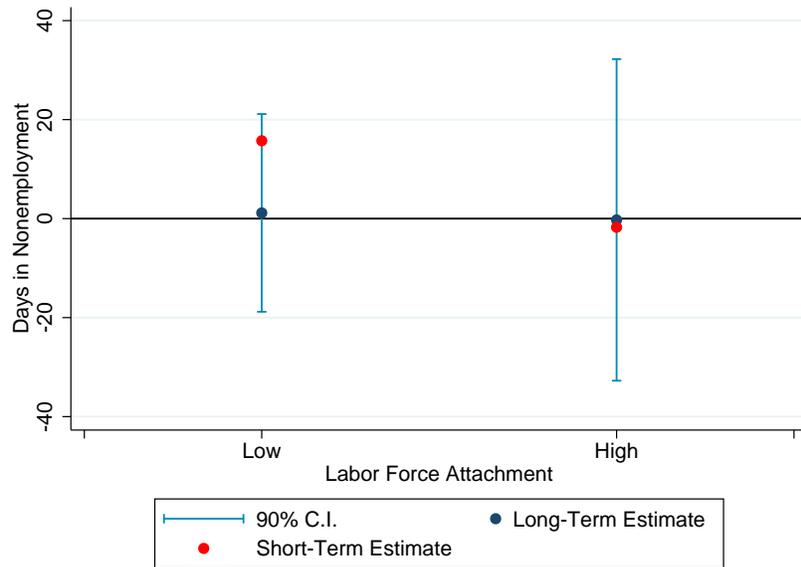


Figure 3.6: Estimates of short and long term effects on days in nonemployment for both high and low labor force attachment samples

Table 3.4: Short and long term effects of extended UI duration for high and low labor force attachment samples

	(1)	(2)	(3)	(4)
	Days UI (Short-term)	Days UI (Long-term)	Nonemployment (Short-term)	Nonemployment (Long-term)
<i>Panel A: low labor force attachment sample</i>				
D_{low}	21.18*** (2.619)	28.44*** (3.425)	15.73** (7.273)	1.154 (12.11)
Observations	32251	32251	32251	32251
R-squared	0.0982	0.0891	0.0891	0.101
<i>Panel B: high labor force attachment sample</i>				
D_{high}	7.557 (5.948)	19.66*** (6.621)	-1.741 (12.59)	-0.262 (19.70)
Observations	19259	19259	19259	19259
R-squared	0.104	0.135	0.0957	0.117
Controls	Yes	Yes	Yes	Yes
Weights	Yes	Yes	Yes	Yes
Short-term	Yes	No	Yes	No
Long-term	No	Yes	No	Yes

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.

3.6 Discussion

A noticeable feature of the results and robustness analyses in the previous section is that only the results at earlier cutoffs seem to be convincing. In the case of the later cutoffs, the number of observations is lower, the outcome variables are noisier, and the results are more sensitive to model specification. For this reason, I abstain from drawing strong inferences from the later cutoffs, and focus the discussion on the earlier ones instead. However, it is important to keep in mind that the lack of robustness at later cutoffs may simply be reflecting a null effect. This argument has theoretical foundation, as discussed earlier, given that the percentage increase in UI duration at these later cutoffs is much lower compared to that at earlier cutoffs. For example, at the last cutoff, the baseline duration of UI benefits is 600 days; if very few people utilize UI up to that amount (the average UI duration for workers in that region is around 270), then an increase of benefit duration up to 660 days is unlikely to have an effect on the average worker.

In terms of short term effects, when focusing on the low labor force attachment sample located at earlier cutoffs, the results show that a UI extension of 60 days results in workers spending on average 21 more days on UI in the short term, or about one third of the increase in benefit duration. For these same workers, there is also a significant increase of 14 days spent in nonemployment, although this effect is somewhat heterogeneous across different cutoffs. The strongest significant effect is observed at the first cutoff, at which extended UI benefits lead workers to spend 23 more days in nonemployment, on average. These results are in line with the literature on short term impacts of UI extensions, although they are on the larger end of the effects that have been previously identified.

In terms of long term effects, for the low labor force attachment sample located at earlier cutoffs, the results show that a UI extension of 60 days results in workers spending on average 28 more days on UI in the long term. The corresponding short term effect is 21 days, indicating that workers spend about seven more days on UI in the long term as a result of the extensions. In

terms of days spent in nonemployment, for the low labor force attachment sample the long term effect of UI extensions is close to zero, whereas the short term effect is positive and around 15 days.²³ Furthermore, when examining the effects at each individual cutoff, we can see that the short term estimate is almost always larger than the long term estimate for these earlier cutoffs (see figure 3.19). These results would seem to suggest that, heading into the long term, workers actually spend *fewer* days in nonemployment as a result of increased benefit duration, which brings the short term effect closer to zero over time. Importantly, the direction of these effects is similar to that in Johannes F. Schmieder et al. (2012). However, it is important to point out that this finding is only suggestive because (1) these long term estimates may not be robust to alternative functional specifications, and (2) these results arise from pooling several cutoffs, and there is noticeable underlying heterogeneity in the corresponding individual point estimates.²⁴

If workers are indeed employed for longer heading into the long term as a result of extended UI benefits, this means that their employment stability will have improved as a result. One way employment stability may improve is if the *match quality* of a worker's new job is higher, which is difficult to measure. Another way employment stability may improve, which is especially relevant in the case of the Spanish labor market, is if the worker transitions to a permanent contract. In fact, Y. F. Rebollo-Sanz and García-Pérez (2015) suggest that UI benefits may encourage transition into a permanent contract for temporary workers. I explore this possible mechanism in table 3.14, while focusing on the third column which includes both control variables and triangular weights. The estimated effects are economically small, in all but one case less than 1% in absolute value, with most quite close to zero. Furthermore, the estimates are statistically insignificant across the board, ruling out large effects in either direction. Therefore, in this analysis I do not find evidence that transition to a permanent contract is a driving force behind the

²³However, the long term effect is imprecisely estimated.

²⁴Even though there is noticeable underlying heterogeneity in the corresponding individual point estimates, as mentioned above the short term estimate is almost always larger than the long term estimate for these earlier cutoffs. Therefore, even if the long term effect is different from zero, a consistent finding seems to be that it is lower than the short term effect. This is interesting on its own and also points to the possibility that, heading into the long term, workers experiencing extended UI duration actually spend fewer days in nonemployment.

discrepancy between the short term and long term estimates.

3.7 Conclusion

In this paper, I contribute to the literatures on both short and long term effects of extended unemployment insurance duration on employment outcomes. I focus on a set of Spanish workers who become unemployed for the first time, which is a particularly young sample. My results on short term effects are in line with previous studies, showing that extended UI duration can affect both days spent on UI and days spent in nonemployment. These effects are robust to alternative specifications. The long term results are less robust, but seem to suggest that workers spend more days on UI in the long term as a result of increased UI duration, and that they possibly spend fewer days in nonemployment.

One caveat of this study that is worth pointing out is that the results may not externalize to the entire Spanish labor force. First, even though the purpose of this study is to study younger workers, the effects may be different for older workers. Second, only workers who claim UI benefits are included in the analysis, so the results may not generalize to individuals who never claim benefits. Third, the RD approach retrieves a treatment effect for individuals around a cutoff, who may not be similar to individuals farther away from the cutoff with different labor force attachment histories.

The results in this paper emphasize the importance of considering long term as well as short term effects of unemployment benefits. More broadly, they emphasize the importance of conducting long term analyses of public programs, as the policy implications may be different from those drawn from short run analyses. Future work on this topic could involve focusing on an even longer time horizon, following young workers affected by these policies over their life-cycle. Furthermore, more research is required to fully understand the effect of UI on individuals' post unemployment work histories, as well as its effect on alternative outcomes which have received

less attention, such as migration.

3.8 Acknowledgement

Chapter 3 is currently being prepared for submission for publication of the material. Ruiz Junco, Pablo; “Short and Long Term Effects of Extended Unemployment Insurance for First Time Claimants.” The dissertation author was the sole investigator and author of this material.

3.9 Appendix

3.9.1 Additional Figures



Figure 3.7: Unemployment in Spain 1990-2015

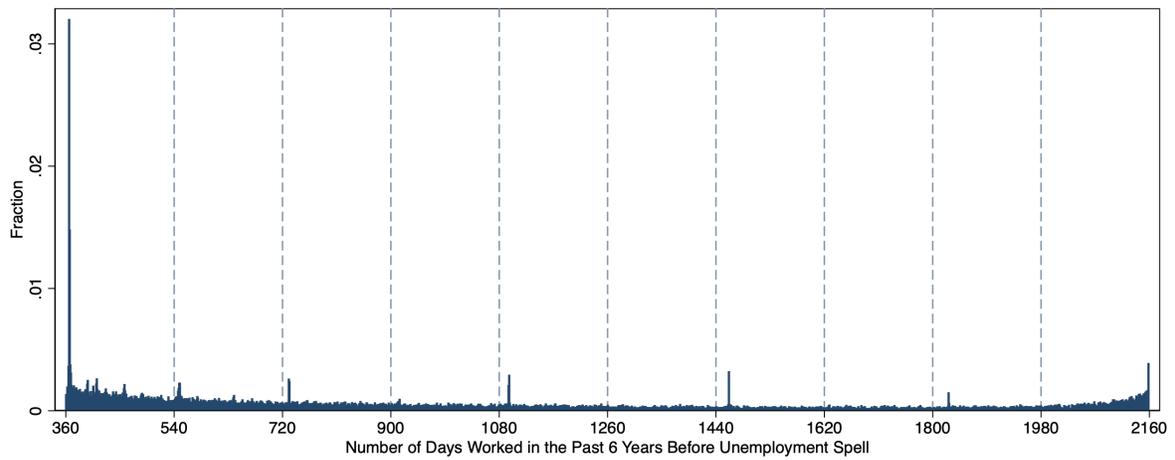


Figure 3.8: Full histogram

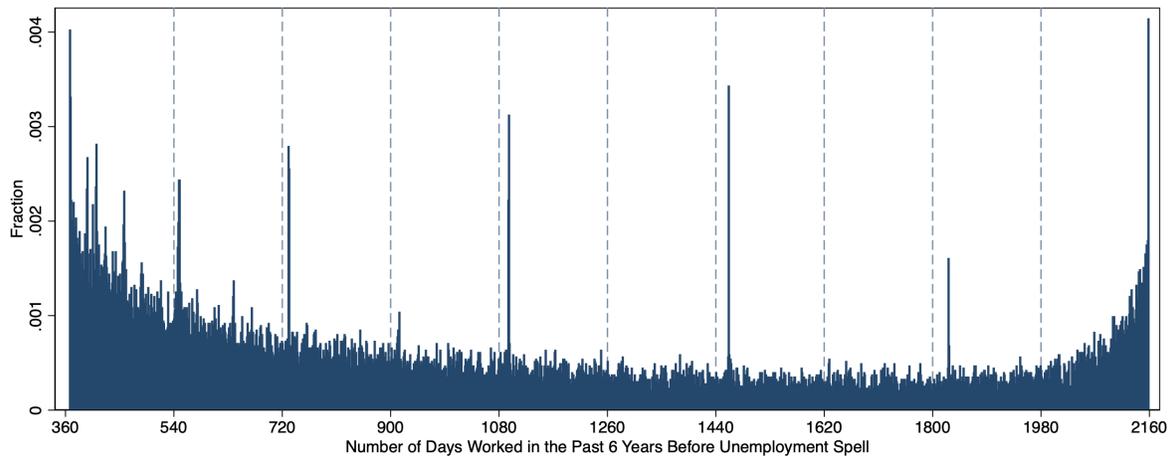


Figure 3.9: Truncated histogram

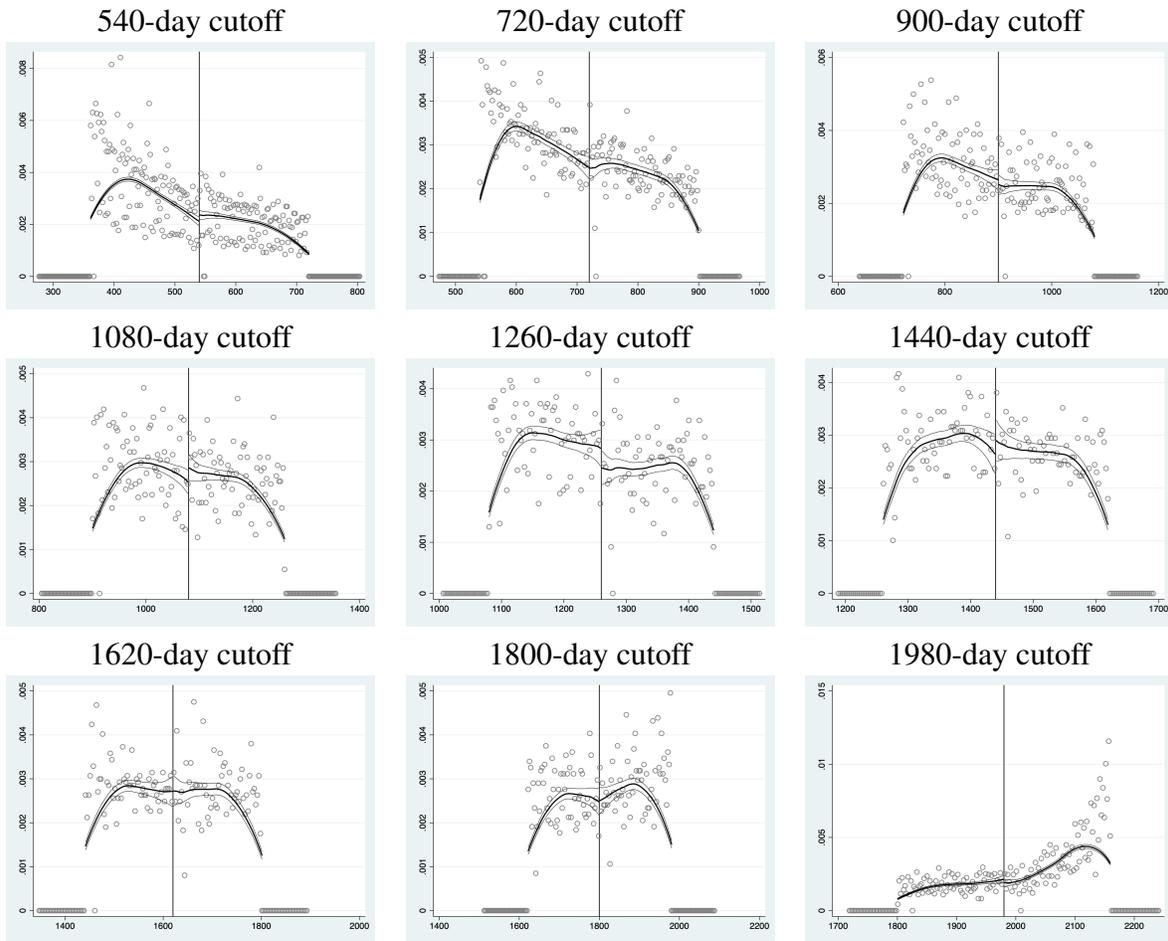


Figure 3.10: Graphical results of McCrary test at different cutoffs

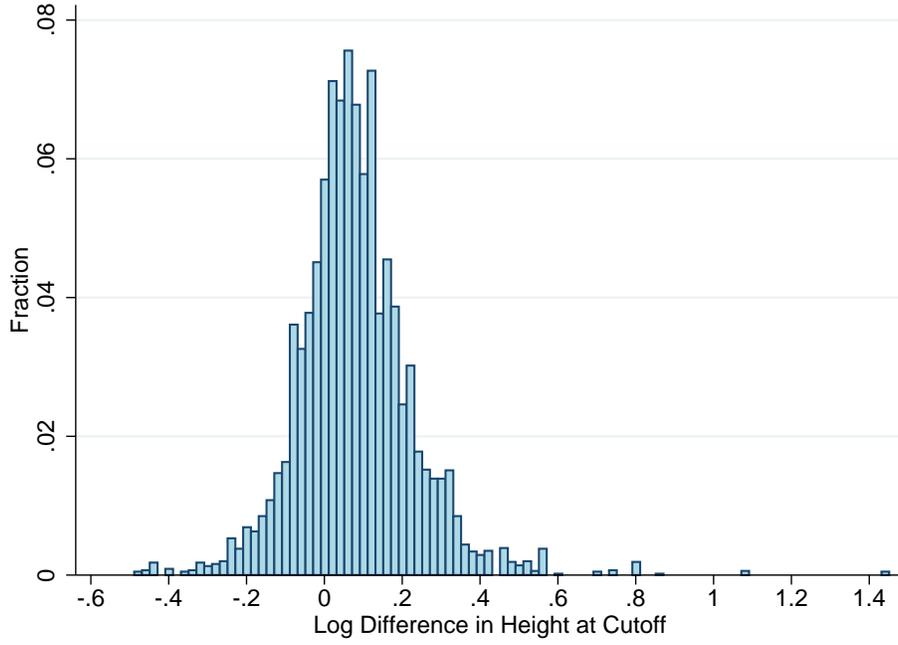


Figure 3.11: Monte Carlo analysis of McCrary test results

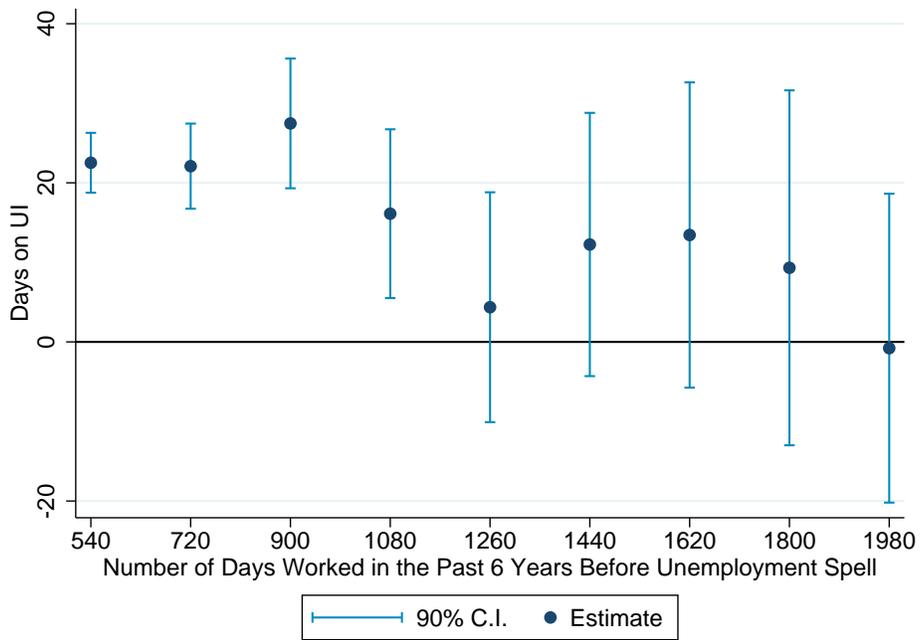


Figure 3.12: Estimates of short term effects on days on UI at different cutoffs

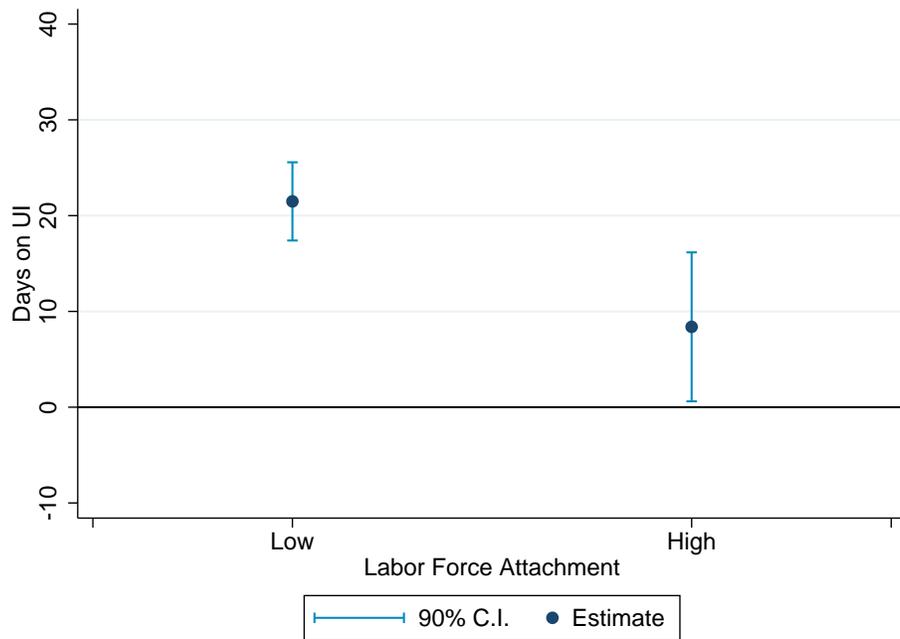


Figure 3.13: Estimates of short term effects on days on UI for both high and low labor force attachment samples

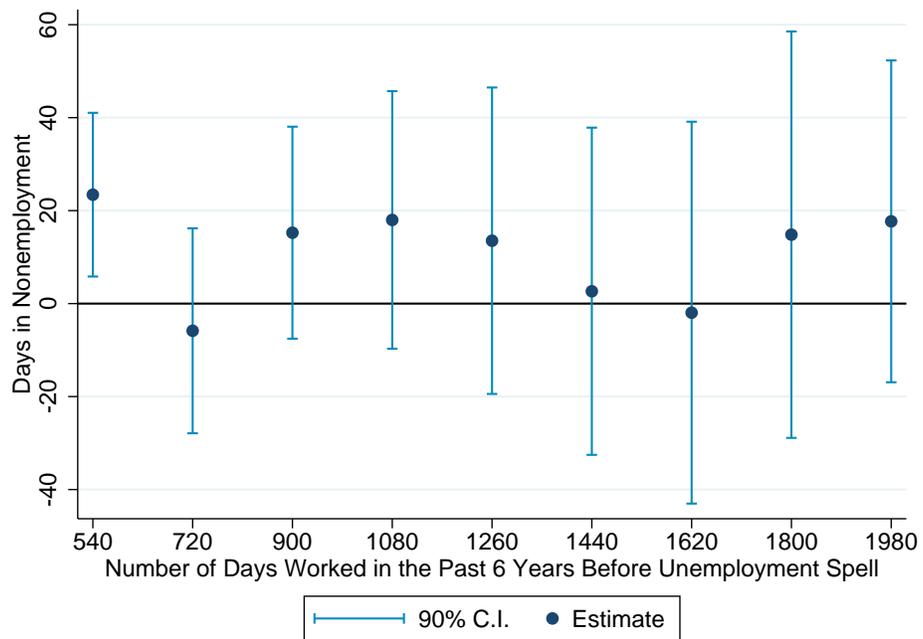


Figure 3.14: Estimates of short term effects on days spent in nonemployment at different cutoffs

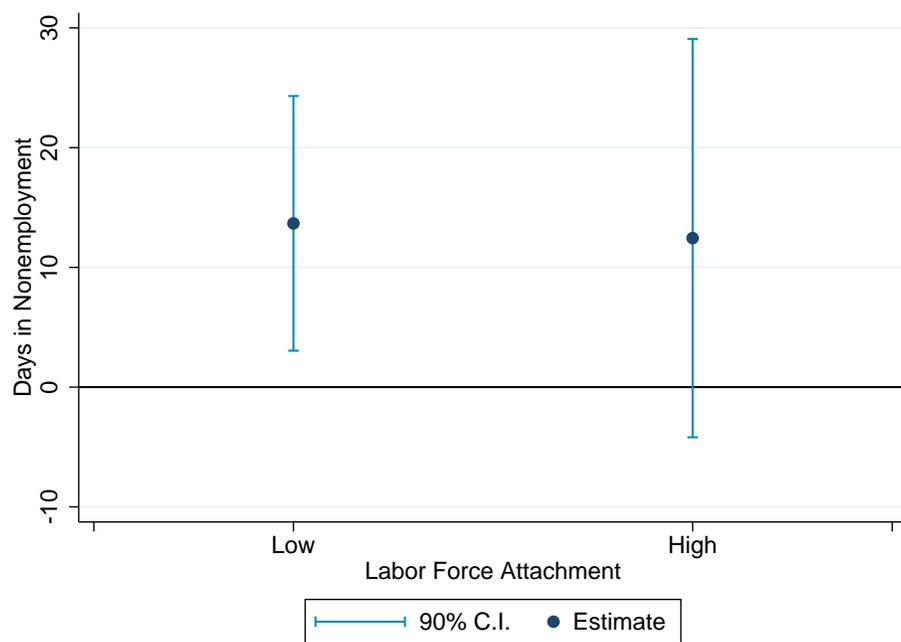
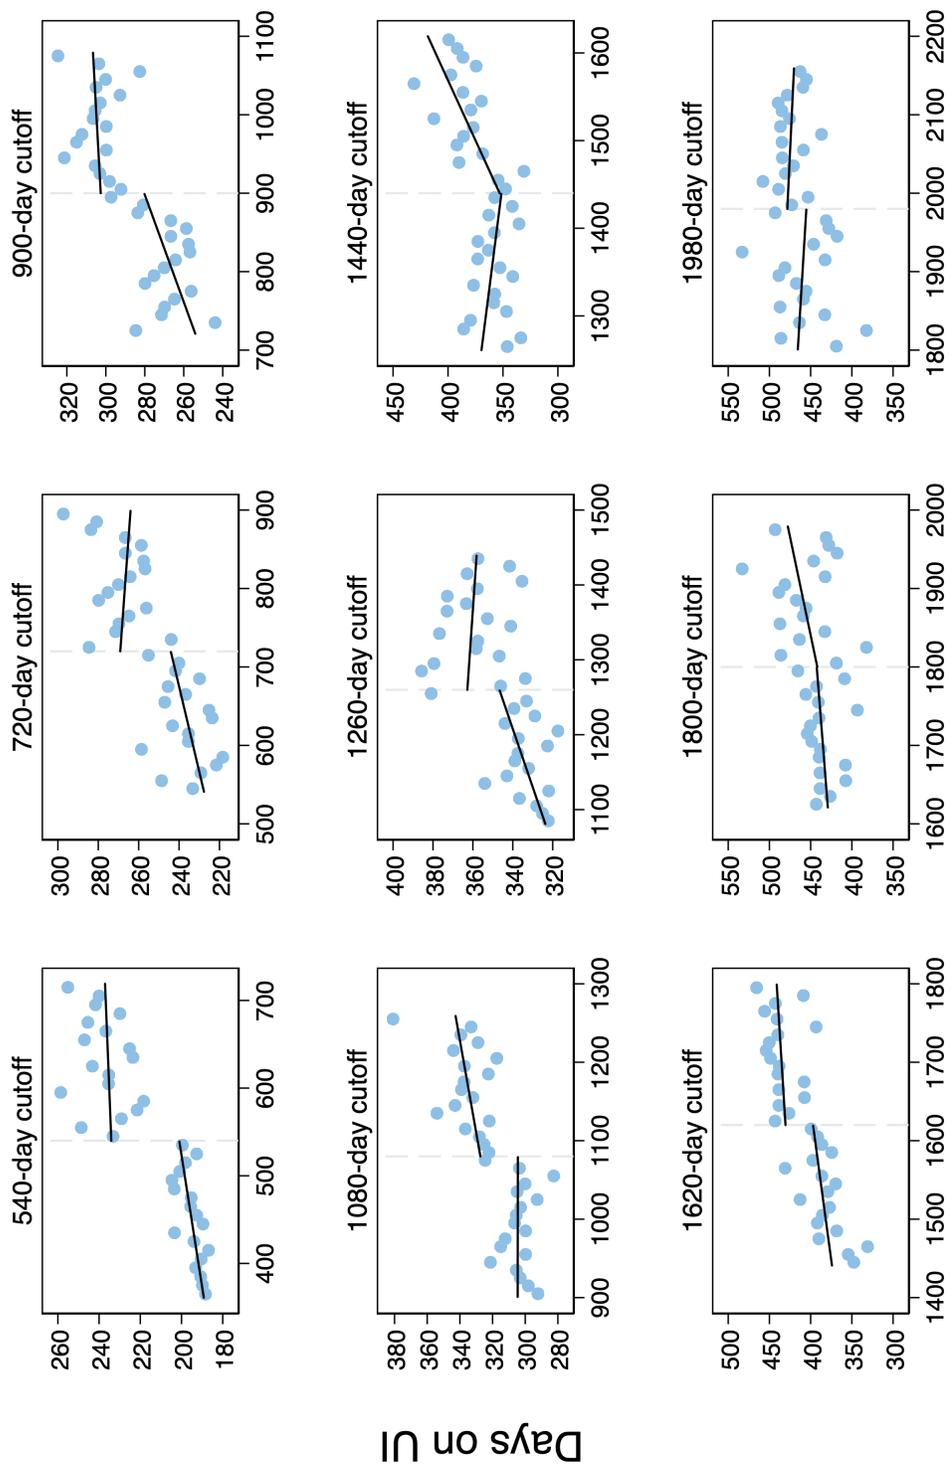
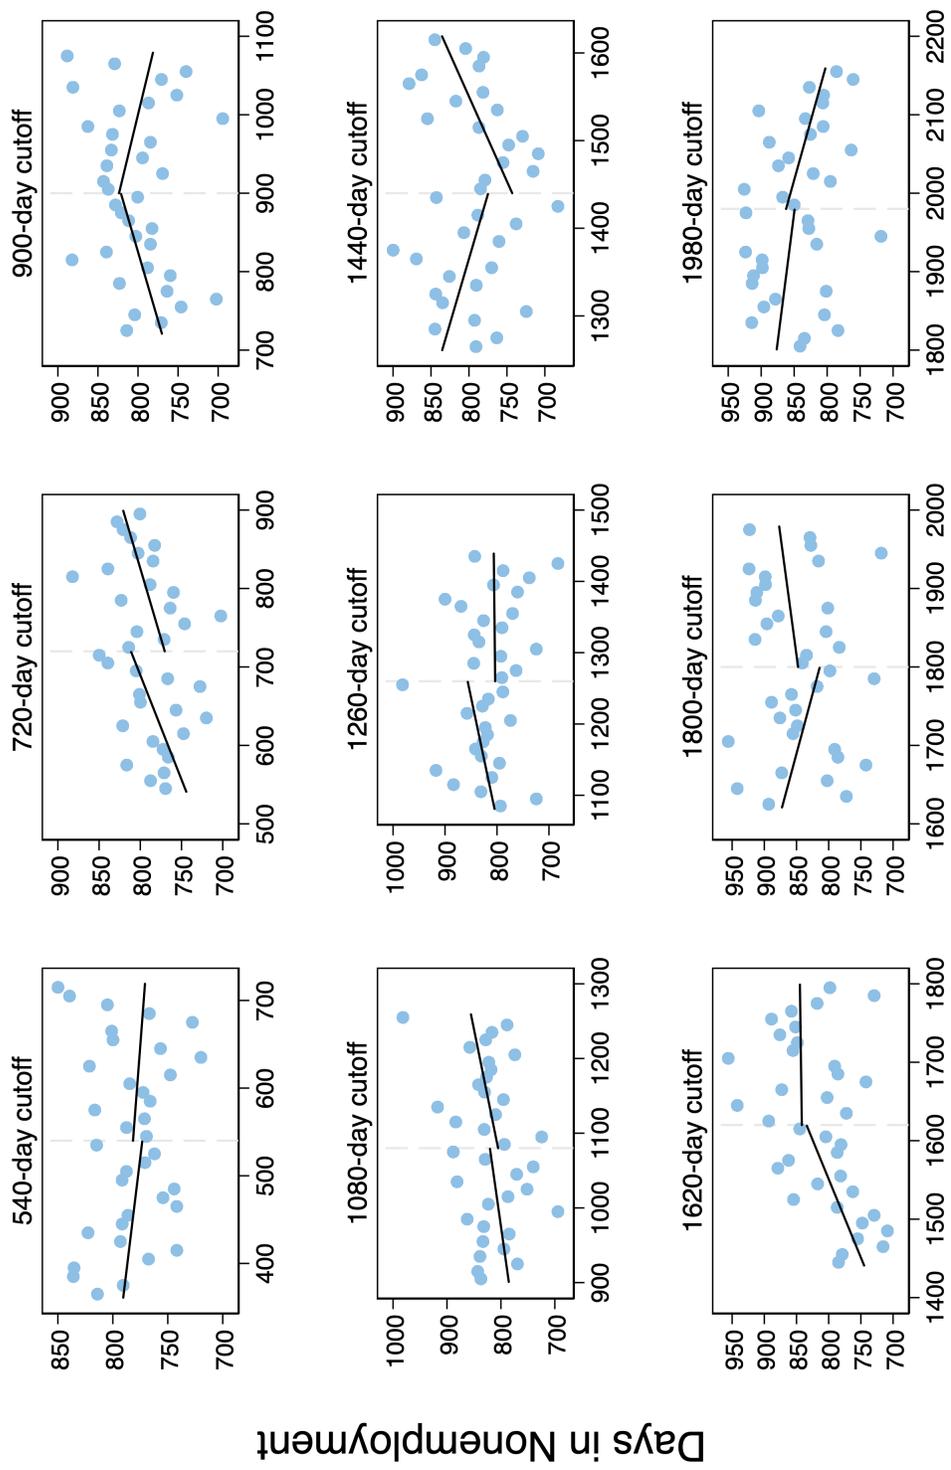


Figure 3.15: Estimates of short term effects on days spent in nonemployment for both high and low labor force attachment samples



Attachment

Figure 3.16: The long term effect of extended UI duration on days spent on UI



Attachment

Figure 3.17: The long term effect of extended UI duration on days in nonemployment

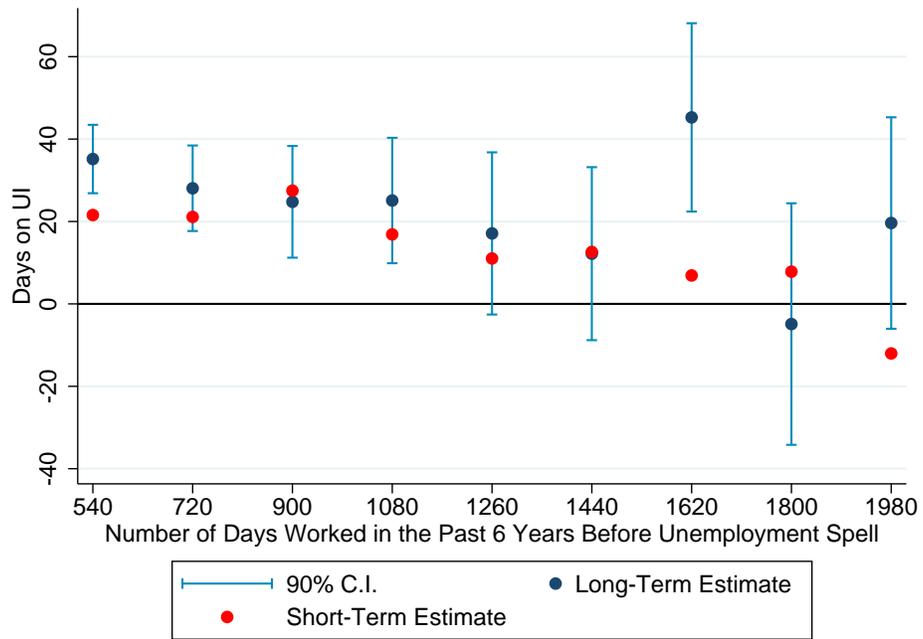


Figure 3.18: Estimates of short and long term effects on days on UI at different cutoffs

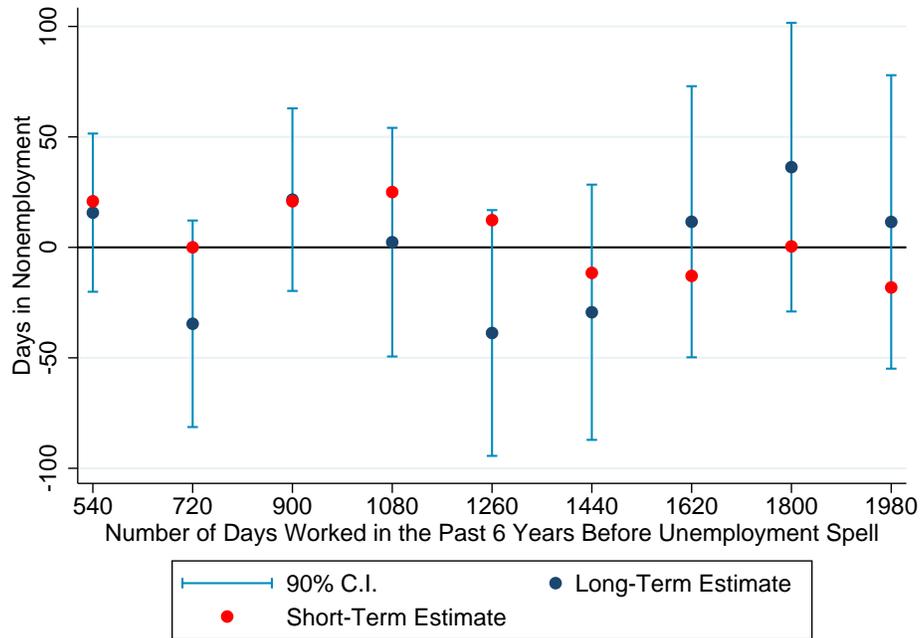


Figure 3.19: Estimates of short and long term effects on days in nonemployment at different cutoffs

3.9.2 Additional Tables

Table 3.5: Occupational category groups

Group name in Spanish	Name in English	Consolidated name
<i>Ingenieros, licenciados y alta dirección</i>	Engineers, graduates, and high-level managers	Skilled
<i>Ingenieros técnicos, peritos y ayudantes</i>	Assistant engineers and qualified specialists	Skilled
<i>Jefes administrativos y de taller</i>	Managers	Skilled
<i>Ayudantes no titulados</i>	Assistants without a post-secondary degree	Staff
<i>Oficiales administrativos</i>	Administrators	Staff
<i>Subalternos</i>	Entry-level and junior employees	Staff
<i>Auxiliares administrativos</i>	Assistant administrators	Staff
<i>Oficiales de primera y segunda</i>	First-tier technical workers	Technical
<i>Oficiales de tercera y especialistas</i>	Second- and third-tier technical workers	Technical
<i>Mayores de 18 años no cualificados</i>	Unskilled workers	Laborer
<i>Trabajadores menores de 18 años</i>	Underage workers	Underage

Notes: original names in Spanish as provided by the SSSA. Translations into English made by the author. Consolidated names based on Miller and Llull (2016).

Table 3.6: Results of McCrary test at different cutoffs

Cutoff	Log Diff. at Cutoff	t-statistic
540	.1099069	1.947959
720	-.0069476	-.0920764
900	-.0490517	-.636777
1080	.1268707	1.645804
1260	-.1466149	-1.510996
1440	.1002719	.9678793
1620	.0017638	.0188002
1800	.0130167	.1450876
1980	-.1175072	-1.313129

Table 3.7: Short term effects for both high and low labor force attachment samples

	(1) Days UI (Low)	(2) Days UI (High)	(3) Nonemployment (Low)	(4) Nonemployment (High)
<i>D</i>	21.49*** (2.477)	8.389* (4.716)	13.67** (6.449)	12.45 (10.09)
Observations	40331	26257	40331	26257
R-squared	0.0989	0.0981	0.0865	0.0873
Controls	Yes	Yes	Yes	Yes
Weights	Yes	Yes	Yes	Yes
Low Sample	Yes	No	Yes	No
High Sample	No	Yes	No	Yes

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.

Table 3.8: The long term effect of extended UI duration on days spent on UI

	p-value of diff.	Short-term	Long-term
D_{540}	.0093	21.57*** (2.518)	35.16*** (5.033)
Obs. (R^2)		12778 (0.102)	12778 (0.102)
D_{720}	.2996	21.11*** (3.765)	28.06*** (6.295)
Obs. (R^2)		8252 (0.103)	8252 (0.097)
D_{900}	.7378	27.48*** (5.755)	24.77*** (8.219)
Obs. (R^2)		6180 (0.101)	6180 (0.109)
D_{1080}	.2639	16.87** (7.293)	25.09*** (9.231)
Obs. (R^2)		5041 (0.112)	5041 (0.127)
D_{1260}	.5648	11.03 (9.278)	17.10 (11.93)
Obs. (R^2)		4246 (0.119)	4246 (0.142)
D_{1440}	.9705	12.61 (10.72)	12.19 (12.73)
Obs. (R^2)		3509 (0.122)	3509 (0.156)
D_{1620}	.0072	6.904 (13.87)	45.26*** (13.86)
Obs. (R^2)		3171 (0.169)	3171 (0.187)
D_{1800}	.3797	7.820 (16.89)	-4.891 (17.78)
Obs. (R^2)		3196 (0.174)	3196 (0.203)
D_{1980}	.0311	-12.02 (15.85)	19.63 (15.56)
Obs. (R^2)		5137 (0.145)	5137 (0.170)
Controls (weights)		Yes (Yes)	Yes (Yes)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.

Table 3.9: The long term effect of extended UI duration on days in nonemployment

	p-value of diff.	Short-term	Long-term
D_{540}	.7938	20.86* (12.25)	15.70 (21.71)
Obs. (R^2)		12778 (0.085)	12778 (0.101)
D_{720}	.1759	0.0324 (14.96)	-34.59 (28.34)
Obs. (R^2)		8252 (0.108)	8252 (0.115)
D_{900}	.9757	20.85 (16.84)	21.60 (25.06)
Obs. (R^2)		6180 (0.100)	6180 (0.132)
D_{1080}	.4219	25.04 (17.79)	2.347 (31.37)
Obs. (R^2)		5041 (0.133)	5041 (0.142)
D_{1260}	.091	12.32 (21.89)	-38.76 (33.72)
Obs. (R^2)		4246 (0.129)	4246 (0.151)
D_{1440}	.5568	-11.56 (22.79)	-29.38 (35.02)
Obs. (R^2)		3509 (0.122)	3509 (0.162)
D_{1620}	.4785	-12.91 (28.32)	11.59 (37.20)
Obs. (R^2)		3171 (0.159)	3171 (0.169)
D_{1800}	.2904	0.432 (30.46)	36.31 (39.62)
Obs. (R^2)		3196 (0.155)	3196 (0.185)
D_{1980}	.3896	-18.15 (26.09)	11.50 (40.27)
Obs. (R^2)		5137 (0.125)	5137 (0.130)
Controls (weights)		Yes (Yes)	Yes (Yes)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.

Table 3.10: Robustness of short term effects on days spent on UI

	(1) Linear	(2) Donut	(3) Quadratic	(4) Local Linear
D_{540}	22.52*** (2.281)	22.57*** (2.346)	19.89*** (3.358)	18.35*** (3.718)
D_{720}	22.09*** (3.245)	22.66*** (3.351)	17.09*** (4.514)	17.60*** (4.478)
D_{900}	27.46*** (4.944)	28.12*** (5.109)	29.51*** (6.902)	29.15*** (5.503)
D_{1080}	16.12** (6.434)	18.68*** (6.330)	8.983 (9.263)	3.539 (10.13)
D_{1260}	4.363 (8.761)	6.216 (8.899)	-13.49 (12.57)	-6.329 (13.60)
D_{1440}	12.25 (10.03)	10.57 (10.47)	25.32* (14.58)	16.04 (13.02)
D_{1620}	13.44 (11.64)	12.19 (12.06)	34.19** (15.99)	33.28* (17.47)
D_{1800}	9.316 (13.53)	14.21 (13.34)	-8.190 (19.23)	9.717 (18.94)
D_{1980}	-0.785 (11.77)	-3.661 (11.78)	-13.76 (17.39)	0.625 (13.12)
Controls	Yes	Yes	Yes	No
Weights	Yes	Yes	Yes	No

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.

Table 3.11: Robustness of short term effects on days spent in nonemployment

	(1) Linear	(2) Donut	(3) Quadratic	(4) Local Linear
D_{540}	23.43** (10.68)	22.62** (11.08)	34.56** (15.93)	18.35*** (3.568)
D_{720}	-5.841 (13.37)	-8.485 (13.67)	-28.54 (20.43)	17.60*** (4.413)
D_{900}	15.25 (13.83)	20.22 (13.89)	37.27* (19.02)	29.15*** (5.481)
D_{1080}	18.00 (16.80)	27.16* (15.91)	2.397 (23.98)	3.539 (10.10)
D_{1260}	13.53 (19.98)	14.27 (20.73)	-15.04 (29.74)	-6.329 (13.52)
D_{1440}	2.656 (21.34)	0.532 (22.04)	14.60 (32.26)	16.04 (13.43)
D_{1620}	-1.951 (24.91)	-6.372 (25.43)	23.70 (36.49)	33.28* (17.27)
D_{1800}	14.83 (26.51)	20.24 (26.71)	-5.626 (40.27)	9.717 (18.47)
D_{1980}	17.70 (20.99)	23.43 (20.77)	1.025 (30.99)	0.625 (13.22)
Controls	Yes	Yes	Yes	No
Weights	Yes	Yes	Yes	No

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.

Table 3.12: Robustness of long term effects on days spent on UI

	(1) Linear	(2) Donut	(3) Quadratic	(4) Local Linear
D_{540}	35.16*** (5.033)	34.26*** (5.153)	40.24*** (6.721)	39.69*** (6.862)
D_{720}	28.06*** (6.295)	26.67*** (6.401)	33.05*** (8.904)	24.06*** (8.709)
D_{900}	24.77*** (8.219)	24.28*** (8.464)	8.875 (12.32)	-5.040 (14.30)
D_{1080}	25.09*** (9.231)	29.73*** (9.000)	13.26 (13.58)	0.579 (16.30)
D_{1260}	17.10 (11.93)	17.64 (12.45)	1.636 (16.85)	-19.34 (23.24)
D_{1440}	12.19 (12.73)	12.70 (12.21)	7.141 (19.05)	-6.735 (16.94)
D_{1620}	45.26*** (13.86)	49.20*** (14.09)	53.76*** (19.43)	37.21** (17.29)
D_{1800}	-4.891 (17.78)	-1.282 (17.97)	-46.20* (25.06)	-13.46 (23.92)
D_{1980}	19.63 (15.56)	16.97 (16.10)	21.72 (21.03)	24.54 (17.61)
Controls	Yes	Yes	Yes	No
Weights	Yes	Yes	Yes	No

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.

Table 3.13: Robustness of long term effects on days spent in nonemployment

	(1) Linear	(2) Donut	(3) Quadratic	(4) Local Linear
D_{540}	15.70 (21.71)	13.62 (22.21)	3.488 (32.24)	39.69*** (7.152)
D_{720}	-34.59 (28.34)	-33.60 (28.64)	-47.80 (43.93)	24.06*** (8.624)
D_{900}	21.60 (25.06)	20.53 (25.76)	56.34 (35.30)	-5.040 (14.45)
D_{1080}	2.347 (31.37)	14.69 (31.69)	-49.79 (44.50)	0.579 (16.31)
D_{1260}	-38.76 (33.72)	-40.40 (33.49)	-84.90 (52.26)	-19.34 (23.36)
D_{1440}	-29.38 (35.02)	-24.76 (36.29)	-1.996 (48.61)	-6.735 (16.97)
D_{1620}	11.59 (37.20)	15.77 (37.32)	31.78 (52.46)	37.21** (17.09)
D_{1800}	36.31 (39.62)	44.37 (39.71)	7.361 (55.94)	-13.46 (24.48)
D_{1980}	11.50 (40.27)	17.23 (40.24)	13.44 (60.16)	24.54 (17.67)
Controls	Yes	Yes	Yes	No
Weights	Yes	Yes	Yes	No

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.

Table 3.14: The effect of extended UI duration on transition to a permanent contract

	(1)	(2)	(3)
D_{540}	0.0112 (0.0121)	0.0122 (0.0121)	-0.00494 (0.0131)
Obs. (R^2)	15947 (0.003)	15947 (0.066)	15947 (0.073)
D_{720}	-0.000521 (0.0165)	0.00353 (0.0158)	0.0149 (0.0176)
Obs. (R^2)	10291 (0.001)	10291 (0.083)	10291 (0.090)
D_{900}	0.00939 (0.0185)	-0.00146 (0.0179)	-0.00263 (0.0192)
Obs. (R^2)	7730 (0.000)	7730 (0.103)	7730 (0.116)
D_{1080}	0.00879 (0.0208)	-0.0000319 (0.0198)	0.00289 (0.0223)
Obs. (R^2)	6363 (0.001)	6363 (0.110)	6363 (0.110)
D_{1260}	0.00788 (0.0239)	0.00849 (0.0220)	0.00858 (0.0239)
Obs. (R^2)	5426 (0.000)	5426 (0.104)	5426 (0.121)
D_{1440}	0.0613** (0.0266)	0.0534** (0.0258)	0.0446 (0.0281)
Obs. (R^2)	4622 (0.002)	4622 (0.105)	4622 (0.114)
D_{1620}	-0.0246 (0.0254)	-0.0189 (0.0257)	-0.00410 (0.0266)
Obs. (R^2)	4338 (0.002)	4338 (0.111)	4338 (0.116)
D_{1800}	-0.00262 (0.0278)	-0.00176 (0.0267)	-0.00980 (0.0297)
Obs. (R^2)	4499 (0.000)	4499 (0.129)	4499 (0.147)
D_{1980}	-0.000664 (0.0227)	-0.000326 (0.0218)	0.0114 (0.0225)
Obs. (R^2)	7372 (0.0000552)	7372 (0.112)	7372 (0.131)
Controls (weights)	No (No)	Yes (No)	Yes (Yes)

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Standard errors clustered at the day (attachment) level. Control variables include gender, year of birth, province where the worker became affiliated to the SSSA, type of affiliation to the SSSA, occupational group, foreign status, year in which unemployment started, age at which unemployment started, and whether the worker was on a temporary contract.

3.9.3 Prediction and Characterization of Workers at Spikes

I examine the characteristics of workers at spikes in tables 3.15 and 3.16.²⁵ In table 3.15, I present the results of five regressions, each regressing a characteristic of interest on an indicator for whether a worker is at a spike and a set of covariates. Importantly, the set of covariates includes dummies for “attachment regions,” the flat areas on the step function governing UI benefit duration. Including these dummies ensures that workers on spikes are only compared to workers that are in the same area of the attachment distribution, since each attachment region covers 180 days.

Table 3.15: Characterizing workers at spikes in terms of predetermined variables

	(1) Female Female	(2) Temp. Contract	(3) Num. Episodes	(4) Num. Employers	(5) Cont. Employed
Full/Half yr.	0.0994*** (0.00820)	0.0984*** (0.00617)	-2.069*** (0.0791)	-0.827*** (0.0242)	0.410*** (0.00822)
Observations	44754	44754	44754	44754	44754
R-squared	0.190	0.196	0.0458	0.136	0.153
Attach. Region	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes

In table 3.16, I run similar regressions but instead I focus on differences in terms of outcome variables for workers at spikes compared to the rest of the sample.

Table 3.17 contains the results of predicting whether a worker is at a spike, using a random forest classifier.²⁶ For this prediction exercise, I use the following variables: number of work episodes, number of employers, gender, whether the worker was on a temporary contract, and a factor variable for the attachment level divided into regions of 100 days each.²⁷ I include two performance metrics for four different values of the classification threshold: the true positive rate

²⁵I consider a worker to be at a spike if they have a work history corresponding to discrete years or half-years.

²⁶For details on this method, see Breiman (2001).

²⁷The inclusion of this factor variable enables the method to predict whether a worker is at a spike for workers with similar levels of labor force attachment.

Table 3.16: Characterizing workers at spikes in terms of outcome variables

	(1) Days on UI	(2) Days in Nonemployment
Full/Half yr.	-2.201 (1.720)	34.25*** (5.945)
Observations	44754	44754
R-squared	0.205	0.0572
Attach. Region	Yes	Yes
Controls	Yes	Yes

(the proportion of actual positives that are correctly identified as such) and the true negative rate (the proportion of actual negatives that are correctly identified as such).²⁸

Table 3.17: Evaluating random forest classifier performance at different threshold values

Threshold value (λ):	0.5	0.25	0.1	0.05
True positive rate	0.44	0.48	0.60	0.63
True negative rate	0.99	0.99	0.96	0.93

In figure 3.20 I plot the receiver operating characteristic (ROC) curve for this classifier. This curve plots the relationship between the true positive rate and one minus the true negative rate, as the classifier threshold varies between zero and one. It is considered desirable for a classifier to have both high true positive and true negative rates, which translates into the line depicted in the ROC curve reaching for the upper-left corner.

²⁸The true positive rate is also known as *sensitivity*, while the true negative rate is also known as *specificity*.

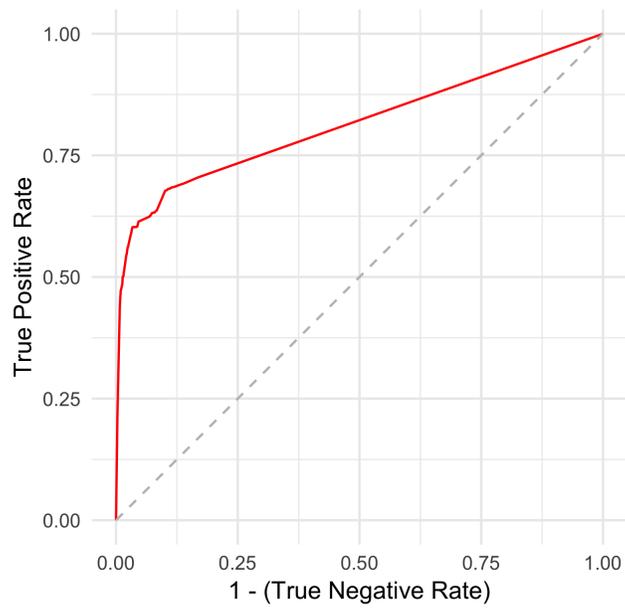


Figure 3.20: ROC curve for spike worker detection

Bibliography

- Aaronson, D. (1998). Using sibling data to estimate the impact of neighborhoods on children's educational outcomes. *The Journal of Human Resources*, 33(4), 915–946. Retrieved from <http://www.jstor.org/stable/146403>
- Agrawala, M., Altonji, J. G., & Mansfield, R. K. (2018). *Quantifying Family, School, and Location Effects in the Presence of Complementarities and Sorting*. Forthcoming, *Journal of Labor Economics*.
- Ainsworth, J. W. (2002). Why does it take a village? the mediation of neighborhood effects on educational achievement. *Social Forces*, 81(1), 117–152. Retrieved from <http://www.jstor.org/stable/3086529>
- Akerlof, K., Maibach, E. W., Fitzgerald, D., Ceden, A. Y., & Neuman, A. (2013). Do people “personally experience” global warming, and if so how, and does it matter? *Global Environmental Change*, 23(1), 81–91.
- Alba-Ramírez, A., Arranz, J. M., & Muñoz-Bullón, F. (2007). Exits from unemployment: Recall or new job. *Labour Economics*, 14(5), 788–810.
- Ammermueller, A. & Pischke, J. (2009). Peer effects in european primary schools: Evidence from the progress in international reading literacy study. *Journal of Labor Economics*, 27(3), 315–348. Retrieved from <http://www.jstor.org/stable/10.1086/603650>
- Angrist, J. D. & Lang, K. (2004). Does school integration generate peer effects? evidence from boston's metco program. *American Economic Review*, 94(5), 1613–1634. doi:10.1257/0002828043052169
- Ansola-behere, S., de Figueiredo, J. M., & Snyder, J., James M. (2003). Why is there so little money in u.s. politics? *Journal of Economic Perspectives*, 17(1), 105–130. doi:10.1257/089533003321164976
- Arceneaux, K. & Stein, R. M. (2006). Who is held responsible when disaster strikes? the attribution of responsibility for a natural disaster in an urban election. *Journal of Urban Affairs*, 28(1), 43–53.

- Arranz, J. M. & Muro Romero, J. (2004). An extra time duration model with application to unemployment duration under benefits in Spain. *Hacienda Pública Española*, (171), 133–157.
- Bagues, M. & Esteve-Volart, B. (2016). Politicians' luck of the draw: Evidence from the Spanish Christmas lottery. *Journal of Political Economy*, 124(5), 1269–1294. doi:10.1086/688178
- Bailey, M., Cao, R., Kuchler, T., Stroebel, J., & Wong, A. (2018). Social connectedness: Measurement, determinants, and effects. *Journal of Economic Perspectives*, 32(3), 259–80. doi:10.1257/jep.32.3.259
- Barber, M. (2016). Donation motivations: Testing theories of access and ideology. *Political Research Quarterly*, 69(1), 148–159. doi:10.1177/1065912915624164
- Bayer, P., Ross, S. L., Topa, G., Journal, S., & December, N. (2008). Place of Work and Place of Residence : Informal Hiring Networks and Labor Market Outcomes Giorgio Topa. *Journal of Political Economy*, 116(6), 1150–1196.
- Bentolila, S., Cahuc, P., Dolado, J. J., Le Barbanchon, T., et al. (2010). Unemployment and temporary jobs in the crisis: Comparing France and Spain. *Documento de trabajo, FEDEA*.
- Betts, J. R. & Zau, A. (2004). *Peer groups and academic achievement: Panel evidence from administrative data*. IZA/SOLE Transatlantic Meeting. Unpublished manuscript.
- Blanchard, O. J. & Diamond, P. (1994). Ranking, unemployment duration, and wages. *The Review of Economic Studies*, 61(3), 417–434.
- Blume, L. E., Brock, W. A., Durlauf, S. N., & Ioannides, Y. M. (2011). Chapter 18 - Identification of Social Interactions. In J. Benhabib, A. Bisin, & M. O. Jackson (Eds.), *Handbook of social economics* (Vol. 1, pp. 853–964). North-Holland. doi:https://doi.org/10.1016/B978-0-444-53707-2.00001-3
- Bonica, A. (2013). Ideology and interests in the political marketplace. *American Journal of Political Science*, 57(2), 294–311. doi:10.1111/ajps.12014
- Bonica, A. (2014). Mapping the ideological marketplace. *American Journal of Political Science*, 58(2), 367–386. doi:10.1111/ajps.12062
- Bonica, A. (2016). Database on ideology, money in politics, and elections: Public version 2.0 [computer file]. *Stanford, CA: Stanford University Libraries*. <https://data.stanford.edu/dime>.
- Boone, J. & van Ours, J. C. (2012). Why is there a spike in the job finding rate at benefit exhaustion? *De Economist*, 160(4), 413–438.
- Bover, O., Arellano, M., & Bentolila, S. (2002). Unemployment duration, benefit duration and the business cycle. *The Economic Journal*, 112(479), 223–265.

- Brännström, L. (2008). Making Their Mark: The Effects of Neighbourhood and Upper Secondary School on Educational Achievement. *European Sociological Review*, 24(4), 463–478. doi:10.1093/esr/jcn013. eprint: <http://oup.prod.sis.lan/esr/article-pdf/24/4/463/1319997/jcn013.pdf>
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
- Brooks-Gunn, J., Duncan, G. J., Klebanov, P. K., & Sealand, N. (1993). Do neighborhoods influence child and adolescent development? *American Journal of Sociology*, 99(2), 353–395. Retrieved from <http://www.jstor.org/stable/2781682>
- Caliendo, M., Tatsiramos, K., & Uhlendorff, A. (2013). Benefit Duration, Unemployment Duration and Job Match Quality: A Regression-Discontinuity Approach. *Journal of Applied Econometrics*, 28(4), 604–627.
- Card, D., Chetty, R., & Weber, A. (2007). *The spike at benefit exhaustion: Leaving the unemployment system or starting a new job?* National Bureau of Economic Research.
- Card, D. & Levine, P. B. (2000). Extended benefits and the duration of ui spells: Evidence from the new jersey extended benefit program. *Journal of public economics*, 78(1), 107–138.
- Case, A. C. & Katz, L. F. (1991). *The company you keep: The effects of family and neighborhood on disadvantaged youths* (Working Paper No. 3705). National Bureau of Economic Research. doi:10.3386/w3705
- Chetty, R. (2008). *Moral hazard vs. liquidity and optimal unemployment insurance*. National Bureau of Economic Research.
- Chetty, R. & Hendren, N. (2016a). *The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects* (Working Paper No. 23001). National Bureau of Economic Research. doi:10.3386/w23001
- Chetty, R. & Hendren, N. (2016b). *The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates* (Working Paper No. 23002). National Bureau of Economic Research. doi:10.3386/w23002
- Currie, J., DellaVigna, S., Moretti, E., & Pathania, V. (2010). The effect of fast food restaurants on obesity and weight gain. *American Economic Journal: Economic Policy*, 2(3), 32–63. doi:10.1257/pol.2.3.32
- Davis, F. L. & Wurth, A. H. (2003). Voting preferences and the environment in the american electorate: The discussion extended. *Society & Natural Resources*, 16(8), 729–740.
- Del Bello, C. L., Patacchini, E., & Zenou, Y. (2015). *Neighborhood Effects in Education* (IZA Discussion Papers No. 8956). Institute for the Study of Labor (IZA). Retrieved from <https://ideas.repec.org/p/iza/izadps/dp8956.html>

- Deryugina, T. (2013). How do people update? the effects of local weather fluctuations on beliefs about global warming. *Climatic Change*, 118(2), 397–416.
- Dolado, J. J., García-Serrano, C., & Jimeno, J. F. (2002). Drawing lessons from the boom of temporary jobs in Spain. *The Economic Journal*, 112(480), F270–F295.
- Durlauf, S. (2004). Neighborhood effects. In J. V. Henderson & J. F. Thisse (Eds.), *Handbook of regional and urban economics* (1st ed., Chap. 50, Vol. 4, pp. 2173–2242). Elsevier. Retrieved from <https://EconPapers.repec.org/RePEc:eee:regchp:4-50>
- Egan, P. J. & Mullin, M. (2012). Turning personal experience into political attitudes: The effect of local weather on Americans' perceptions about global warming. *The Journal of Politics*, 74(3), 796–809.
- Ensley, M. J. (2009). Individual campaign contributions and candidate ideology. *Public Choice*, 138(1/2), 221–238. Retrieved from <http://www.jstor.org/stable/40270840>
- Francia, P. L., Green, J. C., Herrnson, P. S., Powell, L. W., & Wilcox, C. (2003). *The financiers of congressional elections: Investors, ideologues, and intimates*. Columbia University Press.
- Galiani, S., Gertler, P. J., & Undurraga, R. (2018). *Aspiration adaptation in resource-constrained environments* (Working Paper No. 24264). National Bureau of Economic Research. Retrieved from <http://www.nber.org/papers/w24264>
- Gallagher, J. (2014). Learning about an infrequent event: Evidence from flood insurance take-up in the United States. *American Economic Journal: Applied Economics*, 6(3), 206–33.
- Galster, G., Santiago, A., Stack, L., & Cutsinger, J. (2016). Neighborhood effects on secondary school performance of Latino and African American youth: Evidence from a natural experiment in Denver. *Journal of Urban Economics*, 93(100), 30–48. doi:10.1016/j.jue.2016.02.004
- García-Pérez, J. I. [J Ignacio] & Muñoz-Bullón, F. (2011). Transitions into permanent employment in Spain: An empirical analysis for young workers. *British Journal of Industrial Relations*, 49(1), 103–143.
- García-Pérez, J. I. [José Ignacio] & Rebollo-Sanz, Y. (2005). Wage changes through job mobility in Europe: A multinomial endogenous switching approach. *Labour Economics*, 12(4), 531–555.
- Gaspar, J. T. & Reeves, A. (2011). Make it rain? Retrospection and the attentive electorate in the context of natural disasters. *American Journal of Political Science*, 55(2), 340–355.
- Geolytics Inc. (2012). Neighborhood change database (NCDB) 2010: Tract data for 1970-80-90-00-10. <<http://www.geolytics.com/USCensus,Neighborhood-Change-Database-1970-2000,Products.asp>>.

- Gibbons, S., Silva, O., & Weinhardt, F. (2013). Everybody Needs Good Neighbours? Evidence from Students' Outcomes in England. *Economic Journal*, 123(April 2013), 831–874.
- Gomez, B. T., Hansford, T. G., & Krause, G. A. (2007). The republicans should pray for rain: Weather, turnout, and voting in us presidential elections. *Journal of Politics*, 69(3), 649–663.
- Gonzalez-Rozada, M. & Ruffo, H. (2014). *The Effects of Unemployment Insurance Under High Informality: Evidence from Argentina*. Working Paper.
- Gonzalez-Rozada, M. & Ruffo, H. (2016). Optimal unemployment benefits in the presence of informal labor markets. *Labour Economics*.
- Goux, D. & Maurin, E. (2007). Close neighbours matter: Neighbourhood effects on early performance at school. *The Economic Journal*, 117(523), 1193–1215. doi:10.1111/j.1468-0297.2007.02079.x. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0297.2007.02079.x>
- Graff Zivin and Neidell. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1), 1–26. doi:10.1086/671766. eprint: <https://doi.org/10.1086/671766>
- Graham, B. S. (2008). Identifying social interactions through conditional variance restrictions. *Econometrica*, 76(3), 643–660. doi:10.1111/j.1468-0262.2008.00850.x. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0262.2008.00850.x>
- Graham, B. S. (2018). Identifying and estimating neighborhood effects. *Journal of Economic Literature*, 56(2), 450–500. doi:10.1257/jel.20160854
- Grinblatt, M., Keloharju, M., & Ikäheimo, S. (2008). Social Influence and Consumption: Evidence from the Automobile Purchases of Neighbors. *Review of Economics and Statistics*, 90(4), 735–753. doi:10.1162/rest.90.4.735
- Guber, D. L. (2001). Voting preferences and the environment in the american electorate. *Society & Natural Resources*, 14(6), 455–469.
- Hamilton, L. C. & Stampone, M. D. (2013). Blowin' in the wind: Short-term weather and belief in anthropogenic climate change. *Weather, Climate, and Society*, 5(2), 112–119.
- Hanushek, E. A., Kain, J. F., Markman, J. M., & Rivkin, S. G. (2003). Does peer ability affect student achievement? *Journal of Applied Econometrics*, 18(5), 527–544. doi:10.1002/jae.741
- Hastie, T., Tibshirani, R., & Friedman, J. H. (2009). *The elements of statistical learning: Data mining, inference, and prediction, 2nd edition*. Springer series in statistics. Springer. Retrieved from <http://www.worldcat.org/oclc/300478243>

- Healy, A. & Malhotra, N. (2009). Myopic voters and natural disaster policy. *American Political Science Review*, 103(3), 387–406.
- Herrnstadt, E. & Muehlegger, E. (2014). Weather, salience of climate change and congressional voting. *Journal of Environmental Economics and Management*, 68(3), 435–448.
- Howe, P. D., Mildenerger, M., Marlon, J. R., & Leiserowitz, A. (2015). Geographic variation in opinions on climate change at state and local scales in the usa. *Nature Climate Change*, 5(6), 596.
- Hoxby, C. (2000). *Peer effects in the classroom: Learning from gender and race variation* (Working Paper No. 7867). National Bureau of Economic Research. doi:10.3386/w7867
- Imbens, G. & Kalyanaraman, K. (2012). Optimal bandwidth choice for the regression discontinuity estimator. *The Review of Economic Studies*, 79(3), 933–959. Retrieved from <http://www.jstor.org/stable/23261375>
- IPCC. (2013). *Climate change 2013: The physical science basis. contribution of working group i to the fifth assessment report of the intergovernmental panel on climate change*. Cambridge, United Kingdom: Cambridge University Press. doi:10.1017/CBO9781107415324
- Jachimowicz, J., Menges, J., & Galinsky, A. D. (2017). *How wind speed affects voting decisions (may 31, 2017)*. Columbia Business School Research Paper No. 16-82. Available at SSRN: <https://ssrn.com/abstract=2868054> or <http://dx.doi.org/10.2139/ssrn.2868054>.
- Joireman, J., Truelove, H. B., & Duell, B. (2010). Effect of outdoor temperature, heat primes and anchoring on belief in global warming. *Journal of Environmental Psychology*, 30(4), 358–367.
- Karpf, D. (2013). The internet and american political campaigns. In *The Forum* (Vol. 11, No. 3, pp. 413–428). De Gruyter.
- Katz, L. F. & Meyer, B. D. [Bruce D]. (1990). The impact of the potential duration of unemployment benefits on the duration of unemployment. *Journal of public economics*, 41(1), 45–72.
- Kauppinen, T. M. (2008). Schools as Mediators of Neighbourhood Effects on Choice Between Vocational and Academic Tracks of Secondary Education in Helsinki. *European Sociological Review*, 24(3), 379–391. doi:10.1093/esr/jcn016. eprint: <http://oup.prod.sis.lan/esr/article-pdf/24/3/379/1471996/jcn016.pdf>
- Kling, J. R. [Jeffrey R], Liebman, J. B., Katz, L. F., & Sanbonmatsu, L. (2004). *Moving to Opportunity and Tranquility: Neighborhood Effects on Adult Economic Self-Sufficiency and Health from a Randomized Housing Voucher Experiment*. Unpublished manuscript.

- Kling, J. R. [Jeffrey R.], Ludwig, J., & Katz, L. F. (2005). Neighborhood effects on crime for female and male youth: Evidence from a randomized housing mobility experiment. *Quarterly Journal of Economics*, *120*, 87–130.
- Kling, J., Liebman, J., & Katz, L. (2007). Experimental Analysis of Neighborhood Effects. *Econometrica*, *75*(1), 83–119. doi:10.1111/j.1468-0262.2007.00733.x
- Lalive, R. (2007). Unemployment benefits, unemployment duration, and post-unemployment jobs: A regression discontinuity approach. *The American economic review*, *97*(2), 108–112.
- Lang, C. & Ryder, J. D. (2016). The effect of tropical cyclones on climate change engagement. *Climatic change*, *135*(3-4), 625–638.
- League of Conservation Voters. (2007). National environmental scorecard, first session of the 115th congress. <http://scorecard.lcv.org/sites/scorecard.lcv.org/files/LCV_Scorecard-2017-Full.pdf>.
- Lee, B. A. & Campbell, K. E. (1999). Neighbor Networks of Black and White Americans. In B. Wellman (Ed.), *Networks in the global village: Life in contemporary communities*. Boulder, CO: Westview.
- Lemieux, T. & MacLeod, W. B. (2000). Supply side hysteresis: The case of the canadian unemployment insurance system. *Journal of Public Economics*, *78*(1), 139–170.
- Li, Y., Johnson, E. J., & Zaval, L. (2011). Local warming: Daily temperature change influences belief in global warming. *Psychological science*, *22*(4), 454–459.
- Linden, L. & Rockoff, J. E. (2008). Estimates of the impact of crime risk on property values from Megan's laws. *American Economic Review*, *98*(3), 1103–27. doi:10.1257/aer.98.3.1103
- Ludwig, J., Duncan, G. J., Genetian, L. A., Katz, L. F., Kessler, R. C., Kling, J. R., & Sanbonmatsu, L. (2013). Long-term neighborhood effects on low-income families: Evidence from moving to opportunity. *American Economic Review*, *103*(3), 226–231. doi:10.1257/aer.103.3.226
- Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, *60*(3), 531–542. Retrieved from <http://www.jstor.org/stable/2298123>
- McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, *142*(2), 698–714.
- Meier, A., Schmid, L., & Stutzer, A. (2016). *Rain, emotions and voting for the status quo*. Available at SSRN: <https://ssrn.com/abstract=2868316>.
- Meyer, B. D. [Bruce D.]. (1990). Unemployment insurance and unemployment spells. *Econometrica*, *58*(4), 757–782. doi:10.2307/2938349

- Miller, R. & Llull, J. (2016). Internal migration in dual labor markets. In *2016 meeting papers, society for economic dynamics*.
- Missouri Census Data Center. (2017). MABLE/Geocorr 90/2k/14: Geographic Correspondence Engine. <http://mcdc.missouri.edu/websas/geocorr_index.shtml>.
- Moffitt, R. A. (2000). Policy interventions, low-level equilibria and social interactions. In *Social dynamics* (pp. 45–82). MIT Press.
- Moore, F. C., Obradovich, N., Lehner, F., & Baylis, P. (2019). Rapidly declining remarkability of temperature anomalies may obscure public perception of climate change. *Proceedings of the National Academy of Sciences*, *116*(11), 4905–4910.
- Myers, T. A., Maibach, E. W., Roser-Renouf, C., Akerlof, K., & Leiserowitz, A. A. (2013). The relationship between personal experience and belief in the reality of global warming. *Nature Climate Change*, *3*(4), 343.
- Nieuwenhuis, J. & Hooimeijer, P. (2016). The association between neighbourhoods and educational achievement, a systematic review and meta-analysis. *Journal of Housing and the Built Environment*, *31*(2), 321–347. doi:10.1007/s10901-015-9460-7
- Oreopoulos, P. (2003). The long-run consequences of living in a poor neighborhood. *The Quarterly Journal of Economics*, *118*(4), 1533–1575. Retrieved from <http://www.jstor.org/stable/25053946>
- Owens, A. (2010). Neighborhoods and schools as competing and reinforcing contexts for educational attainment. *Sociology of Education*, *83*(4), 287–311. eprint: <https://doi.org/10.1177/0038040710383519>. Retrieved from <https://doi.org/10.1177/0038040710383519>
- Pierre-Louis, K. (2017). It's cold outside. cue the trump global warming tweet. The New York Times. Available at <www.nytimes.com/2017/12/28/climate/trump-tweet-global-warming.html>.
- Rebollo-Sanz, Y. (2012). Unemployment insurance and job turnover in Spain. *Labour Economics*, *19*(3), 403–426.
- Rebollo-Sanz, Y. F. & García-Pérez, J. I. [J Ignacio]. (2015). Are unemployment benefits harmful to the stability of working careers? The case of Spain. *SERIEs*, *6*(1), 1–41.
- Røed, K. & Zhang, T. (2003). Does unemployment compensation affect unemployment duration? *The Economic Journal*, *113*(484), 190–206.
- Rosenbaum, J. E. (1995). Changing the geography of opportunity by expanding residential choice: Lessons from the gautreaux program. *Housing Policy Debate*, *6*(1), 231–269. doi:10.1080/10511482.1995.9521186. eprint: <https://doi.org/10.1080/10511482.1995.9521186>

- Sacerdote, B. (2011). Peer effects in education: How might they work, how big are they and how much do we know thus far? In E. Hanushek, S. Machin, & L. Woessmann (Eds.), *Handbook of the economics of education* (first, Chap. 04, Vol. 3, pp. 249–277). Elsevier. Retrieved from <https://EconPapers.repec.org/RePEc:eee:educhp:3-04>
- Sanbonmatsu, L., Kling, J. R., Duncan, G. J., & Brooks-Gunn, J. (2006). *Neighborhoods and academic achievement: Results from the moving to opportunity experiment* (Working Paper No. 11909). National Bureau of Economic Research. doi:10.3386/w11909
- Schmieder, J. F. [Johannes F], von Wachter, T., & Bender, S. (2012). *The effects of extended unemployment insurance over the business cycle: Evidence from regression discontinuity estimates over twenty years*. National Bureau of Economic Research.
- Schmieder, J. F. [Johannes F.], von Wachter, T., & Bender, S. (2012). The Long-Term Effects of UI Extensions on Employment. *American Economic Review*, 102(3), 514–19. doi:10.1257/aer.102.3.514
- Schneider, M., Teske, P., & Marschall, M. (2002). *Choosing schools: Consumer choice and the quality of american schools*. Princeton University Press. Retrieved from <https://www.xarg.org/ref/a/0691092834/>
- Scholnick, B. (2013). *Bankruptcy spillovers between close neighbors*. Unpublished manuscript.
- Seguridad Social. (2014). *Muestra Continua de Vidas Laborales Con Datos Fiscales*.
- Sisco, M. R., Bosetti, V., & Weber, E. U. (2017). When do extreme weather events generate attention to climate change? *Climatic change*, 143(1-2), 227–241.
- Snyder, J. M. (1990). Campaign contributions as investments: The u.s. house of representatives, 1980-1986. *Journal of Political Economy*, 98(6), 1195–1227. Retrieved from <http://www.jstor.org/stable/2937755>
- Sovago, S. (2015). *The effect of the UI benefit on labor market outcomes-regression kink evidence from the Netherlands*. Working Paper.
- Spence, A., Poortinga, W., Butler, C., & Pidgeon, N. F. (2011). Perceptions of climate change and willingness to save energy related to flood experience. *Nature climate change*, 1(1), 46.
- Stevens, A. (2001). Despite terrorism, candidates make slow return to fundraising. The Hill. Available at <www.hillnews.com/102401>.
- Sykes, B. & Musterd, S. (2011). Examining neighbourhood and school effects simultaneously: What does the dutch evidence show? *Urban Studies*, 48(7), 1307–1331. doi:10.1177/0042098010371393. eprint: <https://doi.org/10.1177/0042098010371393>

- Tatsiramos, K. (2009). Unemployment insurance in Europe: Unemployment duration and subsequent employment stability. *Journal of the European Economic Association*, 7(6), 1225–1260.
- Turley, R. N. L. (2003). When do neighborhoods matter? the role of race and neighborhood peers. *Social Science Research*, 32(1), 61–79. doi:[https://doi.org/10.1016/S0049-089X\(02\)00013-3](https://doi.org/10.1016/S0049-089X(02)00013-3)
- Tversky, A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *science*, 185(4157), 1124–1131.
- U.S. Census Bureau. (1997). United states census 2000: Participant statistical areas program guidelines. Washington, DC: U.S. Dept. Commerce.
- Van Ours, J. C. & Vodopivec, M. (2008). Does reducing unemployment insurance generosity reduce job match quality? *Journal of Public Economics*, 92(3), 684–695.
- Whitmore, D. (2005). Resource and peer impacts on girls' academic achievement: Evidence from a randomized experiment. *The American Economic Review*, 95(2), 199–203. Retrieved from <http://www.jstor.org/stable/4132816>
- Wilcox, C. (2008). Internet fundraising in 2008: A new model? In *The Forum* (Vol. 6, No. 1). De Gruyter.
- Zaval, L., Keenan, E. A., Johnson, E. J., & Weber, E. U. (2014). How warm days increase belief in global warming. *Nature Climate Change*, 4(2), 143.