From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2023

Labor Market and Political Effects of Structural Economic Changes

Choi, Ji Won

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01cz30pw922

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Labor Market and Political Effects of Structural Economic Changes

Jiwon Choi

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE BY
THE DEPARTMENT OF
Economics
Advisers: Leah Boustan and Ilyana Kuziemko

May 2023

© Copyright by Jiwon Choi, 2023. All rights reserved.

Abstract

This dissertation consists of three chapters in labor economics. A theme throughout these chapters is using empirical analysis to examine how local labor markets, firms, and workers respond to policy and structural economic changes in the US.

Chapter 1 documents the persistent effect of deindustrialization on local economies and their residents, from the decline of the New England textile industry in the 1920s. By comparing workers in towns with more and less initial exposure to the textile industry, I document that workers in textile-heavy towns were less likely to out-migrate. The fixed cost of moving and increased family support may have played a role in the decreased migration. Workers switched to the agricultural sector and faced lower earnings on average. Young workers increased their educational attainment, but the increased schooling did not translate to better labor market outcomes nor increased migration likelihood.

In Chapter 2, Ilyana Kuziemko, Ebonya Washington, Gavin Wright, and I explore the local employment and political effect of NAFTA. We show NAFTA hurt employment in areas exposed to Mexican import competition. While NAFTA-exposed counties were more Democratic during the 1980s up to the early 1990s, they became increasingly Republican after the policy. In the areas most exposed to NAFTA, workers significantly opposed NAFTA, and a significant number of voters with protectionist views moved toward the Republican party between 1992 and 1994.

In Chapter 3, Leah Boustan, David Clingingsmith, and I investigate the effect of computer numerical control (CNC) on employment and labor productivity in the US, which are thought to have replaced semi-skilled manufacturing workers in metal manufacturing starting the 1970s. We build a novel measure of exposure to this automation shock, and we find that industries more exposed to CNC experienced higher labor productivity, while reducing production employment. The relative demand for college-educated workers and abstract tasks rose. Workers in more exposed labor markets adjusted by shifting from metal manu-

facturing to other manufacturing industries. Some workers in exposed industries returned to school. Union members were shielded from job loss.

Acknowledgements

My academic and personal growth at Princeton would not have been possible without all the support, guidance, help, and love I received from the people around me. I would like to express my gratitude to them here, even though several pages would not be enough to fit my feelings.

I am grateful to my advisors, Leah Boustan and Ilyana Kuziemko, for their guidance, wisdom, and patience. They have always been eager to discuss research with me. I have greatly benefited from this environment, where I learned not only knowledge and intuition in economics but also the way of research and rigorous thinking. The projects I have worked on with them reassured my passion for economics research and launched my career as an economist.

Leah has been my advisor since my second year, and I still remember how detailed and careful comments she gave me for the flash paper in her class. Since then, I knew I would want to be a kind and thoughtful economist like her and eventually an economic historian like her. The "boot camps" we had for our paper on Computer Numerical Control tools taught me the joy of collaborative writing and made me become a better writer.

Ilyana has been such a thoughtful and enthusiastic advisor. Working with her on our NAFTA project was one of the most pivotal moments in my graduate career, which helped me regain confidence in myself and reignite my passion for research. I learned from her the values of passion and persistence in conducting research, which I hope to replicate during my career.

I thank Steve Redding for his support and mentorship. He has come to every single presentation of mine and always offered me encouragement and careful feedback. His advice and comments made my work stronger and made me a better researcher. I remember at the end of my third year, he scheduled a meeting for feedback and showed his "two thumbs

up" to cheer me up after a presentation where I felt I was so lost during my Ph.D. His encouragement and kindness were greatly appreciated.

I have been blessed to be part of the Industrial Relations Section and the economics department. I wish to thank Ellora Derenoncourt, Hank Farber, David Lee, Alex Mas, and Richard Rogerson for their advice and comments. I received a great deal of support from Valerie Ching, Indiana Gaffney, Ellen Graf, Laura Hedden, Jeannie Moore, Julie Newton, Gene Piccola, and Patti Tracey. The first chapter would not have been possible without the help of librarians at Princeton and Harvard, Charissa Jefferson, Melissa Murphy, and Heather Oswald. I acknowledge the financial support provided by the Graduate School, the Industrial Relations Section, and the Prize Fellowship in the Social Sciences.

My dissertation would not have been possible without amazing collaborators Ilyana Kuziemko, Ebonya Washington, Gavin Wright, Leah Boustan, and David Clingingsmith. Their insights and dedication always inspire me. I learned so much from my fellow graduate students in the economics department: Lukas Althoff, Emily Battaglia, Carlos Burga, Ana Dutra, Elisa Jácome, Victoria Larsen, Chris Mills, Chris Szerman, Joao Thereze, and Yinuo Zhang. I thank Lukas for being a great job market buddy. I thank Emily for making the coursework so enjoyable, which includes resisting laughing at the unintended Ke$ha reference for two full days of class. Thank you also for sharing all the job market wisdom. I thank Joao Thereze for sharing his intuition in economic theory and being an amazing roommate. In addition, I thank my dear friends who always had faith in me and helped me get through difficult times, including Charlie Murphy, Jisoo Park, Jeannine Van Reeken, Maria Hernandez, and Sarah Shin. I thank Charlie for being my best friend and roommate, being the best coding/digitization/IT support, and teaching me the joy of living with a Goldendoodle. Thank you, Jisoo, Jeannine, Maria, and Sarah, for always keeping me grounded and making me feel proud that I have the best friends I could ask for.

I wish to thank my parents, Wonjong Choi, and Jeongja Kim. Their love, devotion, and support made it possible for me to pursue my passion. My education in the US would not have been possible without the tremendous support I received from my uncle and aunt, Junghyun In, and Jeongsoon Kim. I thank Janet In and Jon Pawlowski for being my parents in the US. Thanks to them, I always felt safe and supported in the US, and I wish to reciprocate their kindness and generosity back to my godson Nicholas. Finally, I wish to thank my husband, Minjae Cho. From the moment we first met, you always inspired me to be a better human being. I have always been in awe of your passion for Physics and Music. Your encouragement helped me achieve things I never thought I could and led my life to its best. Your kind and optimistic soul has made me a nicer, happier person. Thank you, and I love you.

For my parents, Wonjong Jeongja,

my uncle and aunt, Jung Hyun and Jung Soon,

and my love, Minjae.

Contents

Abstract 3

Acknowledgements 5

List of Tables 13

List of Figures 16

1 The Effect of Deindustrialization on Local Economies: Evidence from New England Textile Towns **22**

1.1 Introduction 22

1.2 Historical background 29

1.3 Conceptual framework on worker responses to textile industry decline 31

1.4 Data and measuring exposure to the textile industry decline 33

1.4.1 Measuring the textile industry decline using the initial reliance on the textile industry 33

1.4.2 Measuring the textile industry decline using plant closures 35

1.5 The effect of the textile industry decline on workers' migration 36

1.5.1 Main specification 36

1.5.2 Identification 37

1.5.3 First-stage relationship between the exposure measure and textile share of the labor force 38

1.5.4 Muted population and migration responses to the industrial decline . 39

1.6 The effect of the textile industry decline on labor force participation, industry switching, and education . . . 44
1.6.1 Heterogeneous effects on labor force participation by worker age . . . 44
1.6.2 Increased industry switching to the agricultural sector and decreased occupational income . . . 45
1.6.3 Increased educational attainment in response to the deindustrialization 47
1.7 Conclusion . . . 53
Appendices . . . 77
A Appendix Figures and Tables . . . 77
B County-level event-study analysis . . . 89
C Matching Algorithm and matched population characteristics . . . 93

2 Local Economic and Political Effects of Trade Deals: Evidence from NAFTA 98

2.1 Introduction . . . 98
2.2 Measuring local vulnerability to NAFTA . . . 103
A Construction of our measure of NAFTA exposure . . . 104
B Characteristics of counties by NAFTA vulnerability . . . 106
C Relationship between industry-level NAFTA vulnerability and Mexican imports to the US . . . 107
2.3 Empirical strategy for event-study analysis . . . 108
2.4 Employment results . . . 109
A Main county-level event-study results . . . 110
B Robustness to other trade and currency events . . . 112
C Addressing inference concerns . . . 114
D Related results and additional robustness checks . . . 114
E Results at the individual level . . . 115

2.5 Migration and transfer-program response . . . 116
A Migration . . . 116
B Trade Adjustment Assistance . . . 118
C Disability Insurance . . . 119
D Other outcomes . . . 120
2.6 The political response in areas vulnerable to NAFTA . . . 121
A Opinions of respondents in NAFTA-vulnerable states . . . 123
B County-level event-study results . . . 124
2.7 The political response among individuals against free trade . . . 126
A Evidence from repeated cross-sectional data . . . 126
B Corroborating evidence from panel data . . . 131
C Connecting the economic and political impacts of NAFTA . . . 133
2.8 Conclusion . . . 135
Appendices . . . 149
A Appendix Figures and Tables . . . 149
B Data appendix . . . 189
C Assessing main results with randomization inference . . . 196
D Accounting for industry-level benefits from NAFTA . . . 212

3 Automation after the assembly line: Computerized machine tools, employment and productivity in the United States 217

3.1 Introduction . . . 217
3.2 Historical Context . . . 222
3.3 Construction of CNC exposure measure and data sources . . . 226
A Data sources for outcome variables . . . 229
3.4 Graphical evidence of CNC adoption and economic effects . . . 231
A Adoption of CNC machines . . . 231

B Trends in economic outcomes before and after the introduction of CNC 232
C CNC and economic outcomes residualized by industry and year effects 233
3.5 Industry-level results: Economic outcomes and worker responses 234
A Capital expenditures, productivity and labor share 234
B Total employment and employment shares by education level 235
C Task composition of the workforce 236
D Worker response through union protections 236
E Educational attainment and returning to school 237
3.6 Local labor market results 238
3.7 Conclusions 240
Appendices 254
A Appendix Figures and Tables 255

Bibliography **264**

List of Tables

1.1 Average town characteristics by initial textile presence 58

1.2 Migration mechanism: wealth increases the likelihood for migration 64

1.3 Homeowners with mortgages are less likely to migrate 65

1.4 Increased cohabitation 66

1.5 The effect of exposure to textile industry decline on labor market outcomes using a matched sample of workers 71

1.6 Average town characteristics for closure towns and control towns 74

1.7 The effect of large textile plant closures on educational attainment, income, and migration 76

1A.1 Coverage of the American Official Textile Directory 78

1C.1 Average characteristics of matched male individuals and population in 1900 94

2.1 Pre-NAFTA characteristics of counties, by vulnerability quartile 145

2.2 Approval of NAFTA as a function of state-level NAFTA vulnerability . . 146

2.3 Partisan identity and views toward free trade, 1986-2012 repeated cross-sections 147

2.4 Partisan identity and views toward free trade, 1992-1994 panel data . . . 148

2A.1 Education predicts less protectionist views 186

2A.2 How protectionist views predict approval of Ross Perot, 1992 and 1996 . . 186

2A.3 Simulating presidential elections ex. NAFTA effect 187

2A.4 Partisan identity and views toward NAFTA, 1992-1994 panel data 188

2B.1 Datasets used in Table 2.2 (NAFTA approval by state-level vulnerability) 195
2C.1 Estimates from the first-difference model: Change in log of employment (1986-1993) and simulated county vulnerabilities 200
2C.2 Estimates from the first-difference model with share of college grads (1990) control: Change in log of employment (1986-1993) and simulated county vulnerabilities 201
2C.3 Estimates from the first-difference model: Change in log of employment (1993-2000) and simulated county vulnerabilities 202
2C.4 Estimates from the first-difference model with share of college grads (1990) control: Change in log of employment (1993-2000) and simulated county vulnerabilities 203
2C.5 Estimates from the first-difference model: Change in Republican presidential vote share (1984-1992) and simulated county vulnerabilities 204
2C.6 Estimates from the first-difference model with share of college grads (1990) control: Change in Republican presidential vote share (1984-1992) and simulated county vulnerabilities 205
2C.7 Estimates from the first-difference model: Change in Republican presidential vote share (1992-2000) and simulated county vulnerabilities 206
2C.8 Estimates from the first-difference model with share of college grads (1990) control: Change in Republican presidential vote share (1992-2000) and simulated county vulnerabilities 207
2C.9 First-difference model with the AKM inference procedure: Change in log of county employment over pre-NAFTA control period (1986-1993) 208
2C.10 First-difference model with the AKM inference procedure: Change in log of county employment over treatment period (1993-2000) 209

2C.11 First-difference model with the AKM inference procedure: Change in Republican presidential vote share over pre-NAFTA control period (1984-1992) 210
2C.12 First-difference model with the AKM inference procedure: Change in Republican presidential vote share over treatment period (1992-2000) 211

3.1 The effect of CNC exposure on industry-level economic outcomes, 1968-2007 253
3.2 The effect of CNC exposure on employment by union status, sample 1973-2007 253
3.3 The effect of CNC exposure on the two- and four-year college enrollment, 1976-2007 253

List of Figures

1.1 Share of the labor force in the textile industry by region 55
1.2 Worker responses to deindustrialization 56
1.3 Textile share of the labor force by county in 1900 57
1.4 Average textile share of the labor force by quartiles of initial textile presence in 1900 59
1.5 The effect of exposure to textile industry decline on the textile share of labor force 60
1.6 The effect of exposure to textile industry decline on net in-migration rate . 61
1.7 The effect of exposure to textile industry decline on working-age population 62
1.8 The effect of exposure to textile industry decline on individual-level likelihood to out-migrate 63
1.9 Heterogeneous labor force participation by age: School-going aged population (age 14-18) 67
1.10 Heterogeneous labor force participation by age: Population of age 28-65 . 68
1.11 The effect of exposure to textile industry decline on the share of agricultural labor force 69
1.12 The effect of exposure to textile industry decline on occupation-based earning 70
1.13 Share of the population attending school in 1920 by state-age cohort . . . 72
1.14 Normalized trend of share completing at least the eighth grade 75
1A.1 Share of the national labor force in the textile industry located in each region 77
1A.2 Structure of the Official American Textile Directory 78

1A.3 The effect of exposure to textile industry decline on the log of native working-age population . . . 79
1A.4 The effect of exposure to textile industry decline on labor force participation rate . . . 80
1A.5 The effect of exposure to textile industry decline on labor force participation by age group . . . 81
1A.6 Effect of exposure to textile industry decline on school attendance rate among age 14-18 population . . . 84
1A.7 The effect of exposure to textile industry decline on the share of the population (age 11-14) with the eighth-grade completion . . . 85
1A.8 Effect of exposure to textile industry decline on teacher per capita . . . 86
1A.9 Map of matched control towns and plant closure towns . . . 87
1A.10 Normalized trend of log wage income . . . 88
1B.1 The effect of exposure to textile industry decline on county working-age population . . . 89
1B.2 The effect of exposure to textile industry decline on county labor force participation rate, school-going aged population (age 14-18) . . . 90
1B.3 The effect of exposure to textile industry decline on the county-level share of agricultural labor force . . . 91
1B.4 The effect of exposure to textile industry decline on county average occupation-based earning . . . 92
1C.1 The effect of exposure to textile industry decline on working-age population using the matched male individuals only . . . 95
1C.2 The effect of exposure to textile industry decline on the agricultural share of the labor force using the matched male individuals only . . . 96

1C.3 The effect of exposure to textile industry decline on occupational earnings using the matched male individuals only 97

2.1 Variation in protection and NAFTA exposure, across time and space . . . 137
2.2 Log employment as a function of county NAFTA vulnerability 138
2.3 Log population as a function of county NAFTA vulnerability 139
2.4 Trade Adjustment Assistance petitions per capita as a function of county NAFTA vulnerability 140
2.5 Log of DI applications as a function of county NAFTA vulnerability . . . 141
2.6 Presidential election vote shares as a function of county vulnerability . . . 142
2.7 Party identification (increasing in Republican direction) as a function of views toward trade 143
2.8 Heterogeneity in the shift toward GOP after 1992 among protectionist respondents 144
2A.1 U.S. imports from China and Mexico 149
2A.2 Import values from Mexico by 1990 tariff level 150
2A.3 Relationship (in levels) between Mexican imports to the US and pre-NAFTA tariffs 151
2A.4 Relationship (in logs) between Mexican imports to the US and pre-NAFTA tariffs 152
2A.5 Average log employment for four vulnerability quartiles over time (normalized to zero in 1993) 153
2A.6 Log employment as a function of county NAFTA vulnerability, varying controls 154
2A.7 Log employment as a function of NAFTA vulnerability, robustness to Multi-Fibre-Arrangement phase-out 155

2A.8 Log employment as a function of county NAFTA vulnerability, robustness to automation, off-shoring and immigration . . . 156
2A.9 Employment in textile mills, 1990-2020 . . . 157
2A.10 Employment per capita as a function of NAFTA vulnerability . . . 158
2A.11 Log of CZ employment as a function of CZ-level NAFTA vulnerability . . 159
2A.12 Log of employment as a function of NAFTA vulnerability using imputed CBP cells from Eckert et al. (2021) . . . 160
2A.13 Evolution of log employment as a function of NAFTA vulnerability, separating manufacturing v. other industries . . . 161
2A.14 Log of employment as a function of county NAFTA vulnerability, robustness to Peso crisis . . . 162
2A.15 Log total wage bill as a function of NAFTA vulnerability . . . 163
2A.16 Employment as a function of individual vulnerability, PSID sample . . . 164
2A.17 Log population as a function of county NAFTA vulnerability, varying controls 165
2A.18 Log population as a function of NAFTA vulnerability, controlling for long population changes . . . 166
2A.19 Migration response to NAFTA, adjusting for NAFTA vulnerability of a county's pre-existing migration network . . . 167
2A.20 Trade Adjustment Assistance petitions per capita, by county *Vulnerability* quartile (normalized to zero in 1993) . . . 168
2A.21 Trade Adjustment Assistance petitions per capita as a function of county NAFTA vulnerability, varying controls . . . 169
2A.22 Trade Adjustment Assistance *certifications* per capita as a function of county vulnerability . . . 170
2A.23 Evolution of log employment as a function of county NAFTA vulnerability, for a balanced panel of 766 counties for which we have DI application data 171

2A.24 Log DI applications, raw trends by four vulnerability quartiles (1993 normalized to zero) . . . 172
2A.25 Log of DI applications as a function of county NAFTA vulnerability . . . 173
2A.26 Evolution of log DI final *awards* as a function of county vulnerability . . . 174
2A.27 Log of Unemployment Insurance benefits by county vulnerability, normalized . . . 175
2A.28 Evolution of log UI benefits as a function of county vulnerability . . . 176
2A.29 Log reported SNAP benefits, normalized . . . 177
2A.30 Log reported SNAP benefits as a function of county vulnerability . . . 178
2A.31 Coverage of trade-and-jobs related stories by network nightly news programs 179
2A.32 Republican vote share in Presidential elections, separately by vulnerability quartile (raw means, not normalized) . . . 180
2A.33 Democratic vote share in Presidential elections, separately by vulnerability quartile (raw means, not normalized) . . . 181
2A.34 Presidential election vote shares as a function of county vulnerability, varying controls . . . 182
2A.35 Presidential election vote shares as a function of county vulnerability, additional white and college-share controls . . . 183
2A.36 Republican two-party share of House-election votes, by county NAFTA vulnerability . . . 184
2A.37 Heterogeneity in 1992-1994 shift toward GOP among protectionist respondents . . . 185
2C.1 Empirical distribution of actual industrial tariffs in 1990 . . . 199
2D.1 Log county employment as a function of county vulnerability, robustness to controlling for county export advantage from NAFTA . . . 216

3.1 Relative Value Share of Installed Tools by Seven Metal Manufacturing Industries in 1958 . . . 242
3.2 CNC Share by Machine Tool Type . . . 243
3.3 Cumulative CNC Share by Metal Manufacturing Industries . . . 244
3.4 CNC Adoption and CNC Exposure . . . 245
3.5 CNC Share by Machine Tool Type . . . 246
3.6 Core Economic Outcomes and CNC Shock Residualized Binscatters . . . 247
3.7 CNC-Related Degree and Program Completions . . . 250
3.8 CNC Exposure and Local Labor Markets . . . 251
3A.1 Union Share by Seven Metal Manufacturing Industries . . . 255
3A.2 CNC-Related Patent Counts by Tool and Country . . . 256
3A.3 CNC Share by Machine Tool Type for Italy . . . 257
3A.4 Average Task Scores of Occupations by Seven Metal Manufacturing Industries 258
3A.5 CNC-Related Degree and Program Completions by Degree Type . . . 260
3A.6 Import Penetration by Seven Metal Manufacturing Industries . . . 263

Chapter 1

The Effect of Deindustrialization on Local Economies: Evidence from New England Textile Towns[1]

1.1 Introduction

Deindustrialization, a decline in industrial activity in a region or local economy, is a common process of economic development. During industrialization, agglomeration forces often create local labor markets with concentrated industrial activity, and the industrial boom improves residents' economic well-being. However, industries may face a decline or decide to relocate for various reasons, including an increase in trade (Autor et al., 2013, Batistich and Bond, 2022) and an increase in outsourced production for cheaper costs (Fort et al., 2018,

[1]I thank Leah Boustan and Ilyana Kuziemko for their support and guidance throughout the project. I thank Lukas Althoff, David Autor, Emily Battaglia, Ellora Derenoncourt, Claudia Goldin, Stephanie Hao, Shawn Kantor, Ezra Karger, Chris Meissner, Chris Mills, Stephen Redding, Chris Szerman, Neil Thakral, Linh Tô, Yinuo Zhang for their comments. I thank Heather Oswald and Melissa Murphy at Harvard Baker Library for supporting data access, Minjae Cho for data collection, and Charlie Murphy for data digitization. I benefited from valuable comments from seminar participants at the Princeton IR section student seminar, Princeton Prize Fellowship student seminar, and NBER Development of the American Economy Summer Institute Ph.D. session. I am grateful for the financial support from the Prize Fellowship at Princeton School of Public and International Affairs.

Antràs et al., 2017). When an industry leaves a local labor market that has specialized in that industry, the local labor market experiences a decrease in employment opportunities.[2] Understanding how local labor markets and residents react to deindustrialization, for instance how local population adjusts and economic activities are restructured, has important policy implications for regional economic recovery and growth.

Migration has been a highlighted channel of swift adjustment to local economic shocks under the spatial equilibrium framework. However, recent literature has found muted migration response and lagged local recovery. This pattern is often attributed to modern policy changes that can reduce mobility, such as an increase in social welfare benefits (Notowidigdo, 2020, Coen-Pirani, 2021).

In this paper, I document the persistent local labor market effects of deindustrialization from a setting with limited social insurance. In particular, I examine one of the first significant yet under-explored episodes of deindustrialization in the US, the decline of the New England textile industry in the 1920s. New England's textile industry was once the center of industrial development in the US, the site of one of the country's first factory systems. While the industry employed more than 15 percent of the New England labor force in 1920 and roughly 30 percent of the US textile labor force was in the region, the industry faced a massive disruption starting in the 1920s. In over one decade, between 1920 and 1930, the industry shed roughly twenty percent of its labor force.[3]

The size and speed of the shock were comparable to that of the Rust Belt decline: Michigan experienced roughly a 5 percentage-point decline in the car manufacturing share of the labor force between 1970 and 1980 and a 9 percentage-point drop in the share from

[2]This pattern is observed in many examples of once-booming local labor markets. For instance, the Rust Belt and the Sun Belt manufacturing towns were factory towns with specialized industrial bases that faced a long process of economic decline.

[3]The decline of the New England textile industry occurred in two phases, one during the 1920-1930s and another during the 1950-1960s. In this paper, I focus on the first phase of the decline.

1970 to 1990.[4] New manufacturing industries did not emerge in New England until the 1980s, in a manner similar to the lack of re-industrialization in the Rust Belt.[5]

New England's rapid industrial decline is an empirically advantageous setting for studying long-run worker responses to deindustrialization and outcomes. I link full-count historical censuses using an automated method from Abramitzky et al. (2022) to follow residents over time. This allows me to examine various worker adjustment margins, including migration, labor force participation, education, and industry-switching. I am also able to explore whether workers' responses translated to local economic recovery or improved workers' labor market outcomes.

To examine whether and how the textile industry decline in New England states impacted town-level and individual-level outcomes, I construct a town-level measure of exposure to the textile industry decline using the textile share of the labor force in 1900 and compare towns with more or less exposure using a continuous difference-in-differences framework. The empirical strategy relies on the identifying assumption of parallel trends, which is examined for each response margin using event-study analysis. I also flexibly control for the potential confounders, such as the Great Depression or change in immigration policy that may have affected the economic conditions of textile-heavy towns after 1920.

This paper highlights several patterns of worker responses that provide insights into policies for local economic recovery. First, I document the muted migration response from the local shock, even in the absence of sufficient social insurance: I find that both the working-age population and net in-migration rate were not significantly affected by the exposure to the industrial decline. I have the power to reject even a small population decline.

[4]The Rust Belt states, spanning Indiana, Illinois, Michigan, Missouri, New York, Ohio, Pennsylvania, West Virginia, and Wisconsin, experienced more than a 3 percentage-point decline in the car manufacturing share of the male labor force between 1970 and 1990.

[5]Parts of Connecticut and Massachusetts were home to machinery and metalworking industries since the 1880s that benefited from the war production act in the 1940s, but most parts of New England states went through a continued decline in manufacturing activities until the late twentieth century other than during WWII (Rosenbloom, 1999).

I track male individuals using linked census data and examine the effect of industrial decline on these workers regardless of where they live after the shock. This individual-level analysis isolates worker-level effects from the compositional effects measured from the cross-sectional census data, which may mask the in- and out-migration dynamics. I find that workers who were initially located in a town with a greater share of textile presence are *less* likely to out-migrate from their initial town.

I provide suggestive evidence that the fixed cost of moving may have played a role in the migration decision, as wealth measured by homeownership and mortgage status was positively associated with the migration likelihood. Also, individuals in textile-heavy towns increased cohabitation or living in multigenerational housing, documenting a potential role of social networks (e.g., family support) in weathering the reduced economic opportunities.

Second, I show that young and older residents were both likely to stay in their affected towns but responded differently to the local economic shocks. Young residents of age 14-18 reduced their labor force participation in response to the industrial decline. Child labor was a significant part of manufacturing employment in the 1920s, implying that the decline in labor force participation among young residents may be a positive result for their educational attainment and future labor market outcomes. Adult workers did not significantly reduce their labor force participation.

I document that older workers in towns responded by switching to the agricultural sector, which had on average lower occupational earnings than the textile industry. My findings, together with the limited migration results, suggest if residents do not leave their local labor market experiencing a downturn, place-based policies that aid the transition of workers to different industries or promote diversity in the industrial base will help the local labor market resilience.

Lastly, I explore how younger residents reacted to the local labor market shock. To do so, I build a dataset on major plant closures in New England towns from newly collected

historical data. The dataset is created by first digitizing all of the records on New England plants found in the 1920-1930 Official American Textile Directories. For each plant, I digitize the plant name, the town it is located in, and the types and number of machinery. Each plant is geolocated based on its town of record and assigned to its closest census town. By tracking plants across time, I define major plant closures as the closure of plants that are large relative to the town population and plant closure towns as towns that experienced at least one major plant closure. I then use a propensity score matching procedure to create matched control towns with similar characteristics to plant-closure towns. Using linked census data and matched difference-in-differences design, I exploit the variation in the timing and location of plant closures and the variation of exposure *within* towns based on "critical" age of education to examine individual-level responses to the textile industry decline.

Using the plant-closure design, I find that large plant closures in towns are associated with a higher likelihood of completing the eighth grade for young individuals by 1940.[6] In New England, the eighth grade was the modal grade for individuals choosing between finishing education or entering the workforce in the 1920s. The plant closure did not affect the likelihood of attending or completing the high school degree. The increased schooling reflects the reduced opportunity cost of education from the shock.

The young individuals in closure towns who obtained additional schooling, however, were not more likely to out-migrate from their original towns, and they were more likely to face reduced wage income. The lack of migration response from young individuals affected by plant closure indicates that they were negatively affected by the same deindustrialization shock despite their increased education. The returns to schooling were low in the 1940s (Feigenbaum and Tan, 2020).[7] In the absence of an industry that offers high returns to

[6]This is consistent with the findings from the town-level event-study analysis where I document that exposure to the textile decline was associated with a reduced labor force participation rate and increased school attendance.

[7]Feigenbaum and Tan (2020) estimate that each additional year of schooling raises earnings by 4 percent in 1940, compared to the modern estimates of 6 to 15 percent (Card, 1999).

education and in the absence of migration response, their educational attainment may lead to a decreased level of experience and lower wages conditional on age (Mincer, 1974; Lemieux, 2006).

Related Literature

This paper contributes to the extensive literature on local labor market recovery in two ways. First, this paper provides some of the first evidence of the long-lasting effect of deindustrialization from the early twentieth century. Second, this paper thoroughly explores various mechanisms behind such persistence.

A key prediction from spatial equilibrium models is that local economic shocks dissipate in equilibrium with labor and capital mobility and lead to short-run regional economic recovery. Blanchard and Katz (1992) and Dao et al. (2017) find a transient nature of local economic shocks, but literature on more recent shocks has documented the long-term scarring effects and muted population response from various modern local economic shocks, especially in the US (Autor et al., 2013; Dix-Carneiro and Kovak, 2017; Yagan, 2019a; Pierce and Schott, 2020a; Choi et al., 2021)[8].

One hypothesis for the mixed evidence for migration responses to local shocks is the changing migration response to local shocks over time. The studies finding rapid local recovery examine the shocks in the 1970s while the others study the shock in the 1990s and 2000s, which was a period with decreased mobility and an increase in social benefits that may have affected the worker migration decision. My findings suggest that the pattern of lagged local economic recovery is not a modern phenomenon and document the persistent effect of local shocks in a period without social insurance. This paper further highlights the potential role of wealth on migration and the role of social networks as a worker adjustment

[8]Even in the setting of the 1970s, Hershbein and Stuart (2022) find the lagged local recovery from recessions in the 1970s, similar to my finding that documents the persistence of shock in the US before the 1990s.

margin to shocks, using the linked census data and rich individual-level information from the census, including homeownership, mortgage status, and habitation patterns.

Economists have conventionally focused on migration as a critical channel of mitigating local economic shocks based on spatial equilibrium theories, This paper provides empirical evidence of alternative mechanisms for local labor market adjustment, including educational attainment and industry-switching behavior.

This paper extends our knowledge about the relationship between education and local economic conditions. First, I document that the local bust increased the educational attainment of the residents, consistent with the overall finding in the literature that a negative local economic shock increases educational attainment and a local economic boom decreases schooling (Jaworski, 2014; Black et al., 2005; Charles et al., 2018; and Atkin, 2016).[9] Second, my analysis further tracks groups of individuals with increased educational attainment and examines whether their increased schooling improved the labor market prospects.

The paper contributes to the economics and history literature studying industrialization in the US. Many papers focus on examining the effect of deindustrialization on workers and the local labor market from the decline of the Rust Belts (Hobor, 2013; Alder et al., 2014; Ohanian, 2014). I fill in the gap in the literature by offering a systematic analysis of deindustrialization in New England mill towns, which was an essential part of nineteenth-century industrialization in the US. Examining the deindustrialization episode also contributes to the literature on regional development and convergence in the US (Barro and Sala-i-Martin, 1992; Eriksson et al., 2021), as this deindustrialization was an essential driver of the rise of the Southern economy and the slowdown of the Northeastern economy during the early twentieth century.

[9]Stuart (2022) documents a decline in the 4-year college attainment in response to the Great Recession. A potential distinction between other papers and Stuart (2022) is the cost of education—attending a four-year college is much more costly than attending high school or community college. My result is consistent with the prediction that local bust will increase the take-up in tuition-free education.

1.2 Historical background

The New England states, including Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, and Vermont, played an essential role in the US textile industry development and industrialization in the eighteenth and nineteenth centuries, and the industry remained a necessary part of the New England economy until the early twentieth century. In the late eighteenth century, merchants from England realized that the New England states' geography with rivers and streams was advantageous for building water-powered plants. Some New England towns were often the first form of "factory" or "company" towns in the US, which were one of the prevalent forms of industrialization in the US.

The New England states, starting in Rhode Island and Massachusetts, had one of the first industrialized systems for textile manufacturing, such as the Rhode Island mill system, which employed families, including young children, and the Waltham-Lowell system hired young female workers. Figure 1.1 shows that more than 15 percent of the labor force in the Northeastern region was employed in the textile industry in 1920. Even though the emergence of the steam engine made geography less critical, the agglomeration forces kept the industry in the area for a long time. Appendix Figure 1A.1 plots the share of the textile workforce in each census region, and more than 30 percent of the national textile labor force was employed in the Northeastern part of the country.

The textile industry decline was a large shock to the New England states; over a decade from 1920 to 1930, the industry shed around 5 percentage points of its labor force, as shown in Figure 1.1. The level of the shock geographically varied within the states, creating a more considerable impact on mill towns that were developed around and relied on the textile industry. For instance, Manchester in New Hampshire, a town designed by the Amoskeag Manufacturing Company, started experiencing plant closures in 1922. The mills had once

employed more than 15,000 workers, but all of them closed by 1935, leaving many town residents unemployed (Minchin, 2012a).

Multiple reasons led to the textile industry decline in the New England states. The importance of the southern textile industry rose from the late nineteenth century, as shown in Appendix Figure 1A.1. First, proximity to raw cotton and cheaper labor costs made the South an ideal location for the cotton industry. At the beginning of the rise, New England textile merchants were able to keep their competitiveness by producing higher-quality fabrics. However, with the adoption of newer textile machinery, the Southern manufacturers gradually caught up with the New England producers on the range and quality of goods they produce. Second, with the end of World War I, which reduced the foreign demand for cotton, the New England cotton industry faced a more significant decline (Rosenbloom, 1999; Koistinen, 2002). Other textile industries, such as the woolen industry, experienced relatively less decline, even though the end of foreign demand affected these industries, too (Scully, 1951).

The decline of the textile industry led to some community-level efforts to mitigate the economic impact, but the role of local and state governments was minimal in the 1920s and 1930s. Even though there was an effort to introduce state-level Unemployment Insurance, there was no social safety net nor formal unemployment insurance in the New England states until 1935.[10] During the Great Depression, which coincided with the industrial decline in the 1930s, government expenditure increased with the New Deal policies. Still, the New England states were not the primary beneficiaries of the policies (Rosenbloom, 1999).

Instead, business and financial leaders in the states established the New England Council in 1925 as a response to the decline of the textile industry. The council tried to facilitate financial access to local textile manufacturers and promote the adoption of advanced manage-

[10]New England states were aware of the need for the social safety net, however. New England states attempted to create unemployment statistics in the late eighteenth century. Massachusetts introduced the Unemployment Insurance bill in 1916, which was never passed until the passage of the Social Security Act of 1935.

rial techniques and industrial research to increase product competitiveness. These were not helpful for short-term solutions in preventing mills from closing or mitigating unemployment in the area (Koistinen, 2000).

Labor unions of this era are known to be small and have weak organizational power, but the textile worker unions in some states, including Massachusetts, were able to protect labor-friendly legislations such as hours-of-work laws. The efforts of organized labor did not particularly help the recovery efforts, as these laws kept the labor cost higher in the region (English, 2003; Koistinen, 2007).

1.3 Conceptual framework on worker responses to textile industry decline

Given there was limited government welfare or economic policies to counter the textile industry decline, how the New England local labor markets adjusted during and after the deindustrialization mainly depended on workers' responses to the shock. Figure 1.2 summarizes how workers could respond to the decline of local industry.

First, workers could respond to deindustrialization by migrating out of their local labor markets for better economic opportunities. The migrant workers can increase employment probability and wages in the new local labor markets. The migration decisions can lead to an overall population decline in the original local labor market, decreasing local labor supply and increasing wages, serving as a mechanism for mitigating local economic shocks. This channel has been considered the key mechanism through which local labor markets recover from shocks under the spatial equilibrium models.

Workers, however, can also adjust *within* their initial location and adapt through non-migratory responses. Deindustrialization often decreases local labor demand, decreasing the number of available jobs and reducing local wages. Workers can give up seeking employment

opportunities and detach themselves from the labor market. This will lead to a decline in the labor force participation rate in the local labor market.

Alternatively, if other industries in local labor markets can accommodate exposed workers, workers actively seeking employment may succeed in securing a job in a different sector. Workers in the affected local labor market may switch their industries in response to deindustrialization, and the share of the workforce employed in other sectors will rise in these local labor markets.

For young individuals, deindustrialization can affect their educational attainment decision. Under a simple theoretical framework of education decision, similar to the educational choice model from Atkin (2016), deindustrialization in a local labor market could alter the opportunity cost of schooling and increase the educational attainment of young residents. In this model, the decision for educational attainment depends on the returns to schooling and the opportunity cost of schooling.

The opportunity cost of schooling is determined by the utility difference between working and staying at school for one unit of time. The expected utility from working decreases with the industrial decline, which lowers the opportunity cost of schooling. The utility from schooling includes utility from family support and direct utility from education. The utility is lower for older workers, as their earnings are more likely to be the primary source of household income. The psychological cost (disutility) of schooling or the cost of "going back" to school often increases with age. Thus, young workers may be more likely to choose to increase education in response to local economic decline than their older counterparts. This framework highlights that the educational attainment decision could change with local economic conditions, even without the change in marginal returns to schooling.

In Sections 1.5 and 1.6, I examine how the textile industry's decline affected each adjustment margin in more detail.

1.4 Data and measuring exposure to the textile industry decline

In this section, I define town-level exposure measures to the textile industry decline and discuss their variations. First, I construct a measure of the town's exposure to the textile industry decline using its initial textile presence, which will be used to analyze all adjustment margins. For the educational attainment outcome, I construct an additional exposure measure by identifying whether each town experienced major plant closures using the newly constructed panel of New England textile plants, which allows me to exploit the quasi-experimental variation in location and timing of the plant closures.[11]

1.4.1 Measuring the textile industry decline using the initial reliance on the textile industry

I construct a measure of town-level exposure to the textile industry decline that relies on the variation in the presence of the textile industry in 1900, before the regional industry decline. I use the microdata from the 1900 U.S. Census to construct the measure. I start by defining towns based on the 1329 places in New England states drawn from the Census Place Project (CPP) (Berkes et al., 2022). The places are constructed using census sub-county geographic variables, including city and Minor Civil Divisions (MCDs). MCDs are

[11]There are multiple reasons that I use the plant closure measure only in the education analysis. First, education is the only outcome that can fully utilize the variation in the timing of plant closures. Education has the notion of critical age– there is a notion of when individuals are likely to receive a specific stage or grade of education. This allows for creating a control and treatment group *within* plant closure towns for the educational outcome. At the same time, other response margins, such as migration, have a less clear distinction on what age individuals are more likely to utilize the specific adjustment margin. Second, the educational outcome can be observed using the highest grade completed variable in the 1940 census, no matter when each individual completed the education. Other variables are less precisely estimated as they are constructed using a change over one or two decades. Using migration as an example, if two individuals faced plant closure in their town in 1925 and 1927, they might have migrated by 1940, not being able to know when they migrated at what age. The effect of plant closure on migration would then be less precisely estimated.

administrative units of counties, often representing a civil township or local government. I create town-level economic statistics using crosswalk from Berkes et al., 2022, which maps individuals recorded in the census to these places. The sub-county geographic units allow us to utilize heterogeneity in demographics or economic structure, including distinctions between urban and rural areas within a county (Michaels et al., 2012; Berkes et al., 2022).

For each town, the exposure measure is defined by the share of the labor force working in the textile industry in 1900. The CPP towns are defined from combinations of multiple sub-county geographies and detection of population clusters around the geographies, which makes it hard to delineate town boundaries (e.g., GIS shapefiles). For this reason, even though the exposure measure varies by town defined in CPP, I instead show the county-level geographic variation of the New England textile presence in Figure 1.3. The parts of Connecticut, Massachusetts, and Rhode Island exhibited the highest concentration of textile industry in 1900.

Characteristics of towns by the initial textile presence

Table 1.1 summarizes the average characteristics of towns by terciles of initial textile presence, with the average traits of towns with no textile employment in the first column. Towns in the highest textile share quartile in 1900 tend to have a higher population than towns with lower or no textile presence, as presented in Table 1.1. Textile-heavy towns tend to have a higher labor force participation rate as of 1900, around 4 percentage points more, and a higher share of foreign-born individuals. Half of New England townships reported no textile employment in 1900. These towns tended to be more rural, with a larger percentage of agricultural workers.

1.4.2 Measuring the textile industry decline using plant closures

The textile industry decline can be measured by identifying extensive plant closures. To identify the location and timing of plant closures, I digitize editions of the Official American Textile Directory from 1920 to 1930. The textile directory is an annual listing of textile mills in the United States compiled and published by the Textile World Journal. The directory is similar to typical industry-specific directories. It contains advertisements and detailed plant-level information on location and production capacities, enabling producers and intermediaries in the industry to market their products and facilitate commerce.

Appendix Figure 1A.2 presents an example of a plant-level description in the directory. The detailed plant information includes plant name, location and size of the town or city the plant belongs to, types of materials the plant handles (e.g., cotton, wool, silk), types of products (e.g., hosiery, clothes), and production capacity proxied by the number of textile machinery such as spindles, looms, and water wheels.

Using the data, I track plants that closed and disappeared during the 1920 to 1930 editions. A plant is considered closed in a year if it appeared in the year but not in any year after, or if the directory explicitly stated the plant had closed. In particular, I focus on identifying "large" plant closures, which were closures of a sizeable plant compared to the town population. I define large plants based on the number of machines per town capita and describe any plant with more machines per capita than the highest quartile number of looms or spindles per capita across all towns. Using this definition, I identify the location of 224 large plants and the timing of 84 closures that happened by 1930.

Towns and cities in the textile directory do not align with the census geography, such as census city, county, or Minor Civil Divisions. I geocode towns from the directories to assign latitude and longitude coordinates and assign them to their closest census places (towns) to identify census towns that experienced large plant closures. I define plant closure towns as

towns that experienced at least one major plant closure, and I assign each town its closure year as the earliest year the town experienced a large closure.

Appendix Table 1A.1 summarizes the number of cotton textile plants covered in the 1929 Textile Directory. I also compare the coverage with the 1929 Census of Manufactures. I find that the textile directory covers *more* plants than the Census of Manufactures. The Census of Manufactures covers plants with an annual output of more than 5000 dollars, which may explain the difference in the coverage.[12]

1.5 The effect of the textile industry decline on workers' migration

This section examines how the textile industry's decline affected the town population and individual workers' migration decisions. I start the section by discussing the main specification for the analysis, identifying assumptions, and the first-stage relationship between the initial textile presence and the decline in textile share. Then, I present results on the muted population and decreased out-migration response to the industrial decline. The section discusses potential wealth and social network mechanisms explaining the muted migration response.

1.5.1 Main specification

Using the town-level initial share of the textile labor force in 1900 as the treatment variable, I employ a continuous difference-in-differences strategy and compare changes in economic outcomes between towns with higher textile exposure and towns with lower textile exposure.

[12]The Census of Manufactures is not available between 1890 and 1929, which makes it hard to use the CM as the data source to study the deindustrialization and plant closure episodes in 1920-1930 period.

I use the cross-sectional census data from 1900 to 1940 and estimate the following equation:

$$Y_{mt} = \alpha_m + \gamma_t + \sum_{\tilde{t} \neq 1920} \beta_t \left(\text{Textile Share}_{m,1900}\right) \times \mathbb{1}\left(t = \tilde{t}\right) + \mathbb{X}_m \times \lambda_t + \epsilon_{mt} \tag{1.1}$$

where Y_{mt} is the average outcome in town m in census year t, which includes population, labor force participation, the agricultural share of the labor force, occupational-based earning score, and educational attainment. The specification controls for fixed characteristics of towns and years using town and year fixed effects, α_m and γ_t, respectively. β_{1930} and β_{1940} are our parameters of interest, which are coefficients for the interaction between the initial textile presence of town m in 1900, Textile Share$_{m,1900}$, with indicators for the post-decline census decades, $\mathbb{1}\left(t = \widetilde{t}\right)$.

X_m includes a set of town characteristics which is further described in Section 1.5.2, and λ_t are time-varying coefficients estimated from X_m interacted with census-decade indicators. All specifications include a set of interactions between state and census year indicators to account for differences in town sizes across states due to the administrative nature of Minor Civil Divisions. I cluster standard errors at the town level.

1.5.2 Identification

My empirical strategy relies on the identifying assumption of parallel trends: towns with high and low textile presence had similar trends in the outcome before 1920, the onset of the industrial decline. I examine parallel trend assumptions and address identification concerns in multiple ways. First, in order to provide evidence for parallel trends assumption, I plot coefficients β_t for pre-decline census years in all analysis estimating Equation (1).

Second, secular economic trends may have affected the economic conditions of textile-heavy towns after 1920, potentially causing a violation of the assumption that towns with differential textile presence would have followed similar trends if there were no textile de-

cline. Potential confounders during the post-decline periods include the Great Depression (1929-1939) and changes in immigration policy in the 1920s, which limited immigration flow to the US through country-specific quotas. In the analysis, I include a measure of the Great Depression severity using a county-level change in retail sales during 1929-33 drawn from Fishback et al. (2005) in the analysis as a flexible control interacted with census year indicators, so towns with differential exposure to the Great Depression could have their trends. Similarly, I also include a flexible control of the share of the foreign-born population in 1900 to control for the effect the change in immigration policy could have on the town's economic outcomes.

1.5.3 First-stage relationship between the exposure measure and textile share of the labor force

I start by presenting the first-stage results on whether the textile industry presence in 1900 properly measures how each town was affected by the textile industry decline that began in the 1920s. First, I plot the trend of the average textile share of the labor force for four quartiles of towns based on the textile share of the labor force in 1900 in Figure 1.4. The plot highlights two trends in the textile industry in New England: (1) the textile industry employment is heavily concentrated in towns in the highest quartile of initial textile presence; and (2) the textile industry experienced booms before the decline in the 1920s.

In Figure 1.5, I utilize variation across all towns instead of the four town quartiles to examine the relationship between initial textile industry presence and the textile share of the labor force in towns by estimating Equation (1). The first series plots the β_t coefficient estimates from a specification with state times year fixed effects. The second and third series plot the coefficient estimates from specifications that add to the first series the Great Depression severity measure interacted with year fixed effects and share of the foreign-born population in 1900 interacted with year fixed effects, respectively. In all specifications, I

find that towns with a higher initial percentage of the textile labor force are more likely to experience a decline in the textile share of the workforce between 1920 and 1940, confirming the validity of the measure as a proxy for exposure to the industrial decline. The coefficient of β_{1940}, -0.4, implies that a 10-percentage point increase in initial textile presence, comparable to the 90-10 gap, is associated with a 4-percentage point decline in the share of the labor force working in the textile industry by 1940.

The coefficients for the pre-decline census period, 1900 and 1910, show a rising trend, which is reflective of the boom the New England textile industry experienced during the period, as shown in Figure 1.1. It is important to note that the coefficients on the pre-decline period on this graph indicate the industrial trend in the New England states.

1.5.4 Muted population and migration responses to the industrial decline

Net in-migration rate and log of working-age population

Towns with a larger textile industry presence experienced a decline in labor force participation and earnings following the textile industry decline starting in the 1920s. Workers in these towns may respond to these local shocks by out-migrating to towns with better economic opportunities. The out-migration response of affected towns could lead to a decrease in population.

I examine the migration response by estimating the effect of exposure to industrial decline on a town's net in-migration rate over time in Figure 1.6, which is defined by the change in the log of the working-age population over a decade divided by the log of the initial working-age population. The industry decline has no significant effect on the net in-migration. As discussed in Section 1.5.2, controlling for the share of the foreign-born population is crucial for population analysis. There is no significant relationship between the initial textile industry

presence and the working-age population in the preferred specification with the foreign-born share times year fixed effects.

Figure 1.7 plots the relationship between the industrial decline and town working-age population.[13] The most negative value among the lower end of the confidence intervals for β_{1940} is -0.09. I can reject with 95 % confidence a population decline between 1920 and 1940 larger than 0.9 log points associated with the 90-10 gap in the textile decline exposure measure, a much smaller size of decline compared to the textile industry decline measured in Section 1.5.3.

My preferred unit of geography is towns defined from sub-county geographic units (e.g., MCDs) as I can exploit more variation in the industrial structure within counties. However, as there could be commuting across towns and within counties, counties may also serve as the definition of local labor markets. In Appendix B, I show the event-study analyses using counties as the geographic units. The county-level analyses yield analogous results, even though the coefficients are less precisely estimated as there are 68 counties in New England.

The textile industry decline decreased individual's likelihood of out-migration

The cross-sectional results in Figures 1.7 and 1.6 indicate the population level was not affected by the decline in the textile industry. The results reflect the net effect of the textile industry decline, including the compositional effect. To directly measure the impact on individual workers who experienced the textile decline in a town, I track individuals using linked census data and examine the effects of industrial decline on these workers regardless of where they were after the shock. Then, I examine how the textile presence in individuals' initial location affects their out-migration decisions ten years later. In particular, I estimate the following equation:

[13] Appendix Figure 1A.3 plots the effect of the textile industry decline on the log of the native working-age population and finds a qualitatively similar result as Figure 1.7.

$$Y_{i,m_t,t} = \alpha_{m_{t-10}} + \gamma_t + \sum_{\tilde{t} \neq 1920} \beta_t \left(\text{Textile Share}_{m_{t-10},1900}\right) \times \mathbb{1}\left(t = \tilde{t}\right) + \epsilon_{i,m_{t-10},t} \tag{1.2}$$

where m_{t-10} is the individual's initial town of residence, $\alpha_{m_{t-10}}$ include town fixed effects based on one's initial town and γ_t include year fixed effects. The sample comprises the working-age male population in both census years, between 28 to 65 years old in the destination years.

The automated census linking methods in Abramitzky et al. (2022) are currently able to match men only, as women often change their last names after marriage. Thus, in this analysis, I only match male individuals in New England states recorded in each decadal census from 1900 to 1930 to the next decadal census from 1910 to 1940. The matching method, match rate, and matched population characteristics are further discussed in Appendix C. Appendix Table C.1C.1 shows that the matched male population in New England has very similar features to the overall male population, with close average age, average level of literacy, labor force participation rate, and share in textile labor force with those of all males in New England. The exception is that the matched group is less likely to be foreign-born than the male population.[14] I confirm that the difference in foreign-born share does not drive discrepancy in town-level results and matched individual-level analysis results, by running the town-level event-study analysis using the matched male population in Appendix C.

Figure 1.8 shows the estimation result of Equation (2). The plot summarizes two key results. First, during the pre-shock period between 1900 and 1910, individuals in textile-heavy towns were not more likely to out-migrate than individuals in towns with lower textile reliance, compared to the out-migration differences over one decade between 1910 and 1920. After the textile shock, from 1920 to 1930 and from 1930 to 1940, individuals in textile-heavy

[14] Abramitzky et al. (2022) documents that immigrant groups have lower match rates than native men, as rarer last names with transcription errors can lead to higher false positive rates.

towns were significantly less likely to out-migrate from their initial towns. A 10-percentage point increase in the initial textile exposure is associated with a 1.5-percentage point decrease in migration likelihood by 1930 and a 2.2-percentage point decline in migration by 1940.

Role of wealth and social network in decreased migration

Unlike conventional predictions that workers who faced negative local economic shock will out-migrate from the local labor market, the workers in textile-heavy towns were *less* likely out-migrate from their towns. Here I highlight two mechanisms that may explain the decreased out-migration. First, I present evidence that wealth played a role in migration decisions: lack of wealth reduces the likelihood of migration. Second, I show that individuals increased utilizing the support from social networks with the local economic shocks by increasing cohabitation with their families.

There are several ways wealth could affect the migration decision during an adverse economic shock. First, migration requires a fixed cost of moving, and with the absence of wealth, the ability to afford migration may decrease migration with economic shocks (Bound and Holzer, 2000). Also, the decrease in local labor demand may decrease housing prices. This may increase underwater mortgages, leading to a higher share of the population not being able to afford migration, but this can also affect poorer renters to be "compensated" and decrease the motivation to out-migrate (Notowidigdo, 2020).

In Table 1.2, I examine how the textile shock affected workers' likelihood of migration, conditional on owning a home in the initial decade. I utilize the home-ownership variable in the 1900-1940 census to construct a separate sample of homeowners. Table 1.2 column (1) contains the coefficient estimates of Equation (2), including all matched men. Columns (2) and (3) contain coefficient estimates of Equation (2) for a sample of men who do not own houses and a sample of homeowners, respectively. Homeownership heavily depends on age; the older population is more likely to own housing. As migration likelihood also depends

on workers' age, I include age fixed effects for all of the specifications to control the age profile of homeownership and migration. Columns (2) and (3) of Table 1.2 indicate that homeowners were more likely to migrate from their towns affected by the textile industry decline, suggesting a potential role of wealth. Comparing the 1940 coefficients from Columns (2) and (3), I find that with a 10-percentage point increase in the exposure to textile shock, non-homeowners are 1.2 percentage-point less likely to out-migrate compared to homeowners.

Homeownership may not fully capture individuals' wealth, as homeowners may have mortgage obligations. In Table 1.3, I utilize the mortgage status variable from the census to examine whether homeowners' migration response to the textile shock differs by mortgage status. I find that homeowners with no mortgage obligations were more likely to migrate than homeowners with a mortgage.[15]

With the absence of social safety net, workers facing negative local economic shock may utilize the support from their social network. For instance, an unemployed or displaced worker may weather the income shock by living with their family members. Table 1.4 examines whether the textile shock changed workers' likelihood of cohabitation with family members, measured by whether they lived in a multigenerational-family residence.[16] I find that young workers in towns with a higher textile presence were more likely to live in multigenerational housing in the era of textile decline. A 10-percentage point increase in the initial exposure to the textile decline is associated with a 0.8 percentage-point increase in cohabitation.

[15]Mortgage variable is available only until 1920, which means that I could use the 1900-1910 matched sample as the pre-period, the 1910-1920 matched sample as the baseline, and the 1920-1930 matched sample as the post-period data point.

[16]Multigenerational living is defined as living with at least two generations of family that are not a direct parents-children relationship. For instance, if a worker lives with an aunt/uncle, the worker is in multi-generational housing. If a worker lives in three-generation housing with parents and grandparents, the worker is in multigenerational housing.

1.6 The effect of the textile industry decline on labor force participation, industry switching, and education

In the previous section, I documented the muted migration response from the workers affected by the textile industry decline. This section examines workers' non-migratory responses, including labor force participation, education, and industry switching.

1.6.1 Heterogeneous effects on labor force participation by worker age

The decline of a major industry in towns can lead to a decrease in labor demand and overall level of employment. Still, its impact on town-level labor force participation is ambiguous. On the one hand, if displaced or affected individuals cannot find new employment opportunities or the opportunity costs of other non-occupational activities (e.g., education) decrease, the labor force participation rate will decrease. On the other hand, if the towns offer different employment opportunities and their residents can find new jobs easily, the labor force participation may not be affected.[17]

In this subsection, I examine how the textile industry decline affected workers' labor force participation. I assign individuals in the census to be in the labor force if the respondent has reported an occupational response. The labor force participation rate is then computed as the number of individuals in the labor force divided by the working-age (14-65) population of the town.

[17] Another potential proxy of non-employment is unemployment. However, the employment status variable that separates employment and unemployment (EMPSTAT) is not available in the 1900 and 1920 censuses, so I cannot follow the effect on unemployment using the town-level event-study analysis.

I find a heterogenous effect on labor participation by age groups. The decline in the labor force participation rate is concentrated among young individuals, particularly the school-going age population (age 14-18), as presented in Figure 1.9. A town with a 10-percentage point higher textile share of the labor force experienced a 1.5 to 2 percent decline in labor force participation rate among the school-going age population. In Figure 1.10, I find much flatter and smaller changes in labor force participation among older workers of age 28-65.[18]

1.6.2 Increased industry switching to the agricultural sector and decreased occupational income

In Section 1.6.1, I document a relatively small effect of the textile industry decline on the labor force participation of adult workers. This implies that the older labor force may have reacted to the textile industry's decline by securing new employment opportunities in other sectors.

I find that the towns affected by the industrial decline were more likely to increase the agricultural share of the labor force from 1920 to 1940 in Figure 1.11. This pattern of the rise in different industries in towns and the less pronounced decline in labor force participation among older workers suggests that workers in these towns may have switched their industries to the agricultural sector. Labor reallocation to the agricultural industry is also documented in Boone and Wilse-Samson (2021) during the Great Depression, serving as the informal insurance during this era. In a broader context, this highlights that having a diverse set of industries, or industries to fall back on, may serve as a buffer for industry-level shocks and assist the local labor market recovery.

[18] Appendix Figures 1A.4 and 1A.5 report estimates of the relationship between the initial textile industry presence and labor force participation rates. The overall trend in coefficients mirrors that of Figure 1.4, the first-stage results described in Section 1.5.3. Towns with higher textile presence in 1900 are likely to experience a decline in labor force participation; a 10-percentage point increase in initial textile share is associated with a 1 to 1.5 percent decline in labor force participation rate between 1920 and 1940. Groups of age 19-25, 26-45, and 46-65 demonstrate much smaller magnitudes of labor force participation decline by order of magnitude of ten, as shown in Appendix Figure 1A.5.

Now, I examine how the decline of the textile industry affected town residents' overall earnings. I create a consistent measure of earnings across census years by utilizing the occupational-based income (OCCSCORE) variable in the census. OCCSCORE assigns income scores to each census occupation in the 1950 census occupational classification. I compute the average occupational income score for each town and examine its relationship with the exposure to the textile industry decline in Figure 1.12. Towns with more exposure to the industrial decline experienced a significant decrease in average occupational income scores.

The relatively small labor force participation response from the older labor force combined with the decline in occupational income implies that even though they were able to stay in the labor force without out-migrating, they ultimately took lower-paying positions. The decrease in the occupational income scores is mainly driven by the increase in the agricultural share of the labor force, which has a lower occupational score on average than the textile industry. The agricultural occupational income score is lower partly due to its seasonality, so the decline in the occupational income is due to both the decrease in the length of work and the decrease in wage per unit of time.[19]

As discussed in Section 1.5.4, the effect on the aggregate town-level outcomes may be driven by the compositional effect, so I estimate Equation (2) with the labor force participation, likelihood of working in the agricultural sector, and change in occupational income as outcome variables in Table 1.5. I find that individuals in textile-heavy towns were more likely to work in the agricultural industry and experience a significant decrease in occupational income in the next decade, but their likelihood of staying in the labor force was not significantly associated with exposure to the industrial decline.

[19]Given that the census industry and occupation are defined based on the respondent's job that gives the most source of income, the decline in occupational income score still implies that workers faced an overall decline in earnings.

1.6.3 Increased educational attainment in response to the deindustrialization

I document a sizable decline in labor force participation among the school-going age population in Section 1.6.1. For young individuals making decisions between employment and education, the reduction in employment opportunities from industrial decline could lower the opportunity cost of education, leading to increased educational attainment. The long-run impact of the industrial decline on these young individuals is ambiguous. On the one hand, the reduced employment opportunity may decrease earnings in the short run. On the other hand, the increased level of schooling may translate into increased employment opportunities and earnings in the future.

In this sub-section, I examine how the textile industry's decline affected young individuals' educational attainment and career. I start by showing educational attainment trends in New England states and presenting town-level evidence of increased educational attainment. Then, I further document individual-level responses by exploiting an additional source of variation from the location and timing of textile plant closures and utilizing individuals' age to identify which individuals within the closure towns are affected.

Educational attainment trends in New England states

New England states were pioneers in education and schooling systems in the US in the nineteenth and early twentieth centuries. Modern educational institutions were developed in the mid-to-late nineteenth century in New England, and the expansion of secondary education in the early twentieth century marked the onset of the high school movement in the US (Massachusetts Department of Education, 1930; Goldin and Katz, 2008a). New England states had compulsory schooling laws in effect in the 1910-20s that required school

attendance until 16 years old, except Maine, which needed schooling until 17; however, workers as young as 14 could obtain work permits (Lleras-Muney, 2005).

In Figure 1.13, I plot the share of the population attending school for each age cohort in New England using the 1920 full-count census. More than 80 percent of the population in New England attended school until 13 years of age. There are variations across states on when students start leaving school—Connecticut, Massachusetts, and Rhode Island experience up to 20 percentage point decline in school attendance in the age-14 cohort, which translates to an eighth-grade education. In contrast, young adults in Maine start leaving school at the age of 15, and in New Hampshire and Vermont at the age of 16.

The effect of textile industry decline on town-level educational attainment

If the textile industry decline affected young individuals' educational attainment in New England, we expect town-level average educational attainment to increase in towns with large textile industry presence. I test this hypothesis in Appendix Figure 1A.6 by estimating Equation (1) with the share of school attendance among the school-going aged population as the outcome variable. I find that school attendance increases in textile-heavy towns, in a similar magnitude as the decrease in labor force participation in Figure 1.9. [20,21]

If the decrease in employment opportunities induced the need for public school provision, the increased educational attainment might be explained by the increase in the supply side of education. I ran an analogous town-level analysis using the number of teachers per capita as the outcome variable in Appendix Figure 1A.8. The number of teachers per capita,

[20] The figure presents declining pre-trends from 1900 to 1910, which may be driven by the inconsistent definition of duration of schooling to be considered as "school-attending" across census years. This can categorize individuals in the labor force as school-attending in some census years, making the trends from the inconsistent definition be correlated with the initial textile share of the town labor force.

[21] Alternatively, I can avoid the issues from definitions of school attendance variable is to match individuals from each census to the 1940 census and use the highest grade completed information recorded in 1940. In Appendix Figure 1A.7, I create a town-level share of the eighth-grade completion among the population aged 11-14 and examine its relationship with the initial textile industry share. The caveat of the analysis, however, would be that the 1940 census won't be included in the analysis as the population of age 11-14 in 1940 wouldn't have completed the eighth grade. I find that towns with a higher initial textile presence were associated with a rise in the share of the population with at least an eighth-grade education by 1930.

controlling for the foreign-born share of the population, was not affected by the decline in the textile industry. This finding is consistent with an interpretation that the overall supply of education did not drastically change with the industrial decline and the supply-side effect is not the main driver of the increased educational attainment in these towns.

The effect of plant closures on individual-level educational attainment and career

Earlier I documented town-level evidence of younger individuals increasing school attendance. Now, I ask whether younger individuals were more likely to increase schooling in response to the local economic shock and whether the additional educational attainment further affected future earnings and labor market outcomes. To capture individual-level responses, I exploit two sources of variation: (1) the location and timing of textile plant closures identified from a newly digitized data source; and (2) individuals' location and age in 1920 to identify which individuals were facing a decision between employment and education after the plant closure within the affected towns.

Even if young individuals are in the same affected town, their education and employment decisions may differ based on whether they had to make the decision before or after the textile industry shock in their towns. For instance, if a large textile plant closed in a worker's town when the worker was 12, the worker would have to decide on her educational attainment conditional on the plant closure. However, an older resident in the same town who was 20 years old when the plant closed would have already completed making the schooling decision.

Using individuals matched from the 1920 to 1940 censuses, I employ matched difference-in-differences method and compare young individuals in affected towns with individuals who already made their schooling decisions before the closures after controlling for the birth-cohort educational trend from the matched towns.

Matching algorithm and matched town characteristics

I employ a matching algorithm to create a set of control towns with similar observable characteristics to towns with sizable plant closures. Treated towns are towns that experienced at least one large plant closure between 1920 and 1930, and potential control towns must not have experienced any large plant closure between 1920 and 1930.

I use a propensity score matching algorithm to assign a control town to each treated town based on population and share of textile employment in 1910. Each town with a large plant closure is matched to a non-neighboring control town with the closest propensity score among towns located in a different county but in the same state. Choosing a non-neighboring town can address potential spillover effects from geographic proximity. Finally, control towns are assigned counterfactual closure years based on the year their treated counterparts experienced plant closures. Appendix Figure 1A.9 plots the location of plant closure towns and their matched control towns.

Table 1.6 contains the summary statistics of closure and control towns and the rest of the New England towns. On average, the closure towns had 20 percent of the labor force in the textile industry in 1910, and control towns had 16.5 percent of the labor force in textiles. I find that the closure and control towns have a higher share of the foreign-born population than other New England towns. The closure and control towns are less likely to be rural, employing around 15 percentage points less of the labor force in the agricultural sector.

Empirical strategy for matched difference-in-differences design

My empirical strategy for understanding individuals' education responses is a matched difference-in-differences design. Using individuals in the 1920 census matched to the 1940 census, I examine whether young individuals who faced plant closures during their critical age responded by changing their educational attainment. I compare young individuals in affected towns with individuals who already made the schooling decision before the closures

after controlling for the birth-cohort educational trend in the matched towns. In particular, I estimate the below equation:

$$Y_{i,town,1940} = \beta \cdot \left[\text{Age 11-14 at Closure}_{i,town}\right] + \alpha_c + \delta_{town} + \varepsilon_{town,t} \tag{1.3}$$

where town is i's town of residence in 1920, and c is i's birth cohort. $Age\ 11-14\ at\ Closure_{i,\ town}$ indicates whether individual i who were located in town in 1920 experienced a large plant closure between 10 and 16. Outcome variables include indicators for completing at least eighth or ninth grade, a log of income in 1940, and migration. I employ birth-cohort and town fixed effects.

The sample includes individuals who lived in closure towns or control towns in 1920 and those between 11 and 14 or between 20 and 23 at the time of plant closures. I exclude those between 15 and 19 years of age, as they are close to the exposure age. These individuals, even if they completed schooling at the time of plant closure, may decide to go back to school when they observe decreased employment opportunities and career prospects from the plant closures. Also, there can be measurement errors in either the age variable in the census or the plant closure years. In both cases, this group could be both in the critical exposure and control age groups, so I remove them from the analysis sample.

Effects of plant closure on education, labor market outcomes, and migration

Figure 1.14 plots the normalized trend of the average share of the town population who completed at least the eighth grade for plant closure towns and matched control towns. The graph is organized by the age-at-closure cohorts, and it shows that while the cohorts in both closure and control towns obtained similar educational attainment between 20 to 40 years of age at the time of plant closure, the educational attainment of the closure and control towns diverges for the cohorts of 14 years old or younger.[22]

[22] Appendix Figure 1A.10 shows the analogous graph for the log wage income trend.

Table 1.7 presents the estimated effect of plant closure on young individuals' educational attainment, labor market outcome, and migration. Column (1) shows that young individuals affected by large plant closures likely have a 3 percentage-point increase in eighth-grade completion. I find that this increase happens on the margin of completing at least the eighth grade of schooling, and there was no significant rise in the propensity to complete at least the ninth grade of education, as shown in columns (2).

In Columns (3) and (4), I find that the young individuals affected by plant closures did not significantly differ in employment probability by 1940, and plant closures were associated with a 5 percentage-point decrease in the log of wage income. Column (4) does not include individuals who reported zero wage income, which excludes those who are self-employed. In Column (5), I examine the likelihood of being self-employed separately. I find that the self-employment probability was not affected by the plant closures. Column (6) shows that plant closures did not change the out-migration probability for those who received additional education. These results imply that the plant closure has positively impacted young individuals' human capital accumulation, but the additional schooling did not translate to better labor market outcomes.

There are potential mechanisms that could explain the worsened labor market outcomes. First, the additional educational attainment did not translate into a worker migration response as shown in Column (6). This implies that young workers with additional schooling still faced the negative local economic shock the same way as the others. Also, the amount of additional schooling was one or (at most) two years of schooling, as the critical exposure age was 13 or 14 years of age and the only significant margin of adjustment was completing the eighth grade of education. The returns to schooling were lower than the modern estimates of returns to schooling, which may explain why an additional year of schooling may have not created a large impact on future earnings (Goldin and Katz, 2008a, Feigenbaum and Tan, 2020). Also, if the additional amount of schooling was not large enough to create

an occupational upgrade, the education only delays the labor force entry and decreases the experience profile of workers in the absence of migration response and local industry with higher returns to education. Based on the Mincer earnings model, this implies that workers with additional schooling could experience a decline in their wages (Mincer, 1974; Lemieux, 2006).

1.7 Conclusion

This paper documents new evidence on the persistence of local shocks in the early twentieth century. I find that a severe local economic shock did not induce a significant out-migration response. Individuals with a lower level of wealth were more likely to stay in the affected towns, utilize family and social networks, and increase cohabitation. Instead of out-migration, individuals chose to increase educational attainment, even though the increased schooling did not help workers improve their labor market prospects or migration likelihood. Older workers switched to the agricultural sector but faced a reduction in income. My finding presents evidence that the lagged local recovery may not be a modern phenomenon.

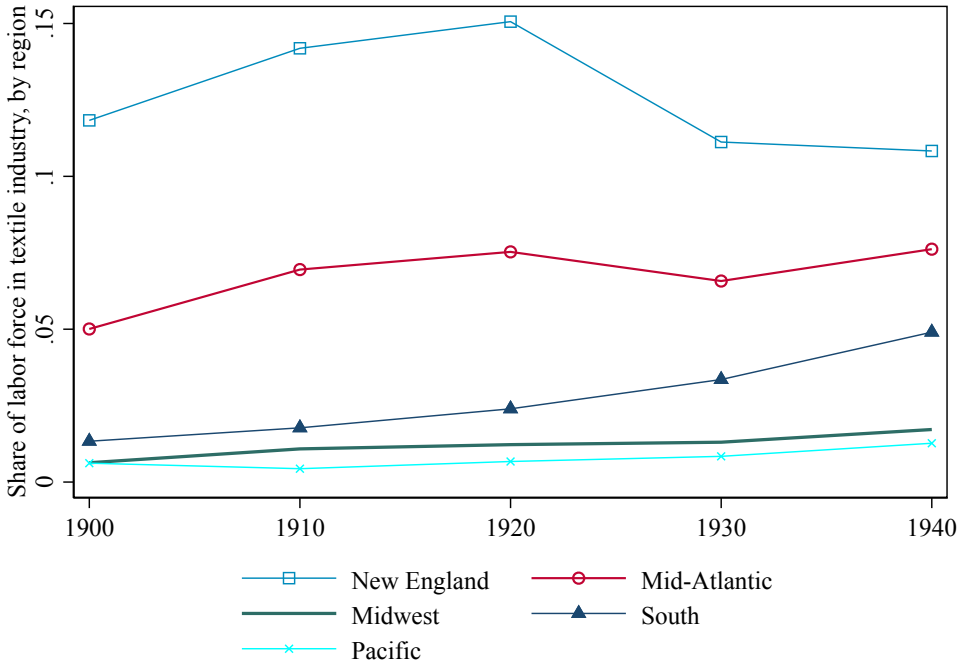
My results speak to several important policy implications. Policies promoting transitions of workers to different industries and diversity of local industrial base can build local labor market resilience. Lastly, even though education is an important margin of worker adjustment to local shocks, in the absence of migration or local industry that offers higher returns to education, the workers may not benefit from the additional education.

The empirical setting of this paper, combined with the comprehensive individual-level data, provides several fruitful avenues for future research. First, while this paper does not focus on the female textile labor force, female workers were an important part of the New England textile industry. Understanding how deindustrialization affected their labor market outcome and marriage and fertility decisions can contribute to our understanding of how local economic conditions shape family formation.

Second, the setup is well-suited to study the effect of deindustrialization on political preferences. While political economy literature documents how economic shocks could elicit political responses in the late twentieth and early twenty-first century, a puzzle often remains about why the affected voters tend to turn rightward even though the redistribution policies are often supported by liberals. Examining how voters in the mill towns responded, in the absence of social safety net, can potentially shed light on this political trend.

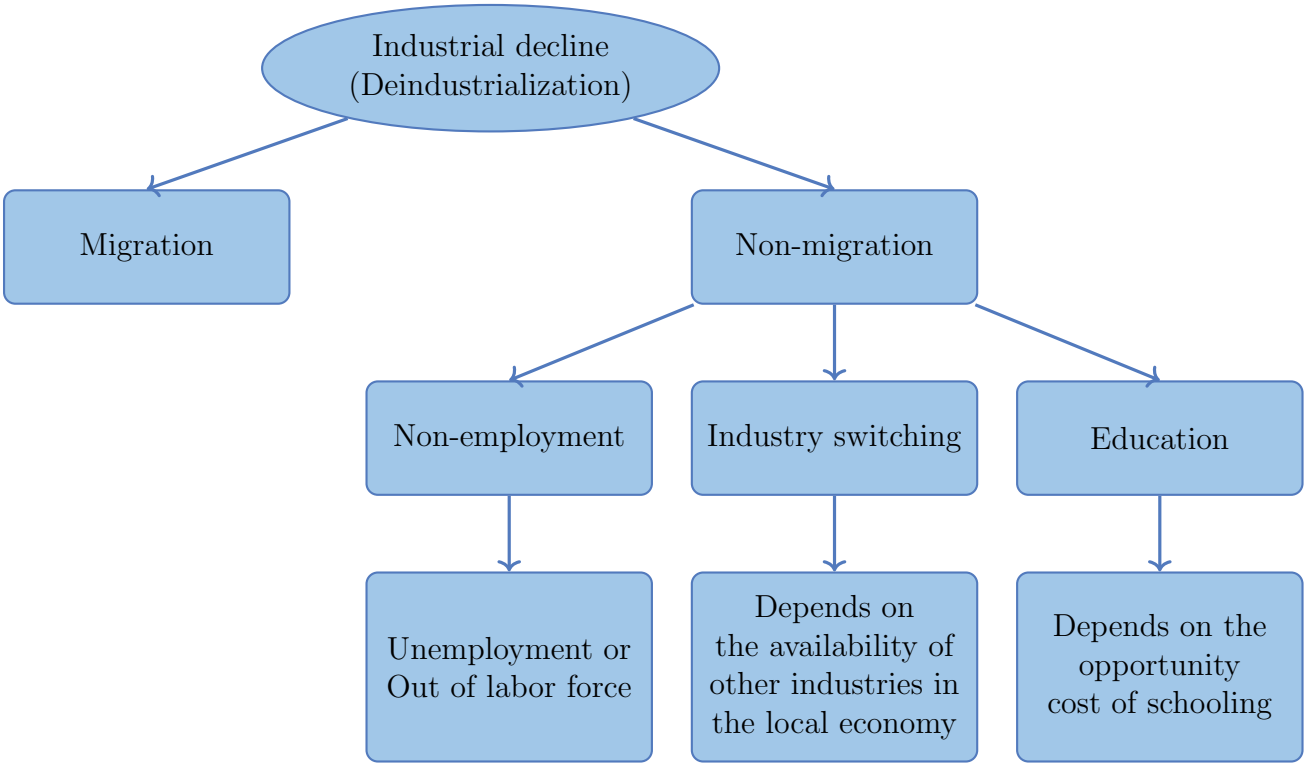
Tables and figures

Figure 1.1: Share of the labor force in the textile industry by region



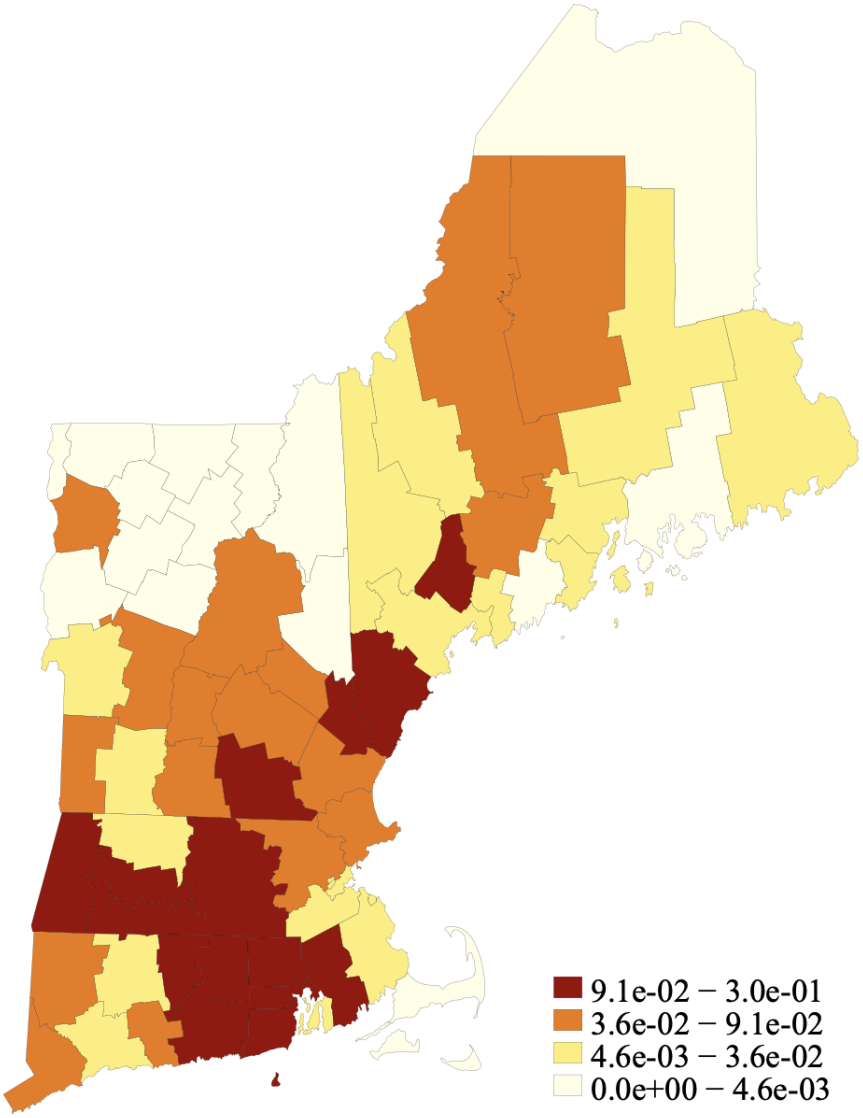
Notes: The figure plots the share of the regional labor force working in the textile industry, computed from the IPUMS 1% census sample. The textile industry is defined as the group of 1950 census industries: (1) Knitting mills, (2) Dyeing and finishing textiles, except knit goods, (3) Carpets, rugs, and other floor coverings, (4) Yarn, thread, and fabric mills, (5) Miscellaneous textile mill products, (6) Apparel and accessories, and (7) Miscellaneous fabricated textile products. The denominator is the number of working-age (14-65) population in each region and area who reported a gainful occupation.

Figure 1.2: Worker responses to deindustrialization



Notes: The figure summarizes worker responses to deindustrialization.

Figure 1.3: Textile share of the labor force by county in 1900



Notes: The figure plots the heatmap of New England counties, whose color intensity indicates which quartile each county belongs to based on the textile share of the county labor force in 1900. The textile industry is defined as the group of 1950 census industries: (1) Knitting mills, (2) Dyeing and finishing textiles, except knit goods, (3) Carpets, rugs, and other floor coverings, (4) Yarn, thread, and fabric mills, (5) Miscellaneous textile mill products, (6) Apparel and accessories, and (7) Miscellaneous fabricated textile products.

Table 1.1: Average town characteristics by initial textile presence

Tercile (lower tercile : smaller textile presence)	No textile presence	1	2	3
Demographics (1900)				
Working-age population	490.961	1726.326	5103.142	4550.736
Avg. socioeconomic index	9.561	11.383	11.664	12.020
Labor force participation rate	0.524	0.531	0.531	0.578
Share foreign-born	0.101	0.133	0.134	0.216
Share literate	0.753	0.779	0.788	0.748
Textile share of labor force	0.000	0.004	0.018	0.203
Agriculture share of labor force	0.517	0.357	0.342	0.196
Manufacturing share (net of textile)	0.074	0.127	0.140	0.122
Number of towns	622	288	281	284

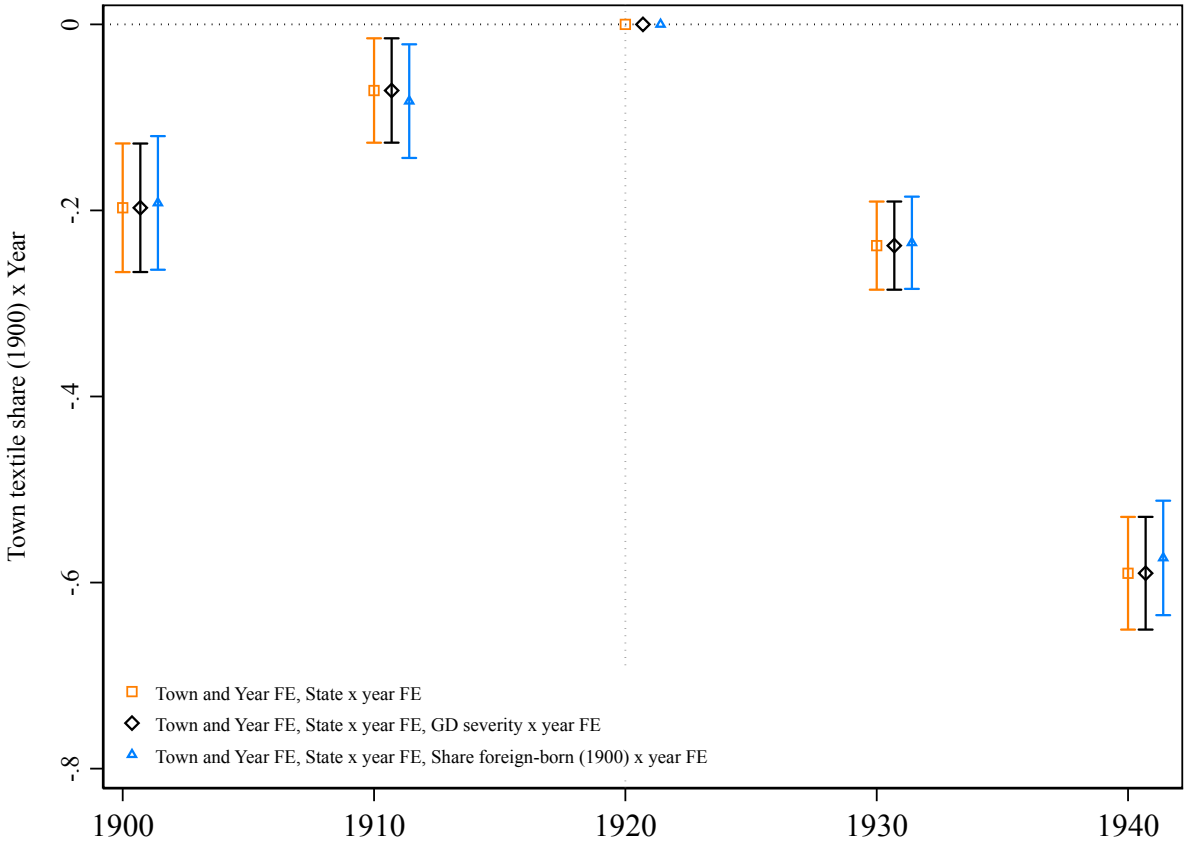
Notes: The table summarizes average town characteristics based on the initial textile presence in 1900. The initial textile presence is defined as the share of the town labor force reported to work in the textile industry. Each demographic characteristic is computed using the 1900 Census, and the average characteristics are computed using the information on the working-age (14-65) population in each town. A working-age worker is defined as participating in the labor force if the worker reports having a gainful occupation. The first column shows the average characteristics of towns with zero textile share of the labor force. The second to fourth columns present the average characteristics of towns in each tercile defined based on the initial textile presence.

Figure 1.4: Average textile share of the labor force by quartiles of initial textile presence in 1900



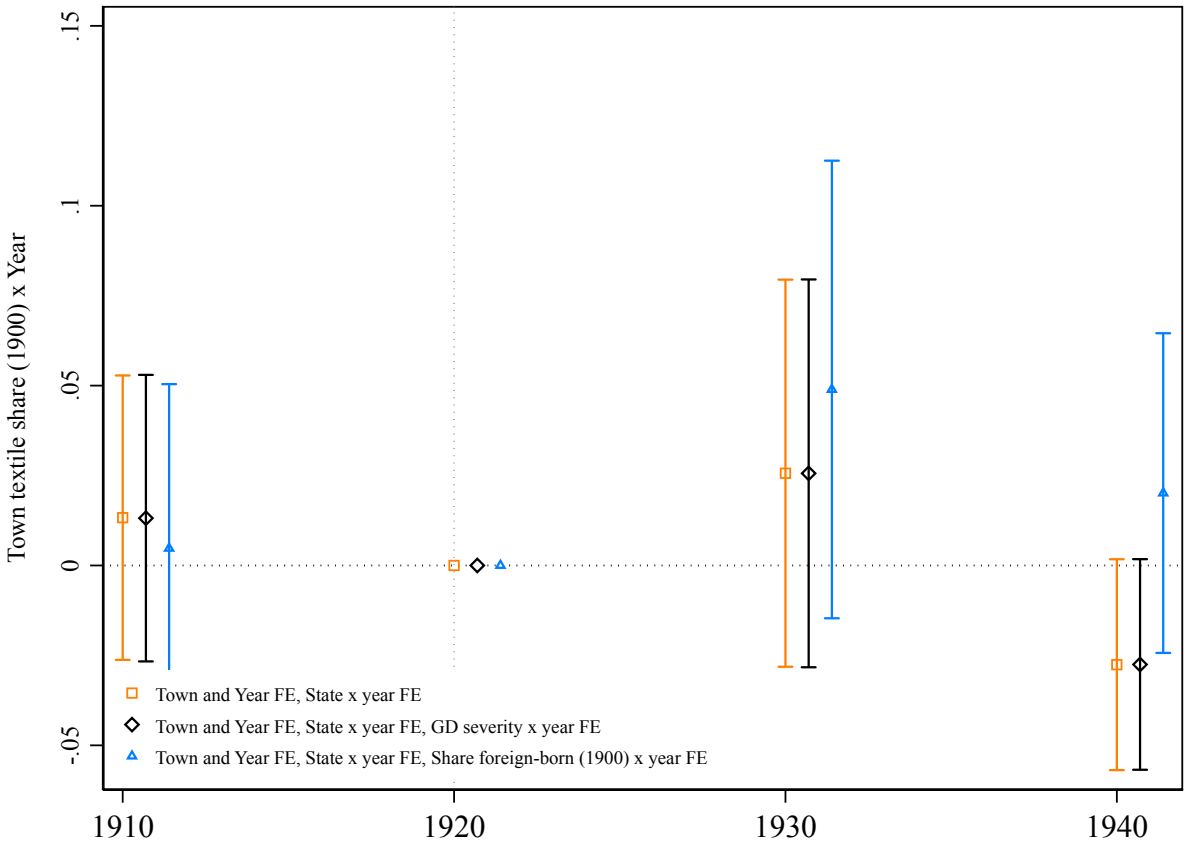
Notes: The figure plots the time series of average textile share of the labor force by quartiles defined by the initial textile share of the labor force in 1900. I compute the textile share of the town labor force by town and year. Each New England town is then assigned to a quartile based on its share of the labor force working in the textile industry in 1900. Finally, I compute the average textile share in each quartile by each census year from 1900 to 1940.

Figure 1.5: The effect of exposure to textile industry decline on the textile share of labor force



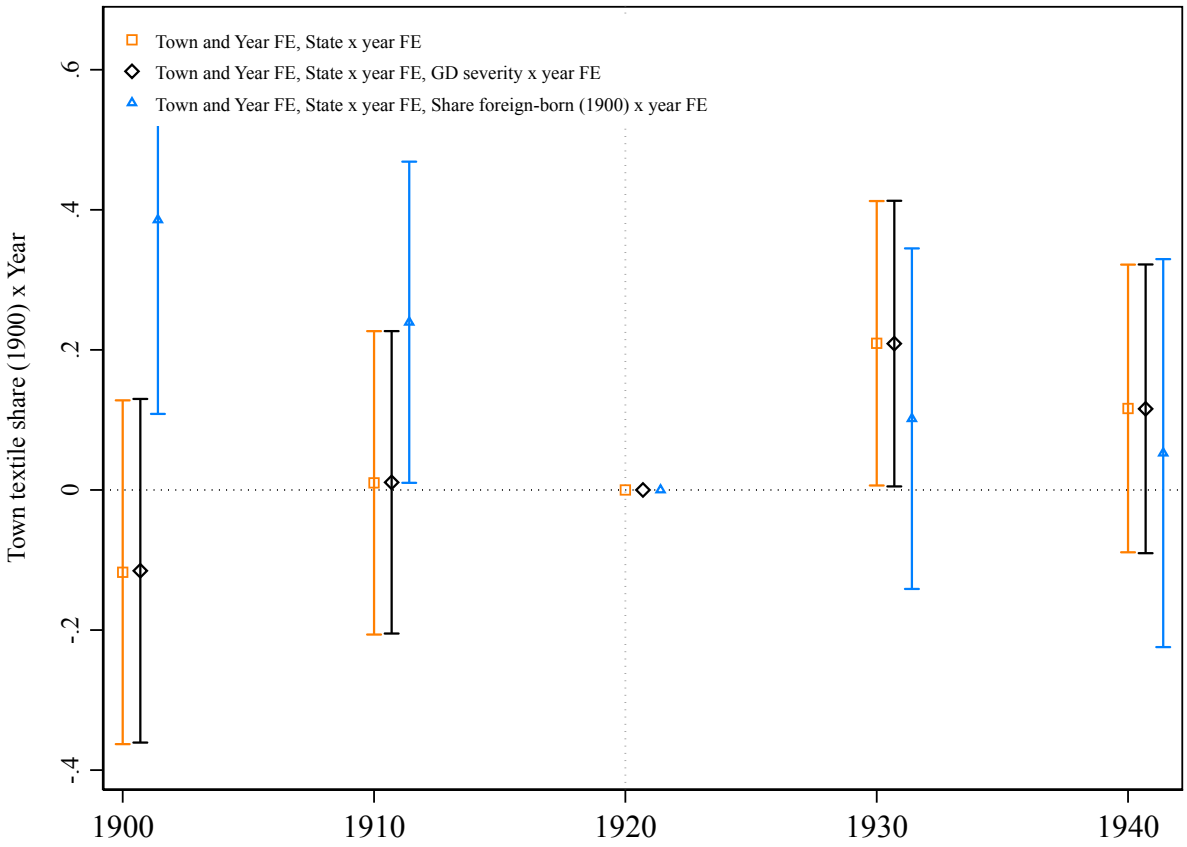
Notes: The figure presents the continuous difference-in-differences analysis results that use the annual textile share of the labor force as the outcome variable. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1.6: The effect of exposure to textile industry decline on net in-migration rate



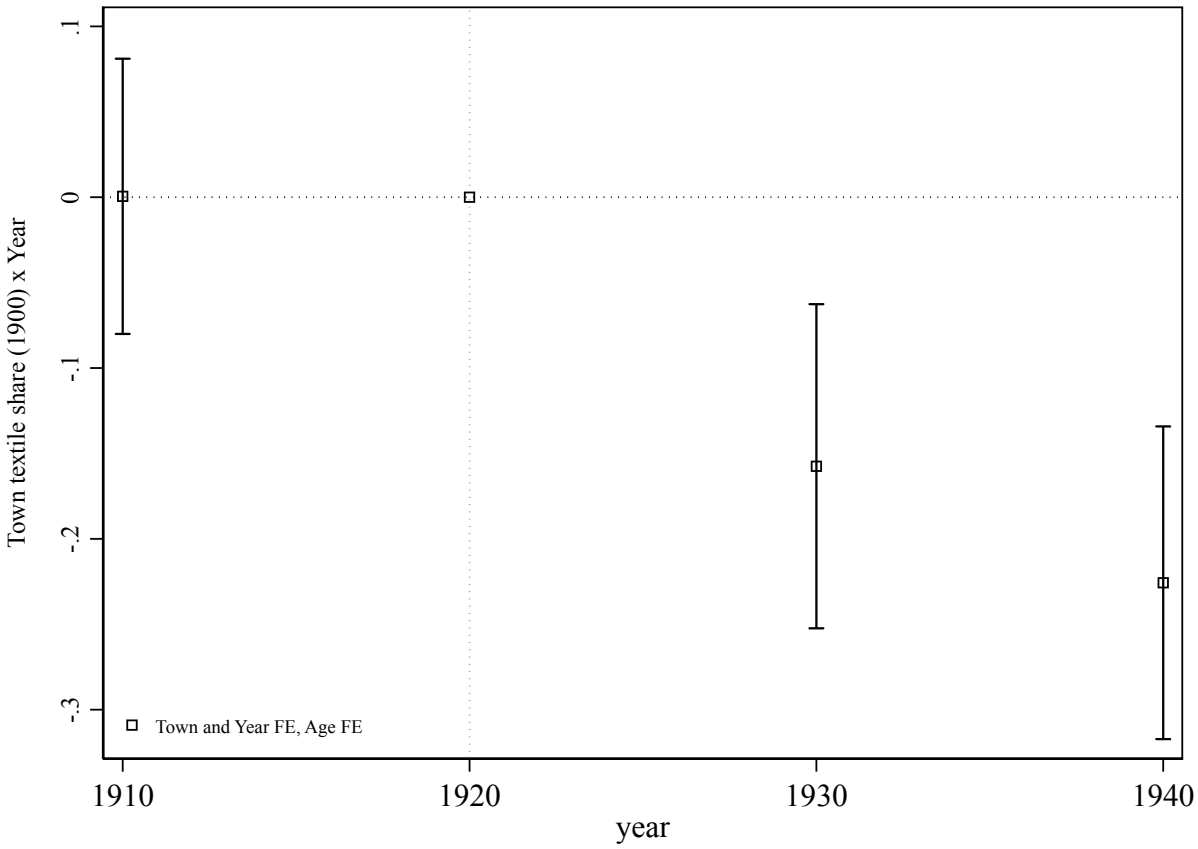
Notes: The figure presents the continuous difference-in-differences analysis results that use the net in-migration rate as the outcome variable. The net in-migration rate in each town is computed by the change in the log of the working-age population over each decade divided by the initial log working-age population. The sample includes each town's information for 1910, 1920, 1930, and 1940, omitting the data point for 1900, as the net rate is computed over a decade. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1.7: The effect of exposure to textile industry decline on working-age population



Notes: The figure presents the continuous difference-in-differences analysis results that use the log of the working-age population as the outcome variable. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state × year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1.8: The effect of exposure to textile industry decline on individual-level likelihood to out-migrate



Notes: The figure presents the estimates of Equation (2), using the individual-level indicator variable for out-migration as the outcome variable. Each individual in 1900, 1910, 1920, and 1930 censuses is matched to the adjacent census years, 1910, 1920, 1930, and 1940. Each individual is assigned an indicator variable of whether the individual moved out of their initial town by the next census decade. Each individual's exposure to the textile decline is defined by the 1900 textile share of the labor force in their original town of residence. The figure plots the coefficient estimates β_t from Equation (2), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects.

Table 1.2: Migration mechanism: wealth increases the likelihood for migration

	Out-migrate (All)	Out-migrate (Non-homeowners)	Out-migrate (Homeowners)
	(1)	(2)	(3)
Textile Share x 1910	0.0005	0.0139	-0.0519
	(0.0411)	(0.0441)	(0.0447)
Textile Share x 1930	-0.1575***	-0.1639***	-0.1289***
	(0.0484)	(0.0523)	(0.0465)
Textile Share x 1940	-0.2257***	-0.2642***	-0.1421***
	(0.0467)	(0.0486)	(0.0481)
Town FE	X	X	X
Year FE	X	X	X
Age FE	X	X	X
Dep. var. mean	0.363	0.421	0.290
Adjusted R^2	0.0495	0.0469	0.0593
N	1441039	799807	641214

Notes: The table presents the estimates of Equation (2), using the individual-level indicator variable for out-migration as the outcome variable. Each individual in the 1900, 1910, 1920, and 1930 censuses is matched to the adjacent census years, 1910, 1920, 1930, and 1940. Each individual is assigned an indicator variable of whether the individual moved out of their initial town by the next census decade. Each individual's exposure to the textile decline is defined by the 1900 textile share of the labor force in their original town of residence. The table presents the coefficient estimates β_t from Equation (2), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. Column (1) uses all matched workers as the sample. Column (2) uses matched workers who do not own their own house as the sample, based on the OWNERSHP variable in each census. Column (3) uses matched workers who reported owning a house as the sample.

Table 1.3: Homeowners with mortgages are less likely to migrate

	Out-migrate (Homeowners)	Out-migrate (No mortgage)	Out-migrate (Mortgage)
	(1)	(2)	(3)
Textile Share x 1910	-0.0467	-0.0432	-0.0376
	(0.0439)	(0.0467)	(0.0498)
Textile Share x 1930	-0.1225***	-0.0844*	-0.1581***
	(0.0439)	(0.0475)	(0.0482)
Town FE	X	X	X
Year FE	X	X	X
Age FE	X	X	X
Dep. var. mean	0.312	0.292	0.333
Adjusted R^2	0.0626	0.0700	0.0579
N	417414	215596	201799

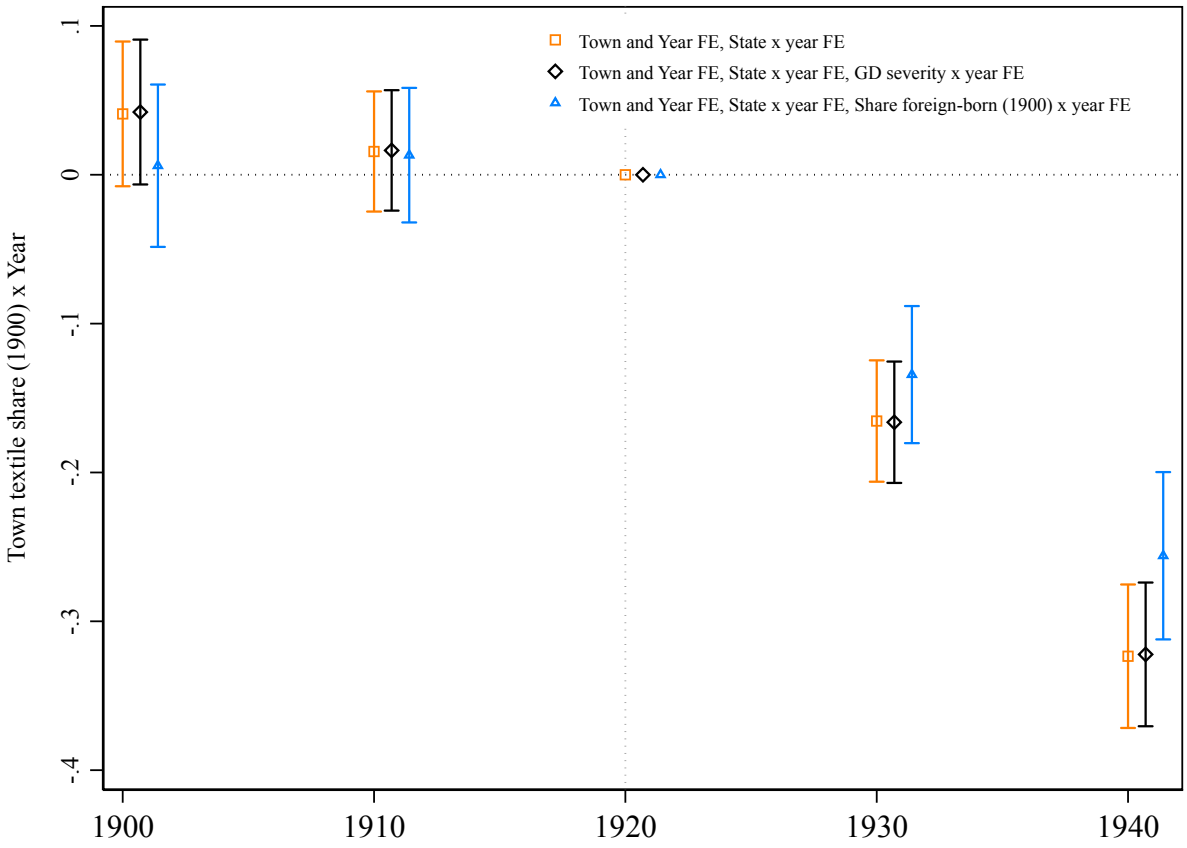
Notes: The table presents the estimates of Equation (2), using the individual-level indicator variable for out-migration as the outcome variable. Each individual in the 1900, 1910, and 1920 censuses is matched to the adjacent census years, 1910, 1920, and 1930. Each individual is assigned an indicator variable of whether the individual moved out of their initial town by the next census decade. Each individual's exposure to the textile decline is defined by the 1900 textile share of the labor force in their original town of residence. The table presents the coefficient estimates β_t from Equation (2), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. Column (1) uses all matched workers who reported owning a house. Column (2) uses matched homeowners with no mortgage obligation as the sample, based on the MORTGAGE variable in the 1900-1920 censuses. Column (3) uses matched homeowners who reported having a mortgage as the sample.

Table 1.4: Increased cohabitation

	Multi-gen. living (All)	Multi-gen. living (28-45)	Multi-gen. living (46-65)
	(1)	(2)	(3)
Textile Share x 1910	0.0030	0.0127	-0.0064
	(0.0142)	(0.0172)	(0.0209)
Textile Share x 1930	0.0134	0.0339**	-0.0196
	(0.0128)	(0.0143)	(0.0174)
Textile Share x 1940	0.0437**	0.0780***	-0.0076
	(0.0188)	(0.0235)	(0.0229)
Town FE	X	X	X
Year FE	X	X	X
Age FE	X	X	X
Dep. var. mean	0.4105	0.3202	0.5301
N	1441039	821359	619664

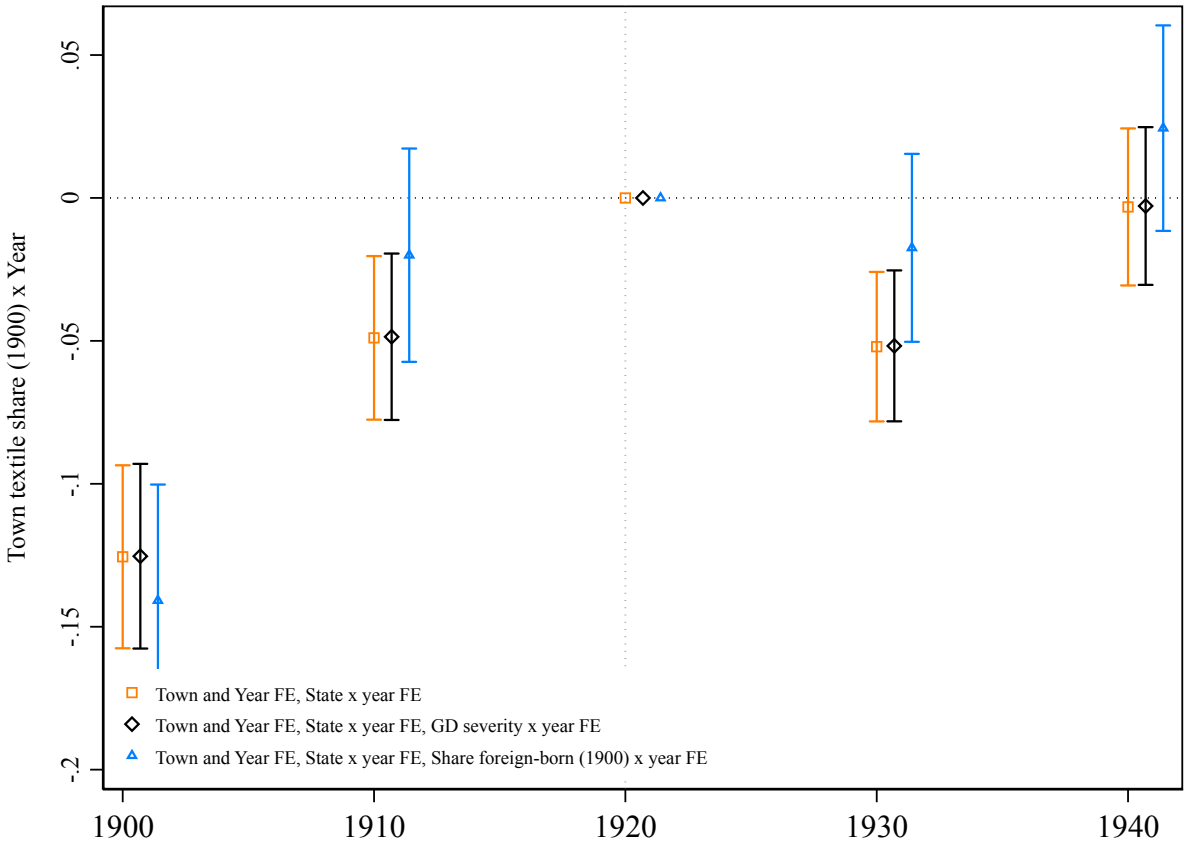
Notes: The table presents the estimates of Equation (2), using the individual-level indicator variable for cohabitation as the outcome variable. Each individual in the 1900, 1910, 1920, and 1930 censuses is matched to the adjacent census years, 1910, 1920, 1930, and 1940. Each individual is assigned an indicator variable whether individuals live with their extended family, indicated by living with more than one generation of the family other than their direct parents or children. Each individual's exposure to the textile decline is defined by the 1900 textile share of the labor force in their original town of residence. The table presents the coefficient estimates β_t from Equation (2), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. Column (1) uses all matched workers who reported owning a house. Column (2) uses matched homeowners with no mortgage obligation as the sample, based on the MORTGAGE variable in each census. Column (3) uses matched homeowners who reported having a mortgage as the sample.

Figure 1.9: Heterogeneous labor force participation by age: School-going aged population (age 14-18)



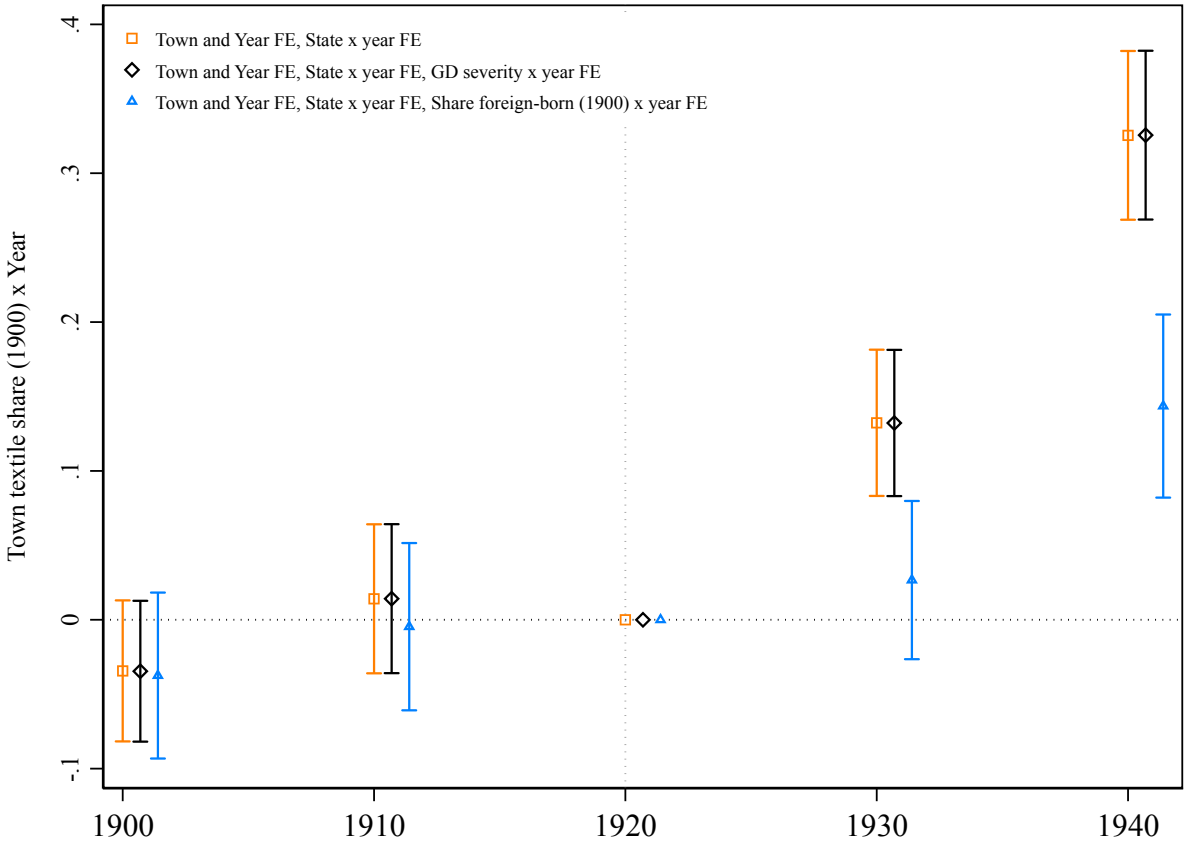
Notes: The figure presents the continuous difference-in-differences analysis results that use the labor force participation rate among the school-going aged population. Individuals with a gainful occupation recorded are considered in the labor force. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state × year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1.10: Heterogeneous labor force participation by age: Population of age 28-65



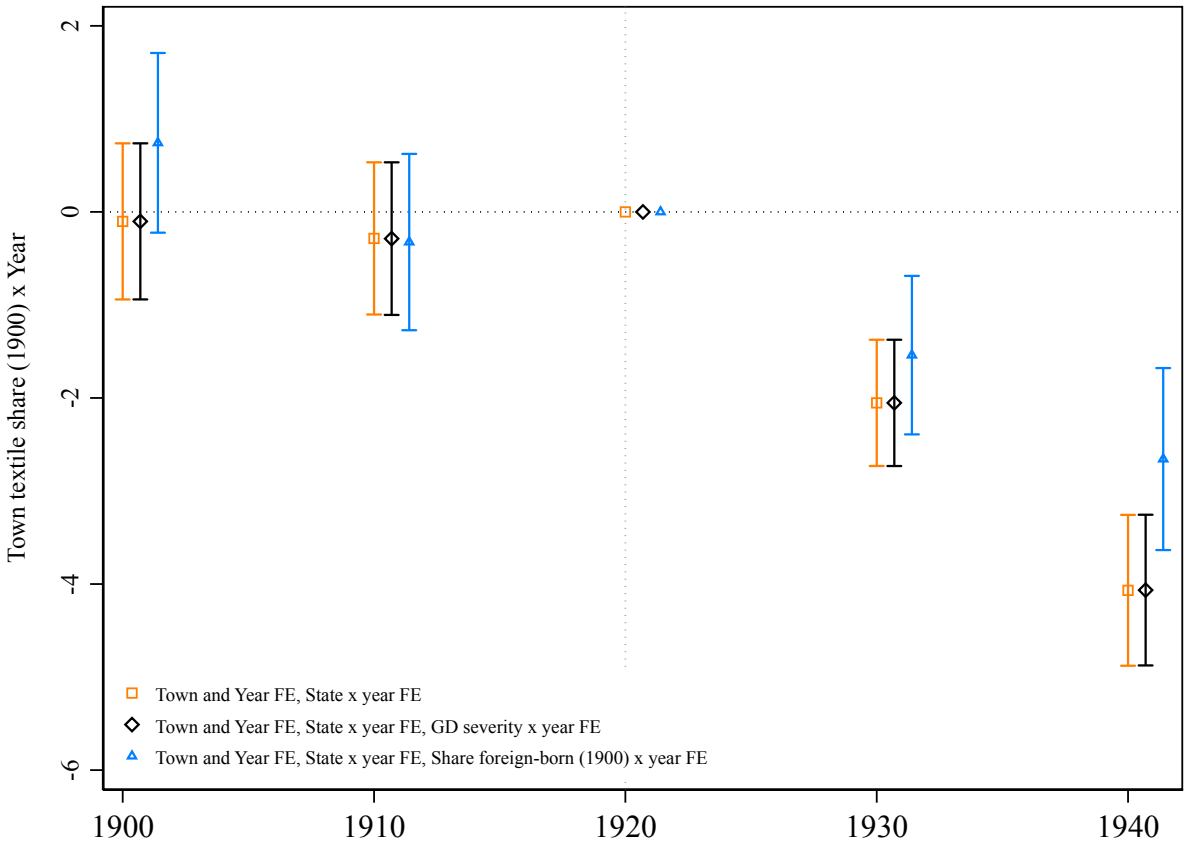
Notes: The figure presents the continuous difference-in-differences analysis results that use the labor force participation rate among the age 28-65. Individuals with gainful occupations are considered in the labor force. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state × year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1.11: The effect of exposure to textile industry decline on the share of agricultural labor force



Notes: The figure presents the continuous difference-in-differences analysis results that use the agricultural share of the town labor force as the outcome variable. The agricultural share of the town labor force is computed by the share of the town labor force who reported having agriculture (code 102 of IND1950) as their industry. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1.12: The effect of exposure to textile industry decline on occupation-based earning



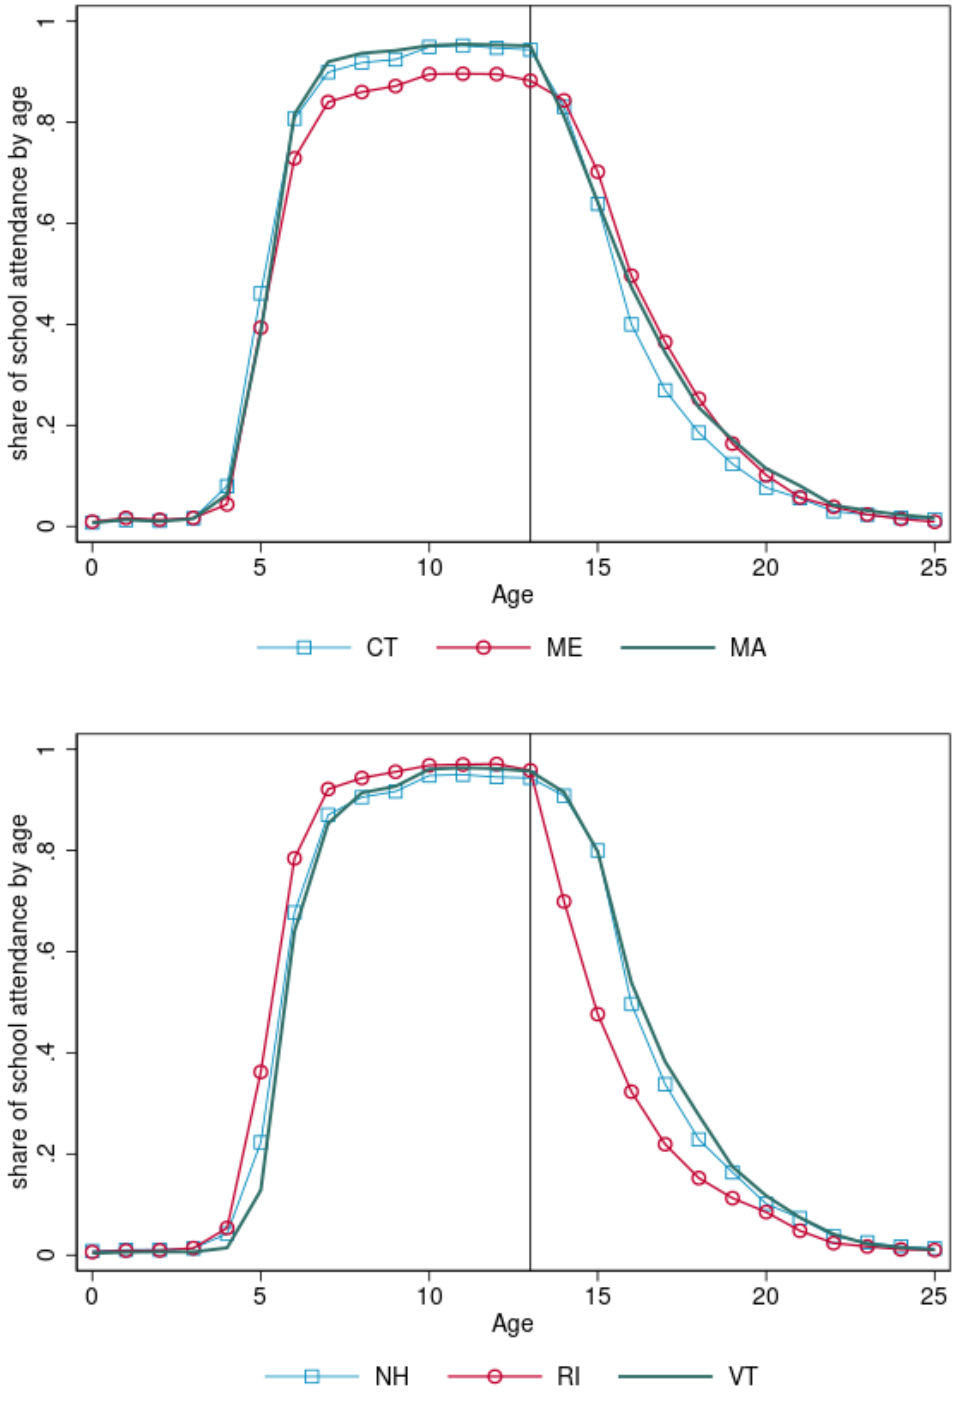
Notes: The figure presents the continuous difference-in-differences analysis results that use the average occupational income score (OCCSCORE) as the outcome variable. The town-level average is computed using the OCCSCORE of town working-age residents who reported having a gainful occupation. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. Standard errors are clustered by town.

Table 1.5: The effect of exposure to textile industry decline on labor market outcomes using a matched sample of workers

	LFPR	Agric. Ind	Change in occscore
	(1)	(2)	(3)
Textile Share x 1910	0.0045	-0.0039	0.4550
	(0.0559)	(0.0115)	(0.3843)
Textile Share x 1930	-0.0224	0.0387***	-0.9743***
	(0.0689)	(0.0111)	(0.3403)
Textile Share x 1940	-0.0497	0.0580***	-0.6073*
	(0.0651)	(0.0177)	(0.3242)
Town FE	X	X	X
Year FE	X	X	X
Dep. var. mean	0.817	0.104	0.955
Adjusted R^2	0.0540	0.2143	0.0079
N	1441039	1441039	907121

Notes: The table presents the coefficient estimates of Equation (2) using the individual-level indicator variables for participating in the labor force, working in the agricultural sector, and change in occupational income score (OCCSCORE) as the outcome variables for columns (1)-(3). Each individual in 1900, 1910, 1920, and 1930 censuses is matched to the adjacent census years, 1910, 1920, 1930, and 1940. Each individual's exposure to the textile decline is defined by the 1900 textile share of the labor force in their original town of residence. The figure plots the coefficient estimates β_t from Equation (2), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects.

Figure 1.13: Share of the population attending school in 1920 by state-age cohort



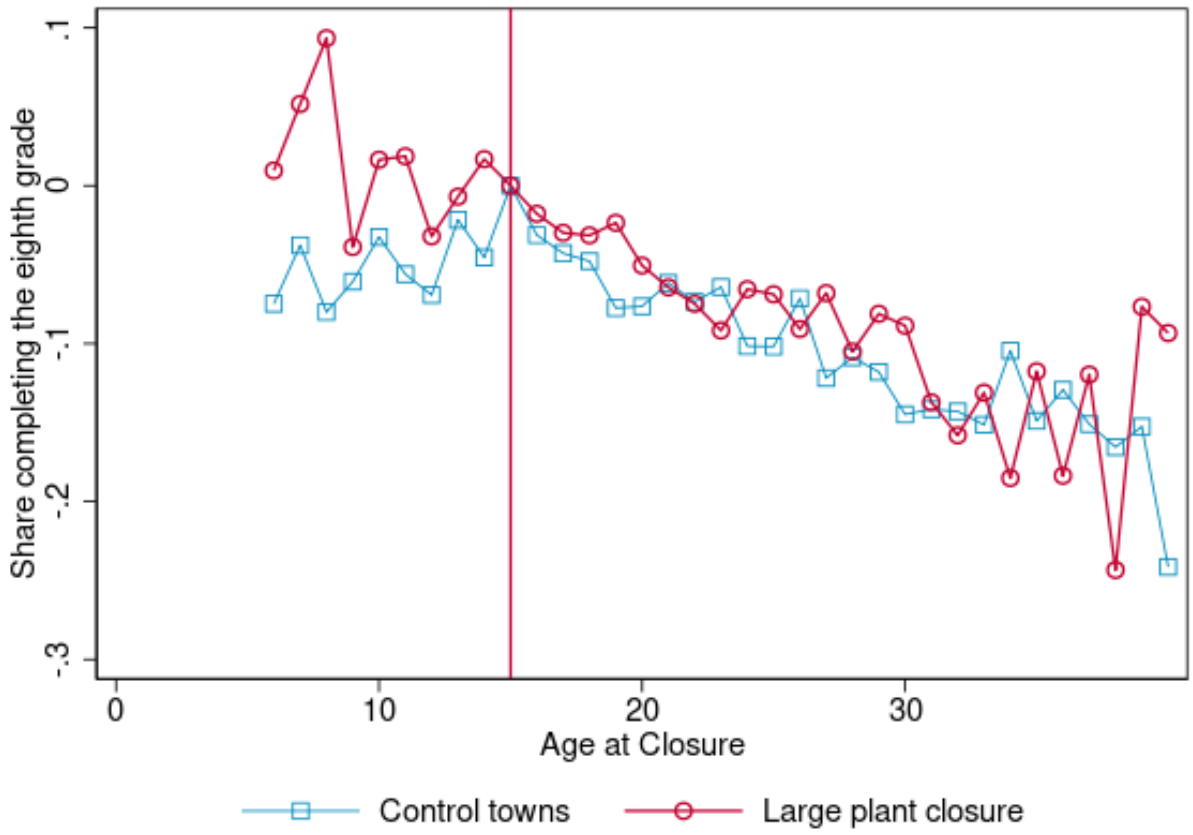
Notes: The figures plot share of each state-age cohort attending school in 1920. Using the school attendance variable (SCHOOL) from the 1920 census, I compute the share of each age cohort in New England states who reported attending school.

Table 1.6: Average town characteristics for closure towns and control towns

	Matched control towns	Large plant closure	Rest of New England towns
Working-age population	4771.889	3895.778	3979.167
Avg. socioeconomic index	10.034	9.810	9.645
Labor force participation rate	0.581	0.577	0.548
Share foreign-born	0.350	0.304	0.217
Share literate	0.913	0.907	0.944
Textile share of labor force	0.165	0.207	0.040
Agriculture share of labor force	0.166	0.199	0.336
Manufacturing share (net of textile)	0.118	0.096	0.110
Number of towns	36	36	687

Notes: The table summarizes average town characteristics for towns that experienced "large" plant closures, their matched control towns, and the rest of the towns in New England states. Large plant closures do not happen in Maine and Vermont, so the sample only includes Connecticut, Massachusetts, New Hampshire, and Rhode Island towns. The demographic characteristics are defined using the 1900 census.

Figure 1.14: Normalized trend of share completing at least the eighth grade



Notes: The graph plots the normalized trend of the eighth-grade completion share among plant closure town and matched control town population who completed at least eighth-grade of education, organized by age-at-closure cohorts. The large plant closure series is computed using the sample of individuals whose residences in 1920 were towns with large plant closures. The control town series is constructed using the sample of individuals located at matched control towns in 1920. The matched control towns are defined using the propensity score matching described in Section 1.6.3. The plant closure year is the earliest year a large plant closed in each closure town, and the closure year is assigned to their matched control towns. Using the closure year, I organize individuals in closure and control towns by their age at the time of the closure year. Then, I compute the average share of the population who completed the eighth grade of education and normalize each series by their average value at the cohort whose ages were 15 during the closure year.

Table 1.7: The effect of large textile plant closures on educational attainment, income, and migration

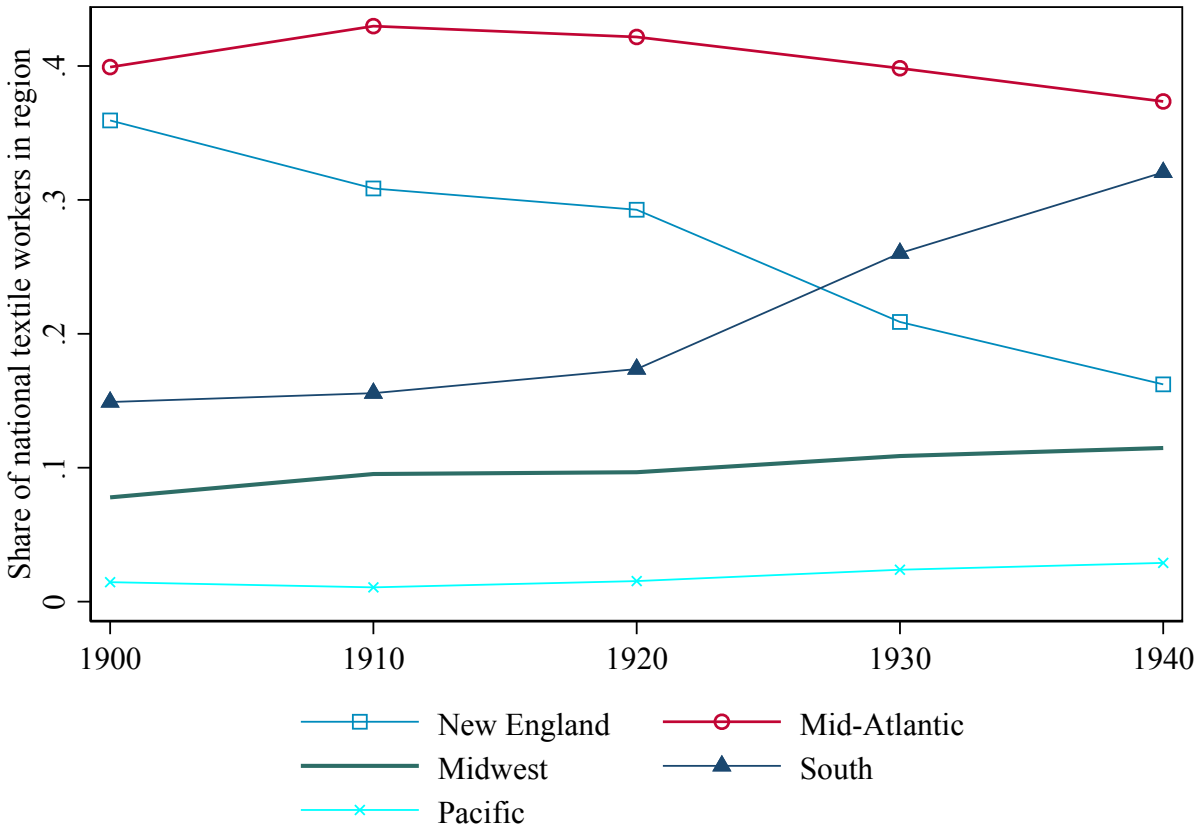
	At least eighth	At least ninth	Employed	Log wage income	Self-employed	Out-migrate
	(1)	(2)	(3)	(4)	(5)	(6)
Affected by plant closure	0.0316*	0.0123	0.0060	-0.0554*	-0.0033	-0.0039
	(0.0179)	(0.0188)	(0.0117)	(0.0285)	(0.0107)	(0.0219)
Birth cohort FE	X	X	X	X	X	X
Town FE	X	X	X	X	X	X
Dep. var. mean	0.811	0.556	0.877	6.938	0.094	0.471
Adjusted R^2	0.0658	0.0783	0.0286	0.0837	0.0172	0.0388
N	11960	11960	11960	9037	10485	11960

Notes: The table presents the analysis result using the matched difference-in-differences specification as described in Equation (3). The sample includes matched individuals (from 1920 to 1940 censuses) in plant closure towns and their match control towns, those whose ages were between 11 and 14, and those between 20 and 23 at the year of plant closure. Individuals affected by plant closures are those who were at the plant closure towns and were aged between 11 and 14 at the time of plant closure in their towns. Each specification includes birth-cohort fixed effects based on age (in the 1920 census) and town fixed effects. The outcome variables are constructed using their matched 1940 census variables. The first and second columns use the individual-level indicator variable, whether one completed at least the eighth and ninth grade of education, as the outcome variables, using the HIGRADE variable. Column (3) uses the indicator variable, whether one was employed as an outcome, using the EMPSTAT variable. Column (4) uses the log of wage income as an outcome, constructed from the INCWAGE variable. Column (5) uses an indicator of self-employment as an outcome using the EMPSTAT variable. Lastly, column (6) uses an out-migration indicator as an outcome, constructed using the state and town information from the 1920 and 1940 censuses.

Appendices

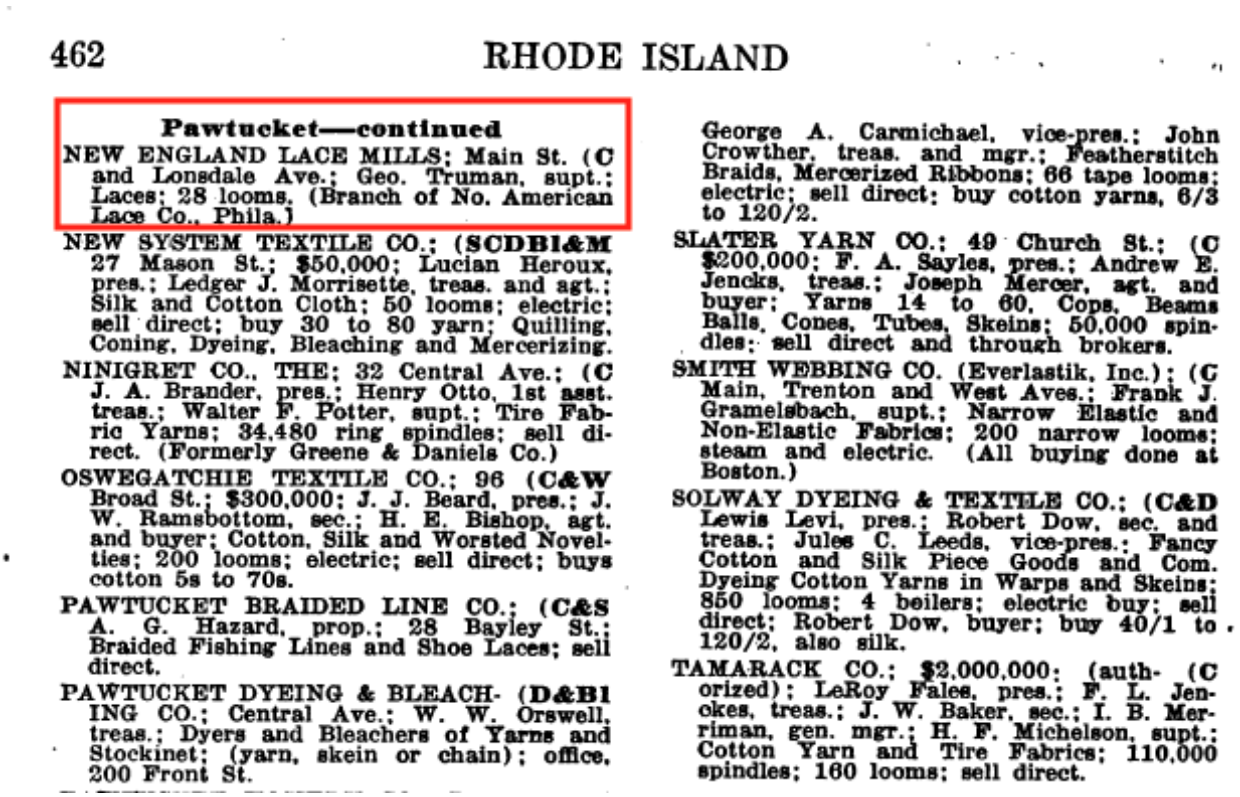
A Appendix Figures and Tables

Figure 1A.1: Share of the national labor force in the textile industry located in each region



Notes: The figure plots the share of the textile labor force located in each census region and area, computed from the IPUMS 1% census sample. The textile industry is defined as the group of 1950 census industries: (1) Knitting mills, (2) Dyeing and finishing textiles, except knit goods, (3) Carpets, rugs, and other floor coverings, (4) Yarn, thread, and fabric mills, (5) Miscellaneous textile mill products, (6) Apparel and accessories, and (7) Miscellaneous fabricated textile products.

Figure 1A.2: Structure of the Official American Textile Directory



462 RHODE ISLAND

Pawtucket—continued

NEW ENGLAND LACE MILLS; Main St. (C and Lonsdale Ave.; Geo. Truman, supt.; Laces; 28 looms. (Branch of No. American Lace Co., Phila.)

NEW SYSTEM TEXTILE CO.; (SCDBI&M 27 Mason St.; $50,000; Lucian Heroux, pres.; Ledger J. Morrisette, treas. and agt.; Silk and Cotton Cloth; 50 looms; electric; sell direct; buy 30 to 80 yarn; Quilling, Coning, Dyeing, Bleaching and Mercerizing.

NINIGRET CO., THE; 32 Central Ave.; (C J. A. Brander, pres.; Henry Otto, 1st asst. treas.; Walter F. Potter, supt.; Tire Fabric Yarns; 34,480 ring spindles; sell direct. (Formerly Greene & Daniels Co.)

OSWEGATCHIE TEXTILE CO.; 96 (C&W Broad St.; $300,000; J. J. Beard, pres.; J. W. Ramsbottom, sec.; H. E. Bishop, agt. and buyer; Cotton, Silk and Worsted Novelties; 200 looms; electric; sell direct; buys cotton 5s to 70s.

PAWTUCKET BRAIDED LINE CO.; (C&S A. G. Hazard, prop.; 28 Bayley St.; Braided Fishing Lines and Shoe Laces; sell direct.

PAWTUCKET DYEING & BLEACH- (D&BI ING CO.; Central Ave.; W. W. Orswell, treas.; Dyers and Bleachers of Yarns and Stockinet; (yarn, skein or chain); office, 200 Front St.

George A. Carmichael, vice-pres.; John Crowther, treas. and mgr.; Featherstitch Braids, Mercerized Ribbons; 66 tape looms; electric; sell direct; buy cotton yarns, 6/3 to 120/2.

SLATER YARN CO.; 49 Church St.; (C $200,000; F. A. Sayles, pres.; Andrew E. Jencks, treas.; Joseph Mercer, agt. and buyer; Yarns 14 to 60, Cops, Beams Balls, Cones, Tubes, Skeins; 50,000 spindles; sell direct and through brokers.

SMITH WEBBING CO. (Everlastik, Inc.); (C Main, Trenton and West Aves.; Frank J. Gramelsbach, supt.; Narrow Elastic and Non-Elastic Fabrics; 200 narrow looms; steam and electric. (All buying done at Boston.)

SOLWAY DYEING & TEXTILE CO.; (C&D Lewis Levi, pres.; Robert Dow, sec. and treas.; Jules C. Leeds, vice-pres.; Fancy Cotton and Silk Piece Goods and Com. Dyeing Cotton Yarns in Warps and Skeins; 850 looms; 4 boilers; electric buy; sell direct; Robert Dow, buyer; buy 40/1 to 120/2, also silk.

TAMARACK CO.; $2,000,000; (auth- (C orized); LeRoy Fales, pres.; F. L. Jenckes, treas.; J. W. Baker, sec.; I. B. Merriman, gen. mgr.; H. F. Michelson, supt.; Cotton Yarn and Tire Fabrics; 110,000 spindles; 160 looms; sell direct.

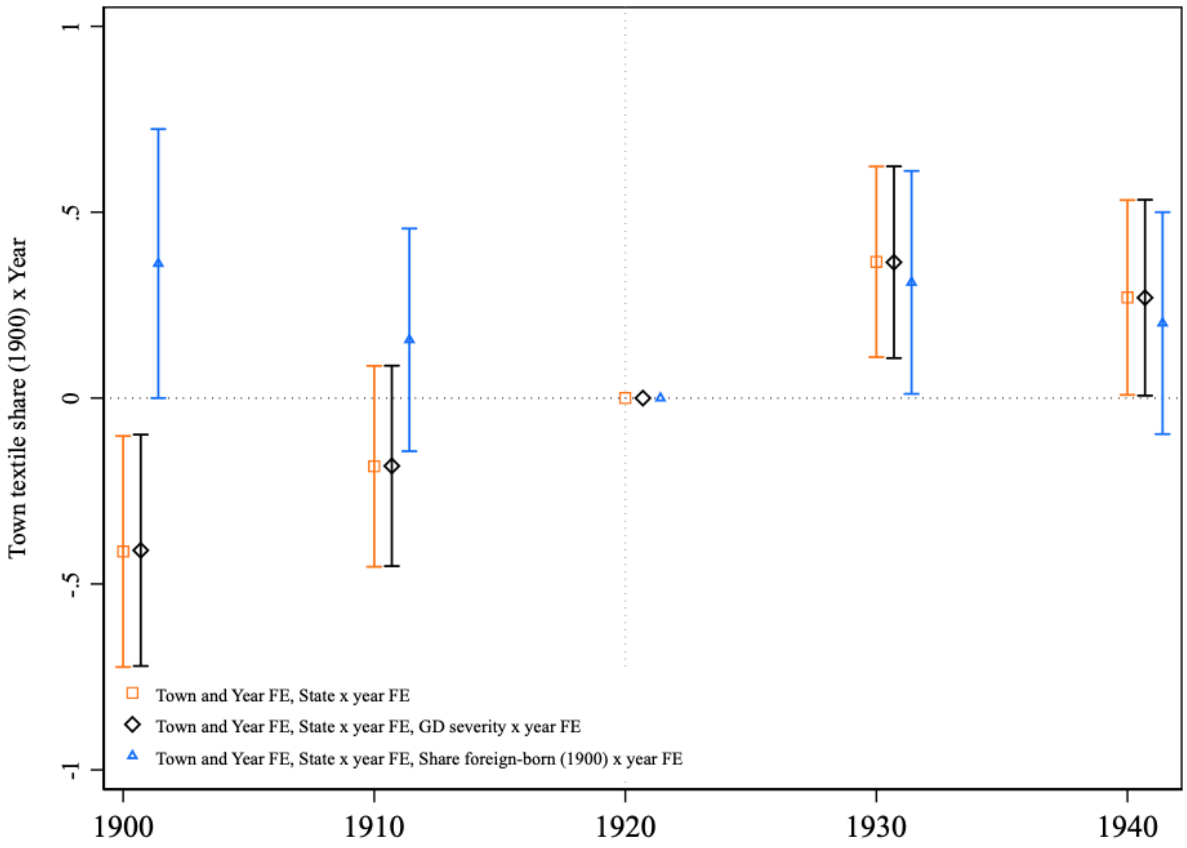
Notes: The figure presents an example of plant-level information contained in the Official American Textile Directory.

Table 1A.1: Coverage of the American Official Textile Directory

States	Textile Directory (1929)	Census of Manuf. (1929)
Connecticut	80	33
Maine	20	15
Massachusetts	241	135
New Hampshire	27	18
Rhode Island	142	55
Vermont	3	3

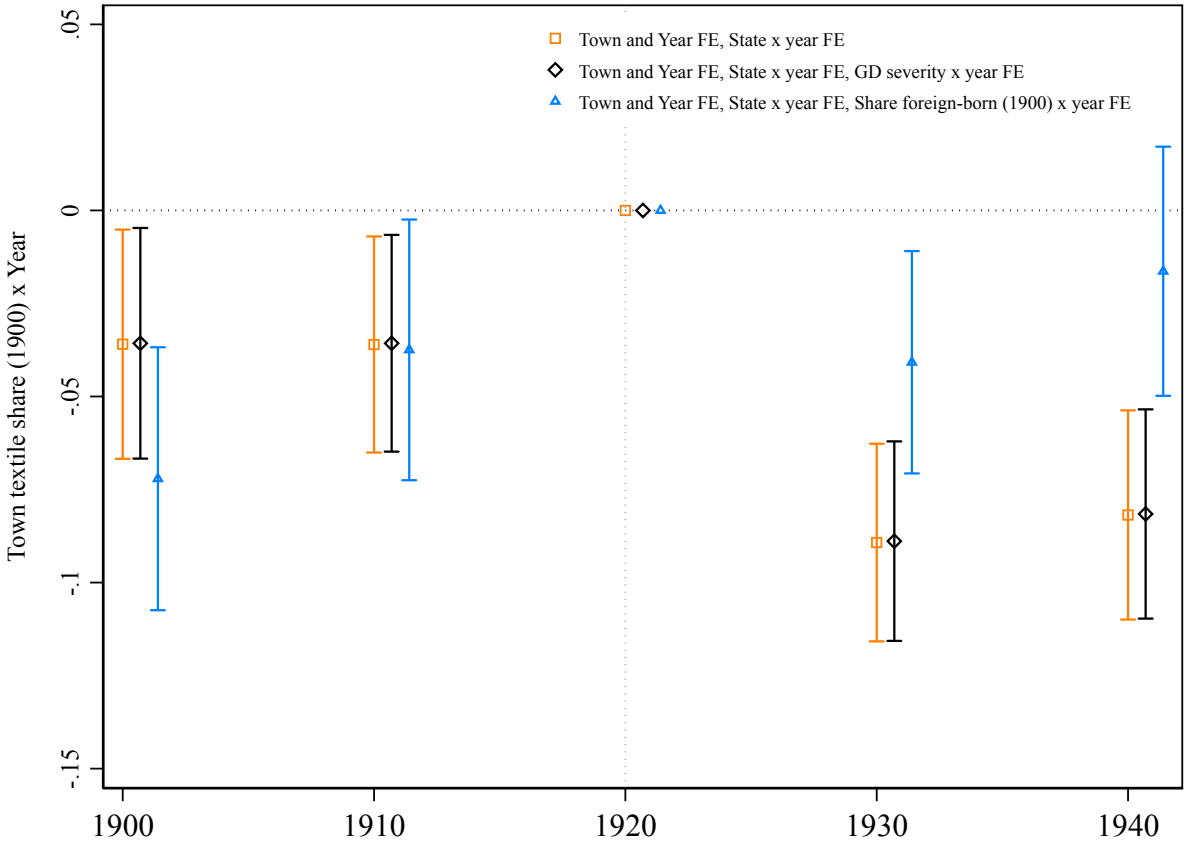
Notes: The table summarizes the number of cotton textile plants by state. The first column presents the number of cotton plants from the 1929 Official American Textile Directory, and the second column presents the statistics from the 1929 Census of Manufacturing digitized by Raff et al. (2015).

Figure 1A.3: The effect of exposure to textile industry decline on the log of native working-age population



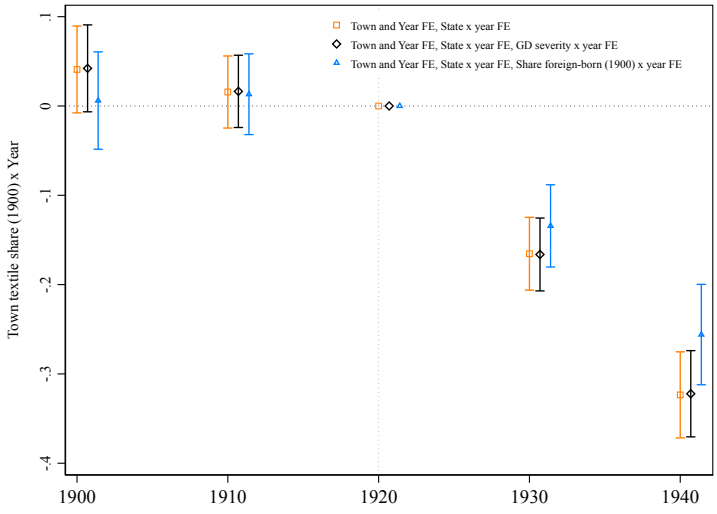
Notes: The figure presents the continuous difference-in-differences analysis results that use the log of native working-age population as the outcome variable. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state × year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1A.4: The effect of exposure to textile industry decline on labor force participation rate

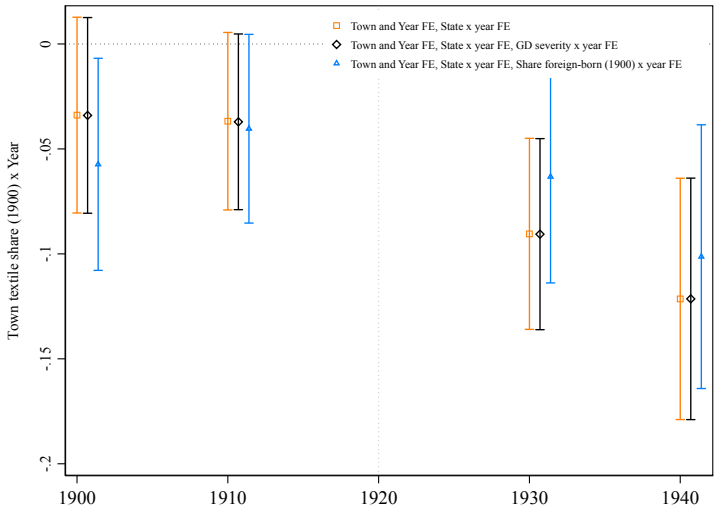


Notes: The figure presents the continuous difference-in-differences analysis results that use the labor force participation rate as the outcome variable. The labor force participation rate is computed by the share of the town's working-age (14-65) population who reported having a gainful occupation in each census year. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1A.5: The effect of exposure to textile industry decline on labor force participation by age group

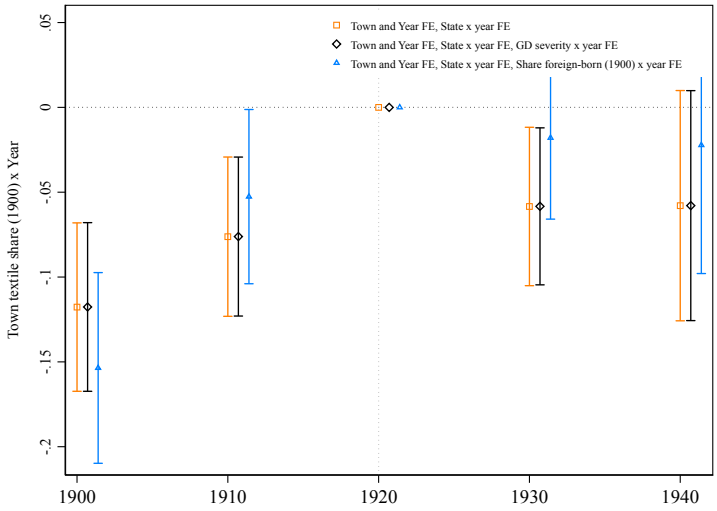


(a) Age 14-18

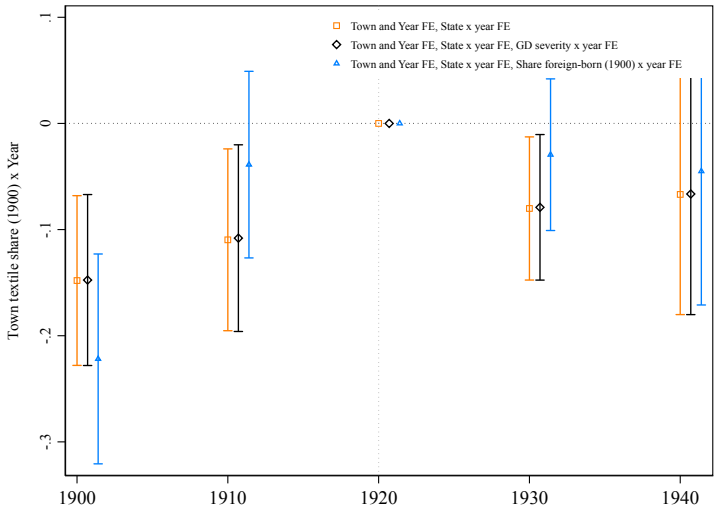


(b) Age 19-25

Notes: The figure presents the continuous difference-in-differences analysis results that use the labor force participation rate in each age group as the outcome variable. The labor force participation rate is computed by the share of four age groups of the town's working-age population who reported having a gainful occupation in each census year. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year



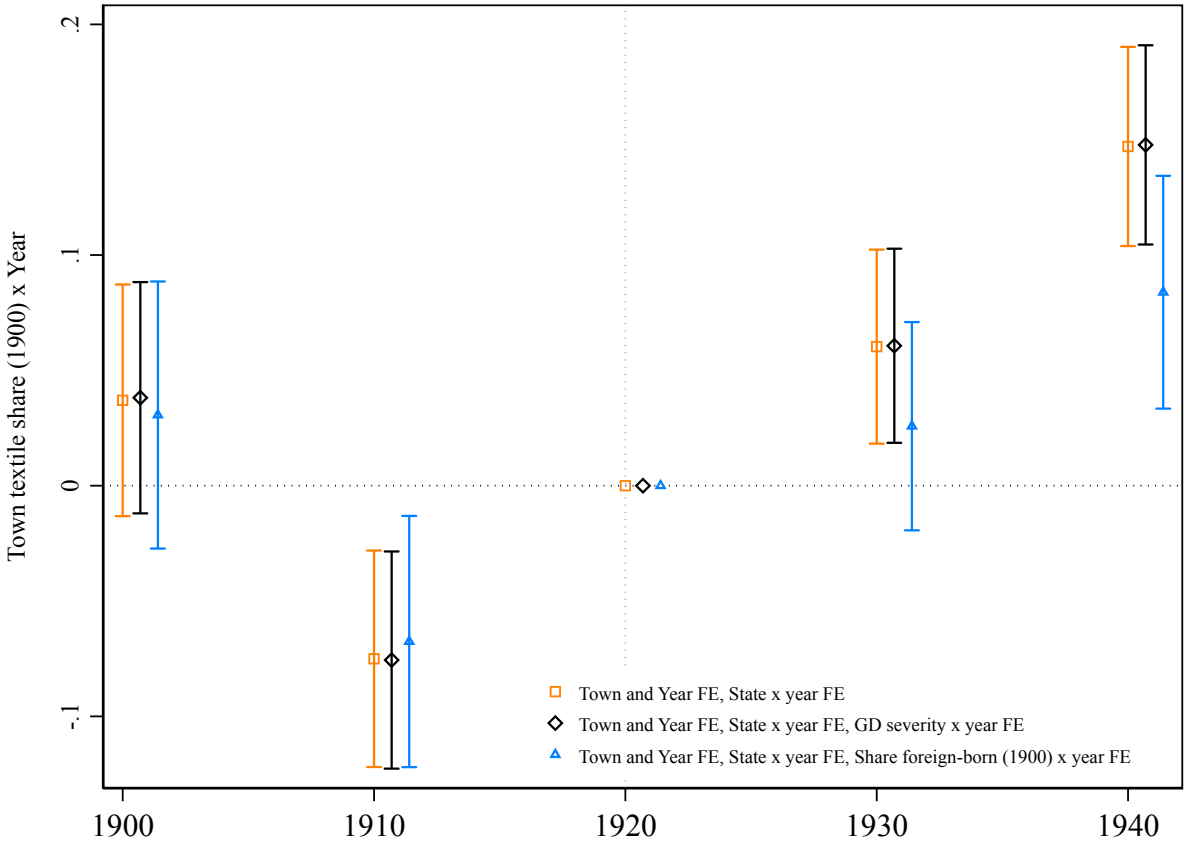
(c) Age 26-45



(d) Age 46-65

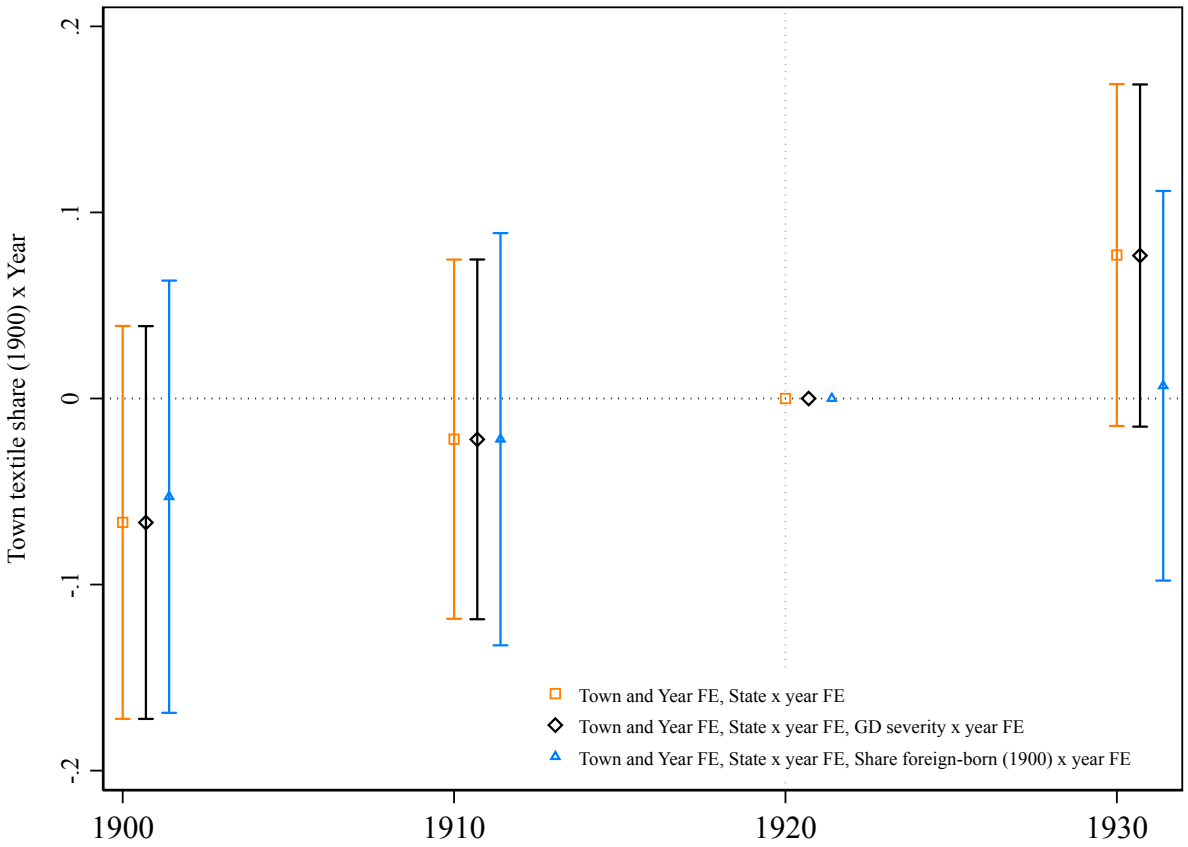
indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators.

Figure 1A.6: Effect of exposure to textile industry decline on school attendance rate among age 14-18 population



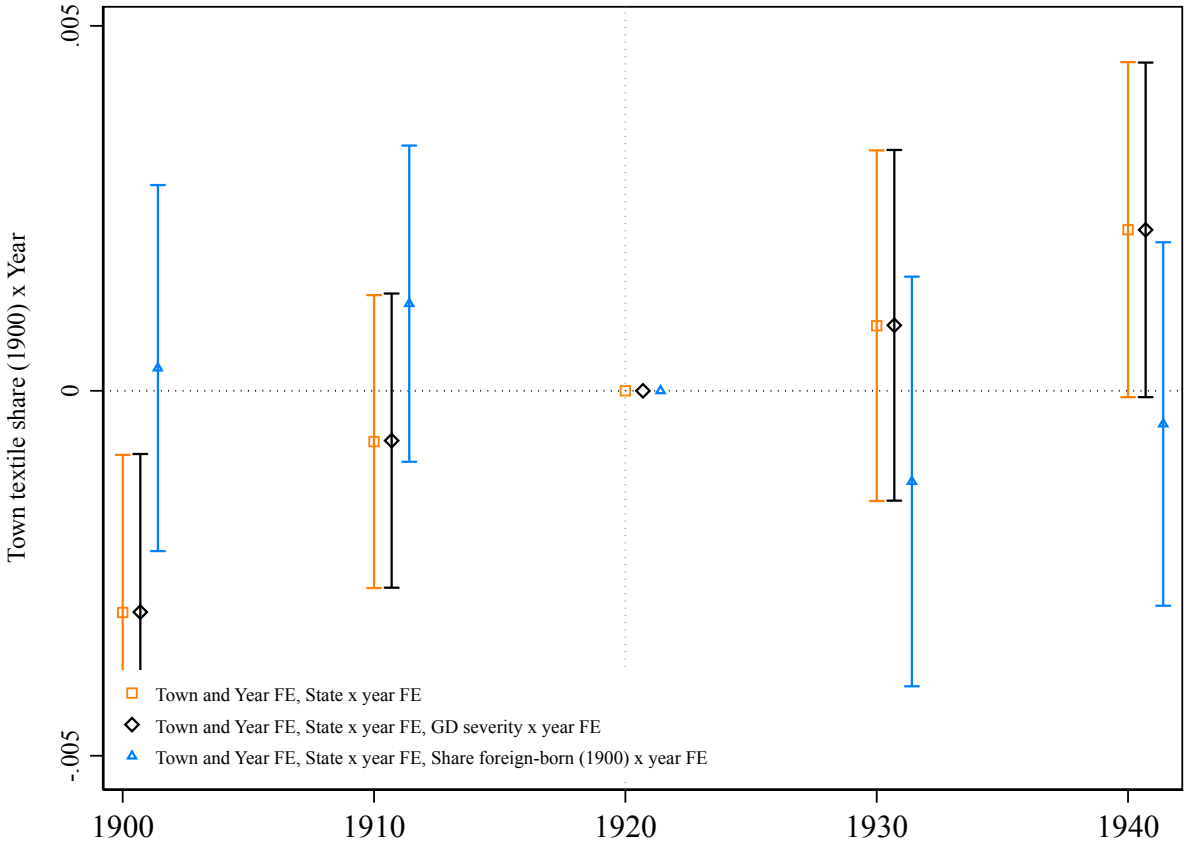
Notes: The figure presents the continuous difference-in-differences analysis results that use the school attendance rate as the outcome variable. The school attendance rate is computed by the share of the town population of 14-18 years of age who reported attending school in each census year. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1A.7: The effect of exposure to textile industry decline on the share of the population (age 11-14) with the eighth-grade completion



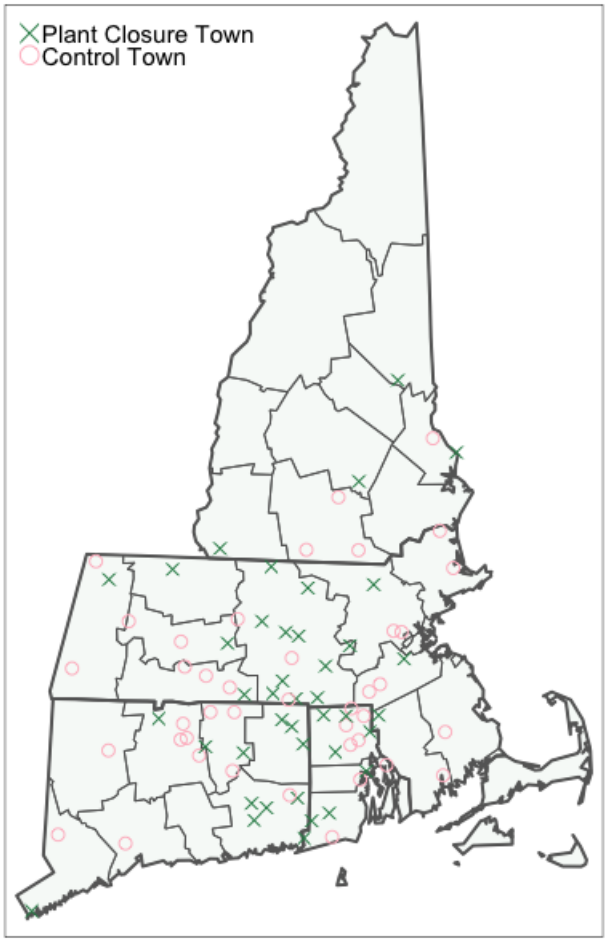
Notes: The figure presents the continuous difference-in-differences analysis results that use the share of the town population aged 11-14 who completed the eighth grade of education. Individuals aged 11-14 in New England towns are matched to the 1940 census to retrieve their completed years of schooling, and I compute the share of individuals whose highest grade of completion is higher than the eighth grade. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1A.8: Effect of exposure to textile industry decline on teacher per capita



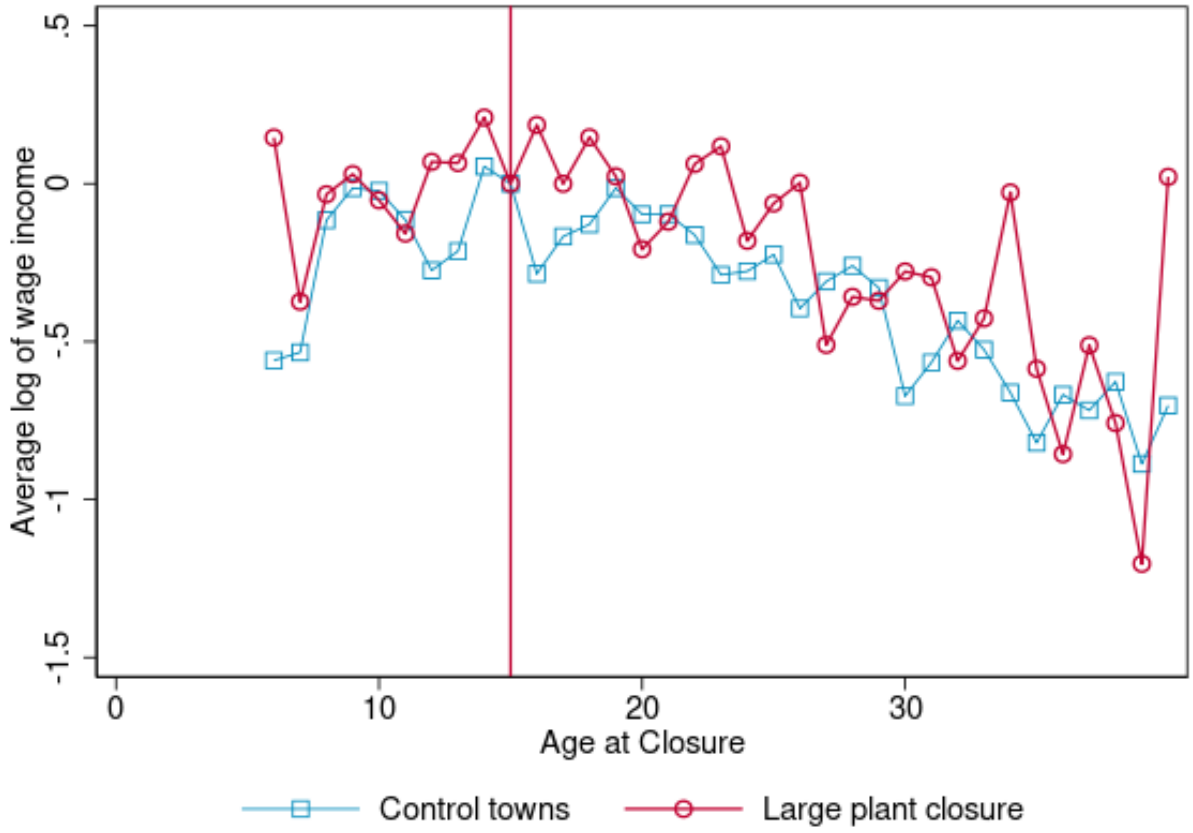
Notes: The figure presents the continuous difference-in-differences analysis results that use the number of teachers per capita as the outcome variable. The teacher per capita is computed by the number of individuals who reported their occupation as teachers divided by the town population. I only include those whose 1950 census occupation codes are 093 "Teachers (n.e.c.)," as most of the teachers are located in this category, as other teacher categories indicate art teachers (e.g., art, music, or dancing teachers). The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state × year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1A.9: Map of matched control towns and plant closure towns



Notes: The figure plots the map of Connecticut, New Hampshire, Massachusetts, and Rhode Island with the location of plant closure towns and their matched control towns.

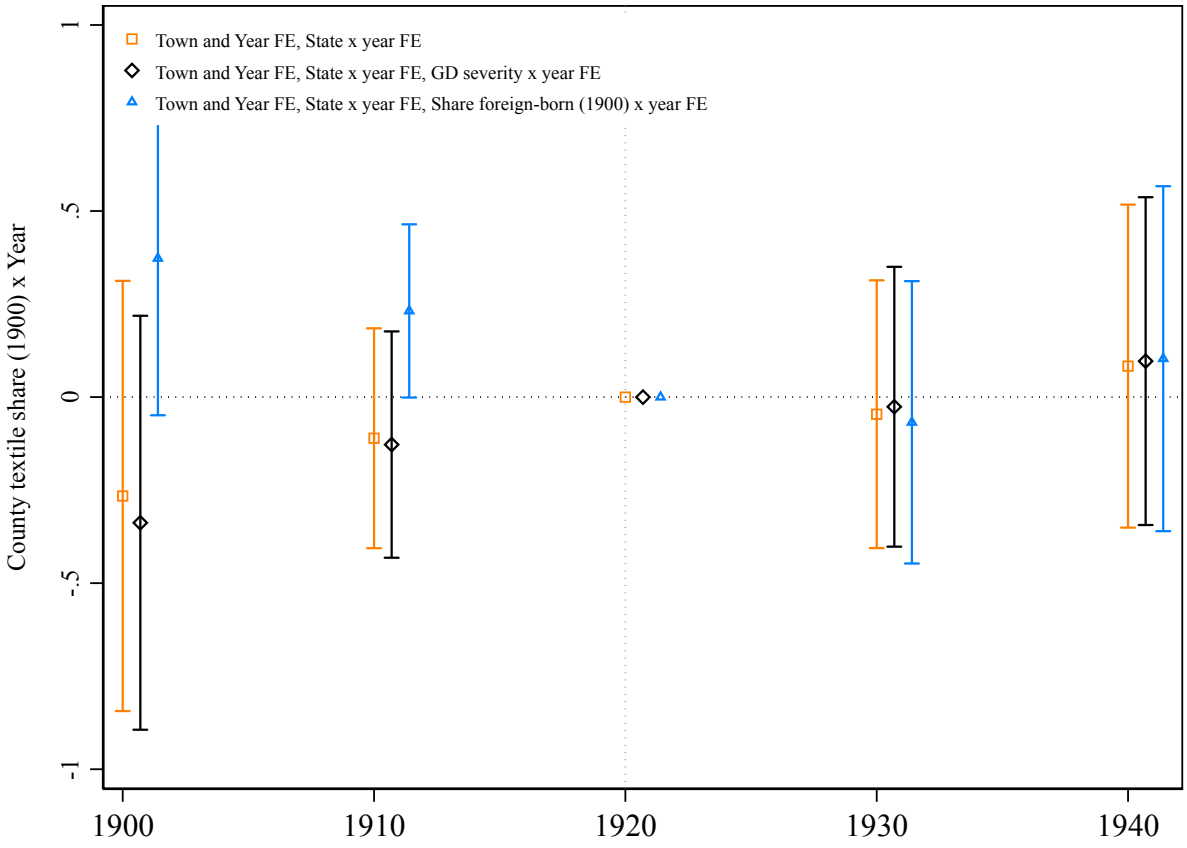
Figure 1A.10: Normalized trend of log wage income



Notes: The graph plots the normalized trend of the average of log wage income among plant closure towns and matched control town populations, organized by age-at-closure cohorts. The large plant closure series is computed using the sample of individuals whose residences in 1920 were towns with large plant closures. The control town series is constructed using the sample of individuals located at matched control towns in 1920. The matched control towns are defined using the propensity score matching described in Section 1.6.3. The plant closure year is the earliest year a large plant closed in each closure town, and the closure year is assigned to their matched control towns. Using the closure year, I organize individuals in closure and control towns by their age at the time of the closure year. Then, I compute the average log wage income and normalize each series by their average value at the cohort ages 15 during the closure year.

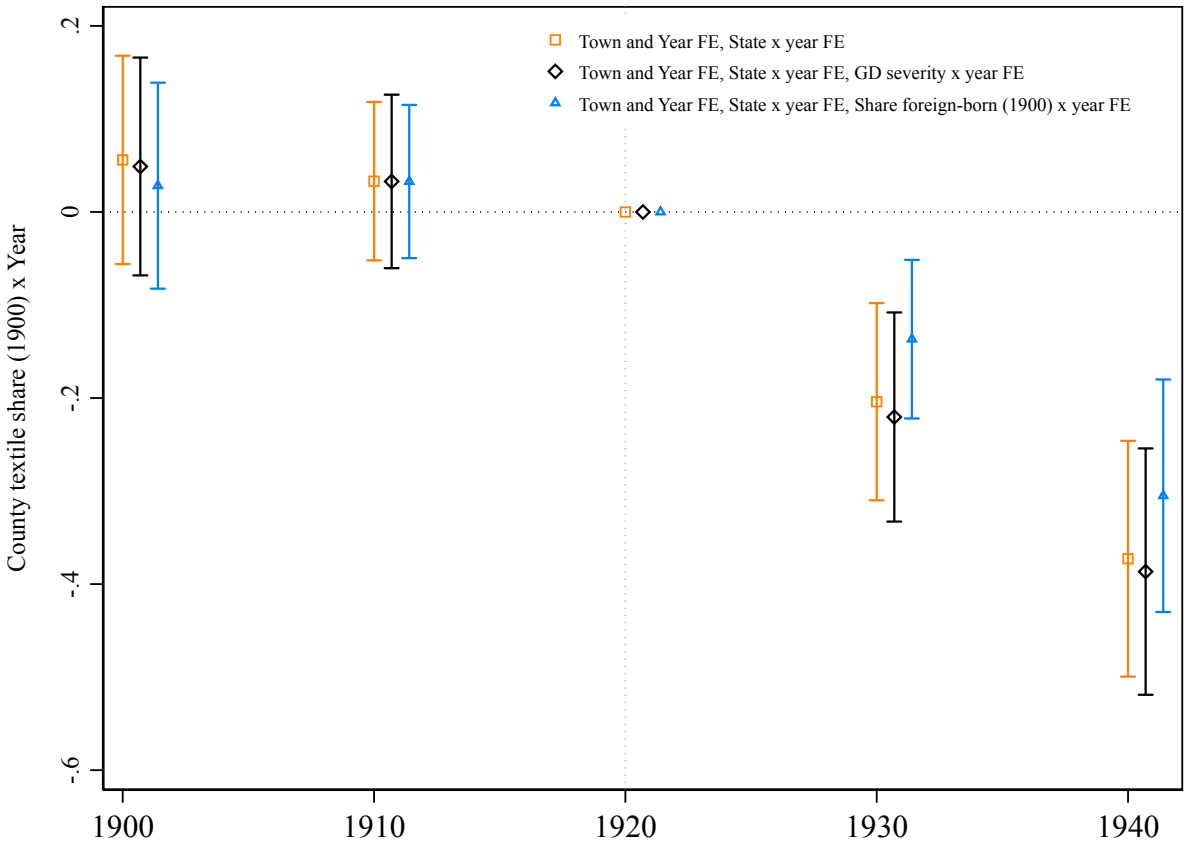
B County-level event-study analysis

Figure 1B.1: The effect of exposure to textile industry decline on county working-age population



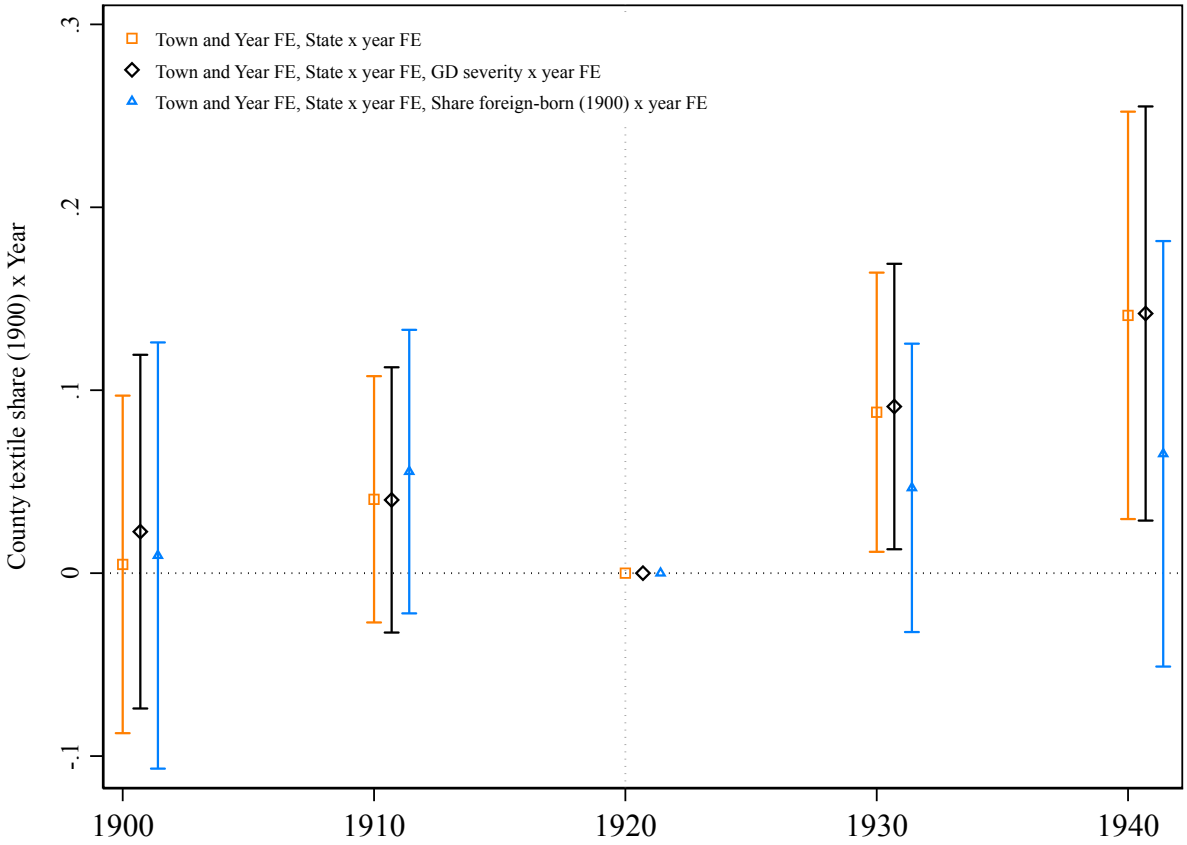
Notes: The figure presents the continuous difference-in-differences analysis results that use the log of the working-age population as the outcome variable. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each county (1900) interacted with census year indicators. The units of geography are counties defined from Berkes et al. (2022). All coefficient series include county and year fixed effects. The first series includes state × year fixed effects, including all interactions between state and census year indicators. The second series includes the county-level measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each county interacted with census year indicators. Standard errors are clustered by counties.

Figure 1B.2: The effect of exposure to textile industry decline on county labor force participation rate, school-going aged population (age 14-18)



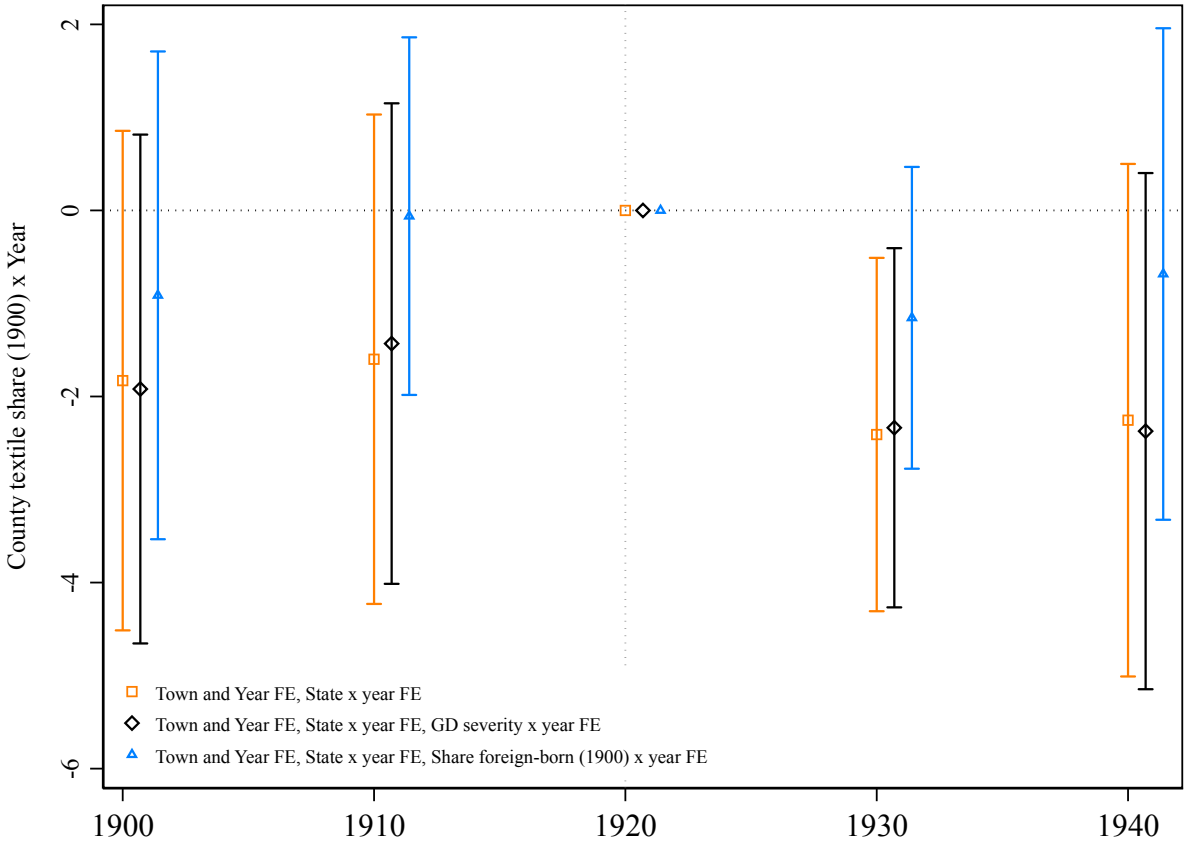
Notes: The figure presents the continuous difference-in-differences analysis results that use the labor force participation rate among the school-going aged population. Individuals with a gainful occupation recorded are considered in the labor force. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each county (1900) interacted with census year indicators. The units of geography are places counties defined from Berkes et al. (2022). All coefficient series include county and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each county interacted with census year indicators. Standard errors are clustered by counties.

Figure 1B.3: The effect of exposure to textile industry decline on the county-level share of agricultural labor force



Notes: The figure presents the continuous difference-in-differences analysis results that use the agricultural share of the county labor force as the outcome variable. The agricultural share of the county labor force is computed by the share of the county labor force who reported having agriculture (code 102 of IND1950) as their industry. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each county (1900) interacted with census year indicators. The units of geography are counties defined from Berkes et al. (2022). All coefficient series include county and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each county interacted with census year indicators. Standard errors are clustered by county.

Figure 1B.4: The effect of exposure to textile industry decline on county average occupation-based earning



Notes: The figure presents the continuous difference-in-differences analysis results that use the average occupational income score (OCCSCORE) as the outcome variable. The county-level average is computed using the OCCSCORE of county working-age residents who reported having a gainful occupation. The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each county (1900) interacted with census year indicators. The units of geography are counties defined from Berkes et al. (2022). All coefficient series include county and year fixed effects. The first series includes state × year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. Standard errors are clustered by county.

C Matching Algorithm and matched population characteristics

I use the automated linking methods from Abramitzky et al. (2022) that link individuals across censuses. The algorithm uses NYSIIS (New York State Identification and Intelligence System) standardized first and last names, ages, and places of birth of individuals to identify unique individuals in each census and link them to other censuses. I use the "ABE_NYSIIS_conservative" method, which requires individuals to be unique within a five-year age bandwidth. The individual-level identification number (histid) crosswalks from each census pair are downloaded from available in https://censuslinkingproject.org/data/.

Appendix Table C.1C.1 shows the match rate and average characteristics of the matched male group in 1900 (matched to 1910) and the male population in 1900. I find that the match rate is 17 percent. The matched individuals have very similar average demographic and economic characteristics compared to the overall male population. The only notable difference is that they are less likely to be foreign-born than the overall population. This difference can be explained by the matching process because immigrants are more likely to have rarer first and last names that are more prone to errors, and this leads to a higher false positive rate (Abramitzky et al., 2022).

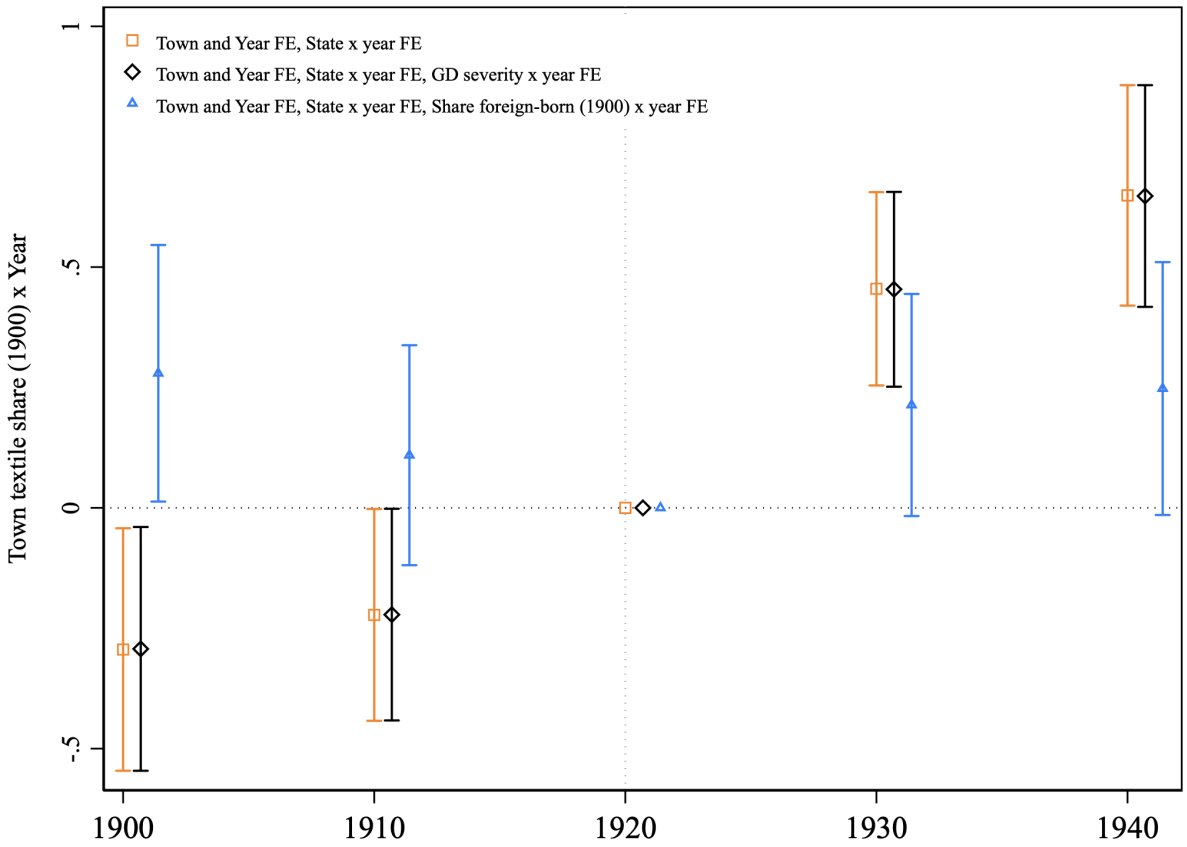
To ensure that this difference does not lead to differences between town-level analysis and matched individual-level analysis, I construct town-level statistics only using the matched individuals from the 1900 to 1940 censuses. I construct 1910, 1920, 1930, 1940 town-level statistics using the matched individuals in the destination census years, who are matched from previous years. For 1900, I use the matched individuals from the 1900 to 1910 census and construct the town-level statistics using their 1900 census information. For 1940Using the town-level data, I run the town-level event-study analysis and find that the results are analogous to the town-level analysis that used the entire town population.

Table 1C.1: Average characteristics of matched male individuals and population in 1900

	Matched (18-55)	All (18-55)	Difference
Average age	33.98	33.82	0.17
Share literate	0.95	0.93	0.02
Labor force participation	0.85	0.85	0.00
Textile share of labor force	0.05	0.06	-0.01
Share foreign-born	0.21	0.36	-0.15
			Match rate
N	261910	1542033	0.17

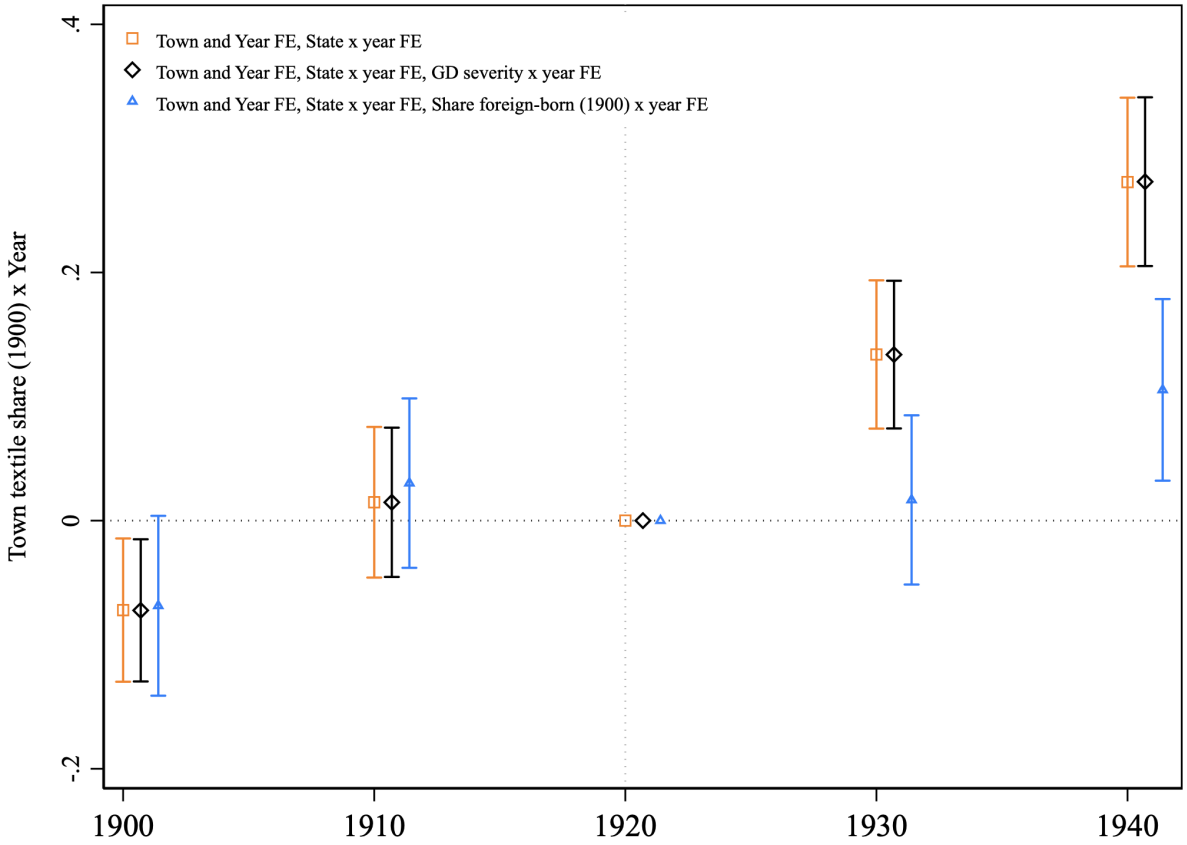
Notes: This table summarizes the average characteristics of the male individuals (age 18-55) used in the matched individual-level analysis, those of all male population (age 18-55), and the difference between the two. The matched individuals are those who are linked from the 1900 to 1910 census, using the "ABE_NYSIIS_conservative" linking method in Abramitzky et al. (2022). 'All (18-55)" column is constructed using all male population in New England from the 1900 census.

Figure 1C.1: The effect of exposure to textile industry decline on working-age population using the matched male individuals only



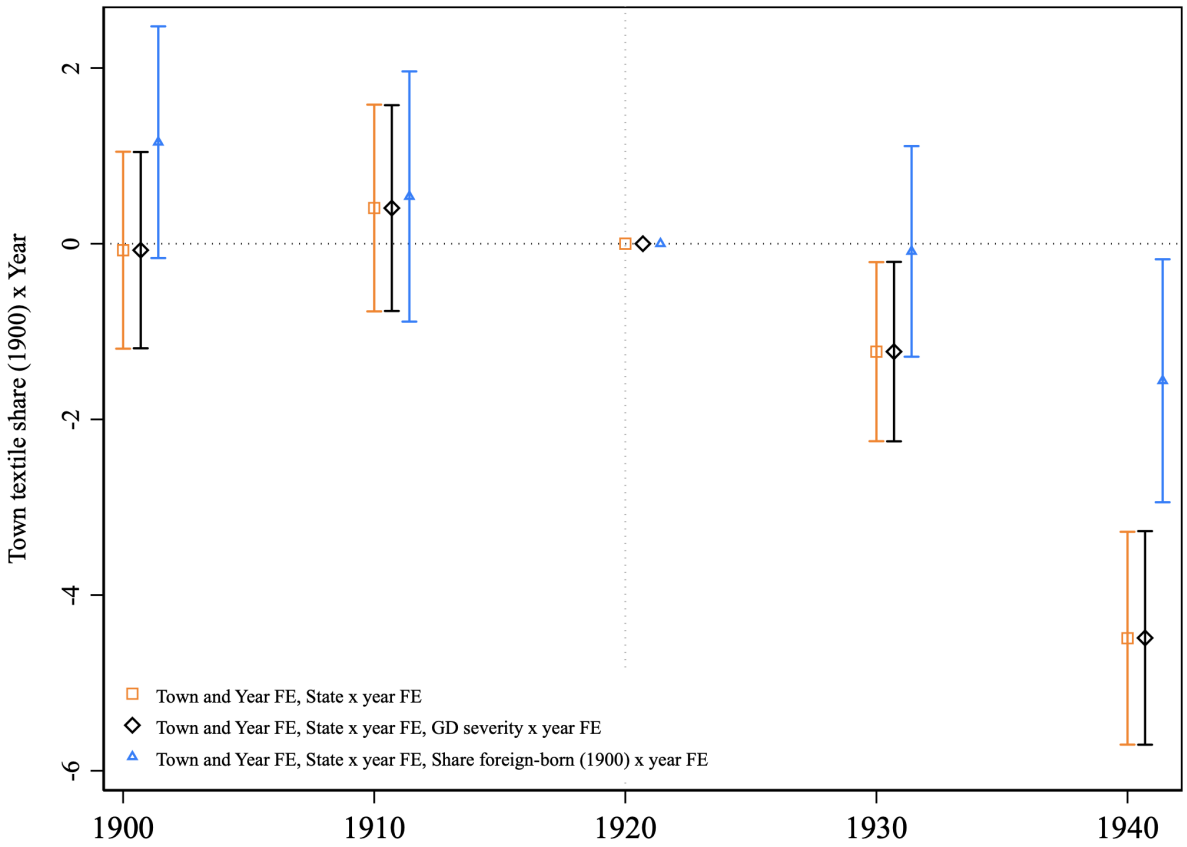
Notes: The figure presents the continuous difference-in-differences analysis results that use the log of the working-age population as the outcome variable, constructed only using the matched male individuals across censuses. The male individuals are matched to the next adjacent census, for instance, from 1900 to 1910, and from 1910 to 1920, using the automated linking method in Abramitzky et al. (2022). The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state × year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1C.2: The effect of exposure to textile industry decline on the agricultural share of the labor force using the matched male individuals only



Notes: The figure presents the continuous difference-in-differences analysis results that use the agricultural share of the labor force as the outcome variable, constructed only using the matched male individuals across censuses. The male individuals are matched to the next adjacent census, for instance, from 1900 to 1910, and from 1910 to 1920, using the automated linking method in Abramitzky et al. (2022). The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state $\times$ year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Figure 1C.3: The effect of exposure to textile industry decline on occupational earnings using the matched male individuals only



Notes: The figure presents the continuous difference-in-differences analysis results that use the average occupational earning score (OCCSCORE) as the outcome variable, constructed only using the matched male individuals across censuses. The male individuals are matched to the next adjacent census, for instance, from 1900 to 1910, and from 1910 to 1920, using the automated linking method in Abramitzky et al. (2022). The figure plots the coefficient estimates β_t from Equation (1), the coefficients of the initial textile share in each town (1900) interacted with census year indicators. The units of geography are places (towns) defined from Berkes et al. (2022). All coefficient series include town and year fixed effects. The first series includes state × year fixed effects, including all interactions between state and census year indicators. The second series includes the measure of Great Depression severity interacted with census year indicators. The third series includes the share of the foreign-born population in 1900 in each town interacted with census year indicators. Standard errors are clustered by town.

Chapter 2

Local Economic and Political Effects of Trade Deals: Evidence from NAFTA[1]

2.1 Introduction

In September of 1993, the Clinton administration released a letter signed by 283 economists, including twelve Nobel laureates, urging Congress to ratify the North American Free Trade Agreement (NAFTA). "The letter is part of a concerted White House campaign to rebut the criticisms of the trade agreement made by Texas billionaire Ross Perot, who has begun spending large amounts of his considerable fortune to promote his view that NAFTA will destroy American jobs," reported the *Los Angeles Times*.[2]

The White House indeed succeeded in passing NAFTA in a close and bi-partisan vote two months later, and it was implemented on January 1st, 1994. A quarter of a century later, it remains controversial. Donald Trump made opposition to NAFTA a key part of his

[1]This chapter was co-authored with Ilyana Kuziemko, Ebonya Washington, and Gavin Wright. We thank seminar participants at Columbia, Cornell, Harvard, LSE, Northwestern, Princeton, Stanford, and UCLA. We also thank conference participants at the Labor Studies meeting at the NBER Summer Institute and the International Trade and Institutions conference at the NBER. David Autor, Jeffry Frieden, Gene Grossman, Larry Katz, Suresh Naidu, Dani Rodrik, and Stephen Redding provided especially helpful feedback. Madhavi Jha, Paola G. Villa Paro, Qianyao Ye, Maria Canals, Arkey Barnett, and Joshua Thomas provided outstanding research assistance.

[2]See"283 Top Economists Back Trade Pact, Letter Shows," *Los Angeles Times*, September 4, 1993, https://www.latimes.com/archives/la-xpm-1993-09-04-mn-31519-story.html.

successful 2016 presidential campaign, claiming in the first presidential debate that "NAFTA is the worst trade deal maybe ever signed anywhere, but certainly ever signed in this country." While his presidential administration ultimately took pro-trade positions, Senator Barack Obama campaigned against NAFTA in the 2008 Democratic primary: "[T]rade deals like NAFTA ship jobs overseas and force parents to compete with their teenagers to work for minimum wage at Wal-Mart."[3]

This paper studies NAFTA's local economic effects as well as its political impact. We begin with event-study analysis showing that NAFTA had a significant, negative effect on employment in counties exposed to Mexican import competition. By 2000, counties in the top quartile of our NAFTA-exposure measure saw a 3–7 log point declines in total employment, relative to the bottom quartile. These results are robust to including *state* × *year* fixed effects as well as flexible controls for manufacturing share of employment, county-level education, the local "China shock" measure from the Autor et al. (2013), and other potential confounds.

The large employment losses do *not* lead to population declines in the counties most vulnerable to NAFTA, and we have the power to reject even small effects. Note that Autor et al. (2013) also find limited migration response to the China shock, so our result deepens the puzzle of why population does not appear to respond to these large, trade-driven employment shocks. By contrast, transfer-programs—Trade Adjustment Assistance (TAA), Disability Insurance (DI), and Supplemental Nutrition Assistance Program (SNAP)—exhibit significant increases in both applications and receipt. However, relative to the size of our estimated employment declines, the magnitude of these effects are small.

[3]For the Trump quote, see the transcript to the first presidential debate in 2016: https://www.washingtonpost.com/news/the-fix/wp/2016/09/26/the-first-trump-clinton-presidential-debate-transcript-annotated/. For the Obama quote, see the transcript to a February 12, 2008 speech: https://www.nytimes.com/2008/02/12/us/politics/12text-obama.html.

NAFTA was a major issue in the 1992 U.S. presidential campaign, with Ross Perot making opposition to it the major motivation for his surprisingly successful third-party campaign.[4] President Bill Clinton eventually won the election and made passage of NAFTA an early goal of his administration, which he accomplished via a close, controversial, and bi-partisan vote in November of 1993. Moreover, his support of NAFTA marked a major switch in Democratic-party policy toward trade, as in the 1970s and 1980s Democrats had been the more protectionist of the two major parties.

The second half of the paper focuses on the political impact of NAFTA and begins by showing that counties most exposed to NAFTA begin to turn away from the Democratic party in the early- and mid-1990s. While the top quartile of NAFTA-exposed counties (many in the upper South) were the most likely to vote Democratic in elections before NAFTA, they begin to trend Republican just as NAFTA is being debated and then implemented. By the early 2000s, they vote as or more Republican than any other quartile. We find analogous results for House elections.

While the Presidential and House election outcomes are at the county level, the rest of our political analysis is at the individual level and thus avoids concerns about ecological fallacy. We begin by showing that in the years since its passage, less than half of Americans approve of NAFTA, and disapproval is especially strong in the areas most vulnerable to it. Second, in repeated cross-section data from the American National Election Surveys (ANES) we show that, in each year of survey data from 1986 to 1992, Democrats enjoy a significant and steady advantage among those with protectionist views, but between 1992 and 1996 a significant number of protectionist voters move toward the GOP and remain there. Finally, in an ANES panel dataset from 1992 to 1994, we can look

[4]In the second presidential debate of 1992, Perot memorably said: "It's pretty simple: If you're paying $12, $13, $14 an hour for factory workers and you can move your factory South of the border, pay a dollar an hour for labor, ..., have no health care,...,no environmental controls, no pollution controls and no retirement, and you don't care about anything but making money, there will be a giant sucking sound going south." Perot captured 19 percent of the popular vote in 1992, making him the most successful third-party candidate since Theodore Roosevelt in 1912.

at the *same voters* over time during this key moment. We indeed find that a significant share of those who in 1992 express protectionist views have moved their party-identification toward the GOP by 1994. These effects are robust to flexibly controlling for a variety of demographic variables as well as views on other political and policy questions.

Our paper contributes to the literature on the local employment effects from exposure to import competition from low-income countries. Shortly after NAFTA's passage, Rodrik (1997a) warned that academics and policy-makers were underestimating the effects of globalization on workers in rich countries. In the U.S. context, the seminal work of Autor et al. (2013) highlighted the large and lasting employment effects of Chinese import competition on exposed U.S. communities. There has been more limited work of this type for NAFTA, an exception being Hakobyan and McLaren (2016a). Like Autor et al. (2013), they use Census data, so focus on longer (ten-year) differences than we do. In particular, they use decennial Census data and model industry-level effects of NAFTA (proxied as changes in earnings by industry from 1990 to 2000) as a function of both 1990 tariff levels and the *change* in tariff levels between 1990 and 2000.[5] Relative to both Hakobyan and McLaren (2016a) and Autor et al. (2013), our use of annual data as opposed to Census microdata (which are available at lower frequency) allow us to visually test for pre-trends and moreover show that breaks in trend are highly correlated in time with NAFTA's implementation.

While there has been limited work on local employment effects of NAFTA, there is a larger literature estimating its *aggregate* welfare effects.[6] These papers have generally

[5] A potentially important issue with including both the *change* in tariff levels from 1990 to 2000 and tariff *levels* in 1990 is that the two are nearly one-for-one (negatively) correlated, as tariffs are mostly stable from 1990 to 1993 and then from 1994 to 2000 almost all tariffs go to zero as a result of NAFTA. Thus, identification is reliant on the relatively small share of industries whose tariffs with Mexico do not go to zero by 2000.

[6] Romalis (2007) uses detailed trade flow and tariff data to estimate import supply and demand elasticities, finding a positive impact on the trade quantities but moderate impact on prices and welfare. Caliendo and Parro (2014) develop a structural general equilibrium model that incorporates the sectoral linkages (e.g., intermediate goods and input-output linkages) and show that NAFTA had a positive impact on U.S.'s welfare by 0.08 percent, while it increased Mexico's welfare by 1.31 percent.

found small (or null) aggregate net welfare effects for the US. Small or null average effects are not inconsistent with the large effects we find in NAFTA-vulnerable counties. In fact, these two results viewed together suggest that the employment and other economic declines we find in NAFTA vulnerable counties are likely *net declines in welfare* (and that workers and residents in non-vulnerable counties enjoyed *net benefits* from NAFTA). Even without evoking the small-on-average results from past papers, our results indicate that NAFTA increased spatial inequality in the US, as the places most vulnerable were *already* lower-income and less educated in the pre-NAFTA period (see, e.g., Moretti, 2012 and Diamond, 2016 on rising spatial inequality in terms of wages and education levels in the US).

We also contribute to a recent literature on the political effects of trade shocks. To date, this literature has found mixed results in the U.S. context. Dorn et al. (2020) find over 2002-2010 a mix of right-ward movement alongside growing polarization in areas facing Chinese import competition. While they conclude that overall the movement is rightward, they also find that exposed areas that begin more Democratic send increasingly *liberal* candidates to Congress. Che et al. (2022) find that over 1990–2000, counties more exposed to competition via the granting of Permanent Normal Trade Relations with China become more likely to vote Democratic. Papers on Germany and France find that greater import competition results in a larger vote share for the far right party (Malgouyres, 2017 and Dippel et al., 2015). In the British case, greater exposure to trade predicts votes for Brexit (Colantone and Stanig, 2018).

By contrast with the literature on the U.S. political response to the China shock, we find a clear shift in the Republican direction in places most exposed to NAFTA and among voters opposed to free trade. We suspect that the difference lies in the political saliency of NAFTA. NAFTA triggered a highly successful third-party presidential campaign in 1992 and remains a politically controversial to this day. NAFTA also involved a president (Bill

Clinton) breaking with the base of his party on a key issue. As we discuss in Section 2.6, NAFTA captured much more attention on network nightly news than did the later easing of trade relations with China.

The paper proceeds as follows. Section 2.2 provides background on NAFTA and describes how we measure local vulnerability to NAFTA. Section 2.3 outlines the empirical strategy for our event-study analysis. Section 2.4 describes the employment results, Section 2.5 the migration and transfer-program results and Sections 2.6 and 2.7 the political results. Section 2.8 concludes.

2.2 Measuring local vulnerability to NAFTA

By 1992, diplomats had hammered out details of an historic agreement to reduce trade barriers across the North American continent, though it awaited approval by the diplomats' governments. As trade between the US and Canada had mostly been tariff-free due to earlier agreements, the debate over NAFTA in the US focused on whether to liberalize trade with Mexico (thus our analysis will focus on import competition with Mexico).

President Clinton signed NAFTA into law in November 1993 after a close vote in both Houses of Congress (we defer the political history of NAFTA to Section 2.6). Many of its provisions went into effect in January of 1994. While growing before NAFTA, Mexican imports enjoy more rapid growth beginning in 1994, especially in industries with pre-NAFTA tariffs (see Appendix Figures 2A.1 and 2A.2).[7] Note that it is not until 2005 that China permanently supplants Mexico as the most important low-wage source of imports to the US. We emphasize this point because much contemporaneous discussion of NAFTA argued trade with Mexico was simply too small in value to have appreciable effects in the US.[8]

[7]Note that Krueger (1999) and Romalis (2007) documented similar trends using 1990s data.

[8]See, e.g., Krugman (2000). Though we should note that even before Autor et al. (2013) there was pushback against the idea that trade with Mexico and other low-wage countries was too small to matter:

A Construction of our measure of NAFTA exposure

Our exposure measure draws heavily from Hakobyan and McLaren (2016a), though we create county-level measures, whereas they examine exposure at the Public-Use Micro-data Area (PUMA) level. In spirit, it is also very similar to that used by Autor et al. (2013): it takes a vector of industry-level measures of exposure to import competition and, for each community, multiplies it by a vector of pre-period industry employment shares.

Following Hakobyan and McLaren (2016a), we begin by creating Mexico's "revealed comparative advantage" (RCA) in a given industry $j \in \mathbf{I}$, using 1990 (pre-NAFTA) data:

$$\text{RCA}^j = \frac{\left(x^{MEX}_{j,1990}/x^{ROW}_{j,1990}\right)}{\left(\sum_i x^{MEX}_{i,1990}/\sum_i x^{ROW}_{i,1990}\right)}. \tag{2.1}$$

In the numerator, $x^{MEX}_{j,1990}$ is the 1990 value of Mexican exports (to all countries except the US) in industry j, $x^{ROW}_{j,1990}$ is the 1990 value of the rest of the world's (ROW) exports (to all countries exluding Mexico and the US) in j. The ratio of the two expressions is roughly equal to Mexico's share of exports in industry j. Of course, the share will be in part driven by Mexico's size. The denominator adjusts for Mexico's *overall* share of all exports, not just those in industry j. Thus, the overall expression in equation (2.1) captures, in 1990, Mexico's relative advantage in producing exports in industry j relative to other industries $i \in \mathbf{I}$. We use data from the UN Comtrade bilateral export series to calculate the RCA for each industry j.[9]

How much a U.S. county is likely to be affected by NAFTA depends on its pre-period reliance on employment from industries with the following two characteristics: (a) Mexico

see Feenstra and Hanson (2000) on aggregation bias in input-output matrices masking the factor content of trade; Feenstra (2010) on how off-shoring of intermediate goods and services means that national statistics will understate the impact of trade; and Rodrik (1997b) that the mere threat of off-shoring to low-wage countries reduces bargaining power and thus wages of domestic low-wage workers even if the jobs ultimately remain in the US.

[9]Note that because we use many different data sources in this paper (and most of them are familiar to labor and trade economists) we do not describe data in detail in the main text. We refer readers to Appendix B for more detail on each data source.

has large RCA in that industry, *and* (b) the industry enjoyed pre-NAFTA tariff protection. These components are reflected in the measure below:

$$\textit{Unscaled vulnerability}_{c,1990} = \frac{\sum_{j=1}^{J} L_{1990}^{cj} \text{RCA}^{j} \tau_{1990}^{j}}{\sum_{j=1}^{J} L_{1990}^{cj} \text{RCA}^{j}}, \tag{2.2}$$

where L_{1990}^{cj} is employment of industry j in county c in year 1990 and τ_{1990}^{j} is the ad-valorem equivalent tariff rate of industry j in 1990.[10] As the units of this variable are not particularly intuitive, in most of the empirical work we divide it by the difference between the average unscaled vulnerability in the most vulnerable quartile and the average unscaled vulnerability in the least vulnerable quartile:

$$Vulnerability_c = \frac{\textit{Unscaled vulnerability}_c}{\mathbb{E}[\textit{Unsc vul}_c \mid c \in \text{ top quartile}] - \mathbb{E}[\textit{Unsc vul}_c \mid c \in \text{ bottom quartile}]}. \tag{2.3}$$

Thus, increasing $Vulnerability$ by one unit is equal to going from the average county in the least vulnerable quartile to the average county in the most vulnerable quartile.

Note that our *Vulnerability* measure uses only *pre-period* measures of both Mexican RCA and community-level industrial composition and thus does not pick up any endogenous reaction to NAFTA itself.[11] Note also it is a constant within county—as we take the τ^j values and employment levels from 1990, it captures how much tariff protection from Mexican RCA a county enjoyed in 1990. As it will serve as the key explanatory variable in our event-study analysis, we essentially ask how this fixed-over-time characteristic covaries with local employment every year of our sample period (so any negative covariance *before* NAFTA's passage in the form of pre-trends would cast doubt on our hypothesis).

[10] County-level employment data come from County Business Patterns (CBP) and tariff data from the U.S. Tariff database from Feenstra et al. (2002). See Appendix B.

[11] Of course, the downside to using pre-NAFTA measures is that the Mexican RCA in a given industry could be endogenous to existing tariffs in 1990 and thus may not fully reflect the true, long-run comparative advantage of Mexico in terms of factor endowments or technology. We suspect that such bias should be small for the short- and medium-run nature of the analysis in our paper.

While our τ^j_{1990} measure is taken from the specific year of 1990 and thus are by construction constant across time, the τ^j_t values can change over time. Figure 2.1 panel (a) shows, separately by quartile of 1990 vulnerability, how the protection measure in equation (2.2) changes if we allow the τ^j_t to follow their *actual* course over time (all other variables in the expression are kept at their 1990 levels, so the value of the four series in 1990 is in fact the average county vulnerability measure, as defined in equation 2.2, for the four groups). Before 1993, there is little change, as tariff rates were largely stable in this pre-NAFTA period. Between 1993 and 1995, there is a large decline in protection, as indeed half of all tariffs on Mexican goods went to zero in the first year after NAFTA's January 1994 implementation.[12] By 2000, even the most protected quartile of counties by the 1990 measure have essentially zero tariff protection.

Because (i) tariffs change very little between 1990 and 1993 and (ii) most tariffs go to zero between 1994 and 2000, there is an extremely high correlation between 1990 tariffs and the 1990 to 2000 *change* in tariffs. Thus, "protection" from Mexican import competition in 1990 is essentially the same as "vulnerability" or "exposure" to NAFTA and we use these expressions interchangeably.

B Characteristics of counties by NAFTA vulnerability

While Figure 2.1 panel (a) shows how tariff protection changed over time, panel (b) shows how protection in 1990 (and thus vulnerability to NAFTA) varies geographically. NAFTA most affected low-wage, labor-intensive manufacturing industries—textiles, apparel, shoes—which were concentrated in the South. But there are high-vulnerability pockets in most states.

[12] See U.S. Information Agency (1998), p. 25. One claim in the 1993 letter signed by economists and circulated by the Clinton administration was that the tariff reductions would be too gradual to create employment losses, but in fact those reductions were mostly complete within the first two years.

Even before NAFTA, those living in the most vulnerable quartile of counties were the least educated and had lowest per capita income, as we show in Table 2.1. They were also the most reliant on manufacturing employment and thus it is not surprising to see a positive correlation with the China-shock measure of Autor et al. (2013). As the most vulnerable quartile is disproportionately Southern, it is also less white, as African-Americans have always disproportionately lived in the South. The differences in pre-NAFTA characteristics by exposure quartile highlights the importance of flexibly controlling for these attributes in order to isolate the effects of NAFTA from secular changes such as skill-biased technological change (Goldin and Katz, 2008b) or the China shock.

Importantly, Table 2.1 also shows that the most NAFTA-vulnerable counties are the least supportive of Republican candidates in 1980s House and Presidential elections. By the 1980s, the Democrats had slowed their Civil-Rights-era losses in the South and the region remained majority Democrat during this period (in part due to the post-1960s entry of Black Democrats) and the party still enjoyed strong support in Northern industrial areas.

C Relationship between industry-level NAFTA vulnerability and Mexican imports to the US

While our *Vulnerability* variable is similar in spirit to the ADH measure, one departure is that we focus on statutory changes in tariff protection instead of changes in actual import penetration. We view this modification as preferable, as actual imports are potentially endogenous to U.S. demand (Autor et al., 2013 themselves note this concern, and thus use Chinese imports to *other* rich countries as an instrument in many specifications).

While we prefer to relate employment changes to statutory changes in tariffs instead of actual import penetration, here we use $industry \times year$ data to show that pre-NAFTA tariff levels (the large majority of which go to zero in the years immediately following

NAFTA) do indeed predict Mexican import growth after NAFTA's implementation. To demonstrate this "first stage" relationship we estimate:

$$MexImports_{jt}^{US} = \sum_{\tilde{t} \neq 1993} \beta_t \left(Avg\ tariff_{j,1990} \right) \times \mathbb{1}\left(t = \tilde{t} \right) + \gamma_1 MexImports_{jt}^{ROW} + \gamma_2 ROWImports_{jt}^{US} \eta_j + \mu_t + e_{jt},$$

where $Avg.Tariff_j^{1990}$ is the average tariff level on Mexican imports in industry j in 1990 (pre NAFTA), $MexImports_{jt}^{US}$ is Mexican imports to the US in industry j in year t, $MexImports_{jt}^{ROW}$ are Mexican imports in industry j in year t to the rest of the world (ROW), $ROWImports_{jt}^{US}$ are the rest of the world's imports to the US in industry j in year t and e_{jt} is the error term. We interact $Avg.Tariff_j^{1990}$ with a vector of year fixed effects (with 1993 as the omitted group). Including $MexImports_{jt}^{ROW}$ and $ROWImports_{jt}^{US}$ separates the NAFTA-triggered decline in industry j tariffs from, respectively, world demand for Mexican imports in j (not necessarily from the US) and U.S. demand for imports in j (not necessarily from Mexico).

Appendix Figures 2A.3 and 2A.4 document that the estimated β_t coefficients show a clear break in trend in the expected direction after NAFTA's passage (and that this pattern is robust to various data sources and functional forms).

2.3 Empirical strategy for event-study analysis

The next two sections examine local economic outcomes of those counties most exposed to NAFTA relative to other counties. For each outcome (employment, population, DI claims, etc.), we begin by showing trends for four groups of counties: four quartiles based on the NAFTA vulnerability measure that we defined in the previous Section. These trends are based on raw data, unadjusted except for normalization of each quartile to zero at 1993.

While this approach is the most transparent, it is difficult to summarize and to adjust for covariates. We thus turn to a standard event-study approach for the bulk of our analysis, where instead of dividing NAFTA exposure into quartiles we simply use (linearly) the measure in equation (2.3), interacting it with year fixed effects. In particular, we estimate:

$$Y_{ct} = \alpha_c + \gamma_t + \sum_{\tilde{t} \neq 1993} \beta_t \left(\textit{Vulnerability}_{c,1990} \right) \times \mathbb{1} \left(t = \tilde{t} \right) + \lambda X_{ct} + \epsilon_{ct}, \tag{2.4}$$

where Y_{ct} is a given outcome in county c in year t (e.g., log employment); α_c are county fixed effects; γ_t are year fixed effects, $\text{Vulnerability}_{c,1990}$ is c's NAFTA vulnerability; X_{ct} include additional controls (which we vary to probe robustness); and ϵ_{ct} is the error term. The exact sample period depends on the outcome variable due to data availability, but in general we begin in the mid-1980s and end in the mid-2000s. We cluster standard errors by state.

Note that this equation does not directly use the schedule of tariff reductions implied by NAFTA (and plotted earlier in Figure 2.1 (a)). Instead, we allow the 1990 level of tariff protection to have an unrestricted effect in each year, captured by the β_t coefficients. We prefer to take an agnostic approach as to how the effects of tariffs evolve over time. Moreover, allowing unrestricted effects of tariffs before 1994 allows us to test for pre-trends.

2.4 Employment results

In this section we document large employment declines in the counties most exposed to NAFTA, relative to less exposed areas, beginning in the mid-1990s. In the next section, we examine common margins of adjustment to local economic downturn (e.g., out-migration and transfer-program application).

A Main county-level event-study results

We begin by showing raw county data broken up into quartiles based on NAFTA vulnerability, though for the sake of space relegate these results to the Appendix. Appendix Figure 2A.5 shows these results for county-year log employment (all employment data in this section is taken from the Census' County Business Patterns data).[13] The raw data suggest that county groups trended together before 1994, after which time the most exposed counties fell behind the other groups in terms of employment growth.

While the plots of raw data that we show as Appendix figures have the virtue of transparency, more parametric event-study figures can more succinctly show robustness to flexible controls and other specification choices. In general, to avoid clutter, in the main text we show results from estimating only two versions of equation (2.4): including $State \times Year$ fixed effects (first series) and adding in addition the county's 1990s manufacturing share interacted with year fixed effects (second series). However, in the Appendix, we will always show at least three additional versions (often more): a version without any additional controls beyond county and year fixed effects (equivalent to setting the X_{ct} vector to empty in equation 2.4); a version that adds to the $State \times year$ fixed effects controls for CZ-level China shock interacted with year fixed effects; and a version that adds controls for 1990 county college share interacted with year fixed effects.

The first series (controlling for $State \times Year$ fixed effects) of Figure 2.2 shows that pre-1994, changes in employment have little relationship to county-level NAFTA vulnerability. But a clear negative trend takes hold in the years that follow NAFTA's implementation. While there is some modest recovery after the early 2000s, by 2008 there is still a statistically significant and negative relationship between 1990 NAFTA vulnerability and county-level changes in employment relative to the 1993 baseline. Adding

[13] Some of the smallest industry-county cells in the CBPD are imputed to protect confidentiality. Appendix Figure 2A.12 shows our main results are robust to using an alternative small-cell imputation proposed recently by Eckert et al. (2021).

flexible controls for 1990 manufacturing share somewhat reduces the magnitude of the decline but it remains statistically significant until the early 2000s. As shown in Table 2.1, counties that would be more exposed to NAFTA were already different on important dimensions in 1990: for example, they had lower rates of college completion and were more vulnerable to the China shock as well. In Appendix Figure 2A.6 we show that while flexibly controlling for these differences often decreases precision, the pattern of coefficients is largely unchanged relative to the second series in Figure 2.2.[14]

How large are these effects? In the year 2000, the estimates of the coefficient on *Vulnerability* in Figure 2.2 and Appendix Figure 2A.6 range from -0.029 to -0.064. That is, relative to 1993, the average county in the top *Vulnerability* quartile lost roughly three to six log points of employment compared to the average county in the bottom quartile. Appendix Figure 2A.10 shows similarly a relative loss of approximately 2.5-4.5 jobs per 100 county residents by 2000.[15] These magnitudes will be useful to keep in mind when we examine the response of migration and transfer programs in the next section.

Readers might find it surprising that there is not a negative pre-trend in our employment event-study graphs, given that NAFTA-exposed counties depended on manufacturing and the sector has been in long-run decline in the US. Indeed, textile and other NAFTA-vulnerable industry lobbyists often complained—even before NAFTA—that policymakers acted as though these jobs were in "sunset industries" and were thus "not worth saving."[16] But at the time of NAFTA's passage, the apparel and textile industries

[14]In Appendix Figure 2A.13 we also split the outcome variable into manufacturing versus non-manufacturing employment, showing that the response for manufacturing employment was larger. However, comparing the results before and after 1997 is complicated because in 1998 the CBP data change from Standard Industrial Classification (SIC) codes to the North American Industry Classification System (NAICS) codes.

[15]Recall that these employment-to-population ratios are based on employment in the county (regardless of the residence of the worker) from the CBPD divided by working-age population in the county (regardless of workplace, if any, of the resident) so is *not* directly comparable to the typical, survey-based measure of the share of a locality whose residents are employed.

[16]Much of the information provided in this paragraph is taken from Minchin (2012b), a history of the decline of the U.S. textile industry.

alone still employed nearly two million people. Whether via a successful "Made-in-America" campaign pitched toward consumers in the 1980s or other factors, employment decline had also plateaued in these industries in the years leading up to NAFTA. We show in Appendix Figure 2A.9 that employment in textile mills was in fact quite stable in the early 1990s (at half a million workers) before beginning a rapid decline in the middle of the decade, coincident with NAFTA's passage.

B Robustness to other trade and currency events

As noted, an important alternative explanation is that the effects in Figure 2.2 are merely picking up early stages of the China shock. As we show directly in Appendix Figure 2A.6, allowing CZ-level China-shock vulnerability to have arbitrary effects across time has little effect on our coefficients of interest relative to controlling flexibly over time for 1990 manufacturing, suggesting that our year interactions with 1990 manufacturing controls are doing a good job of controlling for the covariances between *Vulnerability*, the China shock and county employment.[17] In fact, controlling flexibly for either the China shock or 1990 manufacturing share makes little difference to our results until the early 2000s. We speculate that in the years immediately following China's accession to the WTO in December 2001, the China shock begins to dominate our NAFTA effect, and thus our estimates from this later period are more sensitive to these controls. While longer-run effects may be more sensitive, NAFTA's estimated effect on employment in the first 8-10 years after its implementation appears quite robust to various controls.

Another potentially confounding trade policy during the 1990s was the phase-out of the Multi-Fibre Arrangement (MFA) quotas (which was announced in the mid 1990s but did not become binding until the early 2000s, see Khandelwal et al., 2013 and Chiron, 2004 for

[17]To make the China-shock robustness test more demanding, we use the IV version of their measure, as it happens to be more highly correlated with our NAFTA-exposure measure. We use their IV defined for a period between 1990 and 2000, instead of between 2000 and 2007, as this period coincides more with the timing of NAFTA.

further detail and analysis). Appendix Figure 2A.7 shows robustness of our results to flexibly controlling for the MFA phase-out.

A sudden devaluation of the Mexican Peso took place in December 1994. The devaluation made Mexican goods relatively cheaper in the US and a natural concern is that it could have caused some of the employment effects that we attribute to NAFTA.[18] If the devaluation (and not NAFTA) caused the local employment effects, then we should observe them in counties whose employment is reliant on *any* industry for which Mexico is a strong exporter, *regardless of the industry's 1990 tariff levels*. Many SIC four-digit industries either had no or low tariffs on Mexican imports to the US and thus were unaffected by NAFTA, so we should be able to separate the two hypotheses. In Appendix Figure 2A.14, we replicate our employment results from Figure 2.2 but include as additional controls a *non-tariff-weighted* measure of vulnerability—that is, the expression in equation (2.2) but *excluding* the 1990 tariff industry tariff levels τ_j^{1990}—interacted with each year. Our results barely change, suggesting that the patterns we find in our main Figure 2.2 are driven by the decline in *tariffs*, not a more universal change in relative price levels between the two countries. Nonetheless, that the Peso crisis followed so closely after NAFTA passage suggests caution in generalizing the employment results from NAFTA to other trade agreements.[19]

While not directly related to trade policy, in Appendix Figure 2A.8 we also show robustness to controls for a county's 1990 share foreign born as well as how dependent the

[18] Our read of the literature is that the devaluation (and the economic turmoil that followed) was triggered by a number of factors: a large capital account deficit funded via short-term loans; a large share of debt held by foreigners; and "euphoria" related to the future prospects of a liberalizing Mexican economy. NAFTA may have played a role in the final factor ("the 'euphoria' was linked to the country being a 'model reformer,' as well as its access to NAFTA and OECD," Griffith-Jones, 1998) and if so then the peso crisis is not a confounder but a mechanism.

[19] Another piece of evidence suggesting that our results are being driven by actual pre-period tariff levels and not a more general relative price decline is the randomization inference exercise described in Section C and presented in Appendix C: we show in that exercise that a blunt "any positive tariff" measure does not pick up the same effects on county employment as using the actual 1990 tariff levels).

county is on jobs from "off-shore-able" industries or industries vulnerable to robots or automation.

C Addressing inference concerns

Recent work has suggested that designs such as ours may over-reject the null hypothesis of no effect. While our clustering by state addresses correlation of errors *within* states, Adao et al. (2019) shows that there is an additional correlation of errors *across* geographic entities with similar baseline industry composition. In Appendix C we show all the main results of the paper are robust to using their proposed correction.

We also develop a related but more demanding randomization-inference test, again detailed further in Appendix C. The distribution of pre-NAFTA industry-level tariffs τ has a mass at zero and then a very long right tail. We retain $\tau = 0$ for all industries that have no tariff against Mexican imports in 1990. We then model the positive tariffs with a fourth-degree polynomial (the actual distribution of positive tariffs and our approximation is in Appendix C). One implication of this procedure is that the mean β of the distribution formed by permuting across the τ distribution need not be zero, because by retaining the actual tariff value when $\tau = 0$ the simulations contain some real information. As we show in the Appendix, our estimated $\hat{\beta}$ is in the extreme tails (often not even in the support) of the distributions of simulated $\hat{\beta}$ values after 1,000 draws.

D Related results and additional robustness checks

So far, we have focused only on the heightened competition some U.S. industries faced due to NAFTA. But industries can benefit from NAFTA in at least two ways: they may rely on *inputs* that are now cheaper or they can export goods to Mexico more competitively due to reciprocal declines in Mexican tariffs on U.S. goods.

In Appendix D, we show that accounting for these potential benefits makes little difference to our employment results. While we go into detail in Appendix D, the key point is that NAFTA mostly helped and hurt the same industries (so that our *Vulnerability* measure largely picks up the *net* effect of NAFTA on local labor markets, accounting for both the local gains and losses). Input-output matrices demonstrate that most industries rely heavily on inputs from other industries in their same two-digit classification. Similarly, the sectors whose tariffs were reduced due to NAFTA were very similar in the US and Mexico, a point highlighted by policy-makers at the time (U.S. International Trade Commission, 1993).

Our main results so far use county-year data. We prefer the county over the CZ as our unit of analysis because CZs can cross states (21 percent of CZs, comprising 28 percent of the population, cross state lines) and in the political analysis especially we would like to control flexibly for any effect of state-wide campaigns. While CZs have the advantage of better capturing labor markets, counties are in fact decent proxies for labor markets as well: in the average county (CZ), 71 (92) percent of residents also work in the county and 78 (93) percent of workers also reside in the county (see Table 1 of Monte et al. (2018)). The employment results for CZs are very similar to those we find at the county level, as shown in Figure 2A.11, though the impact of NAFTA manifests more as a negative break in a positive pre-trend.

Finally, we show that our county level results are largely unchanged when instead of using (log) employment as the outcome variable, we use the (log) wage bill.

E Results at the individual level

Most of the analysis of trade-induced employment effects in the literature are, like our results so far, at the geographic level. Of course, county- or CZ-level results are of interest in their own right as they pick up potential spillover effects to other local industries. But

interpreting these results as informative of the individual-level effect of working in a NAFTA-vulnerable industry is subject to the ecological fallacy.

To more credibly estimate individual-level effects, we turn to the Panel Study of Income Dynamics (PSID). We define an *individual worker i's vulnerability* to NAFTA based on the industry j of their main job in 1990. That is, $\textit{Vulnerability}_{j(i)} = \text{RCA}_j^{1990}\tau_j^{1990}$. We show results in Appendix Figure 2A.16. Reassuringly, there is a clear decline in employment probability after 1993 for workers in more NAFTA-vulnerable industries.

2.5 Migration and transfer-program response

In this section we ask how individuals in NAFTA-vulnerable counties reacted to the large job losses we documented in the previous section. The two margins we focus on are migration and (applications to and receipt of) transfer programs.

A Migration

Economists have long studied how migration responds to local economic shocks. Blanchard et al. (1992) found significant migration responses using data from the 1970s and 1980s. While *employment levels* often never recovered from economic shocks during this period, via the migration channel, *unemployment rates* generally did. But researchers studying more recent local employment shocks have found much smaller migration responses. The large employment effects of the China shock produced no (Autor et al., 2013) or small and delayed (Greenland et al., 2019a) population effects. Similarly, Yagan (2019b) finds no statistically significant effect of the local severity of the Great Recession and out-migration from one's CZ.

We estimate a variety of event-study specifications and find no evidence of out-migration as a response to NAFTA. Figure 2.3 (Appendix Figure 2A.17) is the

analogue to Figure 2.2 (Appendix Figure 2A.6) except that log county population is the variable of interest. In contrast to the log-employment results, which showed a downward trend break in 1994 for all of our specifications, we find a series of null results and if anything the patterns hint at population *stabilization* after several years of decline in the pre-period. We can rule out even small declines. In the *State* × *year* fixed-effects specification, the bottom of the 95% confidence interval suggests a loss of population since 1993 for the top versus bottom quartile counties of 1.37 log points. In the previous section, the parallel specification (using the point-estimate, not the edge of the confidence interval) suggested a 6.4 log-point *employment* decline.

We perform a number of robustness checks to probe this null result. In Appendix Figure 2A.18 we add controls for past population growth (as suggested by Greenland et al., 2019b). In Appendix Figure 2A.19 we account for the fact that for many hard-hit counties, the counties to which residents most often migrated in the pre-period were *also* highly NAFTA vulnerable (as suggested by Borusyak et al., 2022). None of these specifications produce the expected, negative population result. While we do not include them in the interest of space, we find similar null results when we replicate Figure 2.3 for various population subgroups (e.g., working-age population, male and female population, the white and Black population).

We conclude that despite the large employment effects in NAFTA-vulnerable counties after 1993, their population growth tracks the rest of the country. This result echos historians' description of 1990s Southern mill towns after a major textile employer closed. "Workers' attachments to their jobs and communities—which had been so important as they endured the hardships of mill life—now made it harder for them to find opportunities. These workers failed to fulfill economists' predictions of a new, mobile workforce who would rationally relocate to find new jobs" Minchin (2012b).

B Trade Adjustment Assistance

Of course, policy-makers are not completely naive to the possibility of local job losses due to import competition, from NAFTA or other sources. Legislation originating in the 1960s and further expanded in the 1970s created a series of measures collectively known as Trade Adjustment Assistance (TAA). Beyond income support, TAA provides opportunities for training, job search and relocation payments.[20]

TAA application and certification data by county-year is extremely skewed: the majority of observations are zero and a few outliers pull up the mean substantially. Log measures are thus not feasible and we instead begin by estimating per capita applications (dividing by 1990 county population) as the outcome in our usual event-study set-up. Figure 2.4 (and Appendix Figures 2A.20 and 2A.21) provides the results for TAA applications. We find no pre-trends in per capita applications. In the years immediately after NAFTA's implementation, the coefficient on vulnerability is statistically significant and hovers between 0.0005 to 0.001 (increasing from the zero baseline of 1993). While the coefficients remain positive in the 2000s, they are no longer significant. The analogous results for TAA *certifications* are in Appendix Figure 2A.22 and show a very similar pattern.

While these effects appear statistically significant in much of the mid and late 1990s, the economic magnitudes are modest: an increase since 1993 of between 0.05 and 0.10 TAA petitions per 100 workers in the most- versus least-vulnerable county quartiles. Even using the smaller estimates we find of job loss per capita in Appendix Figure 2A.10 (3.73), the ratio of TAA petitions to jobs lost would be roughly 1.3 to 2.7 percent (note that Autor et

[20]To receive TAA benefits, a group of three or more workers must first file a petition with the U.S. Department of Labor's TAA Program within a year of separation from the firm. If the group of workers meets the eligibility criteria, they will be issued a group eligibility certification. Each worker in the group then must make an individual application for TAA benefits through their local American Job Center. Hyman (2018) is one of the few economics papers that studies its efficacy. He uses assignment to investigators with varying leniency and finds that certification leads to short-run benefits that appear to fade within ten years.

al., 2013 also find a very small increase in TAA petitions in areas hit by the China shock). Of course, our definition of "NAFTA related" is an econometric one—county job loss correlated to 1990 county NAFTA vulnerability occurring from 1994 onward, conditional on a large set of controls—whereas the definition used by TAA investigators will be different. But these small effects motivate us to ask whether individuals in NAFTA-affected counties turned to other transfer programs.

C Disability Insurance

At least since Autor and Duggan (2003), economists have studied whether individuals exposed to negative local economic shocks turn to the federal Disability Insurance (DI) program. Several mechanisms might operate (and we do not seek here to adjudicate among them). On the one hand, those with marginal health issues but still capable of gainful employment might turn to DI for income support if work opportunities dry up. So, holding health status constant, lack of jobs could push marginal candidates to apply to DI (what public finance economists would view as moral hazard). On the other hand, job loss could harm health—mental health issues given the link between job search and depression (Krueger et al., 2011) and physical health issues, given loss of employer health insurance. Minchin (2012b) describes loss of employer insurance as one of the biggest concerns of those who lost textile jobs in the 1990s.

To test whether NAFTA led residents of exposed areas to apply to DI, we obtain office-year DI application and award counts, from 1989 to 2008, from the SSA.[21] We use contemporary district office locations to assign zip codes to district offices.[22] We then

[21] We are deeply indebted to Manasi Deshpande for facilitating our access to these data and answering our many questions and to Melissa Kearney for sharing her extract.

[22] While the data do include zip code information for many district offices in *later* years (and thus in principle we do not need to match by office location for these years), to have a consistent matching methodology in all years, we match only by the zip code information we find using the contemporary district office locations. Using this methodology, we are unable to match to counties those district offices that closed before 2009, the earliest year to our knowledge that district office locations are available publicly.

match those zip codes to counties based on 1990 geography to create a balanced panel of 762 counties, home to around three-quarters of the U.S. population in 1990.[23]

The event-study analysis shows a clear response to NAFTA along both the applications (Figure 2.5 and Appendix Figures 2A.24 and 2A.25) and the awards (Appendix Figure 2A.26) margins. There are no pre-trends suggesting a pre-NAFTA increase in DI applications in vulnerable counties (if anything, some evidence to the contrary). Applications begin to tick upward in NAFTA-vulnerable counties in 1996 and remain elevated for the remainder of our sample period.

Again, however, the magnitude of applications and awards relative to estimated job losses is modest. Taking the largest coefficients among the various specifications for the year 2000, we estimate that applications (awards) increased by 0.169 (0.116) per 100 residents in the most affected counties relative to the least affected.[24] Assuming roughly three jobs lost per 100 by the year 2000 (on the modest side of the range of estimates in the previous section), we estimate a 5.6 (3.8) increase in DI applications (awards) for every 100 jobs lost.

D Other outcomes

A natural implication of job loss is greater unemployment insurance (UI) payments. As UI eligibility typically lasts only 26 weeks, we would expect that UI payments would rise in the period of active job loss but not remain elevated in the longer run. Appendix Figures 2A.27 and 2A.28 show exactly this pattern.

[23] Appendix Figure 2A.23 shows that our log-employment effects look similar, though somewhat smaller in magnitude, when restricted to these counties as they do in Figure 2.2 for all counties. So any effects on DI applications in this subsample might serve as a lower bound for that on the full sample.

[24] The point estimate for applications (awards) is 0.127 (0.223) in 2000. As we saw earlier that population does not change as a result of NAFTA and taking 1.33 DI applications per 100 in the top quartile of counties (see Table 2.1), we estimate an $0.127 * 1.33 = 0.169$ increase in applications per 100 residents (note that using the top quartile slightly overstates this effect). The analogous calculation for awards is: $0.223 * 0.524 = 0.116$.

The decline in employment also suggests more families may gain SNAP eligibility. We find suggestive evidence of an increase in SNAP receipt in Appendix Figures 2A.29 and 2A.30. We do not emphasize these results more because NAFTA occurs just as states and eventually the federal government embark on major welfare reform, which among other changes ended automatic SNAP enrollment upon enrollment in AFDC.[25]

The evidence in this and the previous section suggest deterioration with respect to a number of important socio-economic indicators in NAFTA-vulnerable counties after 1994. The next section examines whether these developments had any political impact.

2.6 The political response in areas vulnerable to NAFTA

The major U.S. political parties have switched positions on trade policy several times since the nineteenth century. From the post-Civil War Gilded Age through the 1920s, the Democratic Party (concentrated in the South and West) favored lowering tariffs and a championed a progressive federal income tax to make up for lost revenue. Republicans (concentrated in the richer, industrial Northeast) preferred tariffs to protect domestic manufacturing and opposed progressive income taxation. This debate lost its salience during World War II and the early decades of the Cold War, when a bi-partisan consensus maintained that offering favorable trade terms to other countries would reduce the appeal of communism. The US had no industrial rivals at this time, which further diminished the importance of the issue. This consensus broke down in the 1970s as U.S. economic growth slowed and import competition grew, at which point the Democrats emerged as the more

[25]This report articulates how many families eligible for food stamps did not receive them in the years following welfare reform: https://www.brookings.edu/research/welfare-reform-reauthorization-an-overview-of-problems-and-issues/.

protectionist of the two major parties (not a surprise, given their base of union members and blue-collar workers threatened by liberalized trade regimes).[26]

The debate over NAFTA was a major topic in the 1992 and 1994 national elections. While Bill Clinton avoided taking a clear stand on NAFTA during the 1992 presidential campaign, he made passing NAFTA in Congress a major goal of the first year of his administration, against the wishes of his party's base.

By no means are we the first to argue (as we do in this and the next section) that NAFTA led to lasting, negative effects on Democratic identification among regions and demographic groups once loyal to the party. Many historians and political scientists have made this argument, though more in narrative than quantitative terms. In general, a theme of *betrayal* emerges. The traditional base of the party—e.g., union members and other working-class voters—bitterly opposed NAFTA and the Democratic president pushing for it, in what became a highly emotional fight (e.g., anti-NAFTA groups organized candle-light vigils on the White House lawn as the vote in Congress approached). Klinkner (2019) writes: "In a hotly contested and emotional vote, the critics of globalization, led by organized labor and environmental groups, were overcome by NAFTA's supporters, principally corporate lobbyists *and the Clinton administration* [emph. added]." Similarly, Stein (2010) writes: "When it came to measures that the base of his party wanted, Clinton faltered... Clinton had made the NAFTA a priority....and this allowed the Republican opposition to mushroom."

The debate over NAFTA was not limited to academics and policy wonks. As we show in Appendix Figure 2A.31, coverage of stories including the words "trade" and "imports" and "jobs" peaks for all three nightly network news shows in 1992-1993.[27] Vice President

[26]The brief historical summary in this section is drawn mostly from Weisman (2004) and Stein (2010).

[27]Interestingly, even though the easing of trade relations with China had a greater impact in terms of total import value, the topic did not garner much coverage on network news. We speculate that the events of September 11th, 2001 and the resulting U.S. military campaigns in Afghanistan and Iraq crowded out media coverage of trade with China.

Al Gore participated in a nationally televised November 1993 debate against NAFTA-foe Ross Perot, who argued throughout the debate that NAFTA would push blue-collar U.S. jobs to Mexico. The debate set a viewership record for CNN that would stand for two decades (Kornacki, 2018).[28] NAFTA was also the subject of at least two *Saturday Night Live* sketches in 1993.[29]

A final point we emphasize before moving to the empirical work is that NAFTA did not turn out to be a one-time deviation from Democrats' previous protectionist platform. It was instead the start of a lasting pro-globalization shift. At the time of its 1993 passage, "Clinton told the national press that NAFTA ....is 'the symbol of where we want to go in the world' (Minchin, 2017, p. 202)." In addition to NAFTA, Clinton signed other important free-trade measures (most notably the Uruguay Round Agreements Act in 1994 and Permanent Normalized Trade Relations with China in 2000). He made his pro-free-trade position a major talking point of his 1996 re-election campaign.[30]

A Opinions of respondents in NAFTA-vulnerable states

We begin by showing the enduring unpopularity of NAFTA in the places most vulnerable to it. We gather, to the best of our knowledge, all surveys that (a) ask a generic sentiment question regarding NAFTA and (b) include state identifiers. Very few surveys include county identifiers and none that we know claim to be *representative* at the county level, so

[28] Over 38 percent of registered voters reported having watched all or part of the debate, with an additional thirty percent saying they watched at least a "little" or had since heard or read about it. These numbers are from the authors' calculations using November 1993 WSJ/NBC survey data.

[29] See https://www.nbc.com/saturday-night-live/video/mexican-stereotype/n10486 and https://www.nbc.com/saturday-night-live/video/united-we-stand-america/n10497 for the videos.

[30] In the first 1996 presidential debate, he noted: "I've done a lot of things that were controversial. My economic plan, my trade position...Sometimes you just have to do that because you know it's right for the country over the long run." In the second debate, he emphasized that he had opened up the country to trade more than any of his predecessors. "[W]e've had over 200 separate trade agreements in the last four years. By far, the largest number in American history—not just the big ones you read about, but a lot of smaller ones." See https://www.debates.org/voter-education/debate-transcripts/october-6-1996-debate-transcript/ and https://www.debates.org/voter-education/debate-transcripts/october-16-1996-debate-transcript/ for transcripts from these debates.

in this subsection we examine how *state*-level vulnerability to NAFTA predicts residents' views toward the trade agreement. Appendix B provides details on the surveys we include.

Table 2.2 documents a robust, negative relationship between the NAFTA vulnerability of the respondent's state of residence and her approval of NAFTA. Col. (1) shows the coefficient on state-level vulnerability is negative and highly significant when only survey (which subsume time) fixed effects are included as controls. Relative to states in the least vulnerable quartile, support for NAFTA is 3.9 percentage points (roughly ten percent) lower in the most vulnerable states. Columns (2) and (3) adds controls for race, sex, education, income, age and union status. These controls themselves have highly significant effects but in fact the coefficient on *Vulnerability* only becomes more negative and precise. In col. (4) we add nine Census-division fixed effects, which again increases the magnitude of our coefficient of interest. The final column shows robustness to dropping those without a stated opinion on NAFTA.

B County-level event-study results

While the previous subsection showed the enduring unpopularity of NAFTA in vulnerable states, we now turn to our standard event-study analysis (obviously, we could not perform an event-study using NAFTA popularity as the outcome, as no questions on NAFTA *per se* existed in the pre-NAFTA period).

Using our usual county-year level event-study specification, Figure 2.6 shows that NAFTA *Vulnerability* is associated with an increase (decrease) in Republican (Democratic) county-level presidential-election vote share in the mid 1990s, after a generally flat pre-NAFTA trend. We show results for both parties because in the key NAFTA years of 1992 and 1996, Perot's third-party anti-NAFTA campaign means that the Democratic and Republican shares do not sum to close to one, as they typically do.[31] Relative to the 1992

[31] As usual, we show the raw relationships as well as robustness to varying controls in the Appendix (see Appendix Figures to 2A.32 to 2A.34). Given the importance of third-party voting in Presidential elections,

omitted baseline election and using our preferred state-year-fixed-effects specification, by the early 2000s counties in the top quartile of vulnerability have moved in the Republican direction by roughly five percentage points compared to the least exposed quartile. This result is robust to our usual manufacturing controls or using the Democratic share as the outcome.[32] Given recent work by Gethin et al. (2022) showing the movement of less-educated voters away from the Democrats and other center-left parties in rich democracies, we show additional robustness of this result to flexible controls for county-level college share and its interaction with white population share (see Appendix 2A.35).

How large are these effects? To benchmark the political importance of the results in Figure 2.6, in Appendix Table 2A.3 we simulate the 2000-2012 presidential elections after "removing" our estimated NAFTA effects (using the $State \times Year$ specification). We calculate that, without NAFTA, Democrats would have won Florida and New Hampshire in 2000 (29 electoral votes, or a *net shift* of 58 electoral votes from Republicans to Democrats). In 2004, Iowa, New Mexico and Ohio would have switched to the Democratic column (a net shift of 64 electoral votes). Both the 2000 and 2004 elections were very close, so these switches would have changed the national outcome. Given that the Democrats handily won the 2008 and 2012 presidential elections, the NAFTA effect in those elections would have only made the landslides larger, adding Georgia in 2008 and North Carolina in 2012 to the Democratic column. Moreover, we calculate that especially in 2000 and 2004, NAFTA would have put more "states in play," which we define as moving a state the Democrats lost by more than 2.5 percentage points to an under-2.5-percentage-point loss.

especially in 1980 and in 1992-2000, we show in the Appendix results separately for both major parties when presidential vote share is the outcome.

[32] In the Appendix we show that House elections yield similar patterns. Third party voting is less common in Congressional elections, so in the Appendix we only show results for the Republican vote share in House elections. One complication with House elections is that uncontested elections are not uncommon (whereas no presidential election since the early 1800s has gone uncontested). We drop such elections in our House analysis, meaning we have fewer counties in our strictly balanced sample and thus prefer the presidential-election results.

In 2004, a net 72 shift of electoral votes went from a Democratic loss of more than 2.5 percentage points to a loss below that threshold.

2.7 The political response among individuals against free trade

Much of the existing literature on the political reaction to trade policy estimates, as we do in Section B, ecological regressions, with the implicit assumption that the *areas* most affected by trade deals would exhibit the greatest shifts in political outcomes. In this section, we instead focus on *individuals*' self-reported views toward trade and test whether those with protectionist views shift away from the Democrats around the time of NAFTA's passage.

A Evidence from repeated cross-sectional data

In this section, we make heavy use of the ANES. Since 1986 it has asked in most of its surveys a question capturing general protectionist sentiment. In almost all years, the question reads as follows: "Some people have suggested placing new limits on foreign imports in order to protect American jobs. Others say that such limits would raise consumer prices and hurt American exports. Do you favor or oppose placing new limits on imports, or haven't you thought much about this?" We create a *Favor import limits* dummy variable, coded as one if you agree with placing new limits on imports and zero for all others. We will sometimes describe individuals coded as "one" for this dummy variable as having "protectionist views" or being "protectionist."[33] In all years, the ANES asks

[33]Note that the ANES cumulative file codes as missing anyone who says they do not know enough about NAFTA or otherwise do not have an opinion. We thus use the individual survey files, which preserve this detail.

partisan ID, a scale variable from 1-7, increasing in support for the GOP, which we use to measure partisan identity.

In the analysis that follows, we document the reaction of protectionist voters to a Democratic president supporting and signing NAFTA.

Main results

We begin by estimating the following equation, separately by each year t in our sample:

$$Partisan\ scale_i = \beta_t Favor\ import\ limits_i + \gamma \mathbb{X}_i + e_i. \tag{2.5}$$

Note that estimating this equation separately by year t allows the coefficients on the control variables in $\mathbb{X}_i$ to be unrestricted across years.

We plot the resulting β_t coefficients over time in Figure 2.7. The first series shows the coefficient estimates with no controls (an empty vector $\mathbb{X}_i$), so the plotted points are just raw differences between protectionist versus other voters. We see that pre-NAFTA, those with protectionist views were *less* supportive of the GOP (and the difference could be distinguished from zero in each year of the pre-period). Sometime between 1992 and 1996 (the favor-import-limits question is not asked in 1994) a significant number of protectionist voters moved toward the GOP, so that the raw difference goes to zero.

In the second series, we add standard demographic and socioeconomic controls: race, sex, age, education, and income. The same 1992-1996 shift toward the GOP among protectionist voters remains, even after allowing these characteristics to have their own effect each year. As noted in the previous section, the robustness to flexible education controls is especially important given recent work by Gethin et al. (2022), showing a longer-run trend of less-educated voters away from the Democrats and other center-left parties in Western democracies. The pattern in the second series shows that the

NAFTA-era shift among protectionist voters (who are themselves less educated) can be distinguished from the more general trends of educational sorting across parties.

In the third series, we add controls for trust in the federal government, views (a "thermometer" going from cold to warm) toward African Americans (given the importance of race in U.S. politics), views on abortion, and weekly religious attendance. We choose these controls in particular because they are asked in all or most years of our sample period. The pattern of protectionist voters shifting in the GOP direction after 1992 holds. In the final series, we add a control for views toward immigration levels, which is not asked in the ANES in 1986 or 1988. For the years it is available, controlling for this variable separately by survey year yields coefficients almost identical to those in the third series.

While Figure 2.7 is important in showing that the shift happens between 1992 and 1994, we now turn to differences-in-difference estimation to facilitate more succinct robustness and subsample analysis. That is, we pool all years and model an individual's partisan identity as a function of *Favor limits*, its interaction with *After 1992* and year fixed effects (again varying controls to test robustness), where the interaction term is the variable of interest.

Table 2.3 shows the results from estimating variants of this equation. Col. (1) has no additional controls beyond year fixed effects. Consistent with Figure 2.7, the coefficient on the main effect of *Favors import limits* suggests that from 1986 to 1992, protectionist views pushed against identifying as a Republican. The coefficient on the interaction term is positive and statistically significant, essentially erasing the pre-period effect. To give a sense of its practical significance, the shift is over one-half the size of the partisan gender gap (as estimated in our sample), a key divide in U.S. politics.

Col. (2) adds state fixed effects, which we add with the caveat that the ANES warns users it is not representative at the state level. Col. (3) drops state fixed effects and adds instead the same demographic controls in the second series of Figure 2.7. We do not report

the coefficients to save space, but they are of the expected sign. Note that they indeed add significant explanatory power to the estimation (the R-squared values jumps up by ten percentage points), but if anything they only increase the magnitude on the coefficient of interest. A similar dynamic occurs in col. (4), when we add to the col. (3) specification the controls for other political and social issues included in the third series of Figure 2.7.[34] Col. (5) replicates col. (4) after dropping all observations that respond "don't know" to the *Favor import limits* question.

In column (6), we add *After 1992* interactions with all the controls in cols. (3) and (4), so that these variables, like our protectionist dummy, can have different effects before and after NAFTA. Adding these controls in cols. (3) and (4) significantly increased the R-squared value, consistent with their having large explanatory power *in the cross-section*, but they add only minimal explanatory power in col. (5), suggesting they have limited explanatory power *over changes*. Put differently, those who, say, oppose abortion or distrust government are for the most part already Republican by 1992.[35]

Heterogeneity

Which protectionist voters might have the largest political response to NAFTA?Minchin (2012b) and others have argued that for many Democrats (especially white Democrats) in the 1980s, economic issues such as trade policy were key to their party loyalty because on social issues such as guns, affirmative action and abortion they sided with the GOP. We thus hypothesize that for these voters, the response to NAFTA will be stronger.[36] By contrast, for, say, a Black voter opposed to NAFTA but who is also pro-choice and wary of

[34]Note that we do not add views toward immigration, the extra control in the fourth series of the figure, as it is missing for most of the pre-NAFTA period.

[35]Readers may wonder why we have not yet examined support for Ross Perot, given his importance in the anti-NAFTA movement. As he only ran in the 1992 and 1996 presidential races, focusing on him would limit the sample period relative to using party identification. But Appendix Table 2A.2 indeed shows that, in both years, approval of Perot is significantly higher among protectionist voters.

[36]Frank (2004) makes a similar point about NAFTA, writing: "With Democrats and Republicans having merged on free trade, the issues that remained were abortion and guns" (p. 177).

Republicans on Civil Rights, the Democrats' position on free trade would be just one of many issues that binds them to the party.

To test this idea, we examine our results in a series of sample splits that create mutually exclusive and exhaustive groups. For each subsample, we estimate the specification in col. (4) of Table 2.3. While we cannot examine each of the issues highlighted by past work, we try to proxy many of them with questions in the ANES.

First, we examine our results by race, estimating the col. (4) specification separately for whites versus all others and plotting the resulting coefficients and 95% confidence intervals in Figure 2.8. The white coefficient is larger in magnitude, though given the small size of the non-white group, there is large overlap in the confidence intervals. The large gender gap in modern U.S. politics suggests that white *men* might feel especially at home with the GOP on cultural issues. Indeed, when we split the sample into white men versus all others, the former group exhibits a substantially larger shift toward the GOP among protectionist voters. A similar result holds among whites without a college degree versus all other respondents.

We next show that the protectionist response is somewhat larger in the South than elsewhere, perhaps because the South has more conservative social views or because the South was more vulnerable to NAFTA. The final cuts we examine are along two key cultural markers: opposition to abortion and weekly church attendance. Both of these splits of the data reveal large differences in the responsiveness of protectionist voters after 1992. Among respondents who do not oppose abortion or do not attend church weekly, the "protectionist response" that we propose still exists and pushes in the hypothesized direction, but is much smaller and not always distinguishable from zero.

We conclude from the analysis of repeated cross-sectional data that between 1992 and 1996, voters with protectionist views exhibited a significant and enduring shift rightward.

As hypothesized by historians, this shift was especially pronounced among individuals who already shared cultural positions with the GOP.

B Corroborating evidence from panel data

A possible concern about the repeated cross-sectional analysis is that views on trade could be endogenous to party identification, whereas our analysis in Figure 2.7 and Table 2.3 implicitly assumes that views on trade cause changes in party identification. NAFTA signaled that key Democratic leaders were taking a new position on trade, and thus some loyal Democratic voters may *change their views on trade* to limit cognitive dissonance. For this reason, we turn to a 1992-1994 panel dataset that follow the *same voters* across time (so we can model any partisan shift as a function of pre-NAFTA trade views).

In 1992, the ANES fielded a panel survey as part of their usual election-year surveys. That year, they designate roughly 1,000 respondents for a follow-up survey two years later; about 750 in fact take the survey 1994. As our key explanatory variable, we use the same "do you favor imports question" in 1992 that we use in the repeated-cross-section analysis earlier in this section.[37] However, in this analysis, we ask whether protectionist views in 1992 predict *changes* in party identification away from the Democrats over the next two years as NAFTA is debated, passed, and implemented.

Table 2.4 shows results from estimating variants of the following equation:

$$Moved\,Right_{i,\text{94-92}} = \beta Favor\,Import\,Limits_{i,92} + \gamma \mathbb{X}_{i,92} + e_i, \tag{2.6}$$

where $Moved\,Right_{i,\text{94-92}}$ is a dummy for having moved toward the GOP on the seven-point scale and all other variables are defined as before.

[37] The question is not asked in 1994 (neither in the 1994 follow-up survey nor in the standard ANES 1994 cross-section survey).

Col. (1) shows the results with no additional controls and suggests that those with protectionist views had an eight percentage-point higher likelihood of shifting rightward. This effect increases in magnitude when we drop those without an opinion (col. 2). Col. (3) adds to the col. (1) specification standard demographic controls, which have a negligible effect on the coefficient of interest. Recall from the repeated-cross-section analysis in Table 2.3 that these variables had important predictive power *in the cross-section* but not in explaining *changes* between 1992 and 1996. Consistent with that earlier result, adding these variables in an effort to explain *changes* between 1992 and 1994 increases the R-squared only modestly. In col. (4) we add controls for political views. The coefficients on views toward the government helping Blacks, demand for a generally active government, abortion rights and immigration are all close to zero. Again, these views have strong predictive power in any given year, but by 1992, most people who are, say, opposed to affirmative action are already Republican, so these controls have little ability to explain *changes* from 1992 to 1994.[38]

A nice feature of the panel analysis is that we can control for key issues of the day in 1992, which may not have stood the test of time in order to be asked repeatedly in the ANES but which could correlate with views on trade. In the final column, we control for views about gays in the military and health reform (two controversial issues during Clinton's first term) and Congressional term limits (a key item on the Contract with America developed in 1993 by Newt Gingrich, the soon-to-be Speaker of the House). Interestingly, none of these issues have a significant effect, despite their attention in the media. In the final column, we add state fixed effects (although, again, the ANES warns that its samples are not representative at the state level). Results are unchanged.

[38] An important and interesting exception is weekly religious attendance, which has a coefficient nearly equal to that on protectionist views. The inclusion of this variable has little effect on our coefficient of interest because the two are nearly perfectly uncorrelated ($\rho = -0.0099$). Thus, the religious represent a distinct group moving toward the GOP around the same time, an important reminder that NAFTA is not the only issue triggering potential political realignment during this moment and an interesting topic for future work.

We can replicate this analysis using a 1993 question on support for NAFTA *per se*, instead of our 1992 question on protectionist sentiment. Two issues arise. First, the sample becomes smaller. Second, as the question is asked in the fall of 1993 (the peak of the debate over NAFTA as the vote in Congress neared) it is also much more likely to be endogenous to party identity than our 1992 measure. Nonetheless, in Appendix Table 2A.4 we find very similar results in terms of magnitude, though less precisely estimated.

As we did with the repeated cross-sectional data, Appendix Figure 2A.37 examines heterogeneity in the 1992-1994 response of protectionist respondents. Given the much smaller sample size, we cannot always reject equality of subsample coefficients. But in almost all cases, the coefficient patterns mirror those Figure 2.8 whereby more socially conservative groups also have the largest protectionist response to NAFTA.

C Connecting the economic and political impacts of NAFTA

A natural question is whether the economic impacts documented in Sections 2.4 and 2.5 *caused* the political effects we document in this and the previous section. We believe there are (at least) two channels by which NAFTA affected political outcomes, which we call an "economic-downturn" pathway and an "issue-based" pathway.

NAFTA was signed by a Democratic president and ever since has been unpopular in the states most vulnerable to its provisions (see Table 2.2), so it is not surprising that NAFTA-related local job loss would shift voters away from the Democrats in vulnerable areas. The Presidential (Figure 2.6) and House (Appendix Figure 2A.36) county-level results generally mirror the employment effects (Figure 2.2) in that the shifts are concentrated in NAFTA-vulnerable areas and occur *gradually* after NAFTA's passage. As cited in the introduction, past work has shown that, at least in western democracies, economic downturns—regardless of their triggers—lead voters to support

anti-establishment and right-wing parties, so even voters *unaware* of NAFTA might move toward the GOP in these counties merely as a response to local economic decline.

But the results in this section show an *immediate* political shift among protectionist voters and thus cannot be explained by economic downturn *per se* but rather by what we call an "issue based" pathway. For example, in 1992 NAFTA had not yet been passed, but we see (Appendix Table 2A.2) highly significant support for anti-NAFTA Perot among protectionist voters that year. Similarly, we see an anti-Democrat shift in political identification between 1992 and 1994 among these voters (Table 2.4), even though NAFTA had barely taken effect and the medium-run losses in vulnerable areas documented in Section 2.4 had not yet materialized. For protectionist voters (who until NAFTA's passage were more likely to be Democrats, see Figure 2.7), punishing the president who signed NAFTA in 1993 was reason enough to move away from the Democrats, regardless of whether any economic consequences had yet materialized. Indeed, for protectionist voters motivated by the issue-based pathway, in a counterfactual world where Clinton pushed NAFTA but Congress voted it down, they might have still left the Democrats.

Of course, neatly breaking the mechanisms into these two categories is an over-simplification. In fact, voters reacting to NAFTA in 1992 may be *forecasting* the economic consequences, and in that sense are voting in line with our "economic downturn" pathway. Moreover, institutions can mediate voters' reactions, a topic we have not taken up in this paper. A limitation of the by-county analysis in this and other papers is that, by design, it will obscure any political shifts that have national spillovers. Given the importance of unions to Democratic get-out-the-vote efforts (Feigenbaum et al., 2018), scholars argue that the NAFTA fight may have caused lasting damage to the Democrats' ability to organize nation-wide.[39] At least in the 1980s and 1990s, unions over-represented

[39] "In aggressively pursuing passage of the agreement, the Clinton administration put itself in conflict with organized labor. By attacking one of the Democratic party's most important constituencies, the administration succeeded in further weakening the Democratic coalition and exacerbating the party's organizational decline" (Klinkner, 2019, p. 70). As Minchin (2017) writes, NAFTA left a huge and lasting rift between the

less-educated men (see Farber et al., 2021), so unions' sense of betrayal over NAFTA is potentially another indirect reason that the trade agreement led this demographic to leave the Democrats even in situations where the workers themselves did not lose their jobs. We flag the role of unions and other working-class institutions as potential mediators between policy and voting as a topic for future work.

2.8 Conclusion

In this paper, we provide evidence that NAFTA substantially reduced employment in counties most exposed to Mexican important competition and that these areas (as well as voters opposed to free trade, regardless of geographic residence) turned away from the Democratic party as a result.

The movement of working-class whites away from the Democratic party is one of the most debated topics in U.S. politics. As Piketty (2020) documents, it is part of a larger trend in the rich democracies, as less-educated voters have abandoned the traditional center-left parties in Europe. The 2016 election of Donald Trump in the US, the successful Brexit campaign in the UK and the rise of right-wing populist parties in Europe have prompted a recent debate about whether these events are best explained by "economic dislocation" or "ethnocentrism."

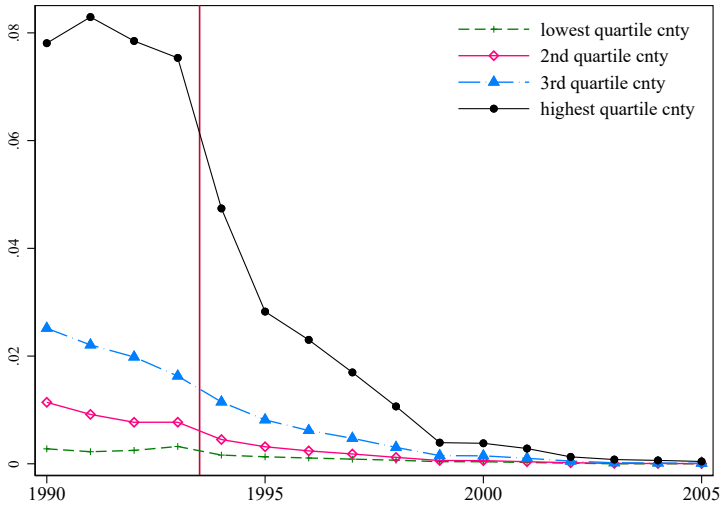
Our results point to an interactive effect between economic dislocation and ethnocentrism or other aspects of social conservatism, at least during the NAFTA era. The trigger for this movement toward the right was an economic one—a trade deal that increased import competition for many low-wage domestic industries, which had been

Democratic party and the labor movement: "For many in the labor movement, the resentment left by the administration's support of NAFTA was long-lasting. Even many years later, some staffers felt that Clinton had betrayed them...'We had just managed to elect a Democratic president, Bill Clinton,' recalled [AFL-CIO economist] Mark Anderson. 'The lousy son of a bitch.' Anderson felt 'terrible' after the NAFTA vote, which he viewed as 'hugely personal' (pp. 203-204)."

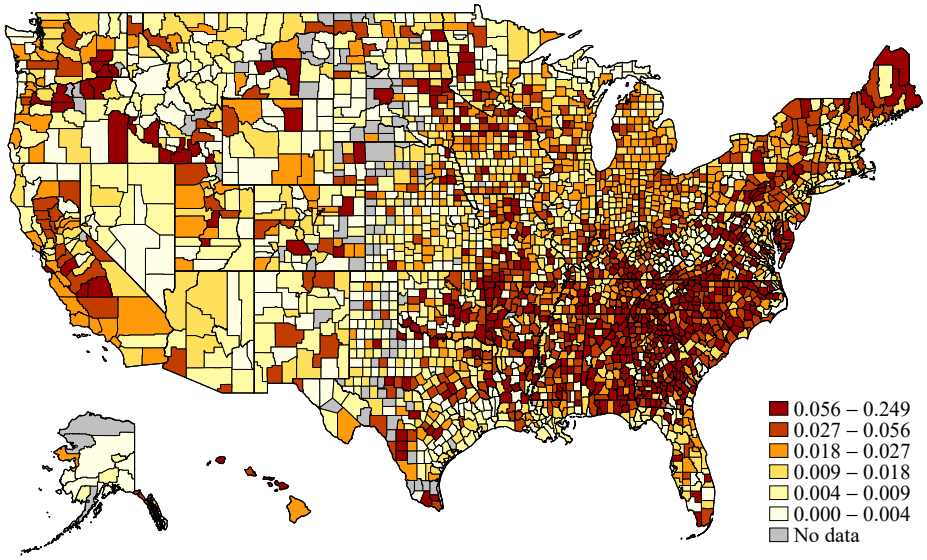
opposed by unions and less-educated voters. But the movement appeared to manifest most significantly among those voters who already had conservative views on many social issues.

Figure 2.1: Variation in protection and NAFTA exposure, across time and space

(a) Protection across time



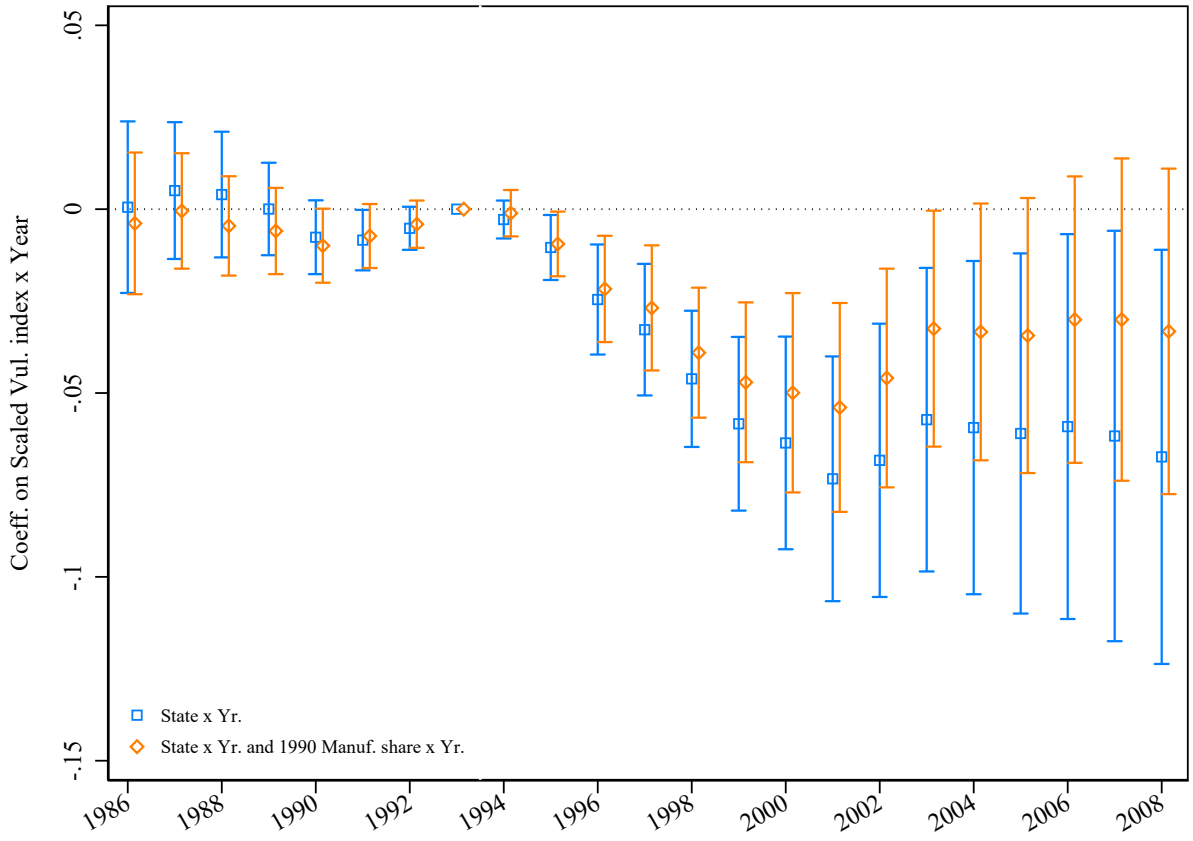
(b) Exposure across counties



Sources: The vulnerability measure is constructed using: ad-valorem equivalent tariffs from Feenstra et al. (2002) and USITC annual tariff data; export series from the UN Comtrade bilateral export series and Hakobyan and McLaren (2016b); and annual county employment by industry from County Business Patterns.

Notes: Panel (*a*) shows the weighted average tariff protection across time by each quartile of 1990 county-level vulnerability. That is, for each county-year, we take baseline (1990) county employment by industry and multiply by Mexico's revealed comparative advantage (RCA) for that industry (in 1990) scaled by τ^t, the U.S. tariff on Mexican goods in that industry in year t. Note that the values of the series in 1990 are in fact the 1990-based county-level *Vulnerability* variable we use in much of the paper, as they use 1990-level tariffs τ^{1990}. Panel (*b*) plots the *Vulnerability* measure across counties. See Section A for more detail on variable construction.

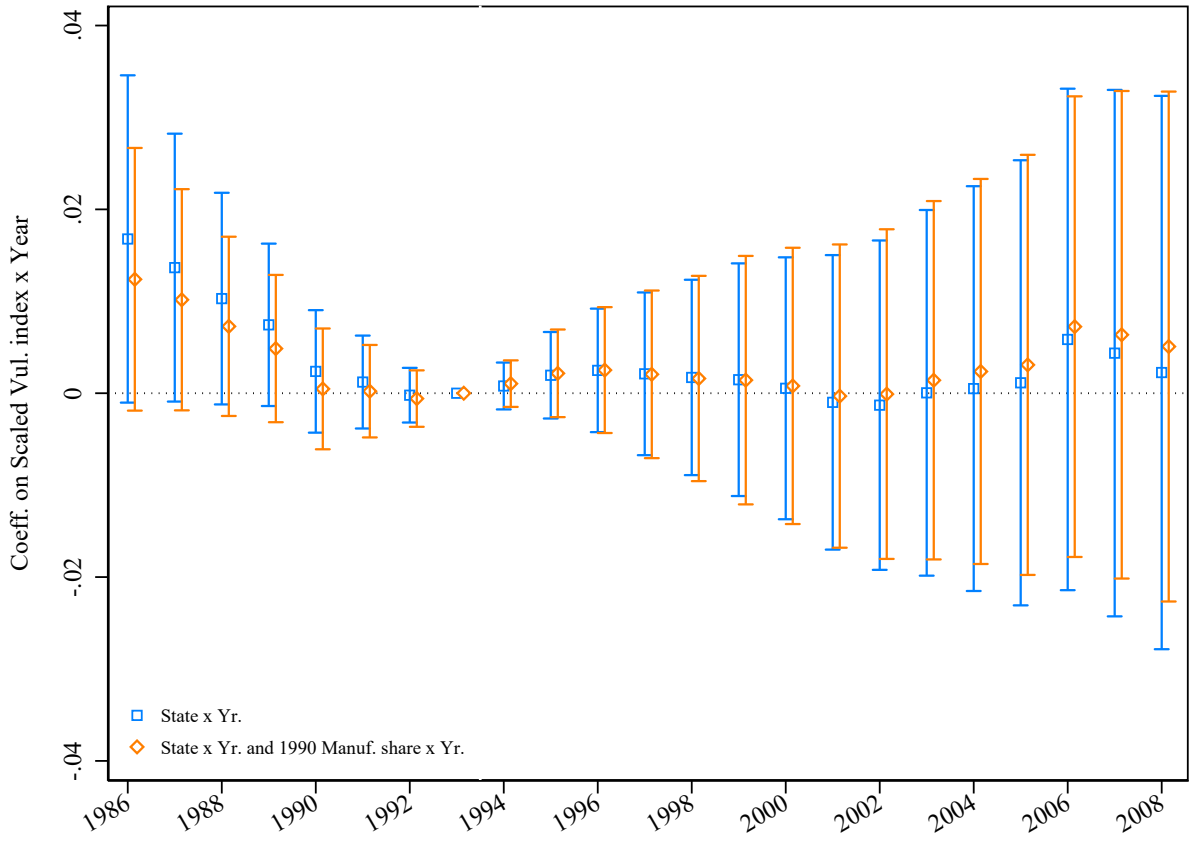
Figure 2.2: Log employment as a function of county NAFTA vulnerability



Sources: The dependent variable is computed from County Business Patterns. See Appendices B and B for more detail.

Notes: The analysis sample is fixed across specifications and strictly balanced, with 2924 counties in each year of the sample. The figure shows the event-study coefficient estimates (and 95% confidence intervals, based on standard errors clustered by state) from different specifications of equation (2.4). Observations are weighted by 1990 county population. We scale the *Vulnerability* variable so that a difference of one unit in scaled *Vulnerability* is equal to the difference between average unscaled *Vulnerability* for the top and bottom quartiles. The first series shows the coefficient estimates from a specification where we control for only county and year fixed effects. The second series adds to this specification *state*×*year* fixed effects. The third specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects. The fourth specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects. The final specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects.

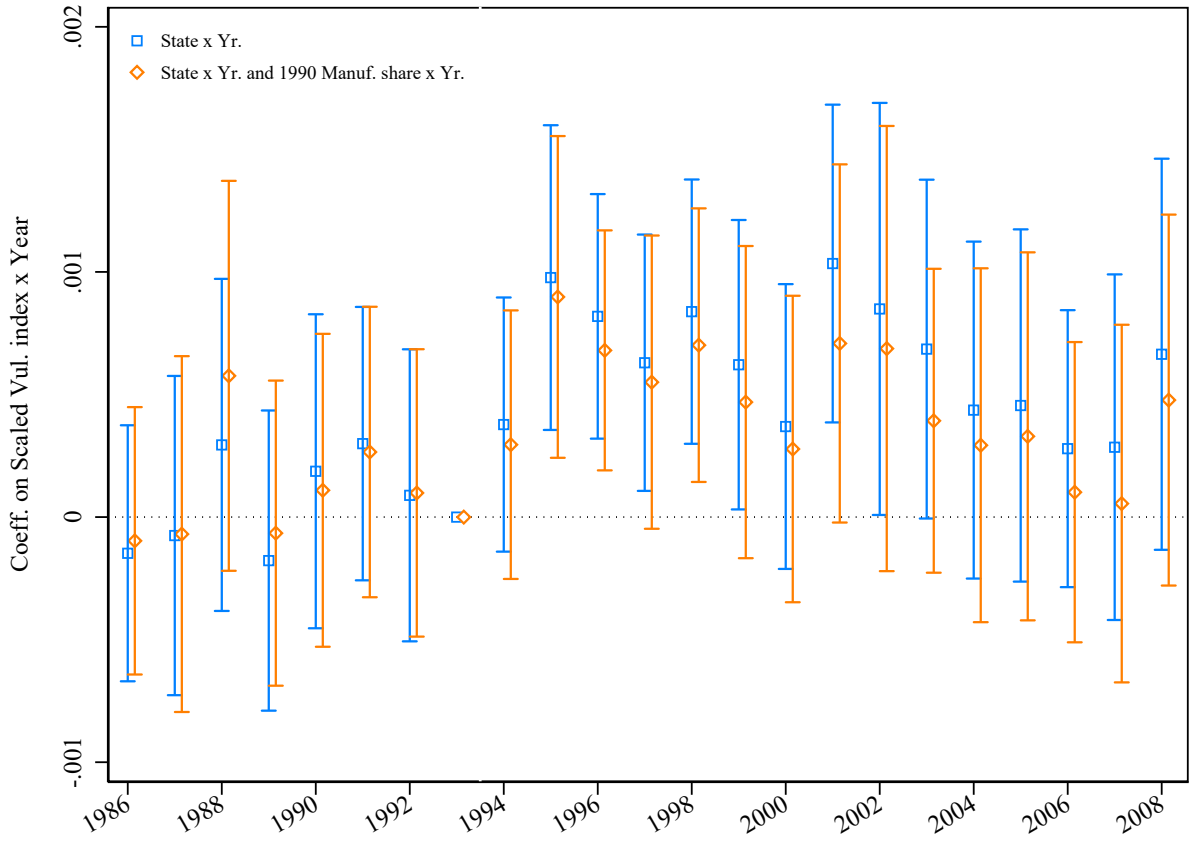
Figure 2.3: Log population as a function of county NAFTA vulnerability



Sources: The dependent variable is taken from the Census Bureau's Population Estimates Program. See Appendix B for more detail.

Notes: The analysis sample is fixed across specifications and strictly balanced, with 2991 counties in each year of the sample. The figure shows the point estimates (and 95% confidence intervals, based on standard errors clustered by state) for the coefficients on *Vulnerability* interacted with year (with 1993 the omitted year). Observations are weighted by 1990 county population. We scale the *Vulnerability* variable so that a difference of one unit in scaled *Vulnerability* is equal to the difference between average unscaled *Vulnerability* for the top and bottom quartiles. The first series controls for county and year fixed effects, as well as *state*×*year* fixed effects. The second series adds to the first controls for 1990 manufacturing share of county employment interacted with year fixed effects.

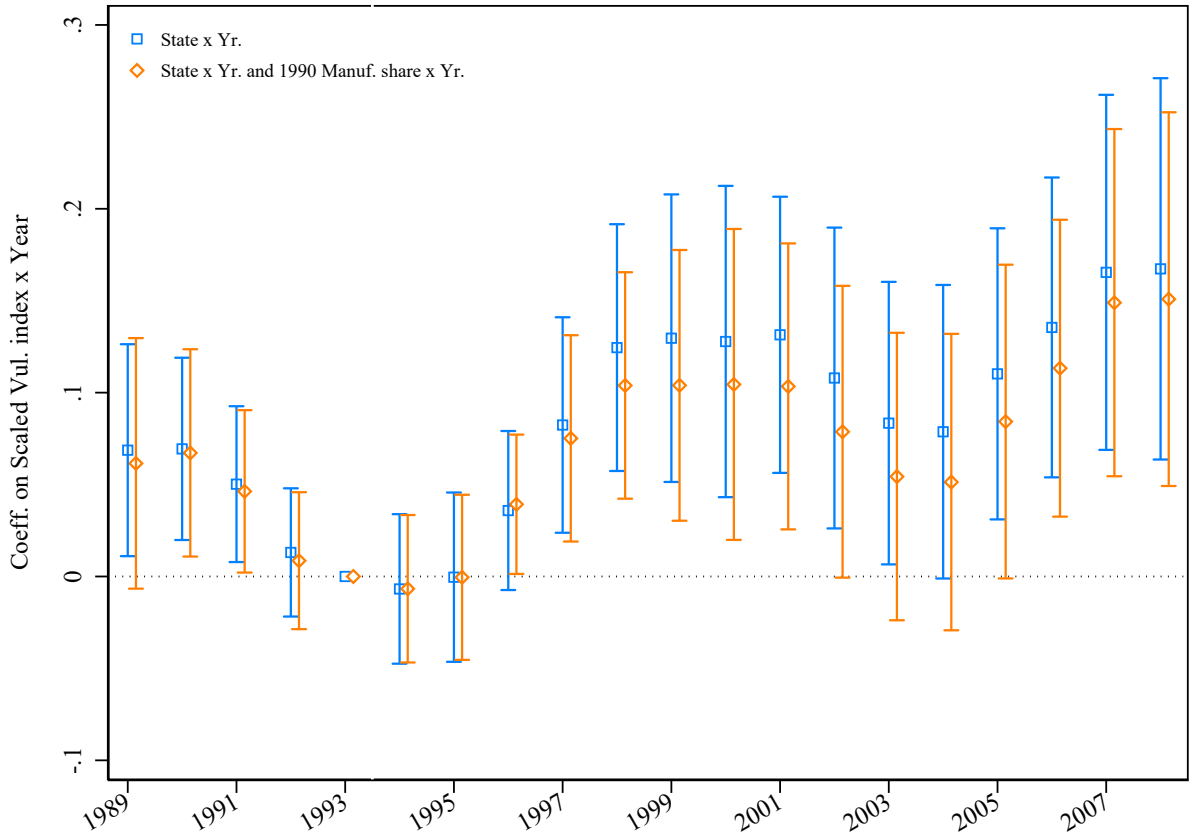
Figure 2.4: Trade Adjustment Assistance petitions per capita as a function of county NAFTA vulnerability



Sources: The dependent variable is taken from the U.S. Department of Labor TAA petition data, divided by 1990 working-age county population. See Appendix B for more detail.

Notes: The analysis sample is fixed across specifications and strictly balanced, with 2914 counties in each year of the sample. The figure shows the point estimates (and 95% confidence intervals, based on standard errors clustered by state) for the coefficients on *Vulnerability* interacted with year (with 1993 the omitted year). Observations are weighted by 1990 county population. We scale the *Vulnerability* variable so that a difference of one unit in scaled *Vulnerability* is equal to the difference between average unscaled *Vulnerability* for the top and bottom quartiles. The first series controls for county and year fixed effects, as well as *state*×*year* fixed effects. The second series adds to the first controls for 1990 manufacturing share of county employment interacted with year fixed effects.

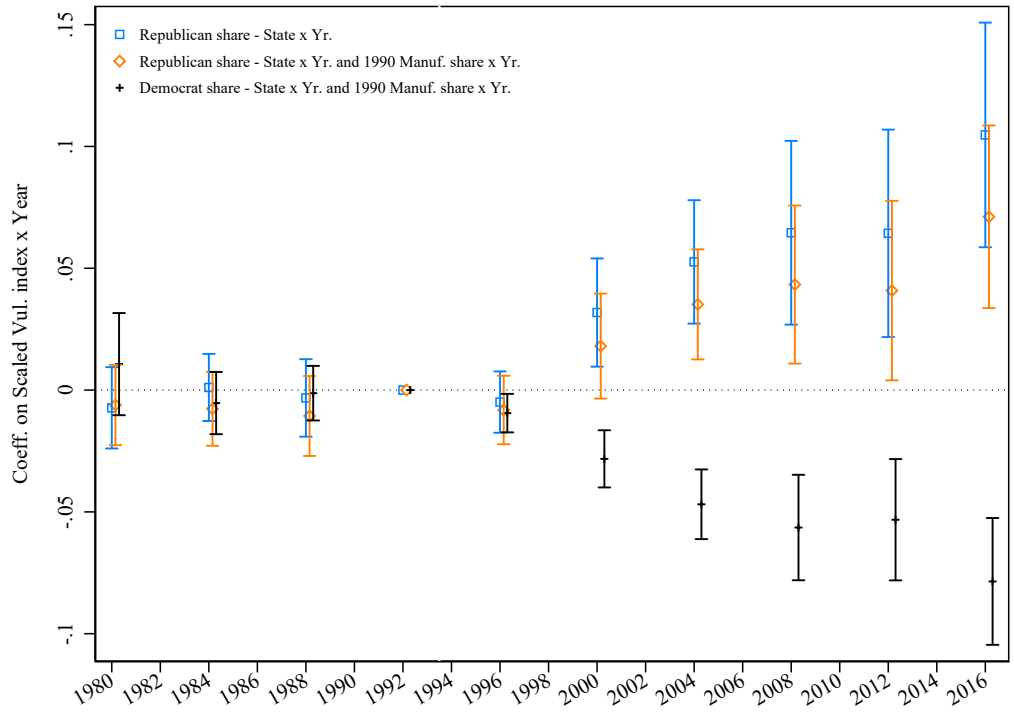
Figure 2.5: Log of DI applications as a function of county NAFTA vulnerability



Sources: The dependent variable is taken from the Social Security Administration (SSA). See Appendix B for more detail.

Notes: The analysis sample is fixed across specifications and strictly balanced, with 762 counties in each year of the sample. As discussed in Section C, we do not have all counties in this analysis, but the 762 counties we have in this balanced-panel analysis account for around three-fourths of the U.S. population. The figure shows the point estimates (and 95% confidence intervals, based on standard errors clustered by state) for the coefficients on *Vulnerability* interacted with year (with 1993 the omitted year). Observations are weighted by 1990 county population. We scale the *Vulnerability* variable so that a difference of one unit in scaled *Vulnerability* is equal to the difference between average unscaled *Vulnerability* for the top and bottom quartiles. The first series controls for county and year fixed effects, as well as *state*×*year* fixed effects. The second series adds to the first controls for 1990 manufacturing share of county employment interacted with year fixed effects. Appendix Figure 2A.26 shows the same analysis for *final awards* of DI.

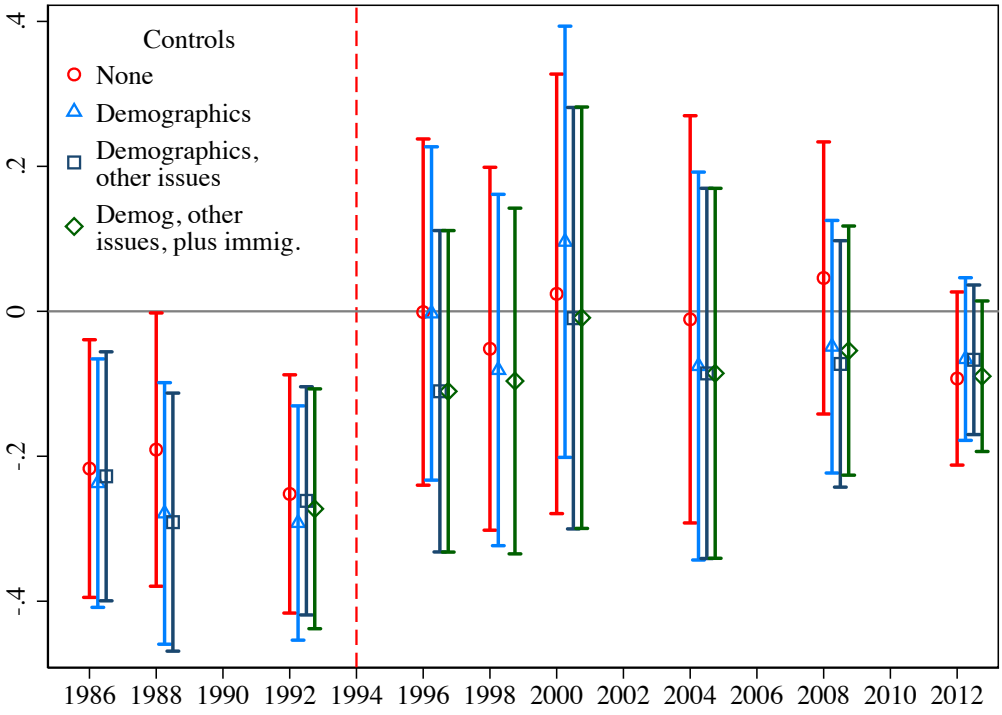
Figure 2.6: Presidential election vote shares as a function of county vulnerability



Sources: The dependent variable is taken from ICPSR general voting data and Dave Leip's Atlas of U.S. Election data.

Notes: Each series includes $N = 2880$ counties in a strictly balanced panel. The figure shows the event-study coefficient estimates (and 95% confidence intervals, based on standard errors clustered by state) from different specifications of equation (2.4), where the Republican (first two series) or Democratic (last series) vote share in presidential elections is the dependent variable. All specifications are weighted by 1990 county population. The first series shows the coefficient estimates from a specification where we control for only county and year fixed effects and $State \times year$ fixed effects. The second series adds 1990 county-level manufacturing share of employment interacted with year fixed effects. The third series is identical to the second except that Democratic presidential share is the outcome.

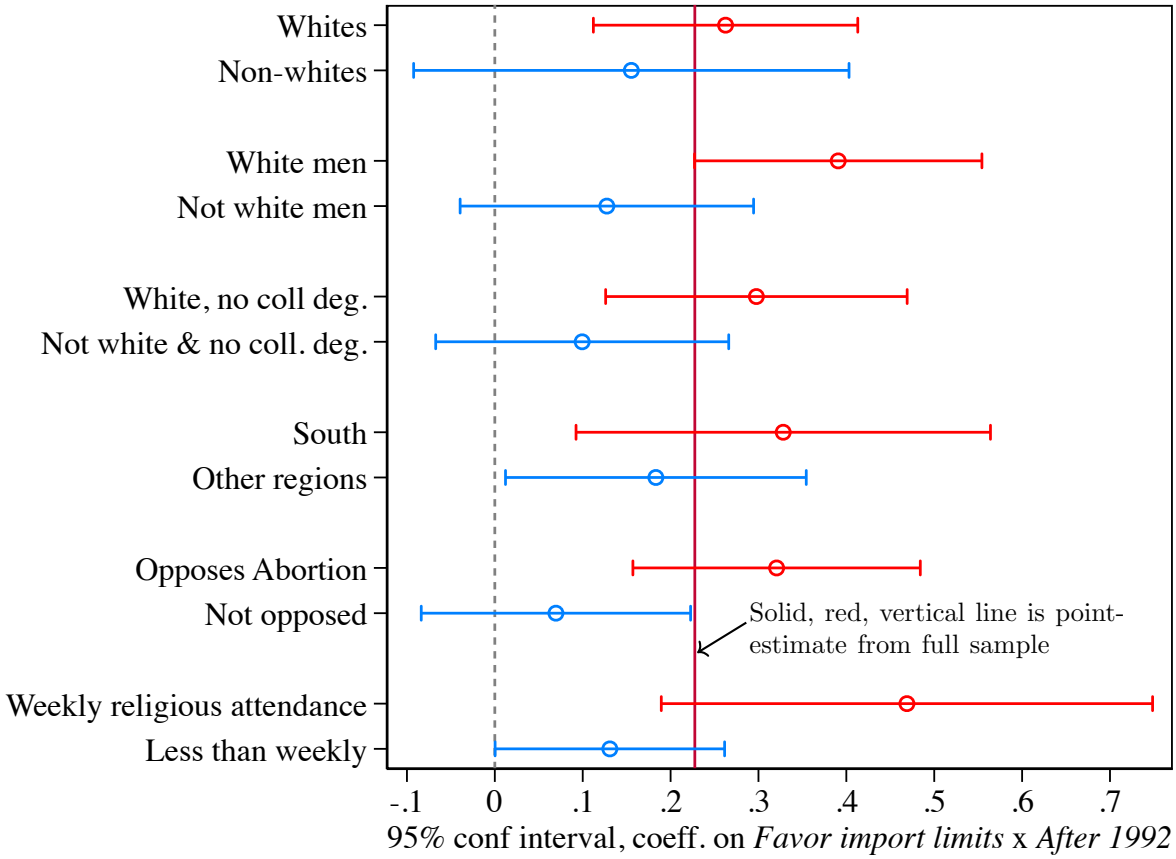
Figure 2.7: Party identification (increasing in Republican direction) as a function of views toward trade



Sources: ANES repeated cross-sectional data. See Appendix B for more detail.

Notes: Separately for each year in our sample period, we regress the party-ID scale (a 1-7 categorical variable, increasing in allegiance to the GOP) on the *Favor Import Limits*$_i$ dummy variable (coded as one if the respondent says that they support additional limits on imports and zero otherwise). The first series includes no other controls (equivalent to raw differences). The second series controls for gender, age, race, education, and family income. The third series adds controls for abortion, trust in government, views toward Blacks and views toward welfare recipients (note that not all of these variables are available in 1998, so the third series is missing that year). The fourth series adds to the third series a control for wanting immigration to increase (this question is not asked in 1986 and 1988). Note that we estimate these regressions separately by year, so the coefficients on the controls are unrestricted across years. We plot 95-percent confidence intervals based on standard errors adjusted for clustering by state of residence.

Figure 2.8: Heterogeneity in the shift toward GOP after 1992 among protectionist respondents



Sources: ANES repeated cross-sectional data, 1986-2012.

Notes: We estimate, for mutually exhaustive and distinct subgroups of the sample, equation (2.5) from the text: $\textit{Party ID}_{it} = \beta^{DD}\textit{Favor limits}_i \times \textit{After 1992}_t + \beta^{main}\textit{Favor limits}_i + \gamma\mathbb{X}_{it} + \mu_t + e_{it}$. *Party ID* is a 1–7 scale variable increasing in Republican-party identification. We use the same control vector $\mathbf{X}_{it}$ as in col. (4) of Table 2.3, namely demographic and political-issue controls. We report the coefficient and 95% confidence interval from the estimate of β^{DD}.

Table 2.1: Pre-NAFTA characteristics of counties, by vulnerability quartile

	Quartile (lower quartile : less vulnerable)			
	1	2	3	4
Demographics (1990)				
Population (in thousands)	32.795	131.218	120.474	42.210
Working-age population (in thousands)	21.093	86.790	79.832	27.217
Household income (in thousands)	23.070	26.145	25.233	21.967
Number of jobs (in thousands)	10.006	53.190	45.879	12.774
Emp-to-Pop ratio	0.349	0.430	0.442	0.396
White share of population	0.905	0.911	0.901	0.844
Manufacturing share of employment	0.080	0.121	0.158	0.168
College-grad share of population	0.131	0.156	0.143	0.111
Political preference (1980-1990)				
Republican House two-party vote share	0.461	0.480	0.472	0.389
Republican Presidential vote share	0.567	0.575	0.567	0.556
Democratic Presidential vote share	0.394	0.386	0.397	0.419
DI and TAA takeup per thousand (1990)				
TAA petition	0.351	0.657	0.934	1.137
TAA certification	0.150	0.270	0.441	0.909
DI application	12.058	9.341	9.767	13.360
DI awards	4.630	3.802	3.988	5.244
NAFTA vulnerability				
Vulnerability based on tariff in 1990	0.003	0.011	0.025	0.078
Exposure to Chinese imports				
ADH (2013) China shock measure (IV)	0.699	0.855	1.185	1.589
Number of counties	757	756	755	755

Sources: County-level demographics are from the Census Population Estimates Program (PEP), House election results are from ICPSR general election data for the United States (1980-1988), pre-NAFTA DI and TAA takeup are computed using the Social Security Administration (SSA) data and the U.S. Department of Labor TAA petition data, and the "China Shock" variable is from Autor et al. (2013) .

Notes: The table contains average county characteristics by quartiles of our NAFTA *Vulnerability* variable. Note that employment-to-population ratio comes from dividing employment in the county (from the CBPD) by population in the county (from the Census) and thus cannot be compared to the usual employment-to-population ratios based on whether *residents* are employed (regardless of where). Note that when we show how the "China shock" in Autor et al. (2013) varies across our four county groups, we use their "IV" version as it is more highly correlated with our NAFTA vulnerability measure.

Table 2.2: Approval of NAFTA as a function of state-level NAFTA vulnerability

	Dept. var: Supports NAFTA				
	(1)	(2)	(3)	(4)	(5)
State-level vulnerability	-0.0387**	-0.0427**	-0.0435***	-0.0445***	-0.0825***
	[0.018]	[0.019]	[0.016]	[0.016]	[0.020]
White			-0.0301**	-0.0210*	-0.0208
			[0.011]	[0.011]	[0.016]
Black			-0.0143	-0.00502	0.0123
			[0.015]	[0.014]	[0.016]
Male			0.0139*	0.0139*	-0.0587***
			[0.008]	[0.008]	[0.009]
No college degree			-0.0697***	-0.0683***	-0.0620***
			[0.008]	[0.008]	[0.011]
Log family income			0.0323***	0.0309***	0.0152**
			[0.007]	[0.007]	[0.008]
Union household			-0.0812***	-0.0757***	-0.104***
			[0.013]	[0.012]	[0.013]
Age / 100			-0.374***	-0.375***	-0.428***
			[0.025]	[0.025]	[0.029]
Dept. var. mean	0.381	0.415	0.415	0.415	0.538
Drop if missing covars	No	Yes	Yes	Yes	Yes
Division FE	No	No	No	Yes	Yes
Drop DK / no opinion	No	No	No	No	Yes
Observations	23297	16143	16143	16143	12431

Sources: Opinion polls from 1993-2015, many of which are from Pew. See Appendix B for survey dates, exact question wording, and other details.

Notes: Survey (which subsume year) fixed effects in all regressions. Col. (1) includes no other controls. Col. (2) replicates the col. (1) specification but on the subsample that has no missing values for the covariates used in subsequent columns. Col. (3) adds the covariates reported in the table. Col. (4) adds Census-division fixed effects. Col. (5) drops respondents who say they do not know enough about NAFTA or do not have an opinion. Standard errors clustered by state. $^{*}p = 0.1,^{**}p = 0.05,^{***}p = 0.01$.

Table 2.3: Partisan identity and views toward free trade, 1986-2012 repeated cross-sections

	Dep't var.: Party ID (1-7, increasing in Republican dir)					
	(1)	(2)	(3)	(4)	(5)	(6)
Favor import limits x	0.182**	0.190**	0.216***	0.228***	0.155**	0.209***
After 1992	[0.0719]	[0.0718]	[0.0699]	[0.0648]	[0.0771]	[0.0653]
Favor import limits	-0.222***	-0.227***	-0.265***	-0.278***	-0.385***	-0.269***
	[0.0706]	[0.0709]	[0.0713]	[0.0708]	[0.0837]	[0.0695]
Dep't var. mean	3.619	3.619	3.620	3.620	3.737	3.620
Controls						
–Demographic	No	No	Yes	Yes	Yes	Yes
–State FE	No	Yes	No	No	No	No
–Issues	No	No	No	Yes	Yes	Yes
–Demogr. x Aft	No	No	No	No	No	Yes
–Issues x Aft	No	No	No	No	No	Yes
Excl. DK	No	No	No	No	Yes	No
R-sq. x 100	0.680	2.787	11.988	16.271	15.462	17.176
Observations	18770	18770	18497	18497	11031	18497

Sources: ANES time-series files (repeated cross-sections), 1986–2012. We include all surveys in this interval that ask the *Favor Import Limits* question (see Section A).

Notes: Year fixed effects are in all regressions. Col. (1) includes no other controls. Col. (2) replicates the col. (1) specification but adds state fixed effects. Col. (3) adds to the col. (1) specification controls for race, gender, education, age, and log of family income. Col. (4) adds to the col. (3) specification views toward abortion, trust in government and feelings towards African-Americans. Col. (5) replicates col. (4) but drops any respondent who says "don't know" in response to the *Favor Import Limits*$_i$ question (they are otherwise coded as zero). Col. (6) adds to col. (4) interactions between *After 1992* and each of the controls in col. (3) and col (4). Standard errors clustered by state. $^{*}p = 0.1,^{**}p = 0.05,^{***}p = 0.01$.

Table 2.4: Partisan identity and views toward free trade, 1992-1994 panel data

	Move in Repub direction dummy x 100					
	(1)	(2)	(3)	(4)	(5)	(6)
Favor import limits	8.304**	9.530**	8.301**	8.065**	8.833**	8.512**
	[3.325]	[4.108]	[3.443]	[3.580]	[3.715]	[4.041]
Minorities sd help self				1.388	1.474	1.600
				[1.071]	[1.037]	[1.026]
Wants active gov't				-0.923	-0.911	-1.820
				[1.121]	[1.281]	[1.439]
Support abortion				-1.770	-1.089	-1.254
				[1.872]	[2.159]	[2.281]
Attend church weekly				7.772**	8.211**	6.940*
				[3.632]	[3.819]	[3.937]
Favors increased immigr.				0.222	-2.946	-4.598
				[5.899]	[6.579]	[7.056]
Oppose gays in military					3.408	2.745
					[7.246]	[7.841]
Oppose gov't health care					-0.544	-0.928
					[0.782]	[0.847]
Favor term limits					-6.178*	-5.557
					[3.620]	[4.138]
Dept. var. mean	26.52	26.76	26.49	26.49	26.54	26.54
Ex. DK	No	Yes	No	No	No	No
Demog. covars	No	No	Yes	Yes	Yes	Yes
State FE	No	No	No	No	No	Yes
R-squared	0.00887	0.0104	0.0388	0.0607	0.0660	0.103
Observations	739	553	736	736	731	731

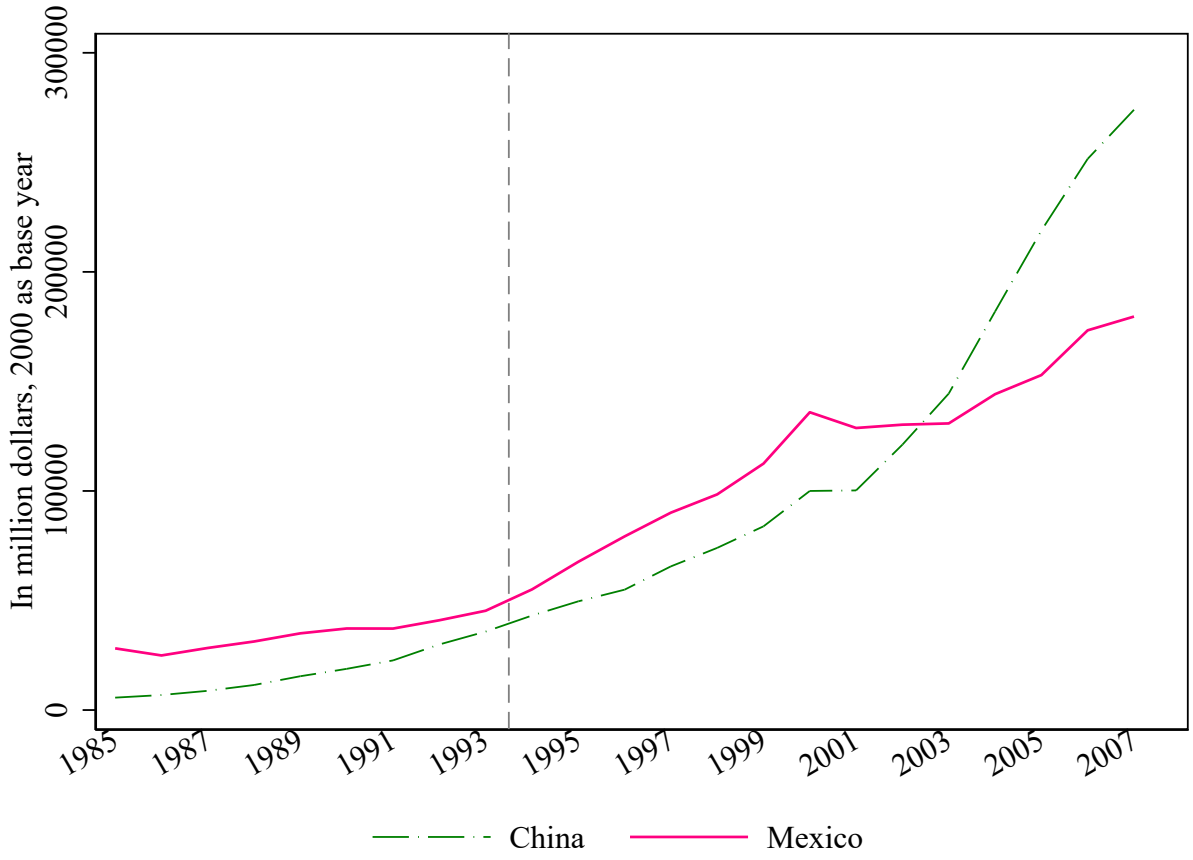
Sources: ANES panel data, 1992-1994.

Notes: The dependent variable is a dummy (multiplied by 100) for whether the respondent moved in the GOP direction in the 1-7 partisan identity scale. All explanatory variables were asked in 1992. "Ex. DK" means that respondents who did not have an opinion are dropped (they are otherwise coded as zero). Demographic controls include race, gender, education, age, log family income, and urbanicity. Standard errors clustered by state. $^{*}p = 0.1,^{**}p = 0.05,^{***}p = 0.01$.

Appendices

A Appendix Figures and Tables

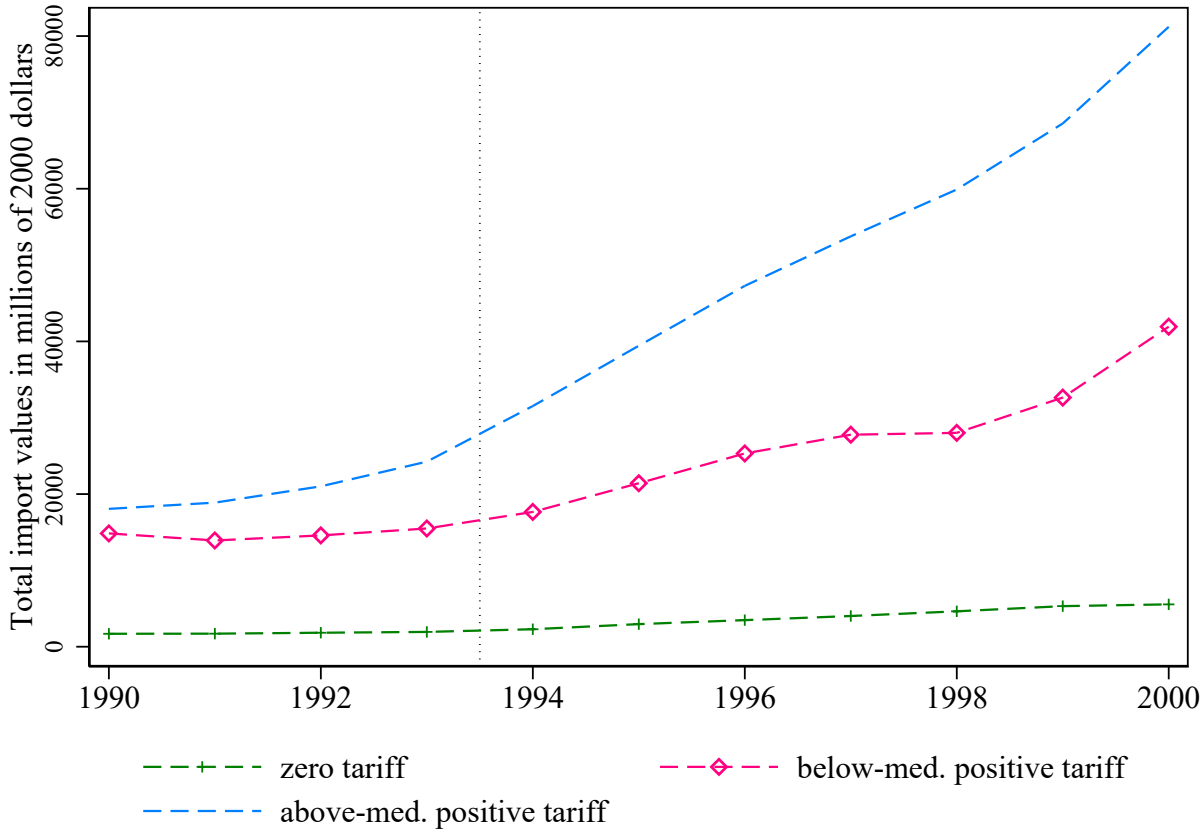
Figure 2A.1: U.S. imports from China and Mexico



Sources: Federal Reserve Economic Data Series (FRED).

Notes: The figure contains the time series of the value of goods imported by the US, based on the custom basis from China and Mexico. The import values are inflation-adjusted using the quarterly-level personal consumption expenditures available from FRED.

Figure 2A.2: Import values from Mexico by 1990 tariff level

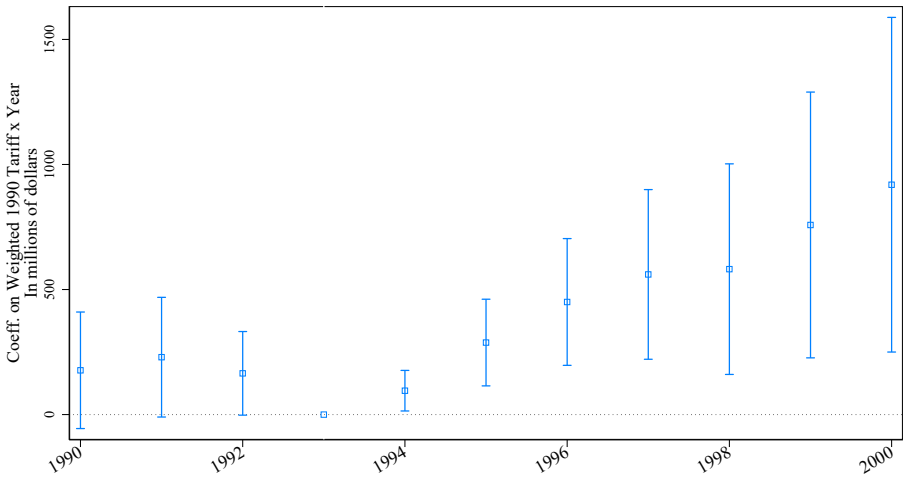


Sources: The import values are taken from the U.S. International Trade Commission (USITC).

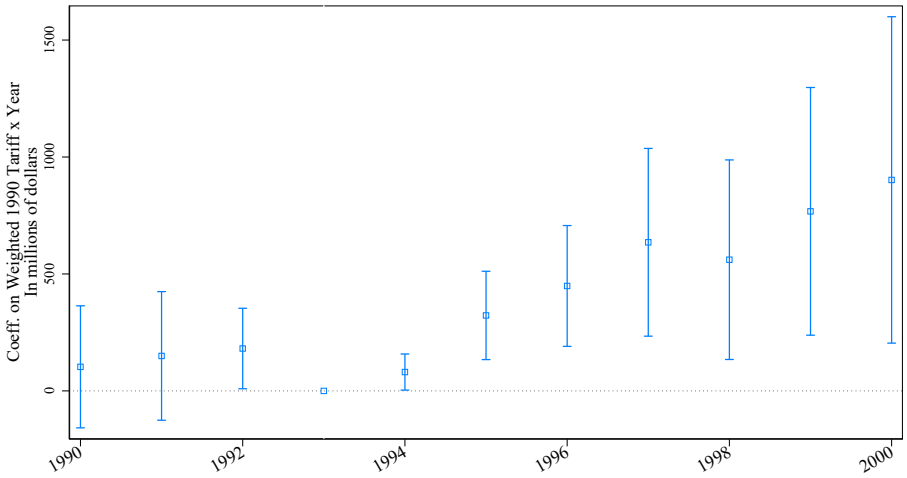
Notes: The figure shows the time series of average import values from Mexico by industries with zero tariff, industries with below-median tariff, and industries with above-median tariff based on 1990 industry-level tariff.

Figure 2A.3: Relationship (in levels) between Mexican imports to the US and pre-NAFTA tariffs

(a) U.S. ITC data



(b) U.N. data



Sources: U.S. International Trade Commission (panel a) and UN Comtrade data (panel b)

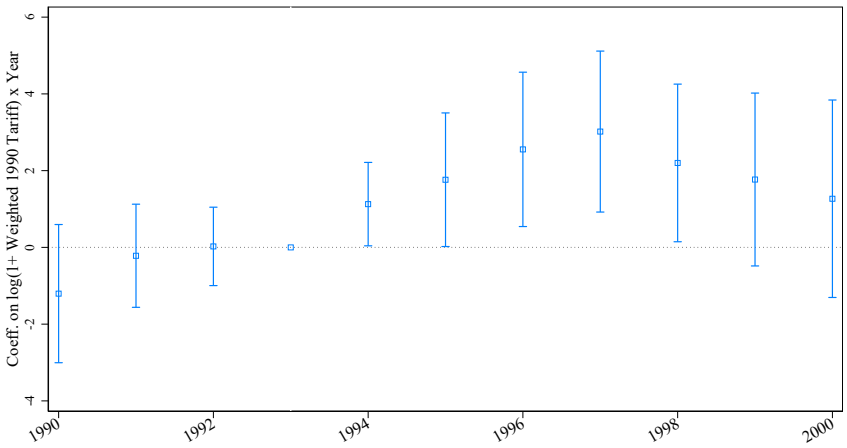
Notes: The figure shows the coefficients β^t from the following regression:

$$MexImports_{jt}^{US} = \beta^t Avg.Tariff_j^{1990} + \gamma_1 MexImports_{jt}^{ROW} + \gamma_2 ROWImports_{jt}^{US} + \eta_j + \mu_t + e_{jt},$$

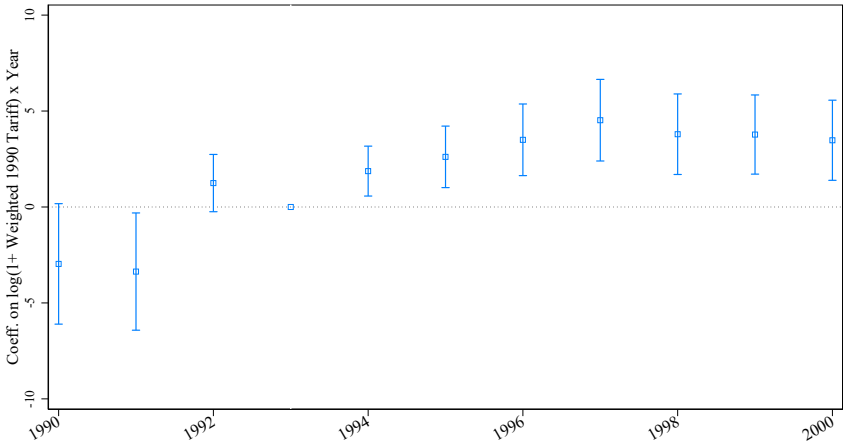
where $Avg.Tariff_j^{1990}$ is the weighted average tariff at the SIC four-digit industry level in 1990, $MexImports_{jt}^{US}$ are Mexican imports to the US for SIC four-digit industry j in year t, $MexImports_{jt}^{ROW}$ is Mexican imports to the rest of the world (ROW) for industry j in year t, $ROWImports_{jt}^{US}$ is the rest of the world's imports to the US for industry j in year t, and η_j and μ_t are industry and year fixed effects, respectively. All import values are in millions of current USD. The 95-percent confidence intervals are based on standard errors clustered at the industry level.

Figure 2A.4: Relationship (in logs) between Mexican imports to the US and pre-NAFTA tariffs

(a) U.S. ITC data



(b) U.N. data

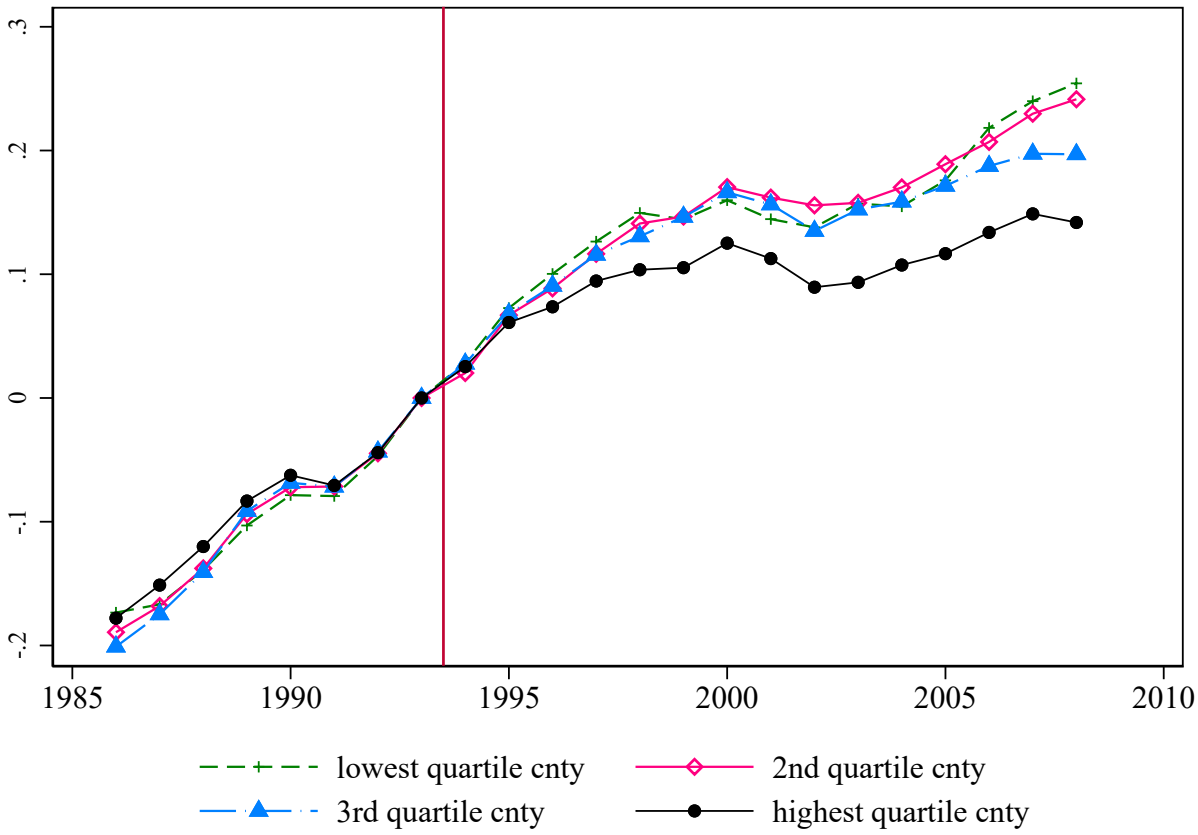


Sources: U.S. International Trade Commission (panel a) and UN Comtrade data (panel b)
Notes: The figures shows the coefficients β^t from the following regression:

$$\log(MexImports)_{jt}^{US} = \beta^t \log(1 + Avg.Tariff_j^{1990}) + \gamma_1 \log(MexImports_{jt}^{ROW}) + \gamma_2 \log(ROWImports)_{jt}^{US}) + \eta_j + \mu_t + e_{jt},$$

where $Avg.Tariff_j^{1990}$ is the weighted average tariff at the SIC four-digit industry level in 1990, $MexImports_{jt}^{US}$ are Mexican imports to the US for SIC four-digit industry j in year t, $MexImports_{jt}^{ROW}$ is Mexican imports to the rest of the world (ROW) for industry j in year t, $ROWImports_{jt}^{US}$ is the rest of the world's imports to the US for industry j in year t, and η_j and μ_t are industry and year fixed effects, respectively. All import values are in millions of current USD. The 95-percent confidence intervals are based on standard errors clustered at the industry level. Note that we use $\log(1 + Avg.Tariff_j^{1990})$ as our variable of interest because over forty percent of four-digit industry groups have zero average tariff in 1990.

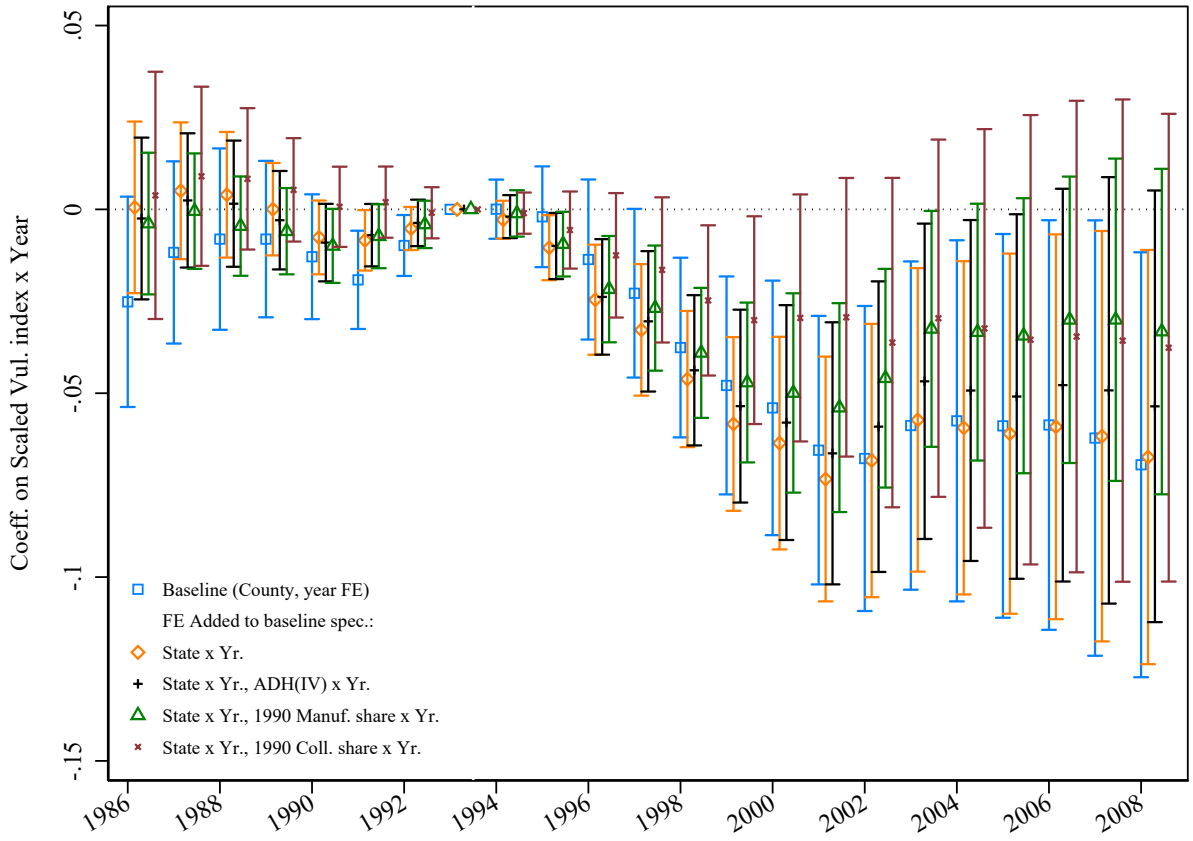
Figure 2A.5: Average log employment for four vulnerability quartiles over time (normalized to zero in 1993)



Sources: The dependent variable is derived from the County Business Patterns (CBP). See Appendices B and B for more detail.

Notes: The figure shows log of total employment trends from 1986 to 2008, separately by 1990 county vulnerability quartiles. Log of total employment is computed using the CBPD. We do not weight and other than normalizing to zero in 1993, the data plotted are simply raw annual means within the quartiles.

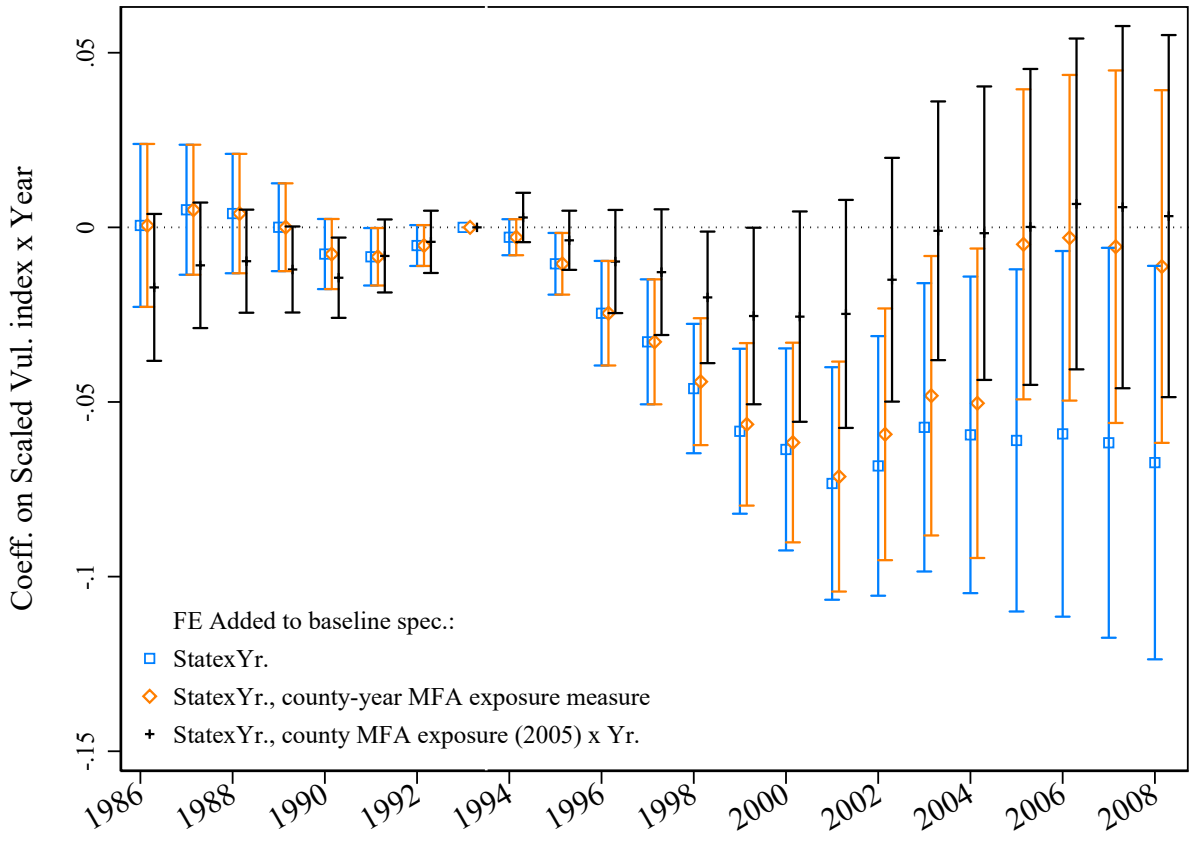
Figure 2A.6: Log employment as a function of county NAFTA vulnerability, varying controls



Sources: The dependent variable is computed from County Business Patterns. See Appendices B and B for more detail.

Notes: This figure is analogous to Figure 2.2 but shows additional specifications after varying the controls. The first series includes only county and year fixed effects. The second series adds to this specification *state*×*year* fixed effects and the the third specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects (and are thus identical to the first and second series of Figure 2A.6). The fourth specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects. The final specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects.

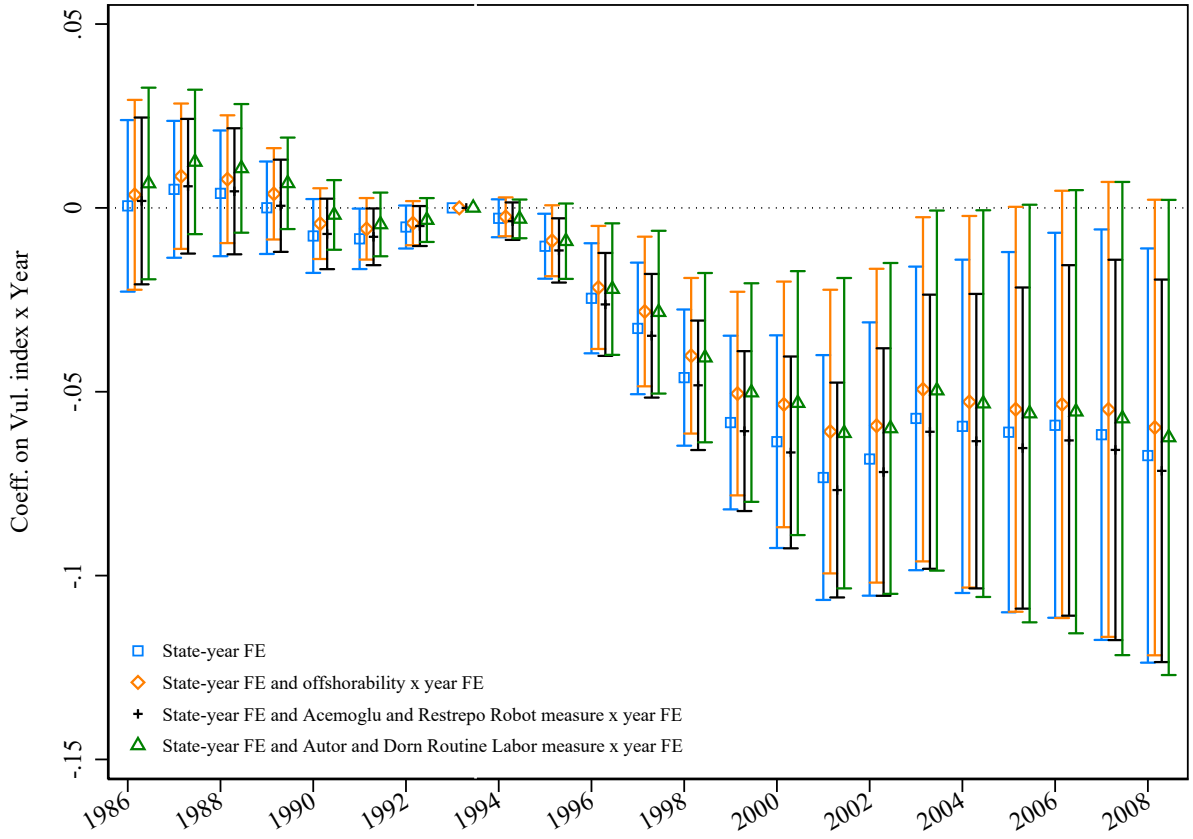
Figure 2A.7: Log employment as a function of NAFTA vulnerability, robustness to Multi-Fibre-Arrangement phase-out



Sources: The dependent variable is derived from the CBP. A county-year-level measure of exposure to the MFA is drawn from Pierce and Schott (2020b), which is based on the approach from Khandelwal et al. (2013).

Notes: $N = 2926$ counties.Under the Uruguay round of the General Agreement on Tariffs and Trade, the quotas under the Multi-Fibre Arrangement (MFA) were put on a phase-out schedule (the final year of the phase-out was announced as 2005). As there was little actual change in the "binding" quotas until the early 2000s, the *contemporaneous* quota fill rate is not a potential confounding variable in our NAFTA analysis, as can we seen by comparing the first and second series of this graph. However, to the extent that agents are perfectly forward-looking, the known end of the quotas by 2005 could potentially cause declines in protected industries many years earlier. We thus draw a county-level MFA vulnerability measure *based on 2005 average quota fill rates* from Pierce and Schott (2020) and interact it with every year in our sample period. That is, we let these future quota fill rates have arbitrary effects in all years. Comparing the second and third series shows that even after flexibly controlling for any forward-looking effects of the MFA phase-out, we still identify a large, negative effect of NAFTA vulnerability on county employment that is statistically significant in the late 1990s.

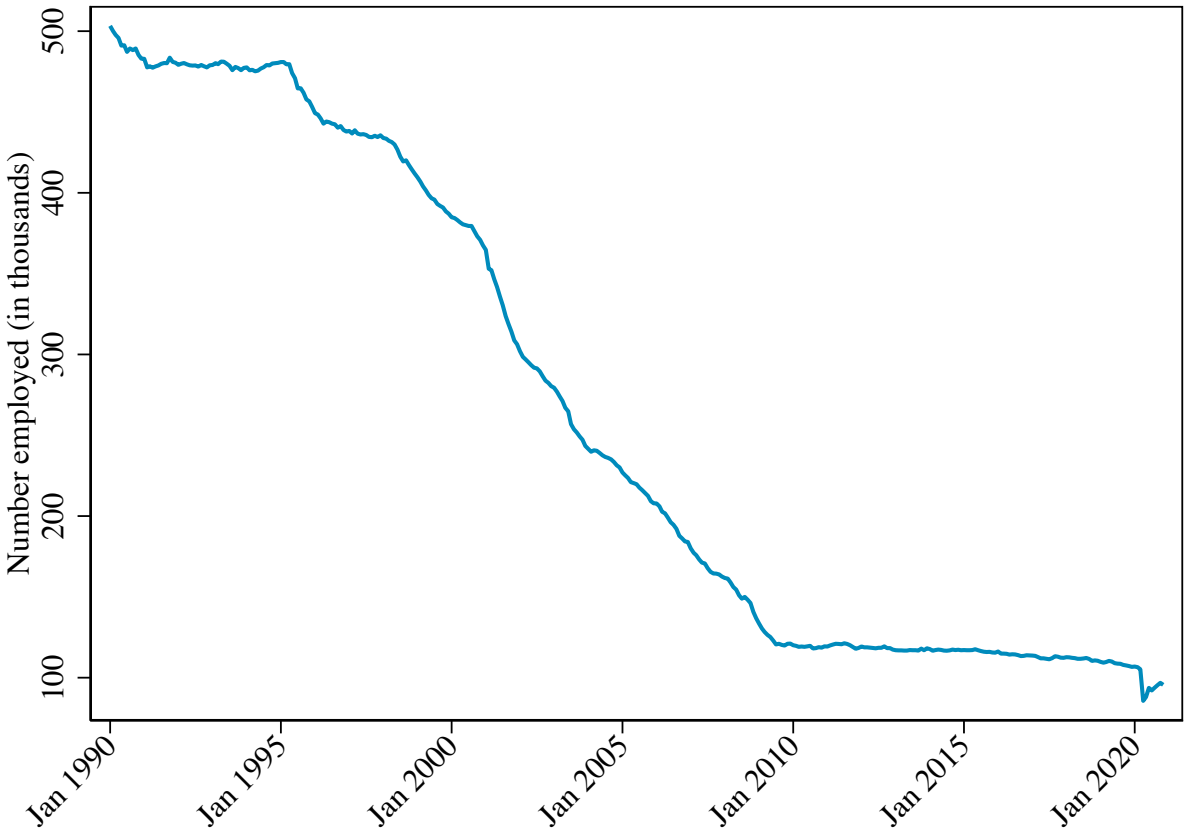
Figure 2A.8: Log employment as a function of county NAFTA vulnerability, robustness to automation, off-shoring and immigration



Sources: The dependent variable is derived from the CBP, and county-level demographics are from the Census PEP.

Notes: Sample contains $N = 2926$ counties for each year of the sample period. This figure extends the analysis in Figure 2.2. The figure shows the event-study coefficient estimates (and 95% confidence intervals, based on standard errors clustered by state) from different specifications of equation (2.4), where log of total employment at the county×year level is the dependent variable. All specifications are weighted by 1990 county population. The first series shows the coefficient estimates from the baseline *state* × *year* specification. The second specification adds to the first specification year fixed effects interacted with CZ-level "offshorability" based on 1980 occupation, as used in Autor et al. (2013). XXX Pls describe the rest with citations as needed.

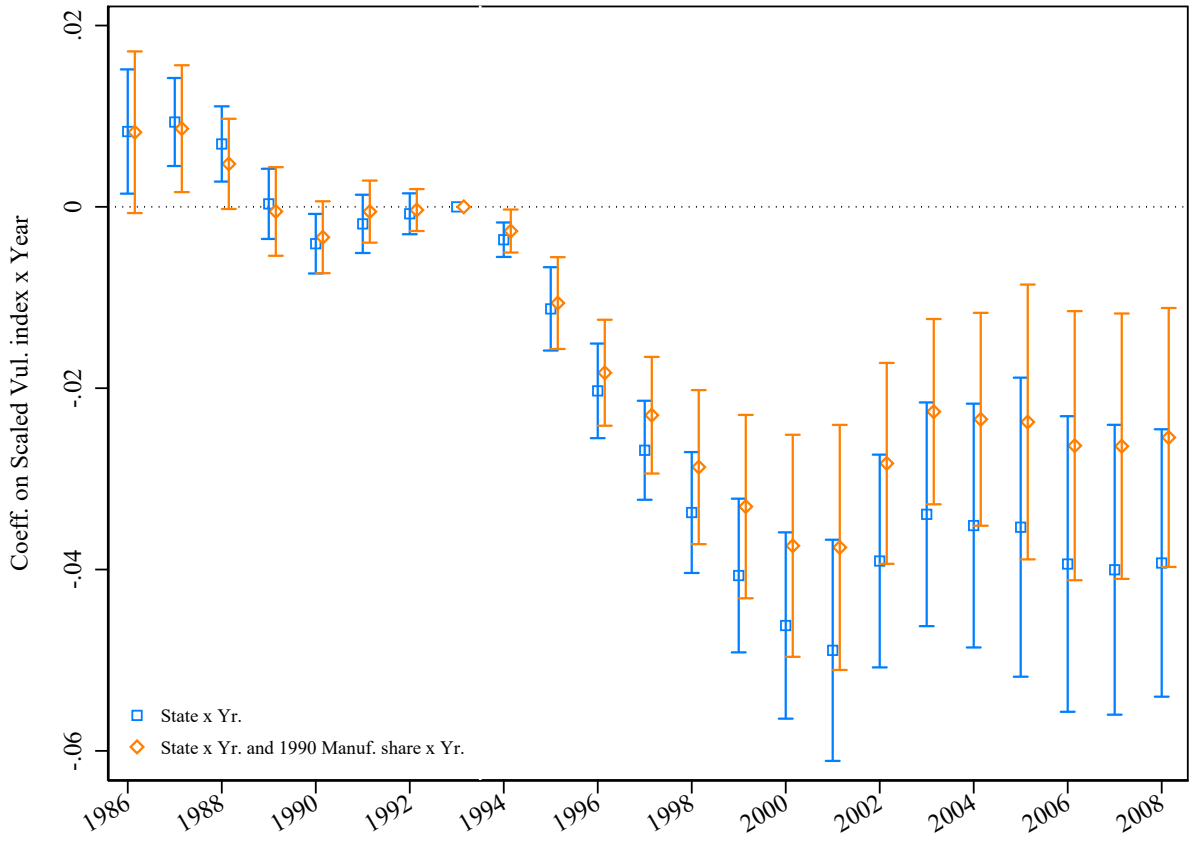
Figure 2A.9: Employment in textile mills, 1990-2020



Sources: U.S. Bureau of Labor Statistics, All Employees, Textile Mills [CES3231300001], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CES3231300001, December 7, 2020.

Notes: The data series provided by FRED begins only in 1990, so we cannot look earlier in time with this data series.

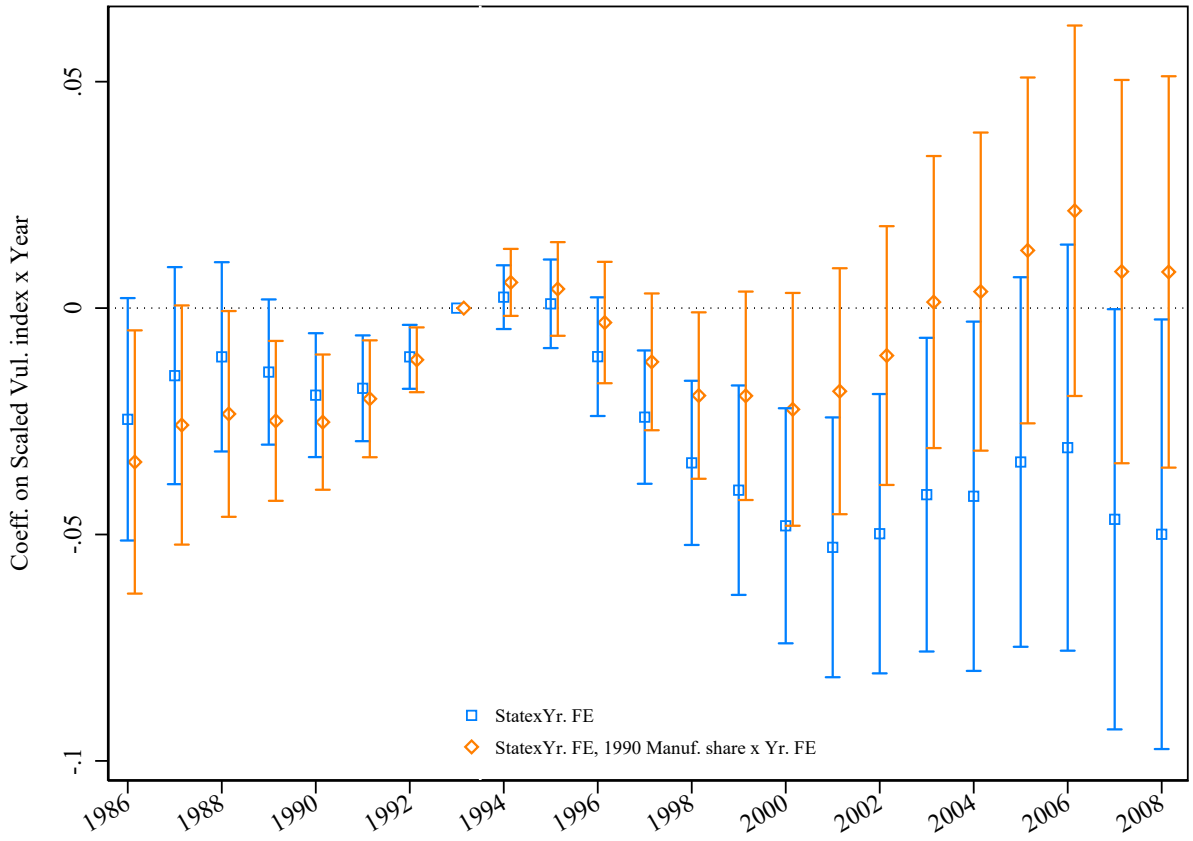
Figure 2A.10: Employment per capita as a function of NAFTA vulnerability



Sources: The dependent variable is derived from the CBP and the census PEP. Note that the denominator is 1990 working-age population.

Notes: $N = 2990$ counties for each year of the sample. The figure is identical to Figure 2.2 except that the outcome variable is per capita employment and not log of total county employment. The figure shows the event-study coefficient estimates (and 95% confidence intervals, based on standard errors clustered by state) from different specifications of equation (2.4), where per capita employment at the county×year level is the dependent variable. All specifications are weighted by 1990 county population. The first series shows the coefficient estimates from a specification where we control for only county and year fixed effects. The specification for the second series adds to this baseline specification *state*×*year* fixed effects. The third specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects. The fourth specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects. The final specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects.

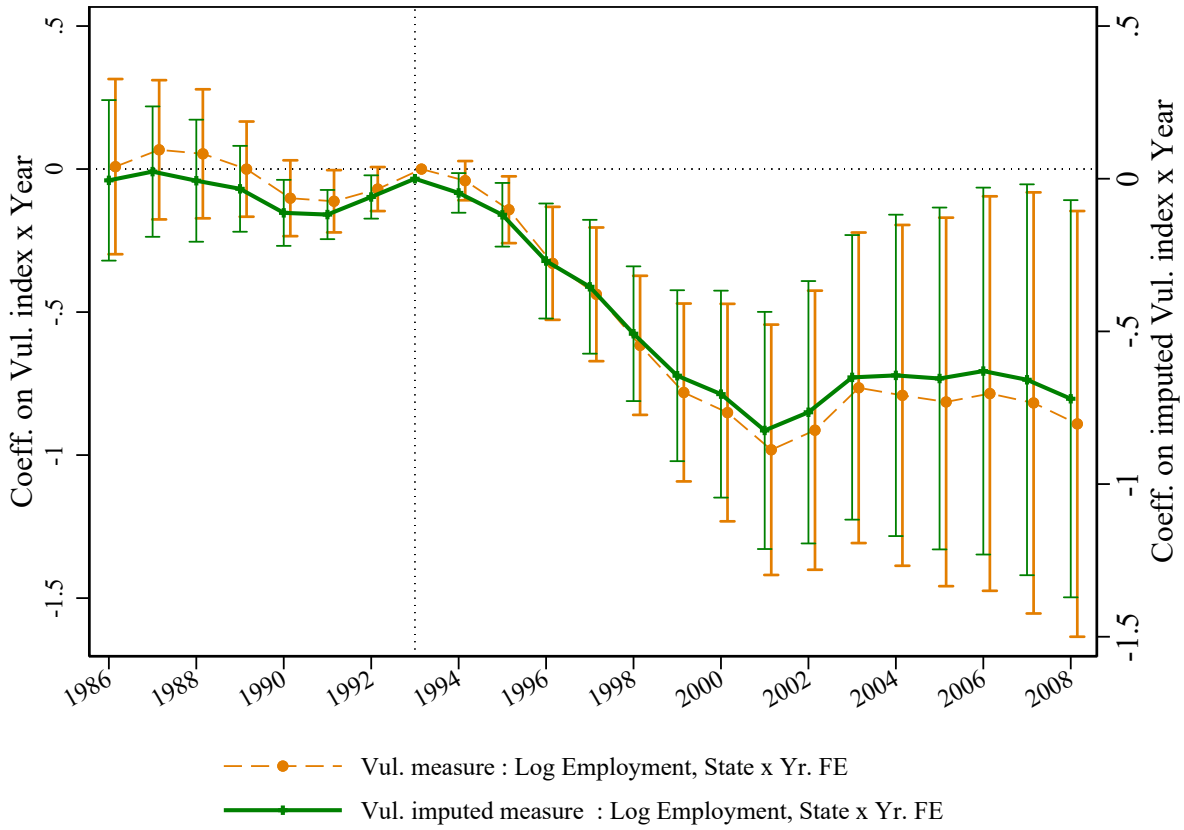
Figure 2A.11: Log of CZ employment as a function of CZ-level NAFTA vulnerability



Sources: The dependent variable is derived from the CBP.

Notes: $N = 705$ CZs for each year of the sample period. This figure is the analogue to Figure 2.2 but at the CZ, not county, level. The figure shows the event-study coefficient estimates (and 95% confidence intervals, based on standard errors clustered by CZ) from different specifications of equation (2.4), where log employment at the CZ×year level is the dependent variable. All specifications are weighted by 1990 CZ population. The first series shows the coefficient estimates from a specification where we control for only CZ and year fixed effects. The specification for the second series adds to this baseline specification *state*×*year* fixed effects, where CZs are assigned to state using David Dorn's CZ to state crosswalk. Whenever a CZ crosses more than one state, the CZ is assigned to a state with the largest share of CZ's population. The third specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects. The fourth specification adds to the second specification 1990 CZ-level manufacturing share of employment interacted with year fixed effects. The final specification adds to the second specification 1990 CZ-level college-graduate population share interacted with year fixed effects.

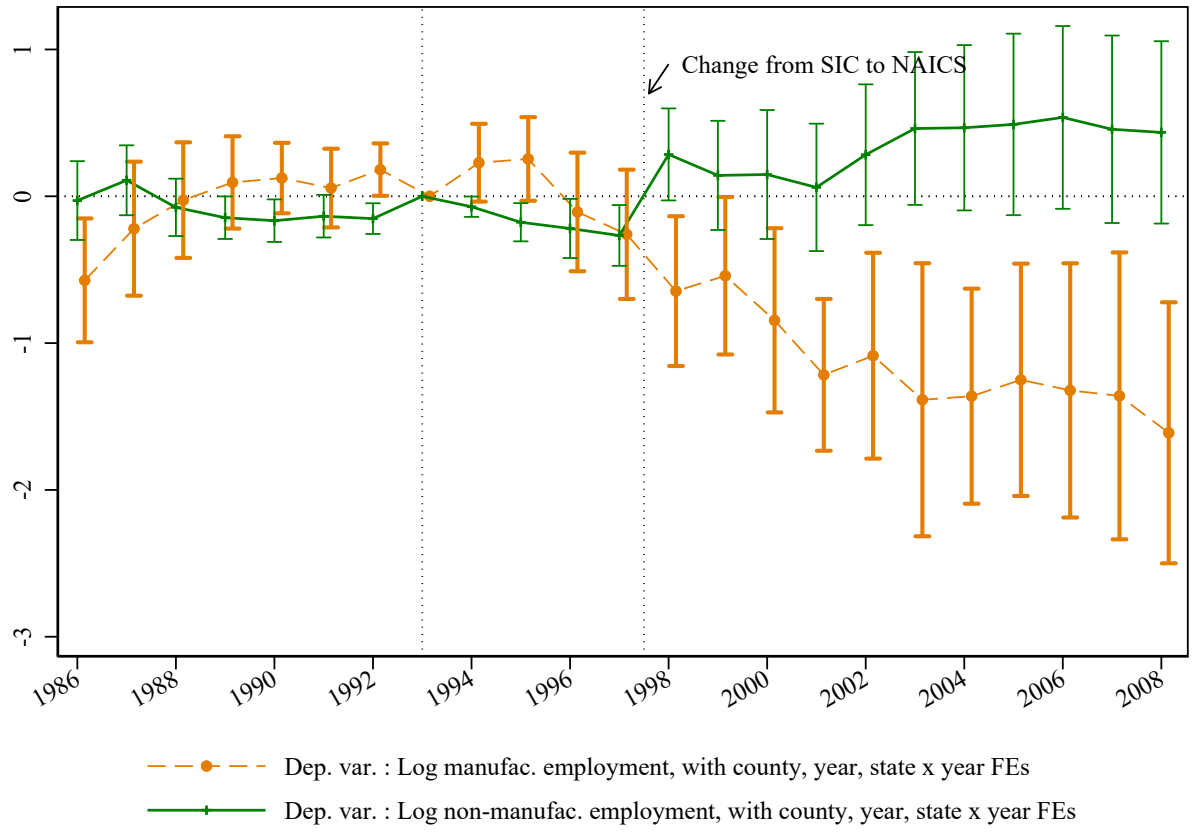
Figure 2A.12: Log of employment as a function of NAFTA vulnerability using imputed CBP cells from Eckert et al. (2021)



Sources: The dependent variable is derived from the CBP.

Notes: $N = 2822$ counties. The figure shows the event-study coefficient estimates (plus 95%-confidence intervals, based on standard errors clustered by state) from specifications of equation (2.4), where log of county employment is the dependent variable. The first series uses our baseline vulnerability measure as the main independent variable, and the specification includes county and year fixed effect and $state \times year$ fixed effects. The second series uses the vulnerability measure using the imputed county-industry cells proposed by Eckert et al. (2021), and the specification includes county and year fixed effect and $state \times year$ fixed effects.

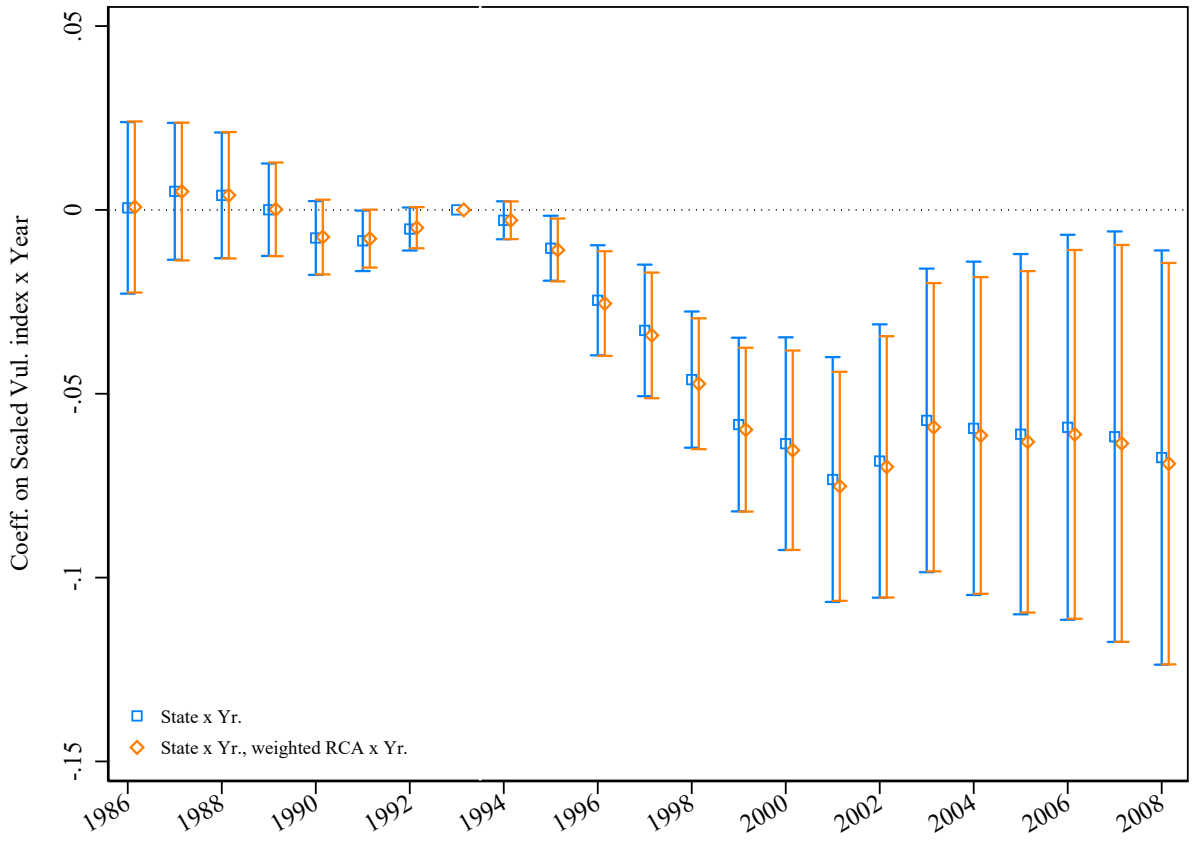
Figure 2A.13: Evolution of log employment as a function of NAFTA vulnerability, separating manufacturing v. other industries



Sources: The dependent variable and the codes to categorize manufacturing industries are derived from the CBP.

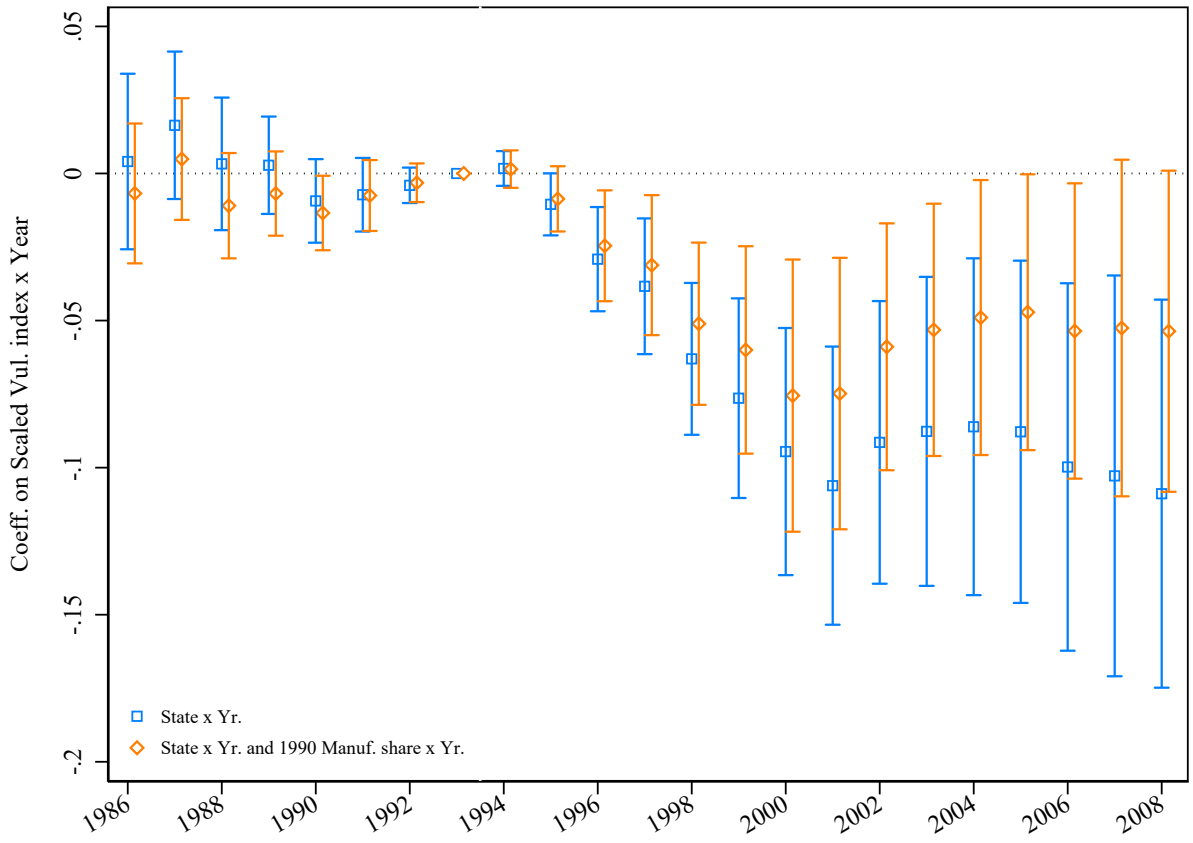
Notes: $N = 2960$ counties. The figure shows the event-study coefficient estimates (plus 95%-confidence intervals, based on standard errors clustered by state) from specifications of equation (2.4), where log of total manufacturing employment and log of total non-manufacturing employment at the county×year level are the dependent variable for the first and second series, respectively. Both specifications are weighted by 1990 county population, and they include county and year fixed effects and *state*×*year* fixed effects.

Figure 2A.14: Log of employment as a function of county NAFTA vulnerability, robustness to Peso crisis



Notes: $N = 2926$ counties for each year of the sample period. The figure shows the event-study coefficient estimates from different specifications of equation (2.4), where log of total employment is the dependent variable. All specifications are weighted by 1990 county population. In the first series we replicated our usual specification with *state*×*year* fixed effects. The second specification seeks to control for any effect of the 1994-1995 Mexican peso crisis, which we do as follows. The peso crisis made *all* Mexican goods cheaper, regardless of pre-NAFTA tariff status. Thus, we create a vulnerability measure that excludes the pre-NAFTA tariff level and simply weights pre-NAFTA 1990 county employment by its dependence on industries where Mexico has high revealed comparative advantage (as measured in 1990, regardless of tariff level). The second series adds as a control this county-level variable interacted with year fixed effects. Comparing the two series suggests that the estimated effect of NAFTA vulnerability does not change much after flexibly controlling for county-level exposure to the devalued peso.

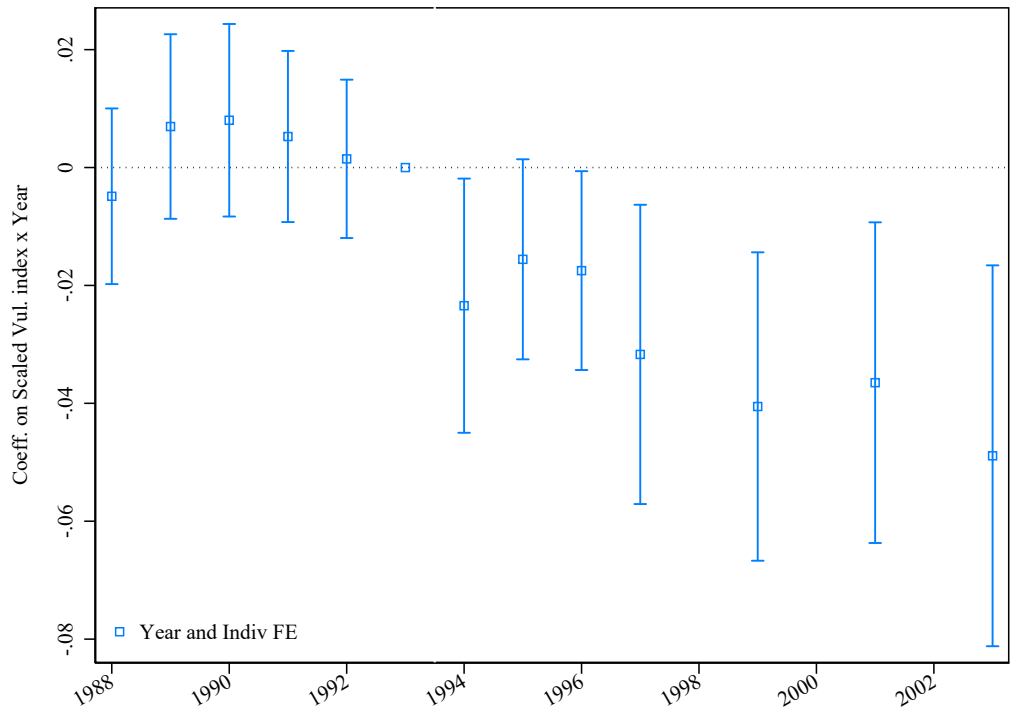
Figure 2A.15: Log total wage bill as a function of NAFTA vulnerability



Sources: CBP data.

Notes: Each series includes $N = 2932$ counties in a strictly balanced panel. The figure is identical to Figure 2.2 except that the outcome variable is total log wage bill instead of total log employment.

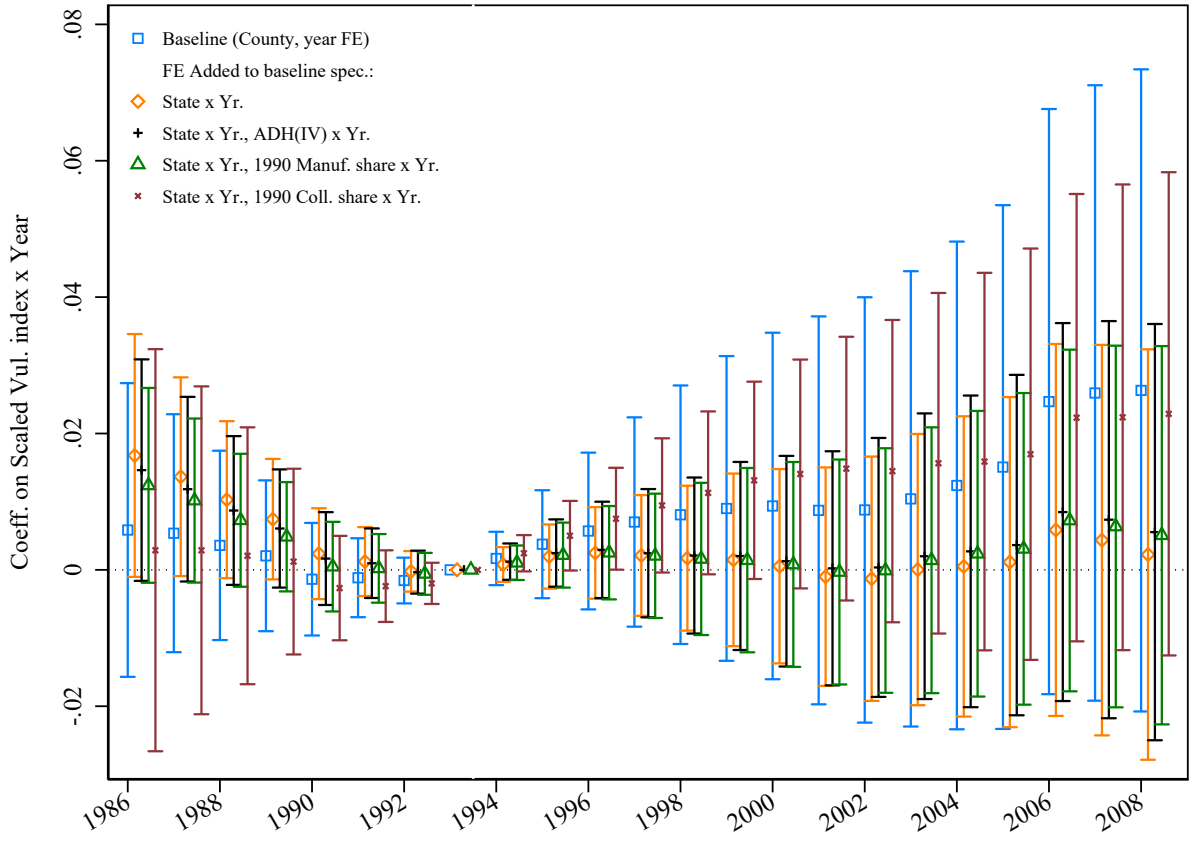
Figure 2A.16: Employment as a function of individual vulnerability, PSID sample



Sources: PSID panel data. See Appendix B for more detail.

Notes: $N = 4373$ individuals. This figure does not use geography to assign vulnerability to NAFTA but instead the individual's industry in a baseline (1988) pre-NAFTA year. We define individual-level i's vulnerability to NAFTA as $Vulnerability_i = RCA_{j(i)} \cdot \tau^{1990}_{j(i)}$, where $j(i)$ is industry j of person i in 1988 (or, if unemployed that year, their most recent industry), $RCA_{j(i)}$ is Mexico's revealed comparative advantage in industry j, and $\tau^{1990}_{j(i)}$ is the U.S. tariff on Mexican imports in industry j in 1990. The specification regressed a dummy variable for being employed in year t on year fixed effects and $Vulnerability_i$ interacted with year and individual fixed effects (and reports the coefficients on these interaction terms).

Figure 2A.17: Log population as a function of county NAFTA vulnerability, varying controls

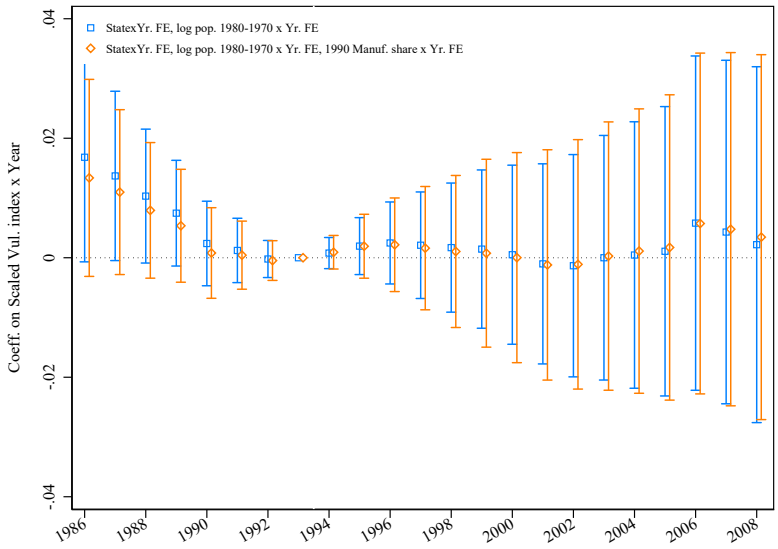


Sources: The dependent variable is taken from the Census Bureau's Population Estimates Program. See Appendix B for more detail.

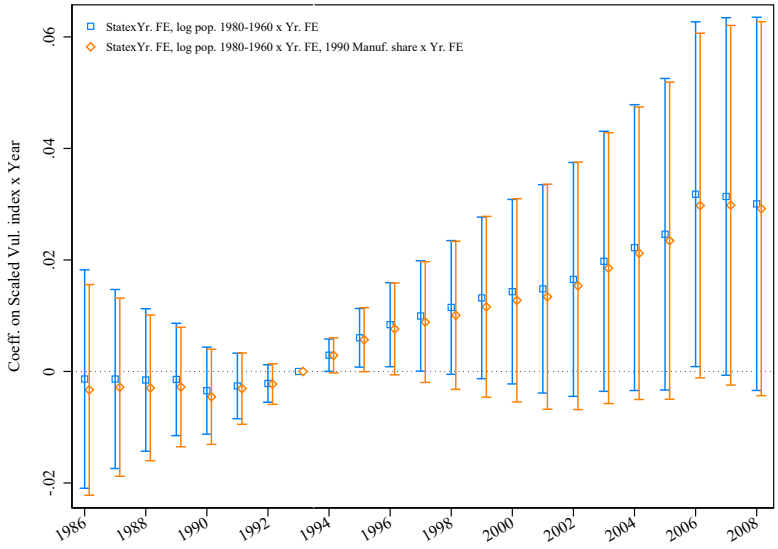
Notes: This figure is analogous to Figure 2.2 but shows additional specifications after varying the controls. The first series includes only county and year fixed effects. The second series adds to this specification *state*×*year* fixed effects and the the third specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects (and are thus identical to the first and second series of Figure 2A.6). The fourth specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects. The final specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects.

Figure 2A.18: Log population as a function of NAFTA vulnerability, controlling for long population changes

(a) Controlling for 1970-1980 county pop change

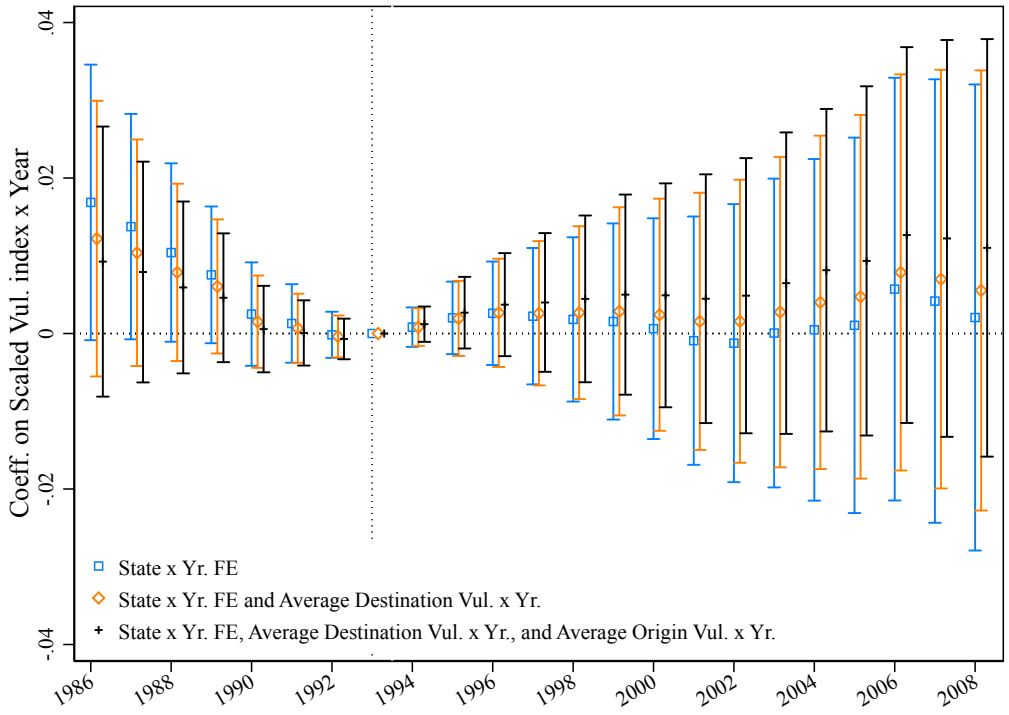


(b) Controlling for 1960-1980 county pop change



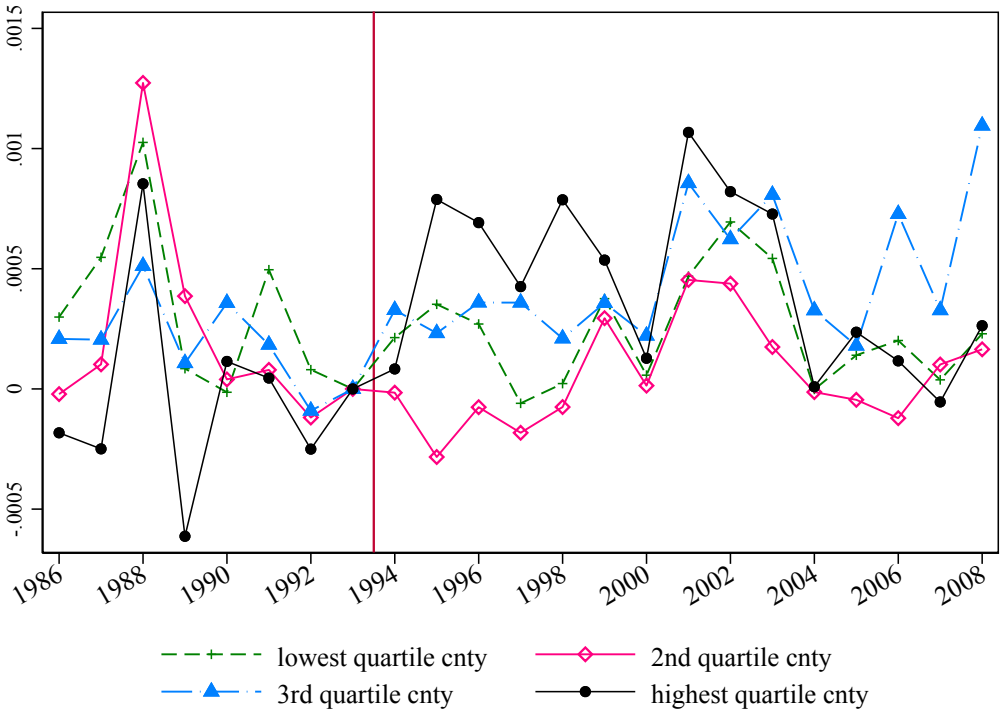
Notes: These specifications mirror those in Figure 2.3 but control for 'long differences' in county population. Panel (*a*) includes in all specifications the 1970-1980 county population change interacted with year fixed effects. Panel (*b*) includes in all specifications the 1960-1970 county population change interacted with year fixed effects.

Figure 2A.19: Migration response to NAFTA, adjusting for NAFTA vulnerability of a county's pre-existing migration network



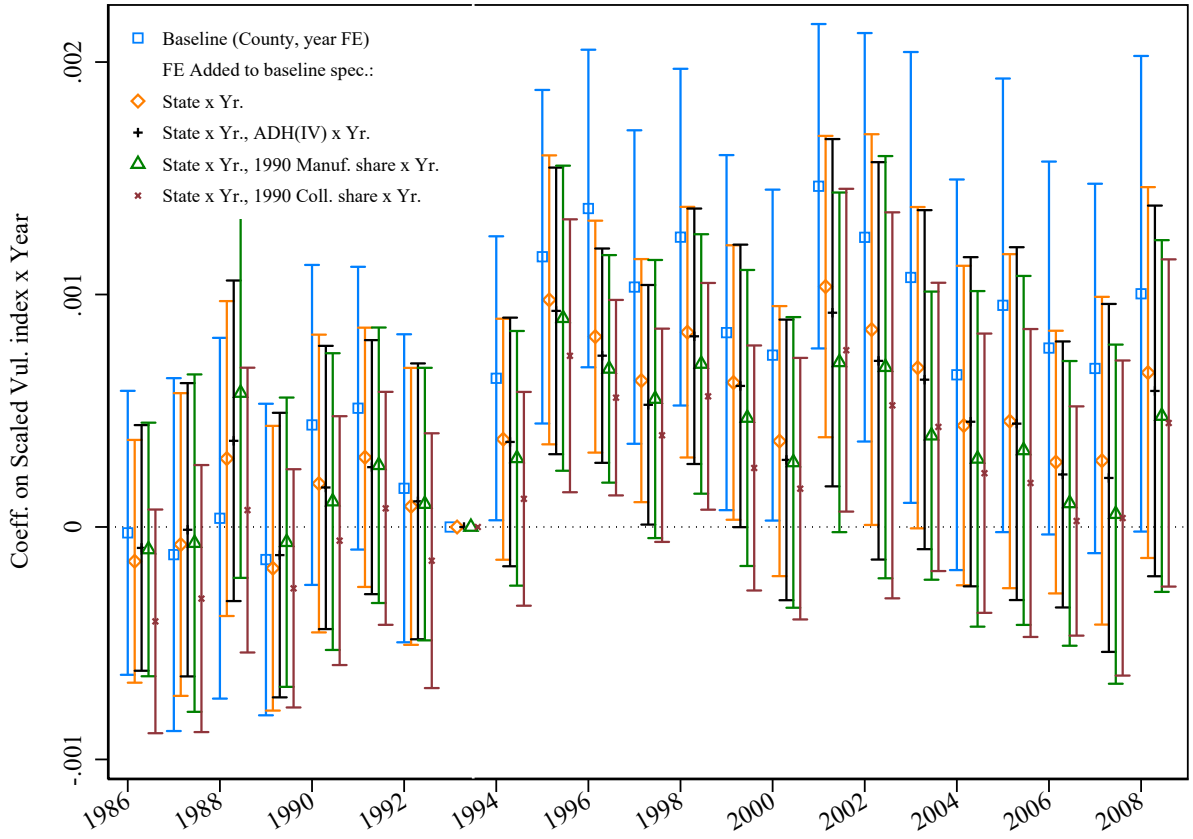
Notes: This figure is an extension of Figure 2.3. The first series replicates the first series of Figure 2.3 but adds the average *Vulnerability* measure of destination counties based on 1990 migration patterns (interacted with year fixed effects). The third series adds the average *Vulnerability* measure of origin counties based on 1990 migration patterns (interacted with year fixed effects).

Figure 2A.20: Trade Adjustment Assistance petitions per capita, by county *Vulnerability* quartile (normalized to zero in 1993)



Notes: This figure uses the same data as in Figure 2.4 but simply shows raw averages (normalized to zero in 1993) by the four quartiles based on county *Vulnerability*.

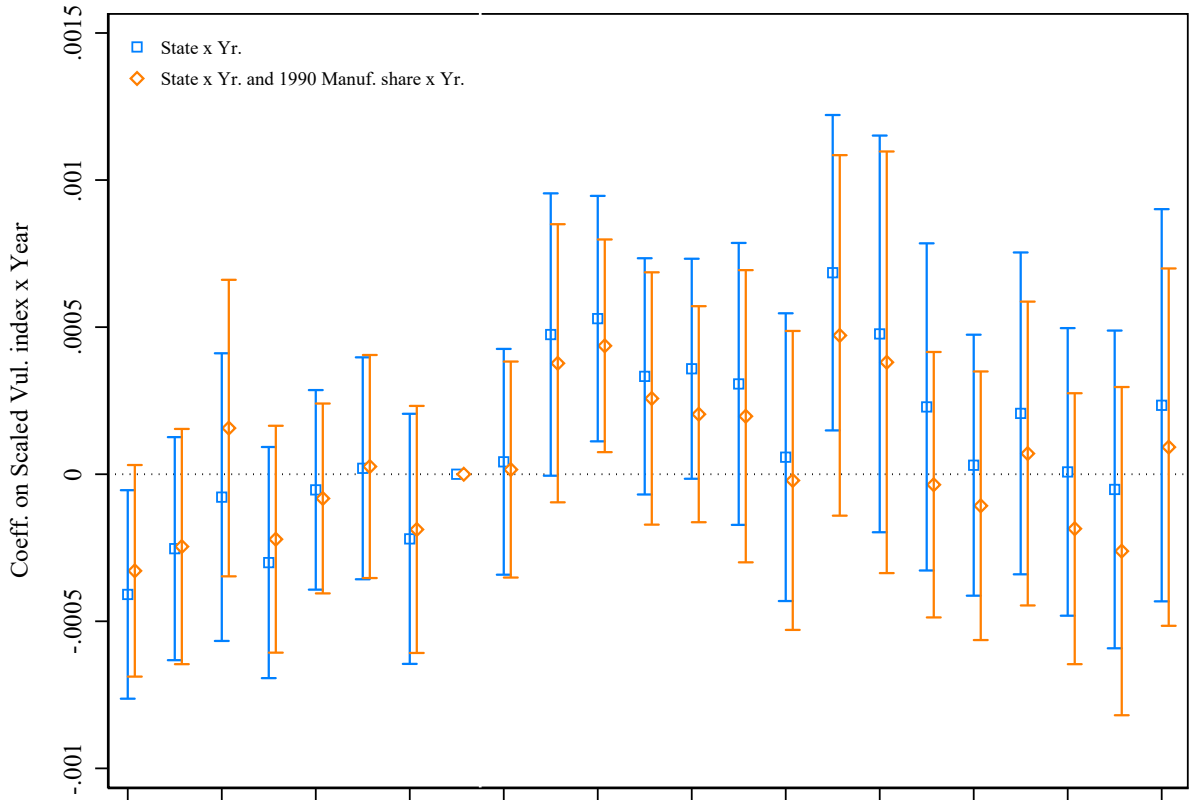
Figure 2A.21: Trade Adjustment Assistance petitions per capita as a function of county NAFTA vulnerability, varying controls



Sources: The dependent variable is taken from the U.S. Department of Labor TAA petition data, divided by 1990 working-age county population. See Appendix B for more detail.

Notes: This figure is analogous to Figure 2.4 but shows additional specifications after varying the controls. The first series includes only county and year fixed effects. The second series adds to this specification *state*×*year* fixed effects and the the third specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects (and are thus analogous to the first and second series of Figure 2A.6). The fourth specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects. The final specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects.

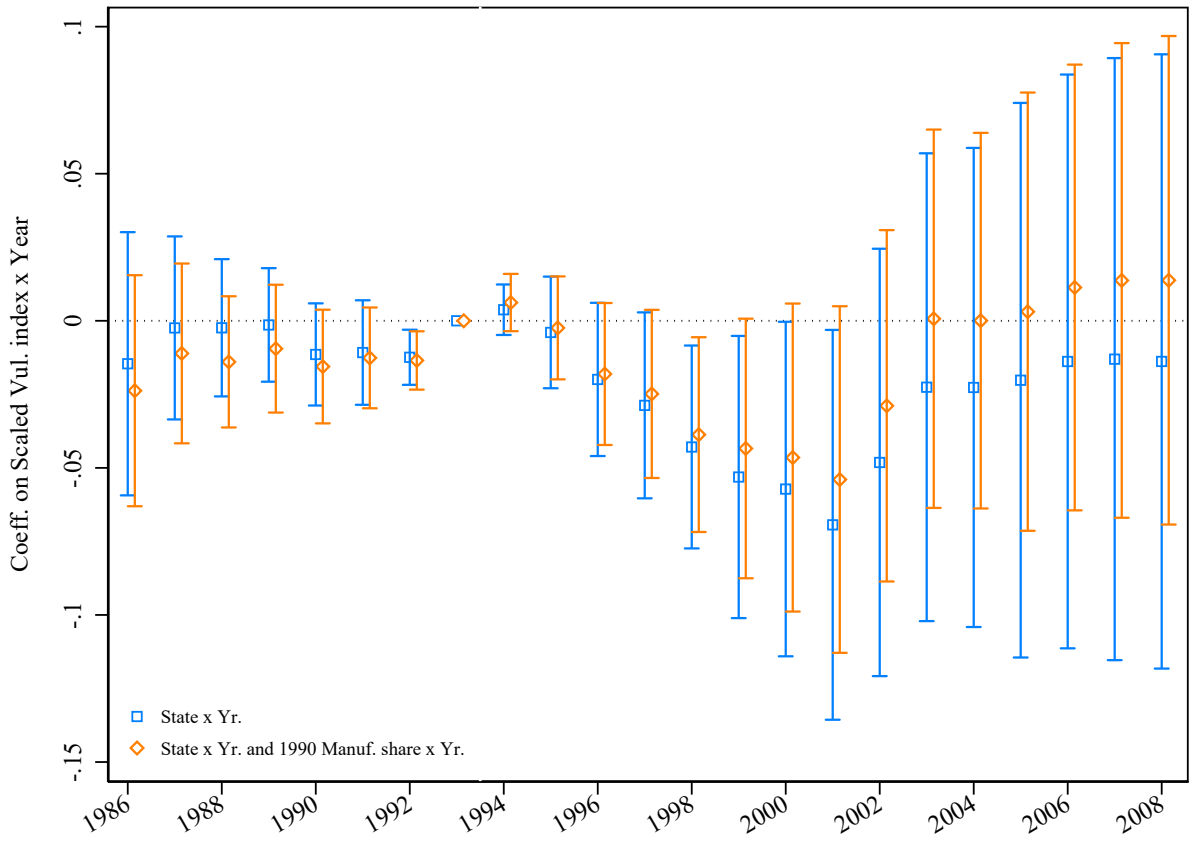
Figure 2A.22: Trade Adjustment Assistance *certifications* per capita as a function of county vulnerability



Sources: The dependent variable is taken from the U.S. Department of Labor TAA petition data. We divide by 1990 county working-age population. See Appendix B for more detail.

Notes: $N = 2914$ counties for each year of the sample period. This figure is identical to Figure 2.4 except that the dependent variable is TAA *certifications* per capita instead of *petitions*. The figure shows the event-study coefficient estimates (and 95% confidence intervals, based on standard errors clustered by state) from different specifications of equation (2.4), where per capita TAA certifications at the county×year level is the dependent variable. All specifications are weighted by 1990 county population. The first series shows the coefficient estimates from a specification where we control for only county and year fixed effects. The specification for the second series adds to this baseline specification *state*×*year* fixed effects. The third specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects. The fourth specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects. The final specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects.

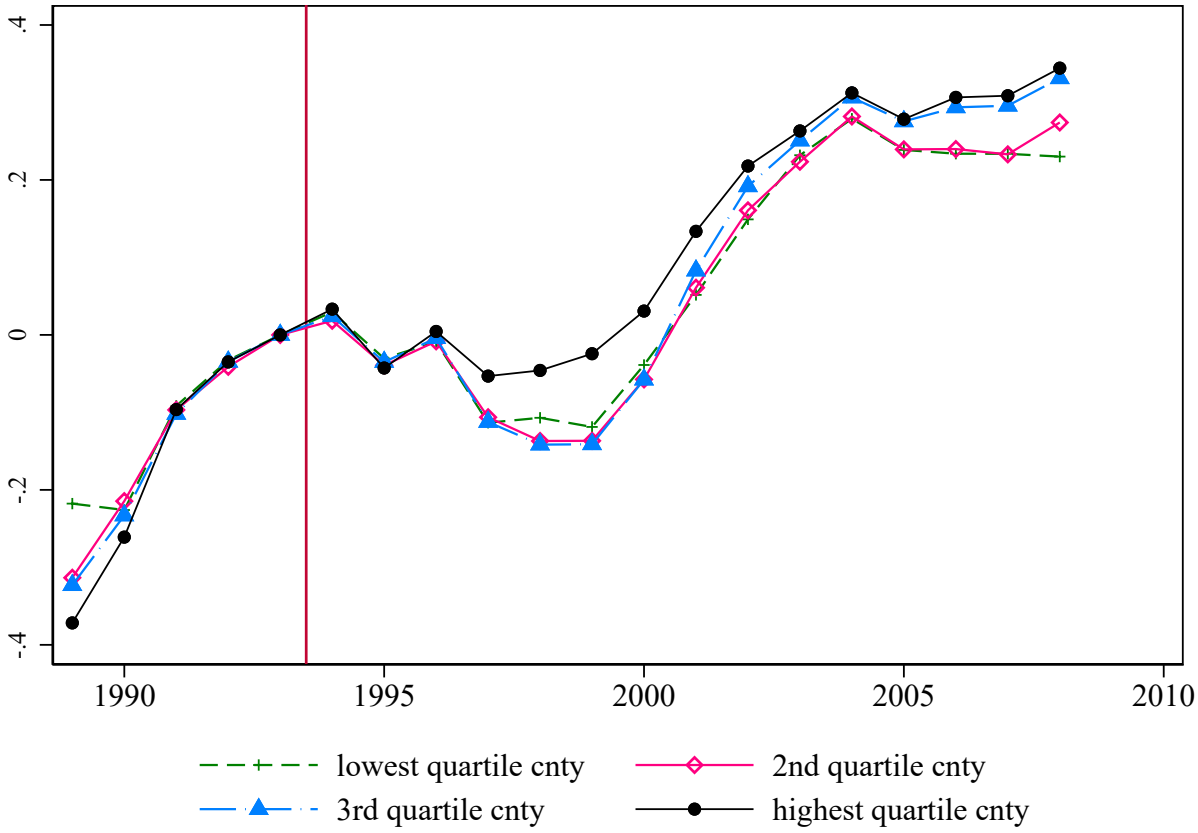
Figure 2A.23: Evolution of log employment as a function of county NAFTA vulnerability, for a balanced panel of 766 counties for which we have DI application data



Sources: The dependent variable is derived from the SSA. See Appendix B for more detail.

Notes: $N = 766$ counties. This figure is identical to Figure 2.2 but is restricted to the 766 counties (which account for around three-fourths of the U.S. population) for which we have DI data. The figure shows the event-study coefficient estimates (and 95% confidence intervals, based on standard errors clustered by state) from different specifications of equation (2.4), where log of total employment at the county×year level is the dependent variable. All specifications are weighted by 1990 county population. The first series shows the coefficient estimates from a specification where we control for only county and year fixed effects. The specification for the second series adds to this baseline specification *state*×*year* fixed effects. The third specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects. The fourth specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects. The final specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects.

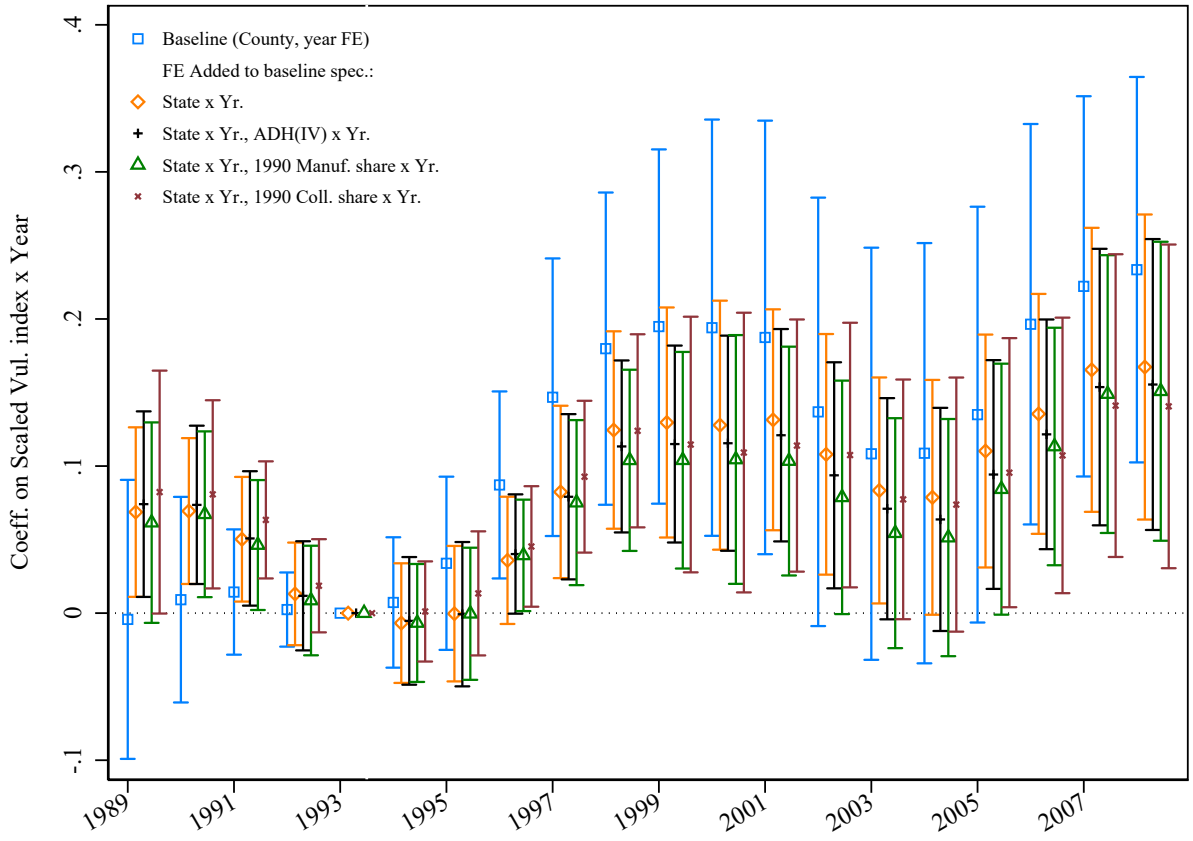
Figure 2A.24: Log DI applications, raw trends by four vulnerability quartiles (1993 normalized to zero)



Sources: The dependent variable is taken from the Social Security Administration (SSA). See Appendix B for more detail.

Notes: The figure shows the log of annual county DI applications by 1990 county vulnerability quartiles. Note that we can only perform this analysis for a subset of counties (see Section C).

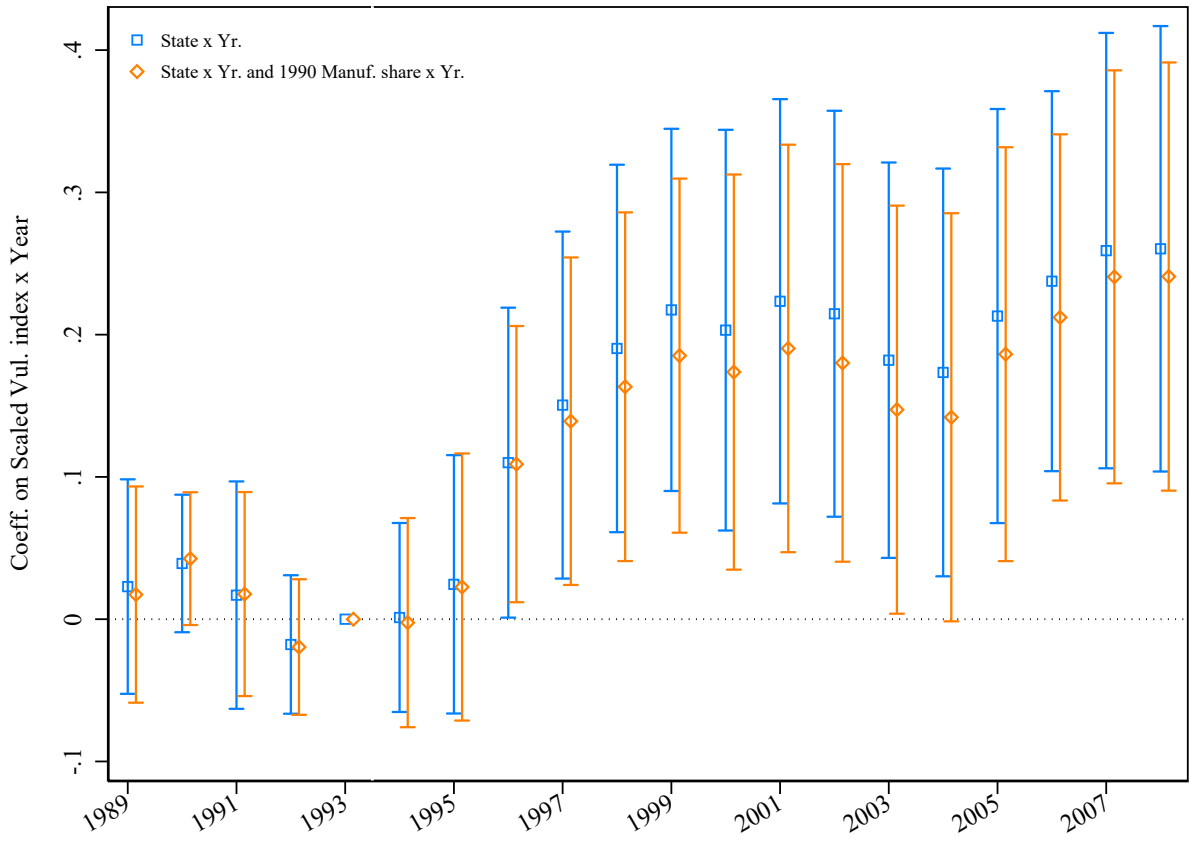
Figure 2A.25: Log of DI applications as a function of county NAFTA vulnerability



Sources: The dependent variable is taken from the Social Security Administration (SSA). See Appendix B for more detail.

Notes: This figure is analogous to Figure 2.5 but shows additional specifications after varying the controls. The first series includes only county and year fixed effects. The second series adds to this specification *state*×*year* fixed effects and the the third specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects (and are thus analogous to the first and second series of Figure 2A.6). The fourth specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects. The final specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects.

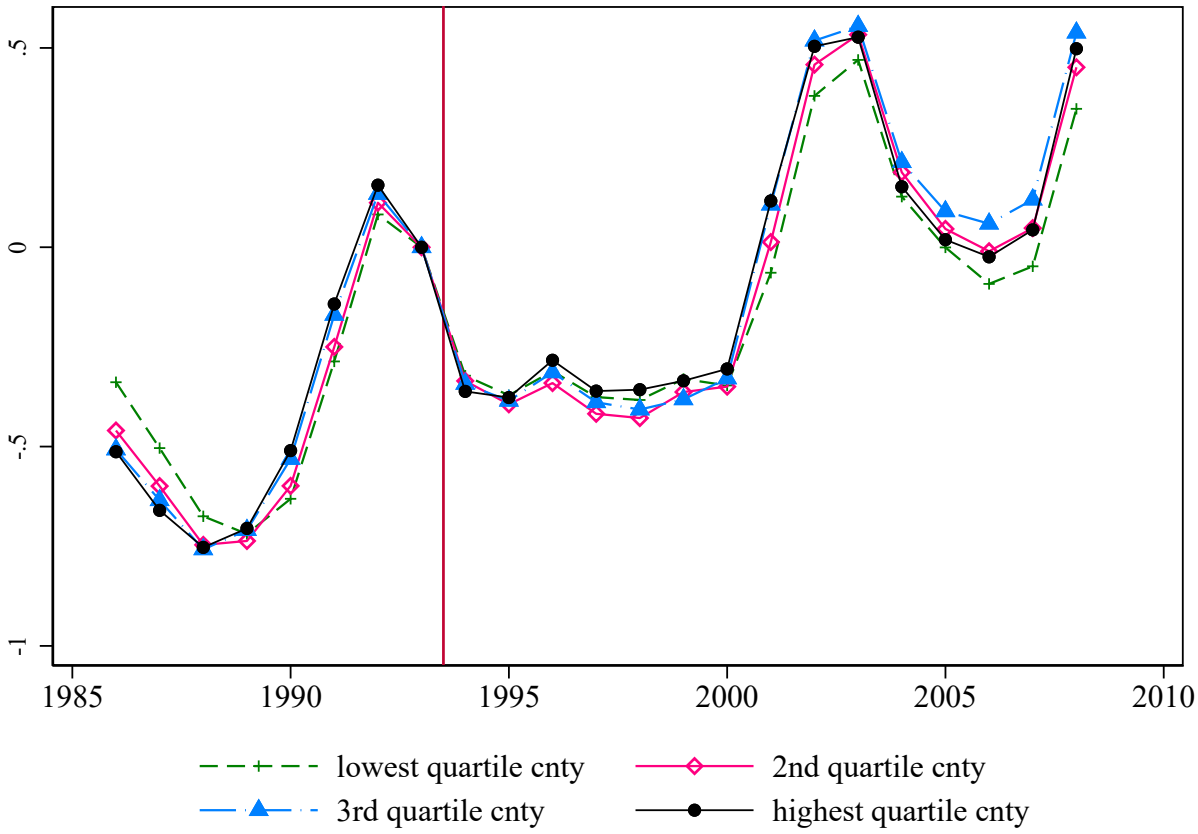
Figure 2A.26: Evolution of log DI final *awards* as a function of county vulnerability



Sources: The dependent variable is taken from the SSA. See Appendix B for more detail.

Notes: $N = 755$ counties. This figure is identical to Figure 2.5 except that the log of final awards instead of applications is the dependent variable. The figure shows the event-study coefficient estimates (and 95% confidence intervals, based on standard errors clustered by state) from different specifications of equation (2.4), where log of Disability Insurance (DI) final awards is the dependent variable. As discussed in Section C, we do not have all counties in this analysis, but the 755 counties we have in this balanced-panel analysis account for around three-fourths of the U.S. population. All specifications are weighted by 1990 county population. The first series shows the coefficient estimates from a specification where we control for only county and year fixed effects. The specification for the second series adds to this baseline specification *state*$\times$*year* fixed effects. The third specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects. The fourth specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects. The final specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects.

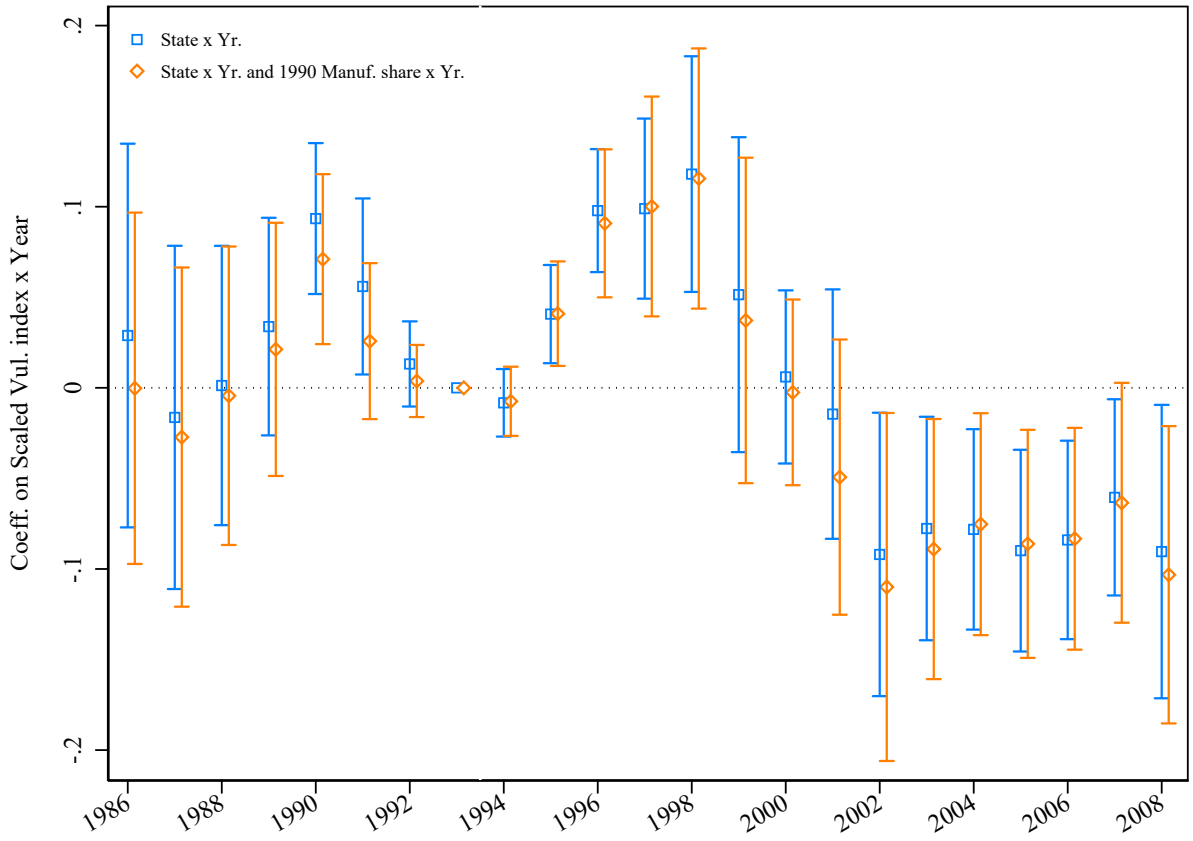
Figure 2A.27: Log of Unemployment Insurance benefits by county vulnerability, normalized



Sources: The U.S. Bureau of Economic Analysis (BEA) personal transfers data

Notes: The log of UI benefits is computed using the annual county-level personal transfers data from the U.S. BEA. The UI benefits in the series includes both state unemployment insurance compensation and other unemployment insurance payments, such as Trade Adjustment Assistance program.

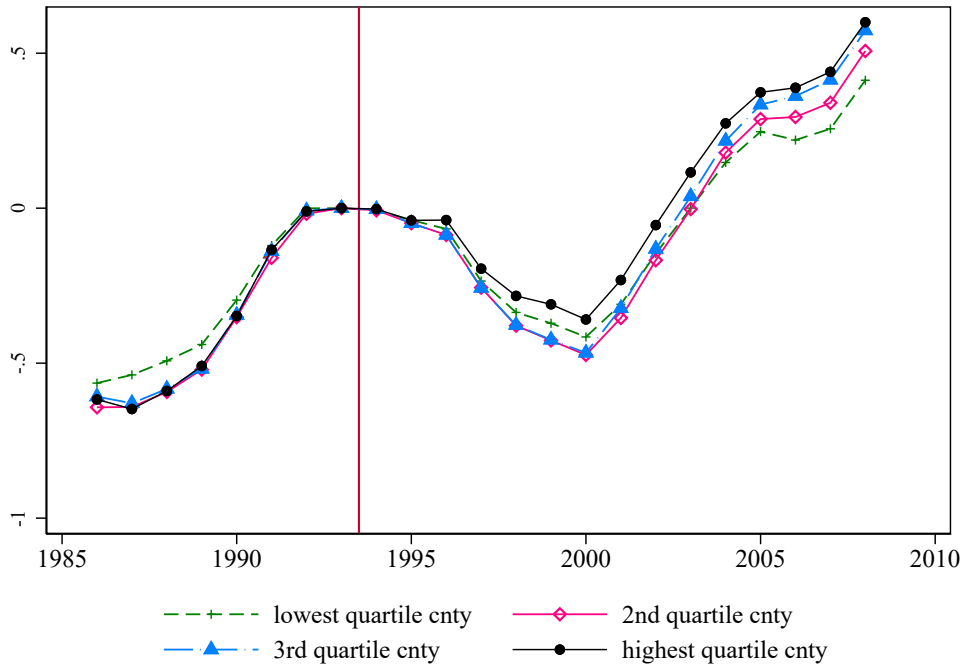
Figure 2A.28: Evolution of log UI benefits as a function of county vulnerability



Sources: The U.S. Bureau of Economic Analysis (BEA) personal transfers data

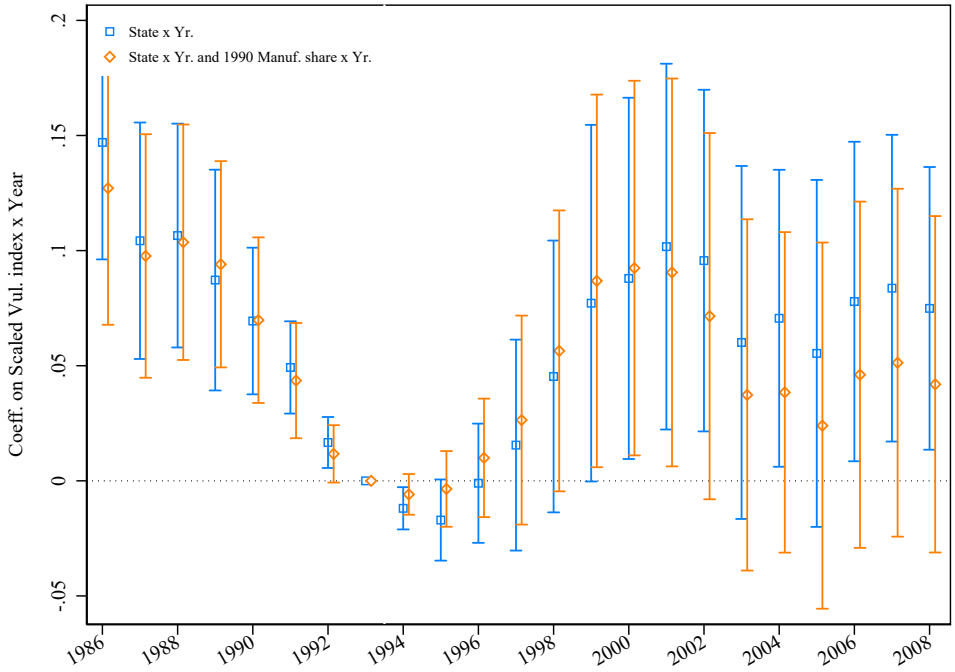
Notes: $N = 2957$ counties for each year of the sample period. The outcome variable is log of total UI benefits in each county. The log of UI benefits is computed using the annual county-level personal transfers data from the U.S. BEA. The UI benefits in the series include both state unemployment insurance compensation and other unemployment insurance payments, such as Trade Adjustment Assistance program. All specifications are weighted by 1990 county population. The first series shows the coefficient estimates from a specification where we control for only county and year fixed effects. The specification for the second series adds to this baseline specification $state \times year$ fixed effects. The third specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects. The fourth specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects. The final specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects. Standard errors clustered at the state level.

Figure 2A.29: Log reported SNAP benefits, normalized



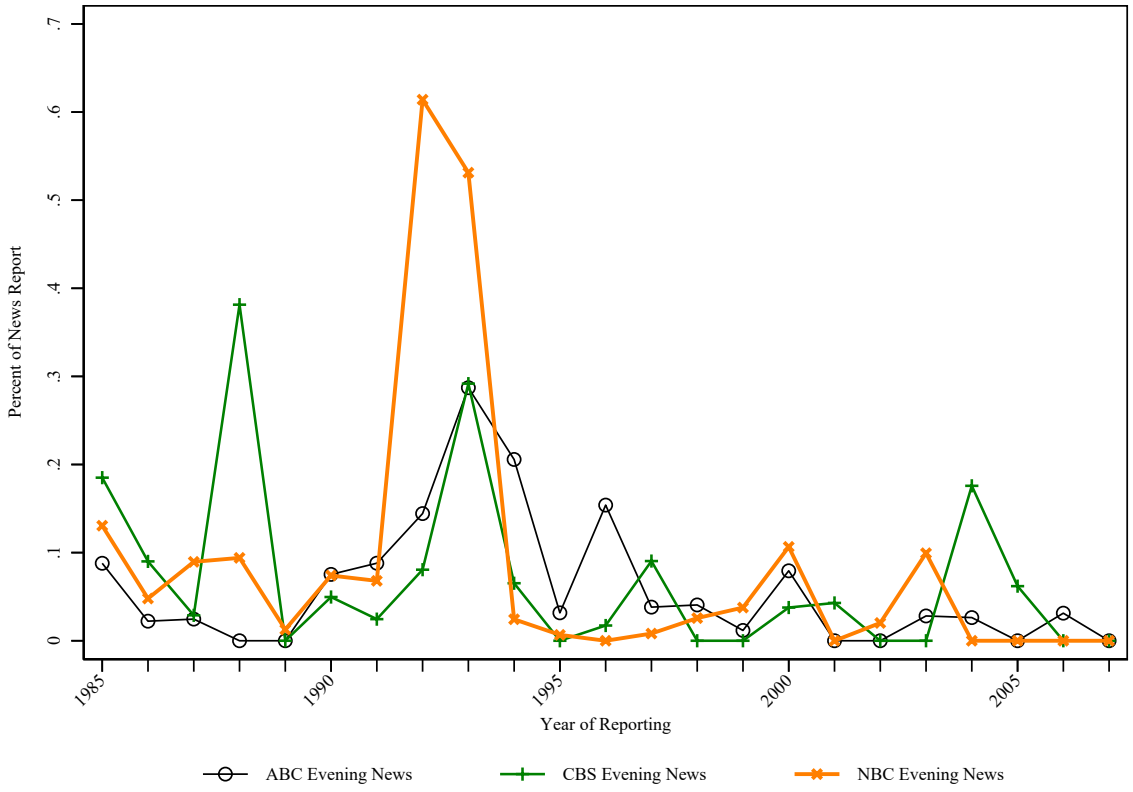
Notes: The log of reported SNAP benefits is computed using the annual county-level personal transfers data from the U.S. BEA. In the data, the SNAP benefits are estimated using the tabulations from the U.S. Department of Agriculture, payments data from state departments of social services, and the Census Bureau's Small Area Income and Poverty Estimates program.

Figure 2A.30: Log reported SNAP benefits as a function of county vulnerability



Notes: $N = 2933$ counties for each year of the sample period. The outcome variable is log of reported SNAP benefits. The log of reported SNAP benefits is computed using the annual county-level personal transfers data from the U.S. BEA. In the data, the SNAP benefits are estimated using the tabulations from the U.S. Department of Agriculture, payments data from state departments of social services, and the Census Bureau's Small Area Income and Poverty Estimates program. All specifications are weighted by 1990 county population. The first series shows the coefficient estimates from a specification where we control for only county and year fixed effects. The specification for the second series adds to this baseline specification *state*×*year* fixed effects. The third specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects. The fourth specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects. The final specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects. Standard errors clustered at the state level.

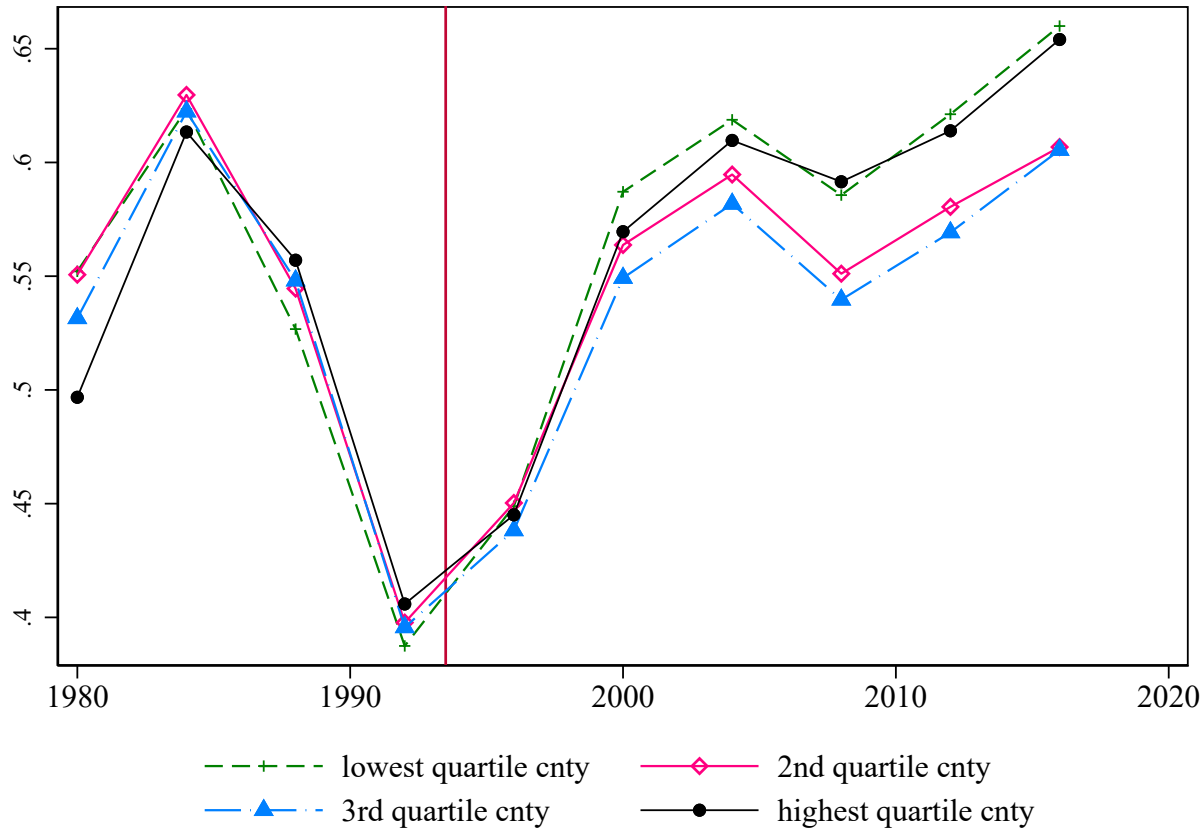
Figure 2A.31: Coverage of trade-and-jobs related stories by network nightly news programs



Sources: Data come from searching The Vanderbilt Television News Archive: `https://tvnews.vanderbilt.edu/search`. See Appendix B for more detail.

Notes: For each year and network, we calculate the share of minutes on the nightly news dedicated to stories that include variants (plurals, capitalizations) the following words: "trade" and "imports" and "jobs" *or* "employment." We exclude any stories (in all years) that include the phrase "trade center" so as not to pick up stories related to the attacks on the World Trade Center buildings.

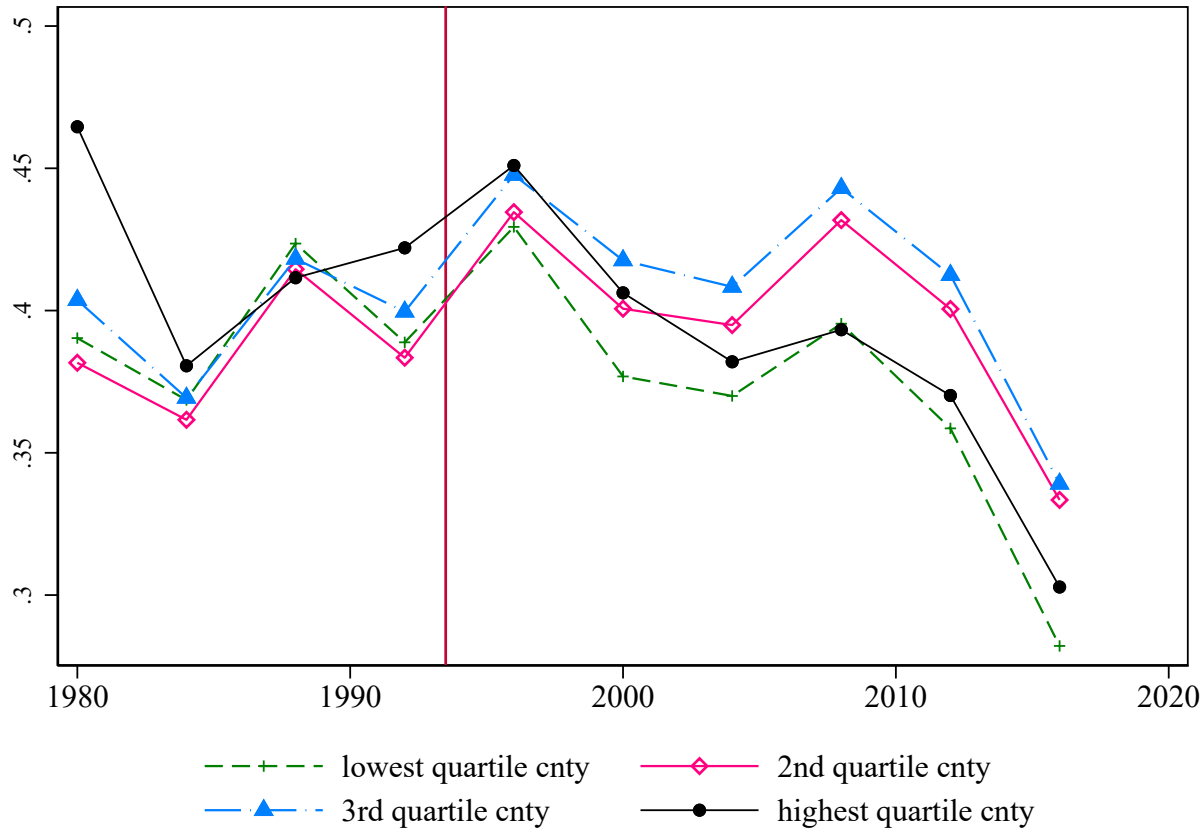
Figure 2A.32: Republican vote share in Presidential elections, separately by vulnerability quartile (raw means, not normalized)



Sources: The dependent variable is computed from ICPSR general election data for the United States (1980-1990) and David Leip's Atlas of U.S. elections (1992-2008).

Notes: The figure shows average two-party Republican House vote share trends from 1980 to 2016 by 1990 county vulnerability quartiles. The two-party vote share is computed using ICPSR general voting data and Dave Leip's Atlas of U.S. Elections data.

Figure 2A.33: Democratic vote share in Presidential elections, separately by vulnerability quartile (raw means, not normalized)

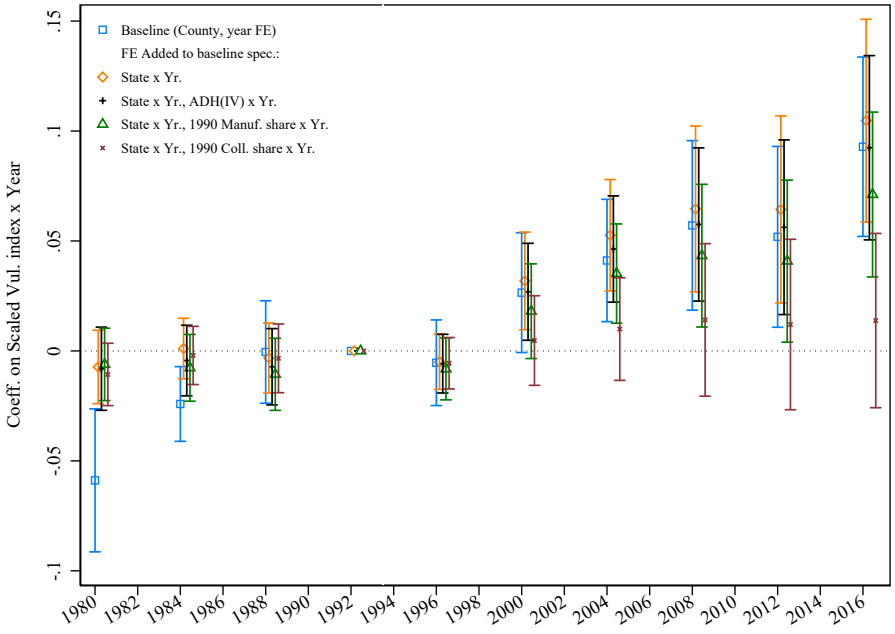


Sources: The dependent variable is computed from ICPSR general election data for the United States (1980-1990) and David Leip's Atlas of U.S. elections (1992-2008).

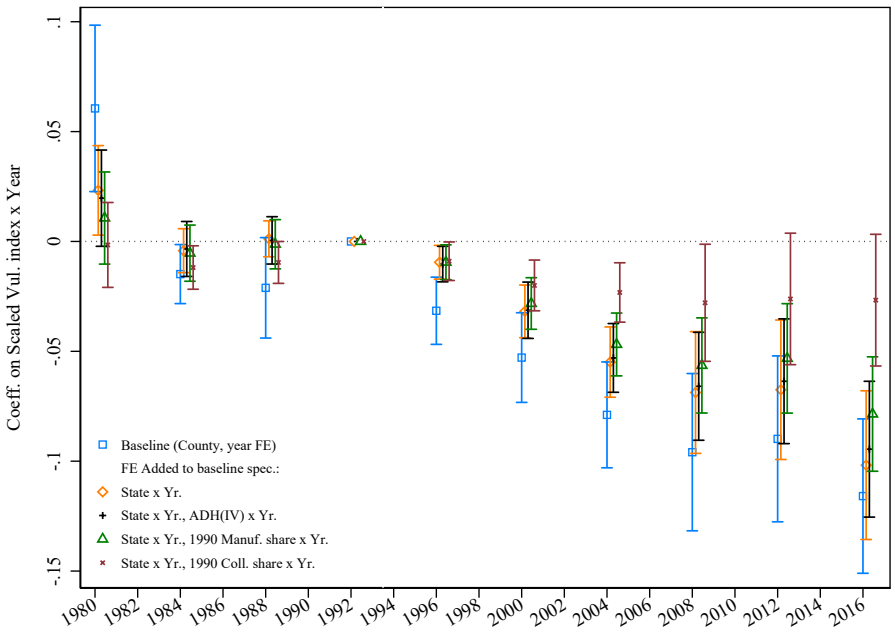
Notes: The figure shows average two-party Republican House vote share trends from 1980 to 2016 by 1990 county vulnerability quartiles. The two-party vote share is computed using ICPSR general voting data and Dave Leip's Atlas of U.S. Elections data.

Figure 2A.34: Presidential election vote shares as a function of county vulnerability, varying controls

(a) Republican share of county votes



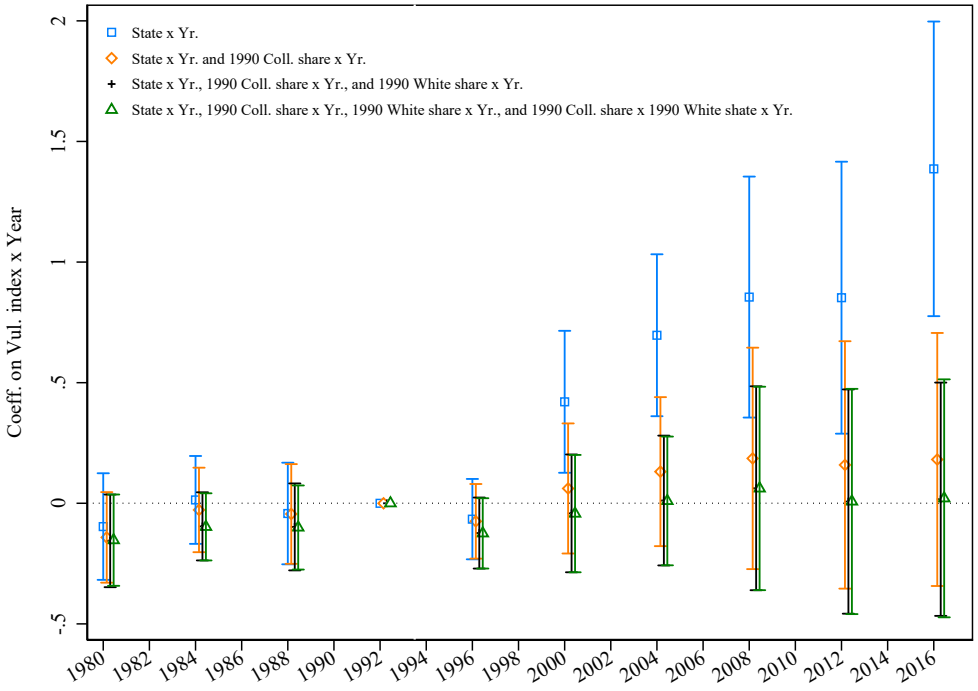
(b) Democratic share of county votes



Sources: The dependent variable is taken from ICPSR general voting data and Dave Leip's Atlas of U.S. Election data. *Notes*: This figure is analogous to Figure 2.6 but shows additional specifications after varying the controls. The first series includes only county and year fixed effects. The second series adds to this specification *state*×*year* fixed effects and the the third specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects (and are thus analogous to the first and second series of Figure 2A.6). The fourth specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects. The final specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects.

Figure 2A.35: Presidential election vote shares as a function of county vulnerability, additional white and college-share controls

(a) Republican share of county votes



(b) Democratic share of county votes

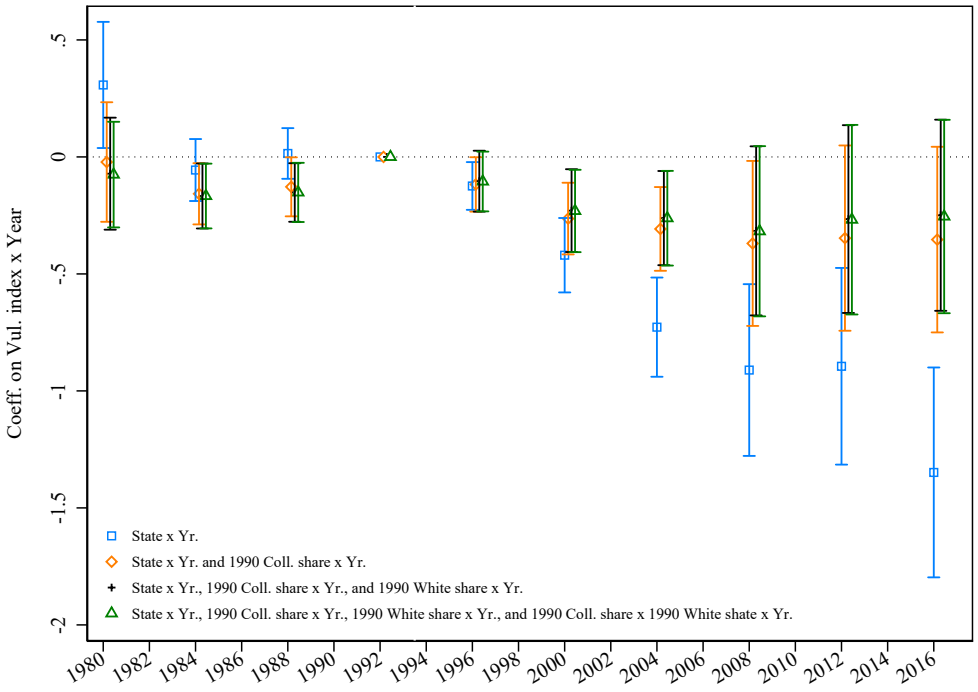
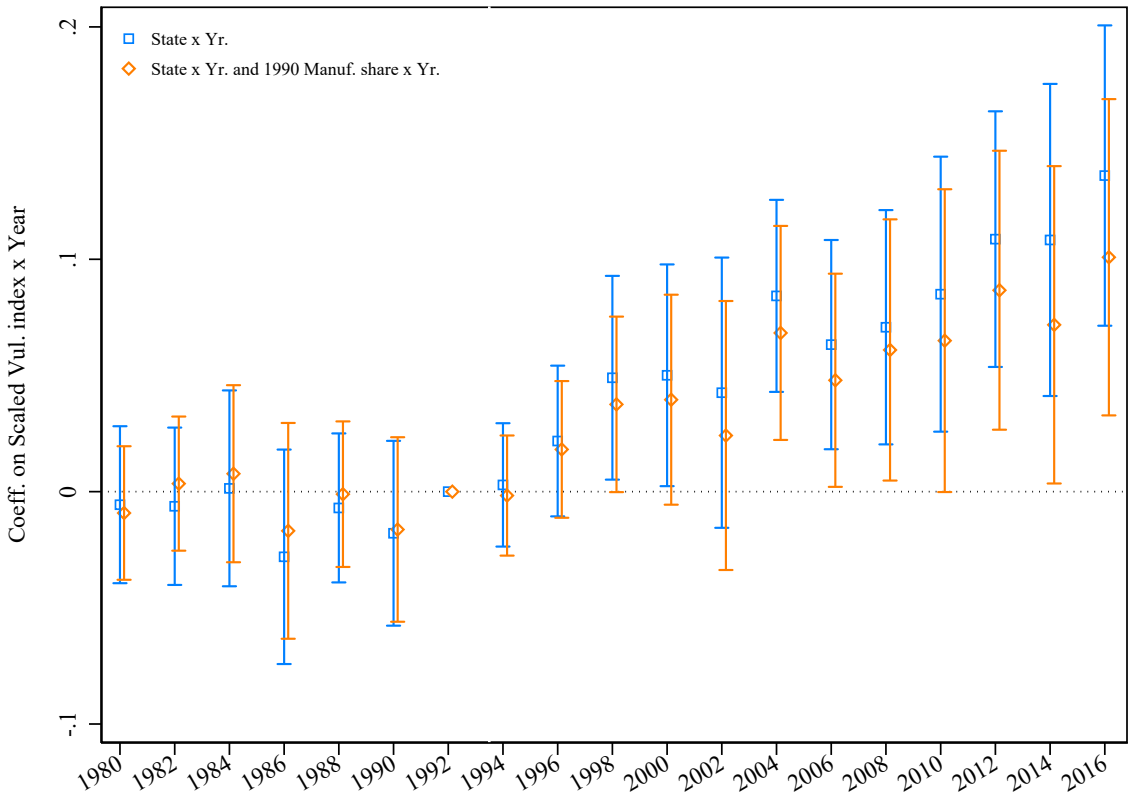


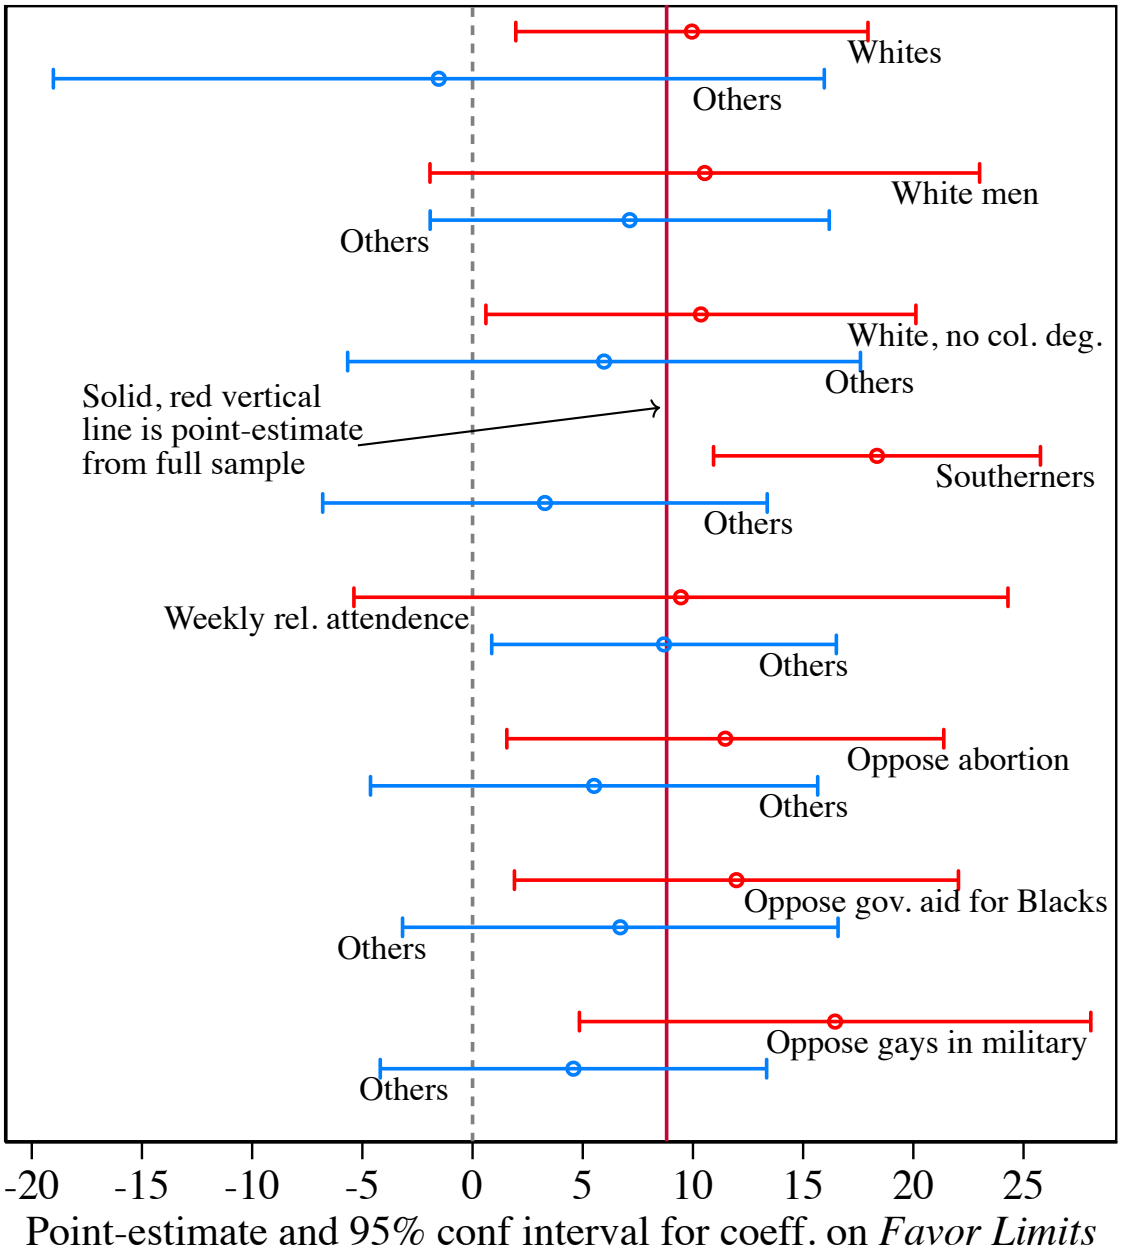
Figure 2A.36: Republican two-party share of House-election votes, by county NAFTA vulnerability



Sources: The dependent variable is computed from ICPSR general election data for the United States (1980-1990) and David Leip's Atlas of U.S. elections (1992-2008). Note that "Republican two-party share" is defined as $\frac{Repub.\,votes}{Repub.\,votes+Dem.\,votes}$ for each county-year. See Appendix B for more detail.

Notes: The analysis sample is fixed across specifications and strictly balanced, with 2461 counties in each year of the sample. The figure shows the event-study coefficient estimates (and 95% confidence intervals, based on standard errors clustered by state) from different specifications of equation (2.4), where the two-party Republican vote share in House elections is the dependent variable. It follows Figure 2.2 though obviously cannot be analyzed annually because elections fall only on even years. All specifications are weighted by 1990 county population. The first series shows the coefficient estimates from a specification where we control for only county and year fixed effects. The specification for the second series adds to this baseline specification $state \times year$ fixed effects. The third specification adds to the second specification CZ-level measure of Chinese import exposure from Autor et al. (2013) interacted with year fixed effects. The fourth specification adds to the second specification 1990 county-level manufacturing share of employment interacted with year fixed effects. The final specification adds to the second specification 1990 county-level college-graduate population share interacted with year fixed effects.

Figure 2A.37: Heterogeneity in 1992-1994 shift toward GOP among protectionist respondents



Sources: ANES panel data, 1992-1994.

Notes: This figure estimates, for mutually exhaustive and distinct subgroups of the sample, equation (2.6) from the text: $Moved\,Right_{i,94\text{-}92} = \beta Favor\,Import\,Limits_{i,92} + \gamma \mathbb{X}_{i,92} + e_i$. It uses the same control vector $\mathbb{X}_{i,92}$ as in col. (4) of Table 2.4, namely demographic and political-issue controls. We report the coefficient and 95% confidence interval (clustered by state) of our estimate of β.

Table 2A.1: Education predicts less protectionist views

	Dept. var: Favor more limits on trade					
	(1)	(2)	(3)	(4)	(5)	(6)
Has BA degree	-0.0823***	-0.0821***	-0.0969***	-0.0973***	-0.0938***	-0.213***
	[0.00970]	[0.0126]	[0.00953]	[0.00894]	[0.00801]	[0.0106]
Some college, no BA degree		0.000420				
		[0.0131]				
Dept v mean	0.382	0.382	0.382	0.382	0.382	0.643
Demog controls	No	No	Yes	Yes	Yes	Yes
Issue controls	No	No	No	Yes	Yes	Yes
State FE	No	No	No	No	Yes	No
Drop DKs	No	No	No	No	No	Yes
Observations	18836	18836	18743	18743	18743	11120

Sources: ANES individual time series files, 1986–2012.

Notes: The dependent variable is a dummy coded as one for respondents who report favoring import limits (and zero otherwise, including no opinion). "Demographic controls" include indicators for white, and male; log of family income and age. "Issue controls" include views toward African-Americans, trust in government, and views toward abortion. $^{*}p = 0.1,^{**}p = 0.05,^{***}p = 0.01$.

Table 2A.2: How protectionist views predict approval of Ross Perot, 1992 and 1996

	Dept. variable: Approves of Perot					
	(1)	(2)	(3)	(4)	(5)	(6)
Favor import limits	0.0628***	0.0586***	0.0597***	0.0602***	0.0574*	0.0555***
	[0.0200]	[0.0191]	[0.0193]	[0.0200]	[0.0287]	[0.0196]
Mean, dept var	0.365	0.360	0.360	0.360	0.456	0.270
Demog. controls?	No	Yes	Yes	Yes	Yes	Yes
Issue controls?	No	No	Yes	Yes	Yes	Yes
State FE?	No	No	No	Yes	No	No
Sample criteria?	None	None	None	None	1992 only	1996 only
Observations	2990	2940	2940	2940	1422	1518

Sources: ANES, 1992 and 1996.

Notes: The dependent variable is a dummy coded as one for respondents who answer "Yes" to the following question: "Is there anything about Mr. Perot that might make you want to vote for him?" (1992) or "Is there anything in particular about MR. PEROT that might make you want to vote FOR him?" (1996). All columns except the final two include year fixed effects. "Demographic controls" include indicators for white, male, and college completion; fixed effects for age rounded to the nearest ten, and log of family income and age. "Issue controls" include views toward African-Americans, trust in government, and views toward abortion. $^{*}p = 0.1,^{**}p = 0.05,^{***}p = 0.01$.

Table 2A.3: Simulating presidential elections ex. NAFTA effect

	2000		2004		2008		2012	
	Flipped	Closer	Flipped	Closer	Flipped	Closer	Flipped	Closer
	FL	AR	IA	CO	GA	MO	NC	GA
	NH	OH	NM	FL		SC		
		TN	OH					
Electoral votes	29	38	32	36	15	19	15	16
Pivotal?	Yes		Yes		No		No	

Sources: Election data from David Leip.
Notes: The counterfactual no-NAFTA Democratic vote count in election year t if calculated as: $Dem\,vote_c^t - \beta_t^D \cdot Vulnerability_c$, where β_t^D is the event-study coefficient estimating the effect of NAFTA vulnerability on Democratic vote share in election year t, using the baseline $state \times year$ specification. We perform a parallel calculation for the Republican vote share. We then aggregate up to the state-election level to determine the counterfactual state winners and electoral-vote allocations. "Flipped" means that the no-NAFTA counterfactual would have *changed* the state's winner in that election year (in all cases, the Democrat would have won instead of the Republican). "Closer" means that under the no-NAFTA counterfactual, the Democrat was within 2.5 percentage points (but would not have won) whereas in reality he lost the state by more that 2.5 percentage points. "Pivotal" indicates that the "flipped" states would have changed the outcome of the entire presidential election.

Table 2A.4: Partisan identity and views toward NAFTA, 1992-1994 panel data

	Move in Repub direction dummy x 100					
	(1)	(2)	(3)	(4)	(5)	(6)
Oppose NAFTA (asked in 1993)	7.777	11.09*	6.428	5.833	6.573	6.506
	[5.095]	[5.853]	[4.617]	[4.683]	[4.784]	[5.068]
Minorities sd help self				2.123*	2.301**	2.127**
				[1.049]	[1.025]	[0.999]
Wants active gov't				-0.282	-0.329	-1.335
				[1.337]	[1.539]	[1.758]
Support abortion				-1.164	-0.506	-0.569
				[1.935]	[2.248]	[2.460]
Attend church weekly				4.526	5.066	2.287
				[3.317]	[3.411]	[3.544]
Favors increased immigr.				-1.766	-5.217	-8.023
				[5.801]	[6.641]	[7.107]
Oppose gays in military					3.661	4.441
					[8.024]	[9.135]
Oppose gov't health care					-0.992	-1.507
					[0.961]	[1.042]
Favor term limits					-5.380	-4.651
					[4.177]	[4.872]
Dept. var. mean	25.93	25.69	25.89	25.77	25.77	25.77
Ex. DK	No	Yes	No	No	No	No
Demog. covars	No	No	Yes	Yes	Yes	Yes
State FE	No	No	No	No	No	Yes
R-squared	0.00489	0.0155	0.0489	0.0685	0.0737	0.123
Observations	621	288	618	617	613	613

Sources: ANES panel data, 1992-1994.

Notes: The dependent variable is a dummy (multiplied by 100) for whether the respondent moved in the GOP direction in the 1-7 partisan identity scale. All explanatory variables were asked in 1992, except for the NAFTA question, which was asked in the fall of 1993. "Excl. DK" means that respondents who did not have an opinion on NAFTA are dropped (they are otherwise coded as zero). Demographic controls include race, gender, education, age, log family income, and urbanicity. Standard errors clustered by state. $^{*}p = 0.1,^{**}p = 0.05,^{***}p = 0.01$.

B Data appendix

Data used to construct the vulnerability measure

Our county-level vulnerability measure is constructed using three components defined prior to NAFTA's implementation (our base year is 1990): (i) average tariff on imports from Mexico by industry, (ii) Mexico's revealed comparative advantage (RCA) by industry, and (iii) industrial composition of each county.

The average tariff on imports from Mexico is drawn from the U.S. Tariff database created by Feenstra et al. (2002).[40] The dataset contains ad-valorem, specific and estimated ad-valorem equivalent (AVE) tariff rates for Most-Favored-Nations (MFNs), Canada, and Mexico by eight-digit Harmonized Tariff Schedule (HTS) industries. Whenever the Mexico-specific tariff rates are not defined for industries, we apply the MFN tariff rates.

We compute Mexico's revealed comparative advantage using the UN Comtrade bilateral export series, available from Hakobyan and McLaren (2016b)'s replication directory. When aggregating eight-digit industries into six-digit industries and computing the weighted average of tariffs, we use USITC import values for each eight-digit industries as the weights.

The UN Comtrade bilateral export series contains the dollar value of exports by origin and destination in six-digit Harmonized Tariff Schedule code. We utilize the County Business Patterns data to compute the industry composition of each county, which is further described in Appendix B.

County Business Patterns

County Business Patterns (CBP) provide county-level economic data by industry, including the number of establishments, employment, and annual payroll. The dataset is based on the week of March 12th every year. We use CBP from 1986 to 2008 to compute county-level

[40] For creating the vulnerability for years other than 1990-2001, we use the USITC annual tariff data. This is not used in the analysis, but to illustrate the change of vulnerability measure over time.

annual employment size and industry employment share in year 1990. [41] County-level industry employment shares in 1990 are computed for each four-digit SIC industry. When we are combining average tariffs of six-digit HTS industries with county-level employment shares of four-digit SIC industries, we use crosswalk from David Dorn's data webpage.[42]

Data on county-level annual population

We utilize county-level annual intercensal population estimates from the Census Bureau's Population Estimates Program (PEP). The PEP calculates population estimates using the most recent census and data on births, deaths, and migration. The county-level estimates are broken down by race, sex, age, and educational attainment. The PEP estimates are obtained from the NBER website (https://data.nber.org/data/census-intercensal-population/) and the U.S. Department of Agriculture Economic Research Service website (https://www.ers.usda.gov/data-products/county-level-data-sets/download-data/).

We use the annual intercensal population estimates as the outcome variable of our population response event-study analysis. We derive county-level pre-NAFTA demographic controls from the PEP, such as share of college-educated population and working-age population defined by population estimates of age 15-64.

[41] We use the published version of CBP county-industry-level employment counts when constructing our vulnerability measure. Eckert et al. (2021) point out that some of these county-industry-level employment cells are missing or imputed when the cell sizes are too small, due to the CBP's disclosure rules. We try to construct the vulnerability measure using a version of the dataset from Eckert et al. (2021), where the missing county-industry cells are imputed, and we end up with a very similar vulnerability measure and analysis results.

[42] The four-digit SIC industries in the crosswalks are the "slightly aggregated" version as used in ADH, but we decide not to add up the adjacent four-digit industries in the county-level industrial composition calculation as they did in the paper. When we use the same aggregation to the industrial composition, we still get a very similar vulnerability measure.

PSID data used for individual-level NAFTA vulnerability analysis

We obtain individual characteristics from the Panel Survey of Income Dynamics (PSID)[43] for 1988 to 2003. The sample consists of 68,814 individuals with information on their demographics (i.e., sex, age, and race), education level, labor status, state of residency, and employment industry code. We follow a sequence of restrictions in order to keep a balanced panel for the relevant variables. First, we use the employment industry code of the individuals in 1990 to match the current NAFTA Vulnerability dataset we have constructed. The 1990 industry code available in the PSID is the 3-digit industry code from the 1970 Census of Population, and because our NAFTA vulnerability dataset includes a crosswalk between the 1970 and 1990 industry codes, we can keep working with the 1990 vulnerability measures seamlessly. Second, when individuals didn't have an industry code associated, we assigned them the value of zero as their vulnerability measure. Finally, the dependent variable is non-missing any year of analysis to guarantee a balanced panel.

Data on county-level annual Trade Adjustment Assistance petition and certifications

We acquire the universe of TAA petition data from 1975 to 2020 from the U.S. Department of Labor. For each petition, the dataset contains information on the name, address, zipcode and industry code of the firm, the product or service that the worker group is engaged with, and the date the investigation starts.[44] We calculate the number of workers included in certified (approved) petitions in a county from 1975 to 2020, based on petitions' institution date.[45] For counties with no petitions filled at a given year, we assign a zero number of affected workers.

[43]The data set is available at https://simba.isr.umich.edu/data/data.aspx

[44]These data also include the date of the petition, which would appear to be a better variable to use to "date" each observation, but it only begins in 1994. However, the gap between petition and investigation is less than a month in the post-1994 data.

[45]We assign all the petition cases to three categories: certification, denial and termination. Termination is not an actual decision but an administrative closing of the case due to petition withdrawal or because the

Data on county-level SNAP benefits

We acquire annual county-level SNAP benefits data from the U.S. Bureau of Economic Analysis (BEA) personal transfers data, available on the BEA website[46]. The series is available from 1969-2020, and we use the data from 1986 to 2008. In the data, the SNAP benefits are estimated using the tabulations from the U.S. Department of Agriculture, payments data from state departments of social services, and the Census Bureau's Small Area Income and Poverty Estimates program.

Data on county-level UI benefits

We acquire annual county-level UI benefits data from the U.S. Bureau of Economic Analysis (BEA) personal transfers data, available on the BEA website[47]. The series is available from 1969-2020, and we use the data from 1986 to 2008. The UI benefits data include both state unemployment insurance compensation and other unemployment insurance payments, such as Trade Adjustment Assistance program. In the data, the county-level benefits are estimated using data from the U.S. Department of Labor, Employment and Training Administration, payments data from the state employment security agencies, and the Local Area Unemployment Statistics program of the Bureau of Labor Statistics (BLS).

Data on county-level annual Disability Insurance approvals

We acquire summary statistics of Social Security Disability Insurance (SSDI) and Supplemental Security Income (SSI) applications and approvals by county and year from the Social Security Administration (SSA). The application-level SSDI and SSI data contain

case is covered by another petition. We therefore only look at the cases that are either certificated or are denied.

[46]The series is available for download at https://apps.bea.gov/regional/downloadzip.cfm

[47]The series is available for download at https://apps.bea.gov/regional/downloadzip.cfm

the application and decision dates, the district office where the application was received, and the zipcode of the district office. The application is collapsed by zipcode and further by county using the 1990 geographic correspondence engine from Missouri Census Data Center.

The zipcode information for each application is not completely populated until mid-1990s, so we focus on the applications with non-missing district office codes where we can locate the zipcode of the district office. We utilize the 2009 and 2019 SSA district office list to create a set of district offices which existed both in 2009 and 2019. We keep the applications from these 1180 district offices and recover the zipcode of each application submitted to the offices. We end up with 755 counties as our "matched district office" sample, which accounts for around 75 percent of U.S. population in 1990. While the method allows us to keep a consistent set of district offices to be included, district offices that opened or closed between 1990 and 2009 are not included in this set.

Data on county-level annual mortality

We construct the number of working-age (age 15-64) deaths by county and year using the National Center for Health Statistics (NCHS) Vital Stats data. The individual-level mortality data contain the cause of death, age, sex, geographic location of each death (e.g., county and state). From 1979 to 1988, the vital stats data are publicly available. For 1989 and onward, the restricted version is acquired from the NCHS as the geographic identifiers are not publicly available.

Data on news coverage

We measure how much NAFTA was covered in the news by computing the share of total news minutes (including commercials, intro, and outro of the news programs) allocated to news with keywords trade, import, and jobs, excluding "Trade Center" in CBS, ABC, and

NBC evening news from 1985 to 2010. We acquire the text data by web-scraping the Vanderbilt's TV news archive (https://tvnews.vanderbilt.edu/). We did not include news sources such as Fox News that were created during our main analysis years.

Data on county-level House election votes

We obtain county-level voter turnout and election votes from ICPSR general election data for the United States (1980-1990) and David Leip's Atlas of U.S. presidential elections (1992-2008). We compute the two-party republican vote share by the share of Republican votes among the votes received by Republican and Democratic candidates, and the republican vote share among the total votes casted in each county.

Survey data on NAFTA favorability

We obtained all publicly available survey data in ICPSR and iPoll that (a) asked a generic sentiment question on NAFTA; (b) contained state identifiers; (c) took place before 2016 (so as not to be affected by the anti-NAFTA presidential campaign of Donald Trump). The datasets always contained information on basic demographics and often union status and family income.

ANES repeated cross-sectional data

We use the *individual* files for each year, *not* the cumulative file that ANES creates for convenience. The individual files have questions that are not included in the cumulative file.

We use every year of data from 1986 to 2012 that includes the *Favor Import Limits* question. In 1990, this question is asked in a different format (a seven-category likert scale instead of a binary yes/no question) so we do not include that year.

Table 2B.1: Datasets used in Table 2.2 (NAFTA approval by state-level vulnerability)

Organization conducting the survey	Date	Sample size
ANES	1993	742
CBS	Oct 1996	1528
Pew	Sep 1997	2000
CNN/Gallup	Aug 1997	481
Pew	Sep 2001	1000
Pew	Dec 2003	553
Pew	Jul 2004	1003
CNN/Gallup	Jan 2004	455
Pew	Mar 2004	1703
Pew	Dec 2004	2000
Newsweek	Feb 2004	1019
Program on International Policy Attitudes	Jun 2005	812
Pew	Oct 2005	1003
Pew	Dec 2006	1502
Pew	Apr 2008	1502
Pew	Mar 2009	2031
Pew	Oct 2009	2000
Monmouth	Oct 2015	1012
CBS/NYT	May 2015	1022

ANES panel data

In 1992 that ANES fielded a small panel data set that followed a subset of the 1992 repeated cross-sectional dataset. We use the provided weights to adjust for attrition.

C Assessing main results with randomization inference

Recent literature has shown that in shift-share regression design, the standard errors could be underestimated even after being clustered by geography (e.g., state) or estimated under heteroskedastic assumptions, since the regression residuals could be correlated across local labor markets with similar industrial compositions (Adão, Kolesár, and Morales 2019; AKM from here on). Our vulnerability measure is an average of industry-level tariff shocks weighted by county industrial composition, so similar downward bias may occur when computing standard errors in our event-study analysis.

In order to check the robustness of our employment and House election results to the above issue, we perform two exercises: (1) we first conduct a placebo exercise where we simulate county-level vulnerability measures by randomly assigning tariffs to each industry drawn from various distributions; and (2) we also apply the AKM inference procedures as in AKM (2019) and compute the standard errors that are robust to cross-sectional correlation across four-digit SIC industries and with three-digit SIC industry clusters.

As the inference discussion in the aforementioned literature is based on a standard regression approach, we modify our main analysis to be the first-difference regression specification below :

$$Y_{c,t} - Y_{c,t-7} = \beta \text{Vulnerability}c^{1990} + \gamma Xc + e_c \tag{2.7}$$

where $Y_{c,t} - Y_{c,t-7}$ is the first-differenced outcome variable such as change in the log of employment and change in two-party republican vote share, $\text{Vulnerability}c^{1990}$ is the vulnerability of county c in 1990, X_c is county-level characteristics including share of college grads in 1990 or a set of state dummy variables (thus state fixed effects).

For the placebo exercise, we use simulated vulnerabilities instead of the actual $\text{Vulnerability}_c^{1990}$, using the simulation procedure described in the following paragraph. For

the AKM inference exercise, we use the actual $\text{Vulnerability}_c^{1990}$ as the main independent variable and report the AKM standard errors along with robust and state-clustered standard errors. We run equation (2.7) separately for the pre-NAFTA and post-NAFTA periods. We define 1986-1993 as the pre-NAFTA period and 1993-2000 as the post-period in our employment analysis, and we use 1984-1992 and 1992-2000 for the pre-NAFTA and post-NAFTA periods in the House election analysis, respectively, so that we are always comparing a midterm election to another midterm election.

For the placebo exercise, we draw simulated tariffs from three distributions: (i) one that follows the empirical distribution of actual industrial tariffs in 1990; (ii) uniform distribution with range [0, 0.4] applied to industries with positive tariffs in 1990; and (iii) uniform distribution with range [0, 0.4] applied to all industries.[48] For (i), we generate the empirical cumulative distribution of all positive tariffs in 1990 by fitting a fifth-degree polynomial, as shown in Appendix figure 2C.1. We take a random draw from U[0, 1] for each sector with a positive tariff and use our modeled CDF to generate a simulated tariff for the sector. We then use this vector of simulated tariffs to generate a simulated $\text{Vulnerability}_c^{1990}$, using the county employment composition in 1990 from the CBP. We construct simulated vulnerability measures from the uniform distribution U[0, 0.4] for (ii) and (iii) in an analogous way.

For each distribution, we repeat the simulation of county vulnerability and estimation of the equation (2.7) 1000 times and create 1000 placebo analysis samples. The outcome variables, change in the log of county employment and change in Republican Presidential vote share, are constructed using the observed data in each county and are identical for all placebo samples. We then compute the estimate of β in equation (2.7) for each sample and report the mean and standard deviation of the estimates across the placebo samples. In

[48] We set the range of uniform distribution 0.4 based on the rough maximum of the actual 1990 tariffs.

order to appropriately apply state fixed effects in our simulated samples, we resample observations in each sample with replacement, using states as the clustering unit.

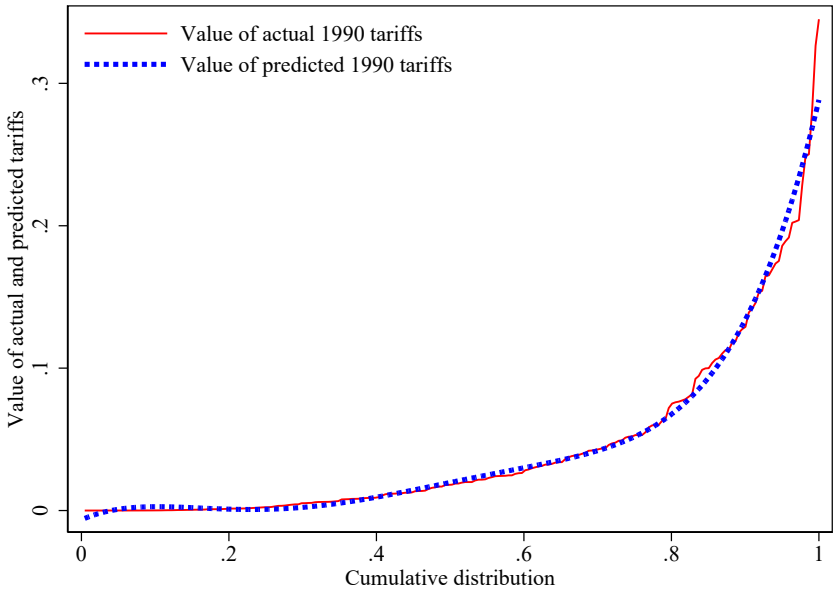
The average and standard deviation of the estimates are reported in Appendix tables 2C.1, 2C.2, 2C.3, and 2C.4 for the employment analysis and Appendix tables 2C.5, 2C.6, 2C.7, and 2C.8 for the House election analysis. The pre-NAFTA employment analysis shows that both original and simulated vulnerabilities have no significant impact on log of county employment in all specifications. The post-period analysis indicates that the original vulnerability measure is associated with the significant decline in log of county employment while simulated vulnerabilities are not significantly associated with the change in county employment. Note that the coefficients on the simulated vulnerabilities using the actual tariff distribution and U[0,0.4] are positive while insignificant, as industries with zero tariffs continue to have zero tariffs in the simulation, the simulated vulnerabilities retain some information on the true industry tariffs and employment shares (e.g., service industries with zero tariffs will continue to have zero tariffs, and counties with high share of service industries will continue to have low simulated vulnerability). When the tariffs are simulated for all industries from U[0, 0.4], the coefficients become more precisely zero.

Similarly, the pre-NAFTA election analysis indicates that both original and simulated vulnerabilities are not associated with the change in county two-party Republican vote share. In post-period analysis, only the original vulnerability measure is associated with a significant increase in the Republican vote share.

In Appendix tables 2C.9 and 2C.10, we report the estimates of equation (2.7) with robust, state-clustered, and the AKM standard errors for the employment analysis. The estimates of Equation (2.7) for the election analysis are reported in Appendix tables 2C.11 and 2C.12. For both employment and election analysis in the pre-NAFTA period, not including state fixed effects overstates the effect of tariff protections on the county employment and Republican vote shares. When implementing state fixed effects, we find

that the pre-period coefficients are not significant, and the post-period effects of vulnerability are significant and robust to allowing for correlation across industry composition of local labor markets, as the coefficients are still significant under the AKM inference exercise.

Figure 2C.1: Empirical distribution of actual industrial tariffs in 1990



Notes: The figure plots the empirical cumulative distribution of positive tariffs of four-digit SIC industries in 1990. The dashed line is a predicted cumulative distribution of positive tariffs, generated by fitting a five-degree polynomial to the empirical CDF.

Table 2C.1: Estimates from the first-difference model: Change in log of employment (1986-1993) and simulated county vulnerabilities

Vulnerability :	Actual	Simulated		
Distribution :	N/A	poly. approx.	uniform [0,0.4]	all industry uniform
Coefficient	-.008	.013	.006	.004
(SE)	(.155)	(.219)	(.08)	(.124)
state FE	yes	yes	yes	yes

Notes: Number of simulation = 1000 times. The table reports the estimates from the first-difference model shown in equation (2.7). The dependent variable is change in log of county employment between 1986 and 1993. The first column uses the original county vulnerability. In column 2, the simulated tariffs for industries with positive tariffs are drawn from the estimated tariff distribution plotted in Appendix figure 2C.1. In column 3, the simulated tariffs for industries with positive tariffs are drawn from U[0, 0.4]. In column 4, the tariffs for all industries are drawn from U[0, 0.4]. The industry-level tariffs are simulated 1000 times and the corresponding county-level vulnerability measure is used as the main regressor for the first-difference analysis. We compute the average and standard deviation of the estimates and report them as the bootstrapped coefficient and standard error. We resample observations with replacement in each sample and use states as clustering units in order to employ state fixed effects. All specifications are weighted by 1990 county population and include state fixed effects.

Table 2C.2: Estimates from the first-difference model with share of college grads (1990) control: Change in log of employment (1986-1993) and simulated county vulnerabilities

Vulnerability :	Actual	Simulated		
Distribution :	N/A	poly. approx.	uniform [0,0.4]	all industry uniform
Initial vul.	-.051	.006	0	.004
	(.221)	(.218)	(.079)	(.121)
state FE	yes	yes	yes	yes
share of colgrad (1990)	yes	yes	yes	yes

Notes: Number of simulation = 1000 times. The table reports the estimates from the first-difference model shown in equation (2.7). The dependent variable is change in log of county employment between 1986 and 1993. The first column uses the original county vulnerability. In column 2, the simulated tariffs for industries with positive tariffs are drawn from the estimated tariff distribution plotted in Appendix figure 2C.1. In column 3, the simulated tariffs for industries with positive tariffs are drawn from U[0, 0.4]. In column 4, the tariffs for all industries are drawn from U[0, 0.4]. The industry-level tariffs are simulated 1000 times and the corresponding county-level vulnerability measure is used as the main regressor for the first-difference analysis. We compute the average and standard deviation of the estimates and report them as the bootstrapped coefficient and standard error. We resample observations with replacement in each sample and use states as clustering units in order to employ state fixed effects. All specifications are weighted by 1990 county population and include share of college grads among 1990 county population as a control and state fixed effects.

Table 2C.3: Estimates from the first-difference model: Change in log of employment (1993-2000) and simulated county vulnerabilities

Vulnerability :	Actual	Simulated		
Distribution :	N/A	poly. approx.	uniform [0,0.4]	all industry uniform
Initial vul.	-.863	-.158	-.114	-.002
	(.194)	(.266)	(.1)	(.129)
state FE	yes	yes	yes	yes

Notes: Number of simulation = 1000 times. The table reports the estimates from the first-difference model shown in equation (2.7). The dependent variable is change in log of county employment between 1993 and 2000. The first column uses the original county vulnerability. In column 2, the simulated tariffs for industries with positive tariffs are drawn from the estimated tariff distribution plotted in Appendix figure 2C.1. In column 3, the simulated tariffs for industries with positive tariffs are drawn from U[0, 0.4]. In column 4, the tariffs for all industries are drawn from U[0, 0.4]. The industry-level tariffs are simulated 1000 times and the corresponding county-level vulnerability measure is used as the main regressor for the first-difference analysis. We compute the average and standard deviation of the estimates and report them as the bootstrapped coefficient and standard error. We resample observations with replacement in each sample and use states as clustering units in order to employ state fixed effects. All specifications are weighted by 1990 county population and include state fixed effects.

Table 2C.4: Estimates from the first-difference model with share of college grads (1990) control: Change in log of employment (1993-2000) and simulated county vulnerabilities

Vulnerability :	Actual	Simulated		
Distribution :	N/A	poly. approx.	uniform [0,0.4]	all industry uniform
Initial vul.	-.412	-.062	-.05	0
	(.223)	(.22)	(.084)	(.109)
state FE	yes	yes	yes	yes
share of colgrad (1990)	yes	yes	yes	yes

Notes: Number of simulation = 1000 times. The table reports the estimates from the first-difference model shown in equation (2.7). The dependent variable is change in log of county employment between 1993 and 2000. The first column uses the original county vulnerability. In column 2, the simulated tariffs for industries with positive tariffs are drawn from the estimated tariff distribution plotted in Appendix figure 2C.1. In column 3, the simulated tariffs for industries with positive tariffs are drawn from U[0, 0.4]. In column 4, the tariffs for all industries are drawn from U[0, 0.4]. The industry-level tariffs are simulated 1000 times and the corresponding county-level vulnerability measure is used as the main regressor for the first-difference analysis. We compute the average and standard deviation of the estimates and report them as the bootstrapped coefficient and standard error. We resample observations with replacement in each sample and use states as clustering units in order to employ state fixed effects. All specifications are weighted by 1990 county population and include share of college grads among 1990 county population as a control and state fixed effects.

Table 2C.5: Estimates from the first-difference model: Change in Republican presidential vote share (1984-1992) and simulated county vulnerabilities

Vulnerability :	Actual	Simulated		
Distribution :	N/A	poly. approx.	uniform [0,0.4]	all industry uniform
Initial vul.	.001	.01	.014	-.001
	(.065)	(.08)	(.029)	(.04)
state FE	yes	yes	yes	yes

Notes: Number of simulation = 1000 times. The table reports the estimates from the first-difference model shown in equation (2.7). The dependent variable is change in county-level Republican Presidential vote share between 1984 and 1992. The first column uses the original county vulnerability. In column 2, the simulated tariffs for industries with positive tariffs are drawn from the estimated tariff distribution plotted in Appendix figure 2C.1. In column 3, the simulated tariffs for industries with positive tariffs are drawn from U[0, 0.4]. In column 4, the tariffs for all industries are drawn from U[0, 0.4]. The industry-level tariffs are simulated 1000 times and the corresponding county-level vulnerability measure is used as the main regressor for the first-difference analysis. We compute the average and standard deviation of the estimates and report them as the bootstrapped coefficient and standard error. We resample observations with replacement in each sample and use states as clustering units in order to employ state fixed effects. All specifications are weighted by 1990 county population and include state fixed effects.

Table 2C.6: Estimates from the first-difference model with share of college grads (1990) control: Change in Republican presidential vote share (1984-1992) and simulated county vulnerabilities

Vulnerability :	Actual	Simulated		
Distribution :	N/A	poly. approx.	uniform [0,0.4]	all industry uniform
Initial vul.	-.032	.004	.011	-.001
	(.047)	(.077)	(.029)	(.039)
state FE	yes	yes	yes	yes
share of colgrad (1990)	yes	yes	yes	yes

Notes: Number of simulation = 1000 times. The table reports the estimates from the first-difference model shown in equation (2.7). The dependent variable is change in county-level Republican Presidential vote share between 1984 and 1992. The first column uses the original county vulnerability. In column 2, the simulated tariffs for industries with positive tariffs are drawn from the estimated tariff distribution plotted in Appendix figure 2C.1. In column 3, the simulated tariffs for industries with positive tariffs are drawn from U[0, 0.4]. In column 4, the tariffs for all industries are drawn from U[0, 0.4]. The industry-level tariffs are simulated 1000 times and the corresponding county-level vulnerability measure is used as the main regressor for the first-difference analysis. We compute the average and standard deviation of the estimates and report them as the bootstrapped coefficient and standard error. We resample observations with replacement in each sample and use states as clustering units in order to employ state fixed effects. All specifications are weighted by 1990 county population and include share of college grads among 1990 county population as a control and state fixed effects.

Table 2C.7: Estimates from the first-difference model: Change in Republican presidential vote share (1992-2000) and simulated county vulnerabilities

Vulnerability :	Actual	Simulated		
Distribution :	N/A	poly. approx.	uniform [0,0.4]	all industry uniform
Initial vul.	.401	.037	.016	-.003
	(.101)	(.131)	(.048)	(.067)
state FE	yes	yes	yes	yes

Notes: Number of simulation = 1000 times. The table reports the estimates from the first-difference model shown in equation (2.7). The dependent variable is change in county-level Republican Presidential vote share between 1992 and 2000. The first column uses the original county vulnerability. In column 2, the simulated tariffs for industries with positive tariffs are drawn from the estimated tariff distribution plotted in Appendix figure 2C.1. In column 3, the simulated tariffs for industries with positive tariffs are drawn from U[0, 0.4]. In column 4, the tariffs for all industries are drawn from U[0, 0.4]. The industry-level tariffs are simulated 1000 times and the corresponding county-level vulnerability measure is used as the main regressor for the first-difference analysis. We compute the average and standard deviation of the estimates and report them as the bootstrapped coefficient and standard error. We resample observations with replacement in each sample and use states as clustering units in order to employ state fixed effects. All specifications are weighted by 1990 county population and include state fixed effects.

Table 2C.8: Estimates from the first-difference model with share of college grads (1990) control: Change in Republican presidential vote share (1992-2000) and simulated county vulnerabilities

Vulnerability :	Actual	Simulated		
Distribution :	N/A	poly. approx.	uniform [0,0.4]	all industry uniform
Initial vul.	.13	-.011	-.018	-.001
	(.086)	(.09)	(.034)	(.047)
state FE	yes	yes	yes	yes
share of colgrad (1990)	yes	yes	yes	yes

Notes: Number of simulation = 1000 times. The table reports the estimates from the first-difference model shown in equation (2.7). The dependent variable is change in county-level Republican Presidential vote share between 1992 and 2000. The first column uses the original county vulnerability. In column 2, the simulated tariffs for industries with positive tariffs are drawn from the estimated tariff distribution plotted in Appendix figure 2C.1. In column 3, the simulated tariffs for industries with positive tariffs are drawn from U[0, 0.4]. In column 4, the tariffs for all industries are drawn from U[0, 0.4]. The industry-level tariffs are simulated 1000 times and the corresponding county-level vulnerability measure is used as the main regressor for the first-difference analysis. We compute the average and standard deviation of the estimates and report them as the bootstrapped coefficient and standard error. We resample observations with replacement in each sample and use states as clustering units in order to employ state fixed effects. All specifications are weighted by 1990 county population and include share of college grads among 1990 county population as a control and state fixed effects.

Table 2C.9: First-difference model with the AKM inference procedure: Change in log of county employment over pre-NAFTA control period (1986-1993)

Dep. Var. :	First-difference in log employment		
Initial vul.	.378	-.001	-.044
robust SE	(.162)	(.14)	(.155)
AKM SE, three-digit SIC cl.	(.204)	(.144)	(.14)
state-clustered SE	(.176)	(.13)	(.208)
pop. weighted	yes	yes	yes
state FE	no	yes	yes
share colgrad (1990)	no	no	yes
Observations	2988	2988	2988

Notes: The table reports the estimates of equation (2.7) with robust (Eicker-Huber-White) standard errors, the AKM standard error with three-digit SIC clusters, and state-clustered standard errors. The dependent variable is change in log of county employment between 1986 and 1993. All specifications are weighted by 1990 county population, and the second column reports the specification with state fixed effects, and the third column reports the specification with share of college grads in 1990 as a control and state fixed effects.

Table 2C.10: First-difference model with the AKM inference procedure: Change in log of county employment over treatment period (1993-2000)

Dep. Var. :	First-difference in log employment		
Initial vul.	-.854	-.993	-.537
robust SE	(.158)	(.139)	(.143)
AKM SE, three-digit SIC cl.	(.184)	(.186)	(.166)
state-clustered SE	(.214)	(.171)	(.192)
pop. weighted	yes	yes	yes
state FE	no	yes	yes
share colgrad (1990)	no	no	yes
Observations	2986	2986	2986

Notes: The table reports the estimates of equation (2.7) with robust (Eicker-Huber-White) standard errors, the AKM standard error with three-digit SIC clusters, and state-clustered standard errors. The dependent variable is change in log of county employment between 1993 and 2000. All specifications are weighted by 1990 county population, and the second column reports the specification with state fixed effects, and the third column reports the specification with share of college grads in 1990 as a control and state fixed effects.

Table 2C.11: First-difference model with the AKM inference procedure: Change in Republican presidential vote share over pre-NAFTA control period (1984-1992)

Dep. Var. :	First-difference in two-party Repub. vote share		
Initial vul.	.175	.016	-.018
robust SE	(.054)	(.049)	(.046)
AKM SE, three-digit SIC cl.	(.051)	(.038)	(.04)
state-clustered SE	(.105)	(.07)	(.053)
pop. weighted	yes	yes	yes
state FE	no	yes	yes
share colgrad (1990)	no	no	yes
Observations	2885	2885	2885

Notes: The table reports the estimates of equation (2.7) with robust (Eicker-Huber-White) standard errors, the AKM standard error with three-digit SIC clusters, and state-clustered standard errors. The dependent variable is change in Republican House vote share between 1984 and 1992. All specifications are weighted by 1990 county population, and the second column reports the specification with state fixed effects, and the third column reports the specification with share of college grads in 1990 as a control and state fixed effects.

Table 2C.12: First-difference model with the AKM inference procedure: Change in Republican presidential vote share over treatment period (1992-2000)

Dep. Var. :	First-difference in two-party Repub. vote share		
Initial vul.	.533	.422	.141
robust SE	(.074)	(.061)	(.061)
AKM SE, three-digit SIC cl.	(.118)	(.073)	(.054)
state-cluster SE	(.125)	(.099)	(.087)
pop. weighted	yes	yes	yes
state FE	no	yes	yes
share colgrad (1990)	no	no	yes
Observations	2885	2885	2885

Notes: The table reports the estimates of equation (2.7) with robust (Eicker-Huber-White) standard errors, the AKM standard error with three-digit SIC clusters, and state-clustered standard errors. The dependent variable is change in Republican House vote share between 1992 and 2000. All specifications are weighted by 1990 county population, and the second column reports the specification with state fixed effects, and the third column reports the specification with share of college grads in 1990 as a control and state fixed effects.

D Accounting for industry-level benefits from NAFTA

Overview

There are two mechanisms by which NAFTA could have benefited U.S. industries and local labor markets. First, NAFTA decreased tariffs U.S. exporters face when U.S. goods are being exported to Mexico, which would have increased the demand of U.S. goods from Mexico (what we will call the *export advantage*). Second, the reduced tariffs on imports from Mexico could decrease the production cost of U.S. goods that use Mexican imports as inputs (what we will call the *input advantage*). We consider the possibility that counties could benefit from NAFTA through these channels, which could also account for changes in the economic condition of local labor markets.

One concern with the analysis in the main part of the paper is omitted-variables bias (from excluding the advantage measures) is causing us to misinterpret the coefficient on *Vulnerability*. In the classic omitted-variables framework, if the export- or input-advantage variables were (a) negatively correlated with *Vulnerability* and (b) positively correlated with county employment, then our estimated coefficient on *Vulnerability* would be negatively biased.

To test this idea, we construct county-level measures of input advantage and export advantage based on county industrial composition in 1990. As we show below, these potential omitted variables are *positively* correlated with our county-level *Vulnerability* measure, suggesting little scope for omitted-variables bias in driving our result and instead that the *Vulnerability* measure is picking up the net effect of NAFTA on local labor markets, including any local benefits.

Constructing county-level measures of export advantage

Our county-level measure of *export advantage* is based on how much the drop in tariffs on U.S. products exported to Mexico can help U.S. industries and thus local labor markets with these industries. Similar to our vulnerability measure, we need three components for constructing the county export advantage measure: (*i*) Industry-level tariffs that are applied to U.S. exports to Mexico prior to NAFTA in year 1993; (*ii*) Revealed Comparative Advantage (RCA) of U.S. industries; and (*iii*) county-level industrial composition prior to NAFTA.[49]

We acquire industry-level Mexican tariff on imports from the US prior to NAFTA in 1993 from López-Córdova (2003).[50] The data contains tariff by four-digit ISIC industries, which we map to SIC codes using ISIC-to-SIC code crosswalk from the Industry Concordance website by Jon Haveman.[51] When aggregating the tariffs by SIC industries, we compute unweighted average tariffs of ISIC industries that correspond to each SIC industry.[52] RCA of U.S. industries is computed in an analogous way to how the Mexican RCA is constructed using Hakobyan and McLaren (2016a)'s replication data (which is from the UN Comtrade bilateral export series) and US HS-level imports and exports from Peter Schott's data webpage.[53] The county-level industrial composition prior to NAFTA is drawn from the 1990 CBP.

[49] We ideally would want tariffs in 1990 to remain consistent with the construction of our *Vulnerability* variable, but our data source described below only has tariff data for year 1993.

[50] We greatly appreciate the generous help from Jose Ernesto López-Córdova and Jose Ramon Morales Arilla in sharing the data.

[51] See https://www.macalester.edu/research/economics/PAGE/HAVEMAN/Trade.Resources/TradeConcordances.html.

[52] We would ideally compute the weighted average tariffs, weighted by the U.S. export values in each four-digit ISIC industry, but we do not have access to such information.

[53] See https://sompks4.github.io/sub_data.html.

Constructing county-level measures of input advantage

Our county-level measure of *input advantage* captures how much the decline in tariffs on Mexican imports to the US could *help* U.S. producers by reducing production input costs (and thus local labor markets that rely on these industries whose input costs decline). We need three components for constructing county-level input advantage measure: (*i*) industry-level weighted average tariffs applied to production inputs prior to NAFTA; (*ii*) Mexican RCA; and (*iii*) county-level industrial composition. The data sources used in (*i*)-(*iii*) are identical to the data sources for computing the vulnerability measure. An additional data source we draw from is the Input-Output matrix of the United States in 1990 from OECD Input-Output database. Input-output (I/O) matrix is a matrix that contains information about what share of production input costs of each industry is spent on each input industry. The OECD I/O matrix has information on 34 industries, which we map to groups of two-digit SIC industries (i.e., each group comprises one or more than one two-digit SIC industries.)

In computing (*i*), we start by computing average tariff of industry groups, weighted by the import value of each industry group. Then we use the input cost composition for each industry group drawn from the I/O matrix to construct a measure of how each industry was affected by tariffs for the imported inputs prior to NAFTA. Computing (*ii*) and (*iii*) is exactly analogous to how we constructed Mexican RCA and county industrial composition for vulnerability measure, but with more aggregated industry groups. When using the input advantage measure in the analysis, we also build the vulnerability measure using same aggregated SIC industry categories for comparability.

Covariance between export- and input-advantage measures and *Vulnerability*

Whether the omission of county-level export or input advantages of NAFTA creates negative bias in our estimates of the coefficient on *Vulnerability* depends on the covariances

of these measures with *Vulnerability*. The correlation between our *Vulnerability* and export advantage measures is 0.15. It is not surprising to see that the export advantage measure is positively associated with the *Vulnerability* measure, as the main export industries in the US are often the main import industries in the US—the United States accounts for a large part of Foreign Direct Investments (FDI) to Mexico in the form of producing intermediate inputs and parts in Mexico and importing them to the US. The main export industries from the US to Mexico include autos and automotive parts, computers and electronics, textile and apparel, and the main import industries from Mexico to the US are also autos and automotive parts, computers and electronics, textiles and apparel, ceramic tile.[54]

The correlation between our county-level input advantage measure and *Vulnerability*, constructed using the aggregation of 2-digit SIC industries, is 0.97.[55] The high correlation is not only because industries heavily rely on own-industry inputs, but also because industry groups must be aggregated to include multiple two-digit SIC industries, making the share of own-industry input even higher.

Is the Vulnerability effect robust to controlling for export advantage?

While the positive correlation between export advantage and *Vulnerability* makes it unlikely that omitted-variables bias is causing the negative coefficients on *Vulnerability* after 1993, we can nonetheless include export-advantage and its interaction with year fixed effects in our standard event-study regressions. We do not perform this exercise for the input-advantage measure given that it is nearly collinear with *Vulnerability*.

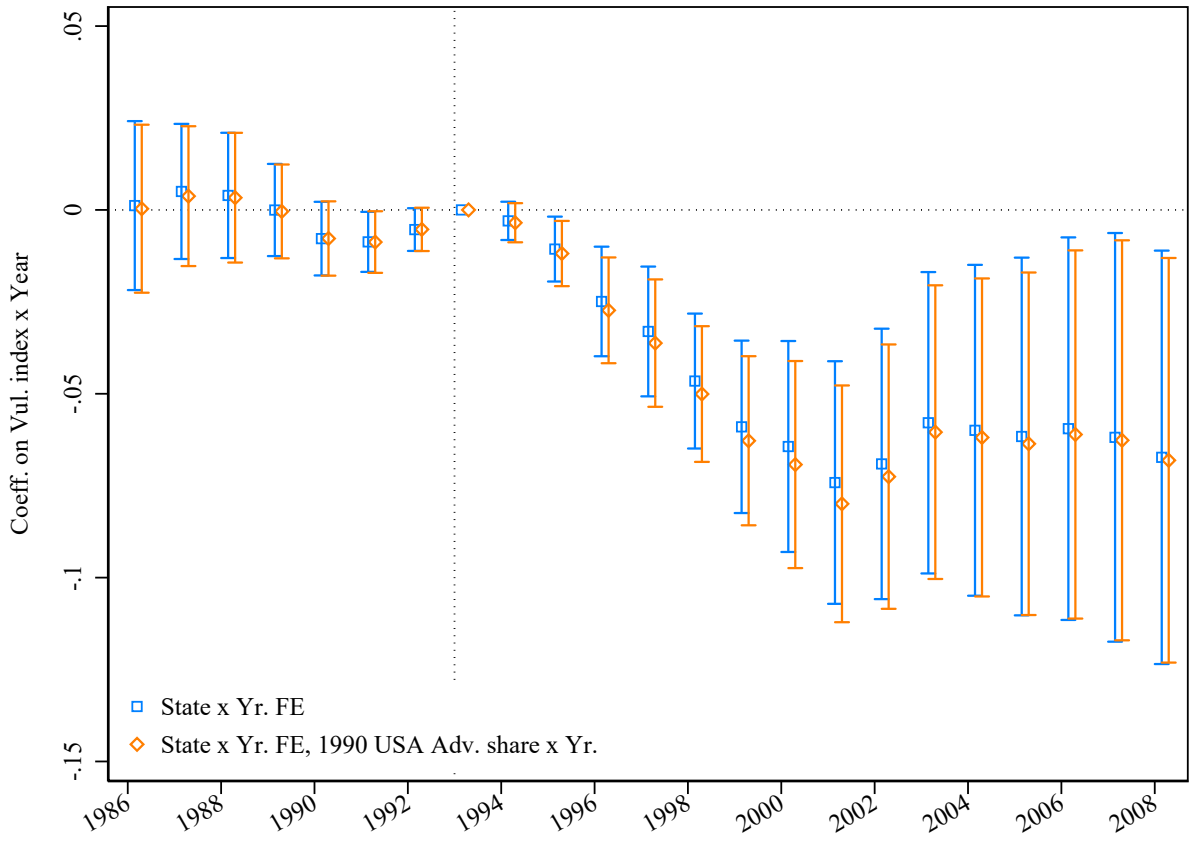
In Appendix Figure 2D.1, we test the robustness of our main log-employment result to including the flexible controls for export advantage. The first series merely reproduces the second series from our main log-employment figure (Figure 2.2 from the main paper). The

[54]This paragraph draws from USITC publication 2596 (1993), which provides a discussion of this strong overlap between import and export industries between the US and Mexico. See `https://www.usitc.gov/publications/332/pub2596.pdf` for the full report.

[55]The correlation between the *Vulnerability* constructed using 4-digit SIC industrial composition and the (two-digit) input-advantage measure is 0.56.

second series adds the export-advantage measure interacted with year fixed effects. The coefficients on *Vulnerability* barely move.

Figure 2D.1: Log county employment as a function of county vulnerability, robustness to controlling for county export advantage from NAFTA



Sources: The dependent variable is log county employment drawn from the CBP. See Appendix B for more detail.

Notes: $N = 2804$ counties. This figure is identical to Figure 2.2 except that we also use county export advantage measure as a flexible control on the second series along with the vulnerability measure as the independent variable. The figure shows the event-study coefficient estimates (and 95% confidence intervals, based on standard errors clustered by state) from different specifications of equation (2.4). All specifications are weighted by 1990 county population. The first series shows the coefficient estimates from a specification where we control for only county and year fixed effects and state×year fixed effects, which is identical to the second series of Figure 2.2. The specification for the second series adds to this baseline specification 1990 county-level export advantage measure interacted with year fixed effects.

Chapter 3

Automation after the assembly line: Computerized machine tools, employment and productivity in the United States[1]

3.1 Introduction

For the past century, manufacturing has been characterized by a continuous (but punctuated) process of automation, whereby new technology enables tasks previously completed by human labor to be accomplished, in whole or in part, with machines. Many technologies have contributed to this transformation – including "computer numerical

[1]This chapter was co-authored with Leah Boustan and David Clingingsmith. We benefited from valuable suggestions from seminar audiences at the Census Bureau, Cleveland Fed, Federal Reserve Board, Johns Hopkins, Rochester, UC-Davis, University of St. Andrews and UVA. We thank David Autor, Henry Farber, Teresa Fort, Susan Helper, Robert Margo, Enrico Moretti and Owen Zidar for useful conversations. We are grateful to the Association for Manufacturing Technology for access to their archival material. Monica Aberg, Camilla Adams, Ethan Bergmann, Kaan Cankat, Nicole Gandre, and Katie Guyot provided excellent research assistance.

control machinery, industrial robots, and artificial intelligence" (Acemoglu and Restrepo, 2019, p. 3).

In this paper, we focus on the productivity and employment effects of computer numerical control (CNC) machine tools, which began to diffuse widely in the 1970s. CNC tools rely on computer programs rather than human operatives to select and perform the tool's physical movements. Like other forms of automation, CNC has the potential to enhance labor productivity. However, CNC may have had differential impacts on workers by skill level, displacing semi-skilled operatives responsible for using hand-operated machine tools, while creating new tasks for high-skilled technicians who could install, program, and fix complex machines.

We find that the diffusion of CNC technology led to rising productivity in manufacturing. In particular, the advent of CNC was associated with growing capital investments, rising labor productivity, and a falling labor share in manufacturing. Although industries that were more exposed to CNC experienced no discernable change in total employment, exposed industries shed many workers on the production floor. Moreover, relative demand shifted away from workers with a high school degree or some college to workers with a college degree. The diffusion of CNC was also associated with growth in occupations associated with abstract tasks and a decline in occupations engaged in routine tasks.

Workers were able to partially adapt to this technology shock in three notable ways. First, workers who belonged to unions were less affected by the shock. The estimated decline in employment was entirely concentrated among non-unionized employees, either because unionized firms adopted the new technology more slowly than their non-union counterparts or because they adjusted their workforce more gradually as new technology arrived.

Second, we find that workers who did not hold a college degree were more likely to go back to school, particularly to earn a four-year degree. We also provide new evidence that colleges and universities expanded their degree offerings related to CNC technology to accommodate this growing interest.

Third, a local labor market analysis shows that, in labor markets that were more exposed to the CNC shock, employment in metal manufacturing declined substantially, but employment in the manufacturing sector as a whole did not fall. This pattern suggests that workers were able to shift from the most exposed industries to other similar types of employment, as in the case of telephone operators studied in Feigenbaum and Gross (2020).

We focus on seven metal manufacturing industries that account for the majority of CNC adoption.[2] Our analysis is based on a novel measure of exposure to CNC technology at the industry-year level. We start with baseline variation in installed machine tools by type – such as lathes, mechanical presses, grinding machines and so on – across industries in 1958, before the spread of CNC. These tools shifted from manual to automatic control in different years and to different extents due to varying engineering challenges across tools and differences in industrial policy in the world's major exporters. When applying our measure to the analysis of local labor markets, we use initial (1970) employment shares by industry and commuting zone (controlling for 1970 employment share in metal manufacturing) to measure local exposure to this automation shock.

We digitized data on machine tool use by industry from the American Machinist Inventory of Metalworking Equipment to calculate tool use in 1958, before the invention of CNC. We also collected annual trade statistics with counts and values of tool exports by type and CNC status to measure differential timing in the switch to CNC by tool. Following a common approach in the literature, we do not rely on production data for the US machine tool industry, which may have been more responsive to changes in demand

[2]Similarly, four industries account for around 70% of robot adoption (Acemoglu and Restrepo, 2019).

among US manufacturers, but instead focus on the three major world exporters – Germany, Japan and Italy. Our completed diffusion measure begins in 1968 and we then end our analysis in 2007 on the eve of the Great Recession and before the uptick in diffusion of industrial robots.

By our measure, all sample industries had very low exposure to CNC in the early 1970s. By 1990, industries like aircraft increased exposure to CNC tools dramatically (up to 40% of its tool base by value), while other industries like motor vehicles were less affected (25% of its tool base). Our exposure measure is closely correlated with actual CNC adoption at the industry level over the three years for which we have industry-level adoption data (1973, 1978, and 1983).

Industries that gained higher exposure to CNC technology in certain periods may have faced other trends, such as rising import competition or other changes to the production process that would have favored high-skilled workers. We provide three pieces of evidence against the idea that our results are driven by other potentially correlated factors. First, we do not find a pre-trend in the workforce composition of industries that were more or less exposed to the later CNC shock before the shock takes place (pre-period = 1968-75). Second, we control for import competition for final products due to growing international trade. Third, our local labor market analysis finds similar employment patterns with a more continuous version of the industry variation.

Our paper is most closely related to classic work on the effect of factory electrification on demand for skill in US manufacturing (Goldin and Katz, 1998; Gray, 2013; Katz and Margo, 2014; Fiszbein et al., 2020). We extend this timeline forward beyond the heyday of mass production employment in the 1940s and 1950s to the automation of the machining process in the 1970s and 1980s. Like these earlier episodes, we find that the diffusion of automated machine tools increased labor productivity. Electrification was associated with

greater employment of low-skilled workers, whereas we find an increase in relative demand for college graduates.[3]

Our findings also contain insights for understanding the effects of industrial robots on the contemporary labor market.[4] Adoption of industrial robots has been associated with falling employment at the industry level in the US, but with null or positive effects on employment in Germany and Japan (Acemoglu and Restrepo, 2020a; Graetz and Michaels, 2018; Dauth et al., 2021; Adachi et al., 2020).[5] Dauth et al. (2021) suggests that higher unionization rates in Germany may explain the null effect of robot technology on dis-employment. Acemoglu and Restrepo (2019) instead point to higher manufacturing wages in Germany and Japan. Our findings are consistent with both explanations. In the early 1970s, as CNC began to diffuse, 45% of workers in the metal industry were covered by a union, compared to around 10% with the spread of robots in 2010 (Appendix Figure 3A.1). CNC also automated mid-skill operatives and machinists involved with producing metal parts, whereas robots automate the lower-skilled assembly of these parts into a final product.[6]

Finally, our study builds on three quantitative studies of CNC adoption, and a large qualitative literature on this era. (We summarize the qualitative literature in the historical context section below.) The earlier quantitative papers compare establishments or plants

[3]Gaggl et al. (2021) and Lewis and Severnini (2020) study the effect of electrification on structural transformation from agriculture to manufacturing and on the agricultural sector in rural areas, respectively.

[4]Both automated machine tools and industrial robots are specific examples of the transformative role of computers in the factory. As such, our paper joins a large literature on the effect of computerization and skill-biased technical change on the workforce (a few touchstone papers include Autor et al., 1998 and Autor et al., 2003).

[5]Related papers compare firms that invest in robot technology (or, more broadly, in "industrial equipment") to competitors who do not. Firms that adopt robots experience rising productivity, output and total employment (Humlum, 2019; Acemoglu and Restrepo, 2020b; Koch et al., 2021; Aghion et al., 2020). Bessen et al. (2019) and Bessen et al. (2020) instead find evidence of displacement, even at the firm level, following automation events in the Netherlands.

[6]CNC "often replaced skilled craftspeople" (Doms et al., 1997, p 260) including "the skilled jobs of setter and set-up operator [and] the semi-skilled job of machine operator" (Keefe, 1991, p. 516). By contrast, robots are primarily used "to perform several manual tasks" in the second stage of the process, "such as welding, painting, assembly, handling materials and packaging" (Acemoglu and Restrepo, 2020a, p. 2189).

that adopt CNC to competitors who do not (Keefe, 1991; Doms et al., 1997; Bartel et al., 2007). Firms that choose to adopt new technology likely differ in other ways (e.g., these firms could be bigger, better managed, more engaged in export, located in different areas, and so on). These differences are hard to control for without the use of detailed event studies. We instead rely on industry-level exposure driven by new technology discoveries, and consider a wider set of productivity outcomes and worker adjustments.[7]

3.2 Historical Context

The punctuated history of automation in the manufacturing sector

This section situates the development of computer numerical control machine tools in the longer-run history of factory automation. Automation has evolved in four main stages. The first step, which was a necessary precursor to all later automation episodes, was the development of interchangeable parts in the 19th century (Hounshell, 1984). With interchangeability, the production of metal goods could be subdivided into two distinct activities: the production of parts from raw metal stock and the assembly of those parts into finished goods in bulk. Interchangeability eliminated the need for skilled "fitters," who adjusted parts to fit together as they were assembled by hand.

The achievement of true interchangeability depended on the advent of modern machine tools that could cut or bend raw metal in ways that were both precise and replicable. Machine tools such as lathes and drills have been in use since the 18th century but underwent rapid improvement and diffusion after 1860 (Holt, 1966, Woodbury, 1972; Atack et al., 2019). These machine tools were operated by skilled machinists who translated

[7]The existing papers on CNC consider different parts of the skill distribution. Doms et al. (1997) find that plants adding computer-assisted production *do not* hire more college graduates. Keefe (1991) documents that machine shops adopting CNC replace skilled occupations like machinists with unskilled occupations like machine tenders. Bartel et al. (2007) show that skill requirements for newly-hired machine operators are higher at firms with CNC, suggesting that there can be some skill upgrading *within* occupation categories.

engineering drawings into precise operations by manipulating the wheels and levers of the machine by hand.

The second step in the history of automation, beginning in metal manufacturing, was the invention of the assembly line, pioneered by the Ford Motor Company in 1908 (Hounshell, 1984). Factory electrification, which diffused rapidly between 1910 and 1930, was important to the development of the assembly line because it allowed for the flexible placement of each machine in the order required to maximize efficiency (Devine, 1983). Previously, machines were powered by a central drive, which limited flexibility and required porters to move parts around the factory. Electrification thus substituted for human labor and was complementary with more high-skilled tasks needed to install and fix machines (Goldin and Katz, 1998).

Automated machine tools, the subject of this paper, were the third step in this automation process. Before this step, semi-skilled operators were required to control the motions of machine tools. Numerical control – as first developed in the late 1950s and then computerized in the mid-1970s – codifies the movements of skilled operators into a program so that a less-skilled operative can execute them automatically. Numerical control required the invention of both computer systems that could execute the programs and servomechanisms that translated the programs into precise physical movements of the tools. Both technologies advanced substantially during World War II, the computer primarily for ballistics computations and servomechanisms for the automatic targeting of guns to ships or airplanes detected on radar (Mindell, 2004). We describe the invention and diffusion process for CNC tools in the next section.

The fourth step in the automation of metal manufacturing has been the use of industrial robots in the assembly of metal products from components parts. Unlike CNC tools, which replaces the fine motor skills of skilled machinists, robots automate the gross motor skills involved with assembly, as well as with "welding, painting... handling

materials and packaging" (Acemoglu and Restrepo, 2020a, p. 2189). Robots began to diffuse widely in the 1990s and have received significant attention from economists.

The invention and diffusion of CNC machine tools

The first numerically controlled machine tool was invented in the United States in the early 1950s at the MIT Servomechanisms Lab, building on the advances in computing and servomechanisms that emerged during World War II. These early tools, developed under contract with the US Air Force, were used for the machining of helicopter rotor components, which required a level of precision that even skilled machine tool operators of the day could not readily attain (Noble, 1986).

The original numerically controlled machine tools built at MIT were too expensive to be commercially viable. The goals of the Air Force, along with the preferences of the scientists involved, resulted in a machine that was extremely precise and of wide capability but also very expensive. While commercialization began in the late 1950s, initial adoption was largely confined to the aircraft industry, where a large share of revenue came from cost-plus contracts with the US government.

The first computer numerical control tools designed for wide commercial applications were developed in Japan in the late 1960s. Japanese tool makers became the dominant producers by the early 1970s, followed by German competitors. Throughout the 1960s, Japanese tool makers – with the support of the Ministry of International Trade and Industry (MITI) – were pursuing lower-cost (and thus less precise) designs that were more suited to Japanese metal manufacturing. American machines used closed-loop feedback mechanisms, in which the location of the cutting edge of a tool was independently measured by sensors. Japanese machines used open-loop systems, which eliminated costly sensors and assumed tools had moved without error. This design was not initially precise enough to use in aircraft manufacture but was suitable for other industries and much cheaper to produce.

In the mid-1970s, microprocessors replaced dedicated hardware modules, a transition marked by the replacement of the term "numerical control" with the alternative "computer numerical control," or CNC. We adopt the term CNC to refer to automated machine tools throughout the paper, even though the earliest periods in our data series are before this transition from NC to CNC. Microprocessors increased the flexibility of CNC tools, lowered production costs directly, and made the addition of more accurate closed-loop controllers cheap. The US machine-tool industry lagged behind Japan's in converting their designs to CNC (Wieandt, 1994).

Innovation in CNC tools was not driven by US demand

The historical record suggests that both the direction and the speed of innovation by Japanese and German tool exporters were driven primarily by their own domestic markets.[8] Japan and Germany specialized in different machine tools – Japan in lathes and Germany in boring machines, for example – as suited their own domestic manufacturing sectors. Japan's small-to-medium sized manufacturing firms created substantial domestic demand for lathe production in the 1960s (Itohisa, 2010). Japan's MITI provided incentives to machine tool makers in the mid-1960s to develop economies of scale in CNC lathes by producing for the domestic market first before promoting exports (Johnson, 1982; Sarathy, 1989). "Exports nevertheless remained of secondary concern to the Japanese industry until it had exploited the domestic market and gained technological leadership in low-cost CNC machine tools" (Collis, 1988). Japan's early expertise in the making of lathes was then persistent as the market shifted to CNC (Collis, 1988).

Expertise in one machine tool did not translate directly into supremacy in others. Indeed, "to develop a lathe required a different design expertise from that needed to develop a grinding machine or a drill" (Collis, 1988). Twenty-five percent of Japan's lathe

[8]The US comprised approximately 40% of the export market for Japan, 13% for Germany, and 10% for Italy. Calculations based on statistics reported in the Economic Handbook of the Machine Tool Industry, various years, and UN COMTRADE.

exports were CNC in 1975, rising to 95% CNC by 1985. By contrast, only 50% of Japanese boring, drilling and grinding machines were CNC by that year.

While Japan specialized in lathes, Germany instead specialized in boring machines, and boring machines were the first to convert to CNC in Germany. These country-specific patterns are confirmed in the US patent records: lathes dominate early patenting by Japanese firms related to CNC technology, while boring machines dominate patenting by German firms (see Appendix Figure 3A.2).

Inherent differences in the difficulty of automation between tool types can also explain some of the temporal patterns in CNC diffusion. For example, grinding is inherently more difficult to automate than milling, drilling, or turning (Collis, 1988). The later and less complete diffusion of CNC in the grinding tool category is consistent with this greater technical challenge.

3.3 Construction of CNC exposure measure and data sources

Industry-year exposure to CNC technology

We construct a measure of exposure to CNC technology that varies by year and by industry. Our measure relies on two sources of variation: the share of each tool type (e.g., lathes, boring machines) in an industry's tool base as of 1958, before the diffusion of CNC tools; and the cumulative value share of exports of new machine tools made up of CNC (rather than hand-operated tools) up to year t from the three major machine tool exporters (Japan, Germany, and Italy). Our exposure measure to CNC machine tools for industry j is thus the cumulative share of CNC tools in the global market as of year t, weighted by baseline tool use in that industry. We use the cumulative CNC share, rather than the annual CNC share. This measure better reflects the transition of the tool stock

toward CNC for a given year. When combined with industry and year fixed effects, the effect is identified by the annual increases the CNC share relative to the industry average.

Our empirical analysis captures the most significant players in the world market. Japan overtook the United States to become the largest producer of CNC lathes in 1975 and served more than 60% of the world market by 1981 (Rendiero, 1985). Germany was the second largest producer with around 20% of the market while the United States was left with only 10%. Italy was the fourth largest producer. We exclude US machine tool production in the analysis, which may have been more responsive to demands from domestic manufacturers. To construct our measure, we begin at the tool level, measuring the cumulative CNC share of exports for each tool type k and exporter i from 1971 up to year t. The CNC share for exporter i by year t can be written:

$$\text{Share CNC}_{i,k,t} = \frac{\sum_{\tau=1971}^{t} X_{i,\tau,k}^{\text{CNC}}}{\sum_{\tau=1971}^{t} X_{i,\tau,k}^{\text{Total}}} \tag{3.1}$$

where $X_{i,\tau,k}^{\text{Total}}$ are the total annual export value of type k machine tools of any mode (hand-operated or CNC) from exporter i to the global market in year τ and $X_{i,\tau,k}^{\text{CNC}}$ are the annual export value of CNC tools of type k from exporter i in year τ.

We then aggregate across exporters to create the tool-level CNC share. To do so, we weight the CNC share for took k from exporter i (equation 1) by exporter i's share of the total export value of tool type k, which can be written $\frac{X_{i,k,t}^{\text{Total}}}{\sum_{\iota} X_{\iota,k,t}^{\text{Total}}}$. The CNC share at the tool-by-year level (weighted across exporters) can be expressed:

$$\text{Share CNC}_{k,t} = \sum_{i} \text{Share CNC}_{i,k,t} \frac{X_{i,k,t}^{\text{Total}}}{\sum_{\iota} X_{\iota,k,t}} \tag{3.2}$$

Finally, we link the CNC share for tool k (equation 2) to industry j by weighing by the 1958 value share of tool k among the tool inventory for that industry $\frac{VT_{k,j,1958}}{\sum_{j} VT_{k,j,1958}}$. Equation 3.3 thus presents our exposure measure to CNC technology at the industry-by-year level:

$$\text{Exposure CNC}_{j,t} = \sum_k \left(\frac{VT_{k,j,1958}}{\sum_k VT_{k,j,1958}} \text{Share CNC}_{k,t} \right) \tag{3.3}$$

In the remainder of the section, we explain the data sources for constructing each component of this measure and illustrate the resulting patterns of variation.

1958 value shares of tool k for industry j: We construct the industry-level measures of tool base in equation 3 from the 1958 American Machinist Inventory of Metalworking Equipment (AMIME). The AMIME contains information on the value of tool inventories for 28 detailed tool types for each metalworking sub-industry.[9] Figure 3.1 demonstrates that there is substantial variation in the intensity of tool use across industries. For example, we find that mechanical presses are relatively heavily used in fabricated metal products, boring machines in aircraft, gear cutting machines in farm machinery, and lathes in precision mechanisms. The figure is organized as a heatmap of tool-type usage for the seven metalworking industries in our analysis, ordered by the amount of variation in use of the tool (standard deviation) across the industries. Cells that are shaded orange reflect greater-than-average use of the tool type relative to other metal manufacturing sub-industries, and purple shading reflects less-than-average use. We report tool use for the 14 tool types included in the export series.

Annual CNC shares by tool k: We collect the export values for each machine tool type by exporter (Japan, Germany, and Italy) from the Economic Handbooks of Machine Tool Industry. The data includes 14 major tool types by CNC status for at least one exporter and is available from 1971 to 2009. We consolidate these 14 tool types to seven categories to reflect differential reporting patterns by exporter. Our measure captures the majority of variation in tool use: the 14 tools included in the trade data represent 78% of the value share of the 1958 tool base. To complete the series, we impute the CNC share to be the

[9]The AMIME contains SIC codes for 16 sub-industries. We aggregate these SIC codes into 7 categories using the 1950 census industry codes to merge in the other variables and outcomes used throughout the project.

lowest CNC share for the given exporter in that year for the 14 tool types that are not in the trade data. Figure 3.2 documents that different tool types shifted toward CNC technology at different times and, ultimately, at different rates. We plot the time series of CNC machine tools as a value share of all machine tool exports for the three major exporters (as in equation 1). For Japan (Panel A), lathe exports reach 50% CNC by 1976, a level only reached by milling machines around 1980, boring and drilling machines around 1984, and grinding machines around 1992. For German exports (Panel B), there is a thirty-year lag between the first tool to reach the 50% CNC mark (boring machines, 1980) and the last tool to do so (mechanical presses, 2010). Comparable Italian export series are reported in Appendix Figure 3A.3. We extend each of these series back to 1968 by assuming that each tool had zero CNC share in 1968, 1969 and 1970 before CNC diffusion truly began.

Figure 3.3 plots our industry-level CNC exposure measure (equation 3), which combines the 1958 tool shares by industry from Figure 3.1 with the annual CNC shifts by tool from Figure 3.2. CNC tools diffused most rapidly for the aircraft and precision mechanism industries and most slowly for fabricated metal and motor vehicles. Aircraft reached a diffusion level of 30% of its tool base by 1986, whereas motor vehicles took another decade to reach this level. The construction of our measure makes it clear why this was so: The aircraft industry was particularly reliant on two types of machines – boring and milling tools – that were early to shift to CNC technology. By contrast, the motor vehicles industry was less likely to use early-adoption tools like lathes and more reliant on late-adoption tools like gear cutting.

A Data sources for outcome variables

We use three data sources to measure productivity and employment effects of CNC technology at the industry or local labor market level. First, we draw on the NBER-CES

Manufacturing Industry database for measures of capital investment, total employment and employment by category (production, non-production), labor productivity and the labor share by industry, originally tabulated by the Census of Manufactures (Bartelsman and Gray, 1996). We map the SIC industry codes used in the database to our seven metal manufacturing categories. Labor productivity is measured as log value added per worker, and the labor share is measured as wage bill divided by value added.

Second, we use the CPS Annual Social and Economic Supplement (ASEC) to count employment by industry and education category (less than high school degree, exactly high school degree, some college, college degree or more) and to compute task scores by industry using the underlying distribution of occupations. In particular, we compute the average abstract, routine, and manual task scores of workers in each industry over time by mapping the occupation-level task scores from Autor et al. (2003) to the worker-level occupation variables.

Third, we use Census data in 1970 to calculate the baseline share of employment by industry and commuting zone. We measure total employment and education shares of the workforce by sector (metal manufacturing and total manufacturing), commuting zone and decade (1970-2000).

To measure worker adjustments to the CNC technology shock, we use three CPS supplements. We collect union status from the CPS May supplement (1973-1983) and then from the CPS Outgoing Rotation Groups (1984-2007) in order to count workers by industry and union sector (member/non-member of a union). We use data on school enrollment from the CPS October Educational supplement (1976-2007). Our sample includes prime-age men who did not already hold a bachelor's degree in the survey year and whose current or most recent employment (within the past five years) was in a metal manufacturing industry.

In the appendix, we report firm responses to the arrival of the CNC technology. We compute the share of workers by industry located in the US South from the CPS ASEC sample. We then use the County Business Patterns (CBP) data to measure changes in the distribution of establishment size by industry over time.

3.4 Graphical evidence of CNC adoption and economic effects

We start by presenting a series of figures illustrating the relationship between industry-level exposure to CNC machine tools and economic outcomes.

A Adoption of CNC machines

Figure 3.4 graphs the correlation between our CNC exposure measure and the adoption of CNC machine tools, adjusting for year and industry fixed effects. We only have direct measures of CNC adoption – or the fraction of the installed tool base comprised of CNC tools – for three years: 1973, 1978, and 1983.[10] Figure 3.4 shows a strong correlation ($\phi = 0.74$) between exposure to CNC and actual adoption in these years, providing validation of our exposure measure.

If we had access to data on CNC tool use by industry at an annual level, we could use our exposure measure as an instrument for actual CNC adoption. Because adoption data is not otherwise available, we instead consider the direct effect of CNC exposure on economic outcomes (an "intent to treat" specification). If the relationship between CNC exposure and actual CNC adoption is less than one-for-one (as Figure 3.4 suggests), then the 'true'

[10]We collected this data from the AMIME in 1973, 1978 and 1983. The AMIME included information on CNC adoption in 1989 but the three-digit breakdown necessary to construct was not included in the published report.

effect of CNC adoption on economic outcomes would be larger than the 'intent to treat' estimates that we report in the paper.

B Trends in economic outcomes before and after the introduction of CNC

Figure 3.5 considers trends in the core economic outcomes in our analysis before and after CNC technology began to diffuse circa 1975. For graphical exposition, we partition the industries into two groups – high and low exposure – based on the intensity of the CNC shock. Motor vehicles, fabricated metals, electronics, and general industrial equipment are classified as low exposure, and farm equipment, aircraft and parts, and precision mechanisms as high exposure (see Figure 3.3).

Panel A presents the logarithm of capital expenditures and the logarithm of total employment for high and low CNC shock industries. Adoption of CNC technology, which requires the purchase and installation of new machine tools, should lead firms to undertake new capital expenditures. We find that investments in high and low shock industries move together from 1960 to 1974, with the exception of a short-lived spike from 1966-1970 driven by the aircraft industry (a high shock industry) during the Vietnam War. After the diffusion of CNC technology begins in the late 1970s, the two series begin to diverge, with high shock industries engaging in more capital investment.

In a task-based model, automation can have positive or negative effects on overall labor demand. On the one hand, a smaller share of the tasks will be allocated to labor as some tasks are automated and shifted to machines. On the other hand, automation can increase productivity, thereby increasing overall employment, and can lead to the creation of new labor-intensive tasks.

There is no obvious time-series relationship between CNC diffusion and total employment. Total employment in high and low shock industries trend together from 1960

to 1980, although high shock industries display a larger Vietnam-era amplitude. After the diffusion of CNC technology and the associated capital investments, high shock industries seem, if anything, to be gaining employment in the early 1980s and then losing employment after 1990. Low shock industries instead experience two periods of falling employment, one from 1980 to 1990 and the other from 2000 to 2010.

Panel B depicts the share of workers by education category for high and low shock industries. Note that the scales for the low shock industries are on the left axes and those for high shock industries are on the right axis to facilitate trend comparison. Following the diffusion of computerized machine tools, less-educated workers who operate machines by hand may lose work in favor of college graduates who can fix and install the new machinery. The figure shows that before the introduction of CNC tools, high school graduates represented 40% of the workforce in both high and low shock industries. From 1975 onward, the share of workers with a high school degree held steady in low shock industries and declined in high shock industries. Particularly after 1980, the high school share was replaced by a growing share of workers with a college degree in high shock industries. High shock industries also hired more high school dropouts from 1975 onward, reflecting a polarization of skill away from mid-skill level to both high- and low-skilled workers.

C CNC and economic outcomes residualized by industry and year effects

Figure 3.6 exploits the full array of industry-level variation in exposure to CNC tools over time, presenting a series of bin-scatter graphs relating each outcome $y_{j,t}$ to CNC exposure, adjusted for industry and year fixed effects. The slope of the graph depicts the coefficient β from the following regression:

$$y_{j,t} = \alpha_j + \gamma_t + \beta \text{Exposure CNC}_{j,t} + \varepsilon_{j,t} \tag{3.4}$$

CNC exposure rises by 40 to 55 percentage points over the 40 years in the data (see Figure 3.3). The binscatter residuals reveal that substantial variation in exposure remains even after the removal of industry and year fixed effects. The residuals have an eight-percentage point range, which is large compared to the overall rise in the CNC measure.

The figure again documents a positive relationship between changes in exposure to CNC technology and growing capital expenditures. There is a mild decline in total employment with exposure to CNC tools, but with substantial variation around the regression line. This negative relationship is hard to observe in the time series shown in Figure 3.5. The final four binscatters offer further evidence that exposure to CNC technology is associated with a shift away from mid-skill workers – including both high school graduates and workers with some college – in favor of both workers with a college degree and workers with less than a high school education.

3.5 Industry-level results: Economic outcomes and worker responses

Table **??** reports coefficients for a series of regressions relating CNC technology and economic outcomes, based on estimates of equation (4). These results are the analog of the bin-scatter graphs presented above, but include a wider set of outcomes. Row 1 contains outcomes from the Census of Manufacturers, and Row 2 presents results from CPS data.

A Capital expenditures, productivity and labor share

Column 1 (Row 1) presents the positive correlation between exposure to CNC and the logarithm of capital expenditures observed in Figure 3.6. A 10-percentage point increase in

CNC exposure, which is the approximate difference between the aircraft and motor vehicles industry in 1990, corresponds to a 29% increase in annual capital expenditures.

Labor productivity should rise with automation as labor is reallocated away from tasks for which it does not have a strong comparative advantage and as the new technology augments labor in the tasks that remain. We measure labor productivity as the logarithm of value-added, controlling for the logarithms of production and non-production workers (results are similar if we instead use the logarithm of value added per worker as an outcome). A 10-percentage point increase in exposure to CNC tools corresponding to an 20% increase in labor productivity (Row 1, Column 2).

The labor share will fall as firms reallocate tasks from labor to capital (displacement) but will rise if a sufficiently large number of new tasks are created (reinstatement). We find that the labor share – measured as wage bill divided by value added – is falling with industry exposure to CNC technology. A 10-percentage point increase in CNC exposure is associated with a 1.6 decline in the share of revenue paid out to labor (Row 1, Column 3).

B Total employment and employment shares by education level

We next turn to the relationship between CNC diffusion and employment. Columns 4-6 (Row 1) show results for total employment, production employment, and non-production employment, respectively. A 10-percentage point increase in CNC exposure is associated with a 4% decline in total employment (not significant) and a 7% decline in production workers.

Row 2 of Table ?? reports a series of outcomes calculated from the Current Population Survey. We start in Columns 1-4 with the share of workers by education category. A 10-percentage point increase in CNC exposure is associated with declining share of employment from two mid-skill groups, including a 9 percentage point decline in the share of workers with a high school degree and a 2 percentage point decline in the share of

workers with some college. The mid-skill share is replaced by a growing share of workers with less than a high school degree (4 percentage points) and workers with a college degree (7 percentage points).

C Task composition of the workforce

The loss of mid-skill workers is associated with a shift in the manufacturing production process away from routine tasks toward abstract and cognitive tasks (Appendix Figure 3A.4).[11] Columns 5-7 (Row 2) document that a 10-percentage point increase in CNC exposure is associated with 0.17 increase in average abstract task ratings, which can account for roughly 20% of the rise in the abstract content of manufacturing employment.

D Worker response through union protections

Workers in the union sector may be less affected by the CNC shock, either because unions blocked firms from adopting this new technology, or because unions bargained with firms to share a portion of the technology rents in the form of job protections. Table ?? reports the relationship between CNC exposure and total employment in the union sector and non-union sector based on worker counts in the CPS. That is, we estimate two versions of equation 4 with total employment as the outcome, one for workers who are members of a union and one for workers who are not members of a union. We find that employment losses after the diffusion of CNC technology is concentrated among non-union jobs, which decline by 7.5% for a 10 percentage point rise in CNC exposure. There is no relationship between CNC exposure and employment in the union sector.

We suspect that unions were able to protect workers against employment losses by bargaining for job protections, rather than by blocking new technology. Following the

[11] Note that occupations can require both routine and abstract work and that the three task measures are not constrained to sum to a fixed number, and so the increase in abstract tasks alongside the decline in routine work is not a mechanical effect.

landmark "Treaty of Detroit" between the United Auto Workers and General Motors in 1950, industrial unions in the United States upheld a shared norm to support technical change in exchange for employment protections and other benefits (Levy, 2021, p. 472-475; Reuther, 1962; Barnard, 1983; Brown, 1997). For example, when General Motors planned to upgrade its Linden, New Jersey plant for the use of computer numerical controlled machine tools in 1984, the union negotiated buyouts and job guarantees but did not oppose the retooling effort (Milkman, Ruth, 1997).

E Educational attainment and returning to school

Another possible action that workers in exposed industries may have taken to insulate themselves from the CNC shock was to go back to school to improve their skills. Table ?? reports estimates of equation 4 using an indicator for returning to school as the outcome.

We find that workers are more likely to re-enroll in higher education as CNC technology diffused through their industry. A 10-percentage point increase in CNC exposure is associated with a 2.4 percentage point increase in total enrollment (column 1), which is derived primarily from increases in full-time enrollment in four-year degree programs (columns 2-5). The average enrollment in higher education in this sample is around 5%.

To support this finding, Figure 3.7 graphs a time series of completed college degrees and certificates related to computerized machine tools from 1967 to the present. The series is based on administrative data from Higher Education General Information Survey (HEGIS, 1967-84) and IPEDS (Integrated Post-Secondary Information Data System (1986-present). We classified a subset of degrees as related to CNC technology – including programs like Machinist/Machine Tool Technologist and Industrial/Manufacturing Engineering (full list in footnote).[12]

[12]Programs coded as CNC-related: programs coded as CNC programs: Automation Engineer Technology/Technician; Computer Numerically Controlled (CNC) Machinist Technology/CNC Machinist; Electromechanical Engineering; Electromechanical Tech./Technician; Industrial/ Manufacturing Engineer-

CNC-related degree programs grew rapidly in two bursts: first, from 1970 to 1980 and then from 2010 and 2020. The first period corresponds to the advent of CNC machine tools, and the second period corresponds to the spread of industrial robots. General degree completions follow a different pattern, growing steadily from 1985 to 2010. Appendix Figure 3A.5 reports time series in degree completion by degree type (certificates, associate degrees and bachelor's degrees).

3.6 Local labor market results

Industries that were more exposed to CNC technology may have experienced other correlated shocks that cannot be fully ruled out with industry-level analysis. By incorporating variation across more than 700 labor markets, we can take advantage of more continuous and granular measures of CNC exposure. This section compares employment outcomes in local labor markets that were differentially exposed to the CNC automation shock in the 1970s and 1980s. Furthermore, estimating the effect of an automation shock on a local labor market allows us to explore another margin of worker adjustment: switching between the affected industries (metal manufacturing) and other sources of employment.

We measure exposure to the CNC shock at the local level for 722 commuting zones by weighting CNC exposure in industry j and year t by the initial share of employment in industry j. We impose that CNC exposure is zero in any non-metal manufacturing industry and calculate exposure in local market m in year t as:

$$\text{Exposure CNC}_{m,t} = \sum_j \left(\frac{\text{EMP}_{m,j,1970}}{\text{EMP}_m, 1970} \text{Exposure CNC}_{j,t} \right) \tag{3.5}$$

where $\text{EMP}_{m,j,1970}$ is a count of workers employed in industry j in 1970 and $\text{EMP}_{m,1970}$ is a count of all workers in market m in 1970. By 2000, exposure to CNC tools varied from

ing; Industrial/Manufacturing Tech./Technician; Machinist/Machine Technologist/Assistant; Mechatronics, Robotics, and Automation Engineering; Robotics Tech./Technician.

0.2% to 14.5% across commuting zones. A portion of this cross-market variation is driven by the size of the metal manufacturing sector. The residual of local CNC exposure after controlling for the metal manufacturing share of employment in 1970 still has a sizeable range (2 percentage points).

We stack data from Census decades (1970-2000) and estimate the relationship between employment counts or education shares ($y_{m,t}$) and market-level exposure to CNC technology.

$$y_{m,t} = \alpha_m + \gamma_t + \beta \text{Exposure CNC}_{m,t} + X'\Gamma + \varepsilon_{m,t} \tag{3.6}$$

We consider this relationship both for workers in metal manufacturing and for the full manufacturing sector. All specifications control for commuting zone (α_m) and decade fixed effects (γ_t). Our preferred specification includes controls for initial 1-digit industry employment shares in a commuting zone interacted with decade to allow for industry-specific employment trends. Our 1-digit classification sub-divides initial employment share in metal and non-metal manufacturing. In this case, we identify β by comparing commuting zones with the same initial share of employment in metal manufacturing but a different initial composition of employment within the metal sector (e.g., a higher share of aircraft employment and a lower share of motor vehicle employment). Our vector of controls (X) also includes a series of initial demographic measures interacted with decade.[13]

Figure 3.8 presents two sets of results at the local labor market level. Panel A presents the relationship between exposure to CNC technology and employment outcomes in the metal industry, while Panel B instead considers the relationship with employment in the broader manufacturing sector. At the commuting zone level, one standard deviation

[13]Interacted controls are log population, education shares (high school dropout, high school graduate, some college, and college graduate or above), race/ethnicity shares (White, Black, Asian, Hispanic), share female, and share over 65.

difference in CNC exposure is 3 percentage points. A one standard deviation increase in exposure to CNC is associated with a 3.5 point decline in the share of the population employed in metal manufacturing, but only a 1 point (not statistically significant) decline in the share of the population employed in the manufacturing sector more broadly, suggesting that workers were able to shift from heavily-affected metal sub-industries to other segments of the manufacturing sector. The employment of men in the metal manufacturing industry is more sensitive to the shock than the employment of women.

There is some evidence of a shift away from mid-skill level workers, but only for workers with some college. A one standard deviation increase in exposure to CNC is associated with a 6 percentage point decline in the share of workers with some college. If anything, this mid-skill share is made up by employment of low-skilled workers, but the estimates are quite noisy in general, given that they rely on subdividing commuting zone population by industry and education level.

3.7 Conclusions

Compared to the past, the modern factory floor is filled with machines and empty of people. Computerized machine tools produce complex parts based on instructions encoded in computer programs, conveyors move parts from station to station, and robots assemble the parts into finished products. Jobs in the factory increasingly require a sophisticated understanding of the programming of machines and often a college degree.

This paper studies one important step in the long process of factory automation: the advent and diffusion of automated machine tools. Automated machine tools began to diffuse widely in the 1970s. At mid-century, machine tools required a semi-skilled machinist to perform operations to specification by hand. New CNC tools replaced these routine operations with detailed computer programs overseen by skilled workers.

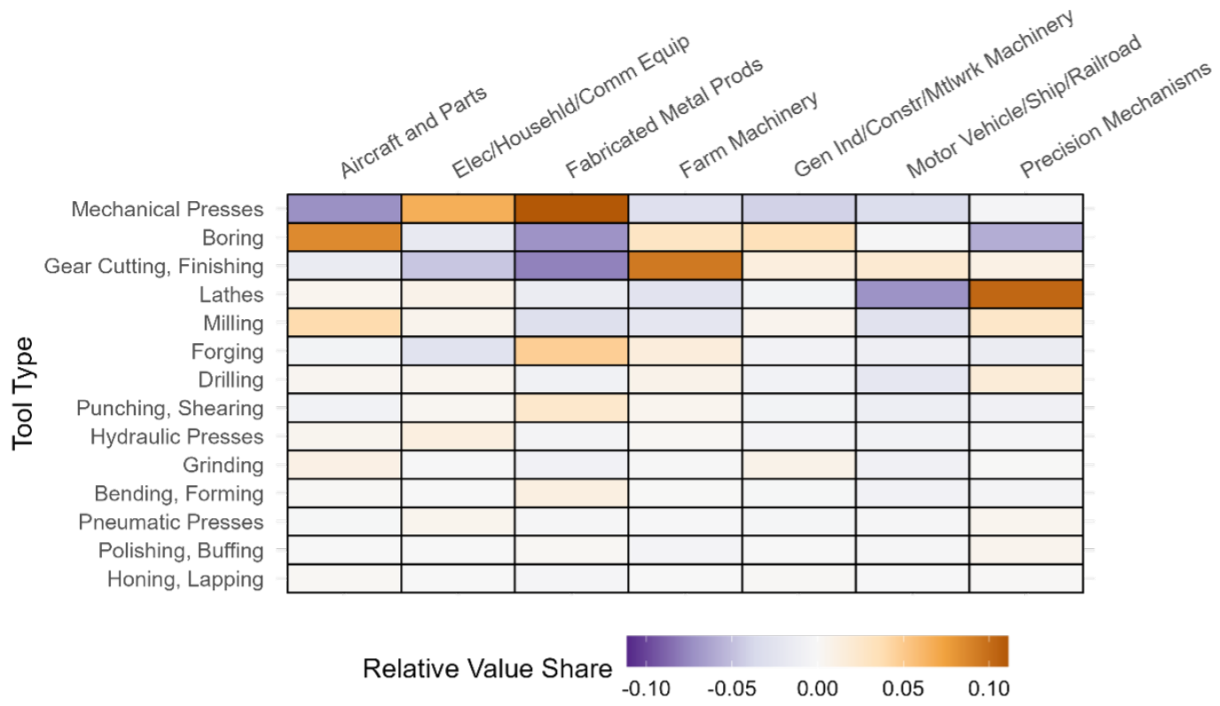
We find that metal manufacturing industries that were more exposed to CNC tools experienced rising labor productivity but some decline in total employment. Labor demand shifted from mid-skill workers (high school graduates and workers with some college) to college graduates. Workers in exposed industries responded to this technology shock in three ways: by shifting to other parts of the manufacturing sector, by returning to school, perhaps to qualify for higher-skilled work, and by enjoying some protections of union status if members of a union.

At least since the initial development of automated machine tools, scholars have speculated that worker protections would constrain the adoption of productivity-enhancing automation (Killingsworth and Heckman, 1986). In countries or settings with stronger worker protections, the spread of industrial robots is not associated with falling employment, and can even lead employment to rise (Adachi et al., 2020; Dauth et al., 2021). Similarly, we document fewer employment losses for low-skilled workers and stronger overall employment growth in the union sector following the spread of CNC technology in the US.

Future work could explore how firms in settings with stronger worker protections respond to automation shocks. For example, are unionized firms in the United States slower to adopt automation technologies, or do they adopt technology while offering job protection to their existing workforce, thus sharing more of the gains from new technology with workers? Do firms react to new technology by shifting production from union to non-union establishments? A firm-level analysis of the diffusion of automation technologies could shed light on these questions and help guide the adjustment process in the future.

Tables and figures

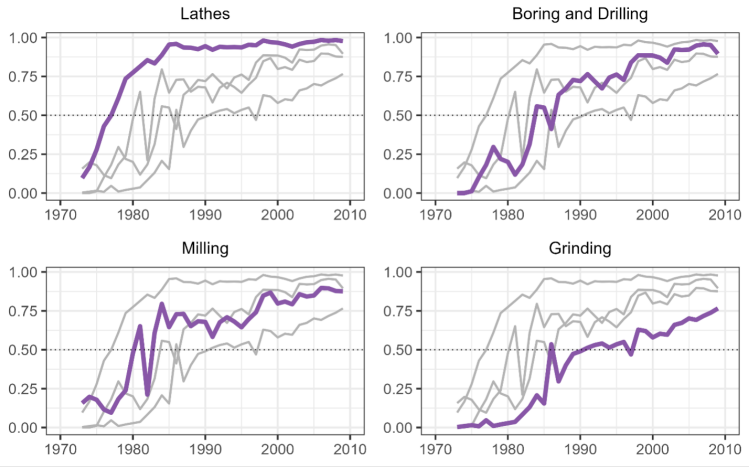
Figure 3.1: Relative Value Share of Installed Tools by Seven Metal Manufacturing Industries in 1958



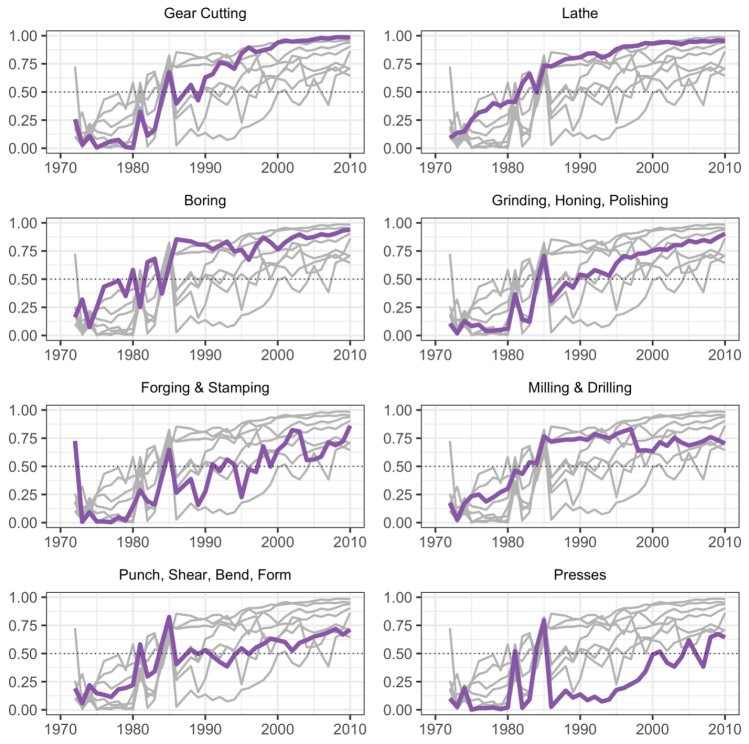
Notes: Cells in the figure show the relative value share of for each type of machine tool by industry. Orange indicates that an industry is more intensive in that tool type than the average industry, while purple indicates it is less intensive than the average. To compute the relative value share, we first compute the value share of installed tool types for each of the seven metal manufacturing industries. We then subtract the mean across industries for each tool type. Data on tool value come from the 1958 American Machinist Inventory of Metalworking Equipment.

Figure 3.2: CNC Share by Machine Tool Type

Panel A: Japan

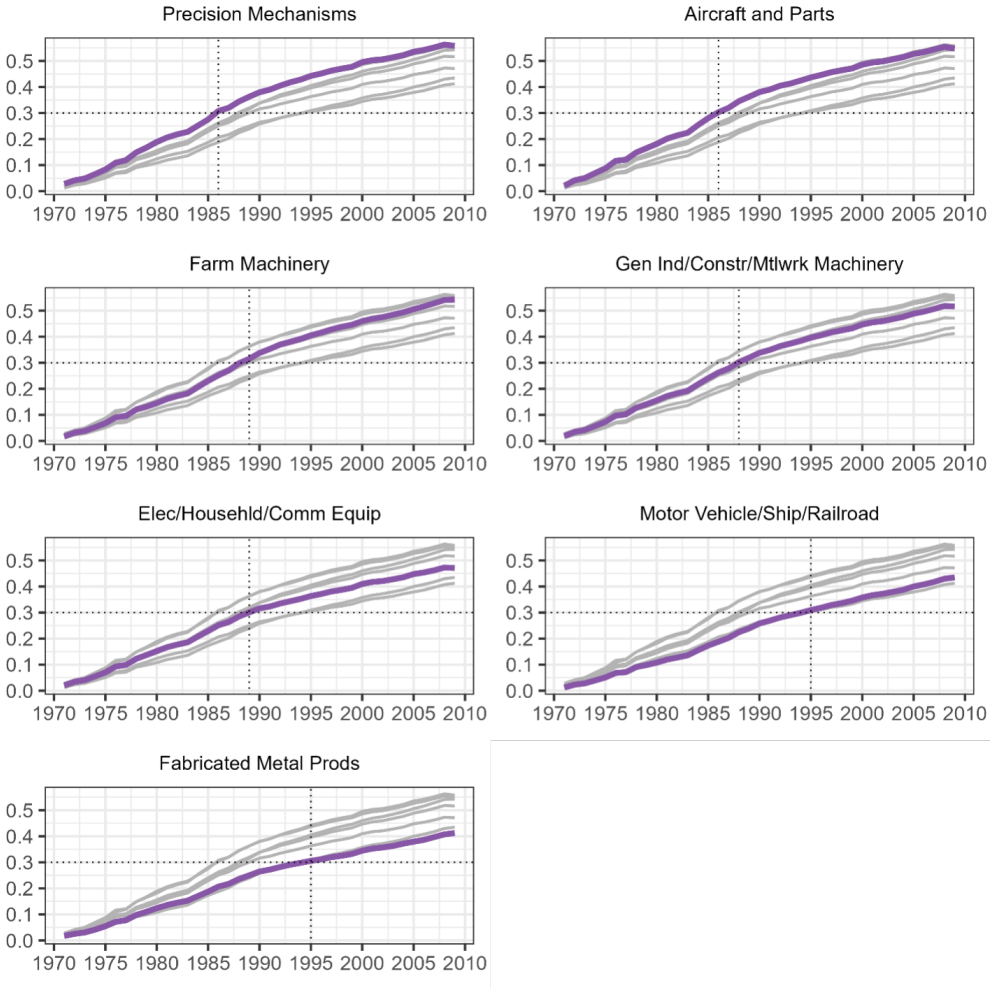


Panel B: Germany



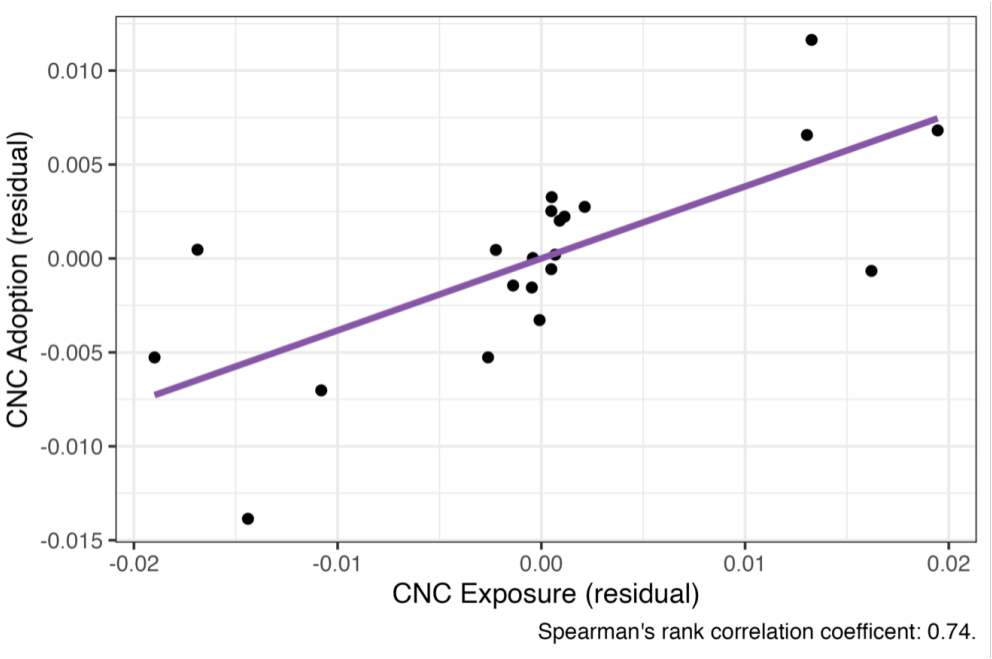
Notes: This figure presents the annual CNC shares of exports by machine tool types for both Japan and Germany. See Appendix Figure 3A.3 for Italy. The data come from volumes of the Economic Handbook of the Machine Tool Industry as described in Section 3.3.

Figure 3.3: Cumulative CNC Share by Metal Manufacturing Industries



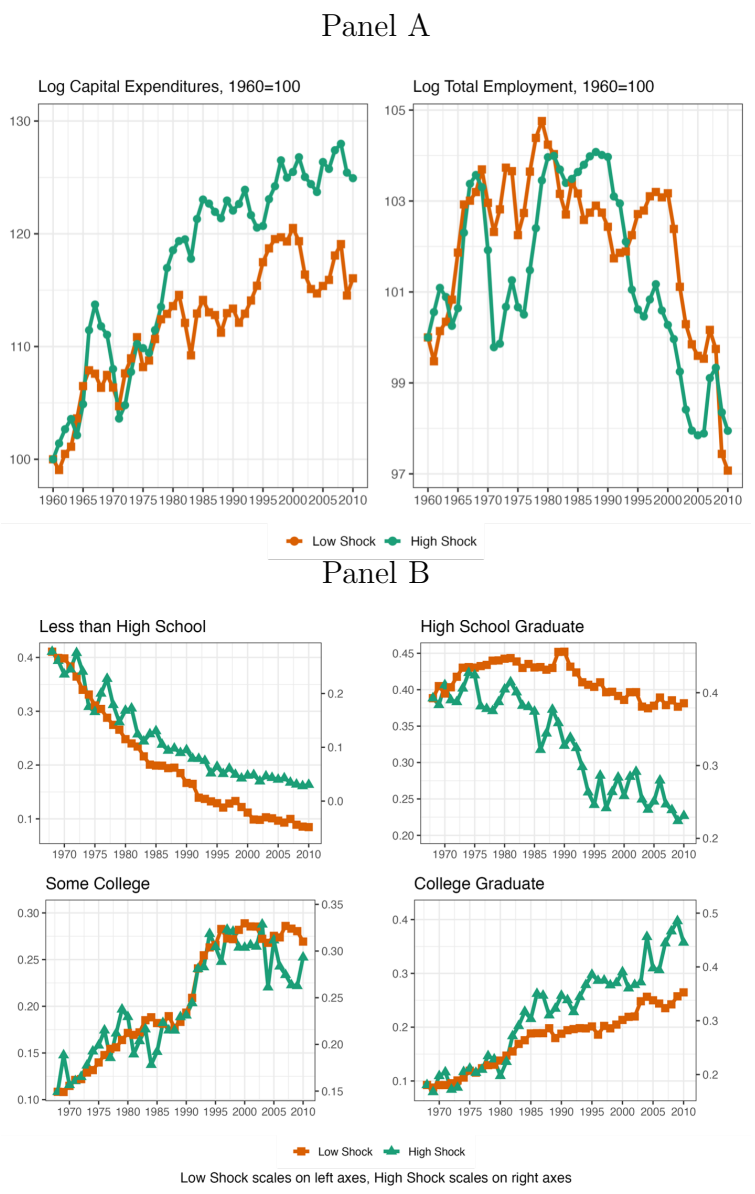
Notes: Cumulative CNC shares for the seven metal manufacturing are computed from the initial tool distribution by industry and the CNC share of exports by tool as described in Section 3.3. Dotted lines indicate the year in which the cumulative share passed 30%. Data underlying the figure come from the 1958 American Machinist Inventory of Metalworking Equipment and volumes of the Economic Handbook of the Machine Tool Industry.

Figure 3.4: CNC Adoption and CNC Exposure



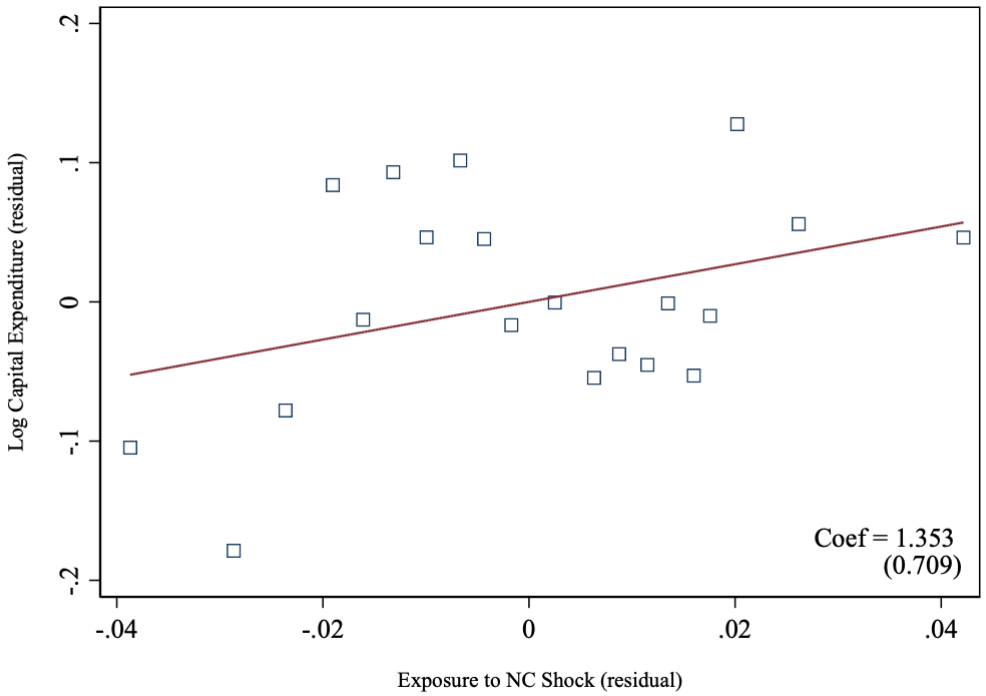
Notes: The figure plots residuals from regressing our CNC exposure measure and CNC adoption measure on industry and year fixed effects. The CNC adoption measure is available for 1973, 1978, and 1983 and we use the values of the exposure measure in those years to construct the figure.

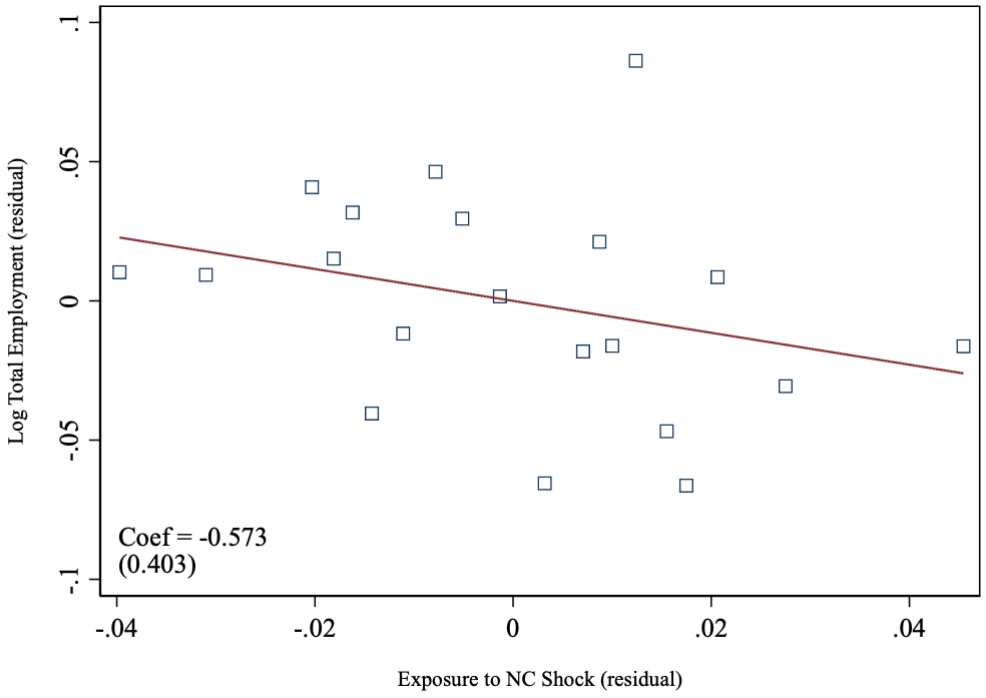
Figure 3.5: CNC Share by Machine Tool Type

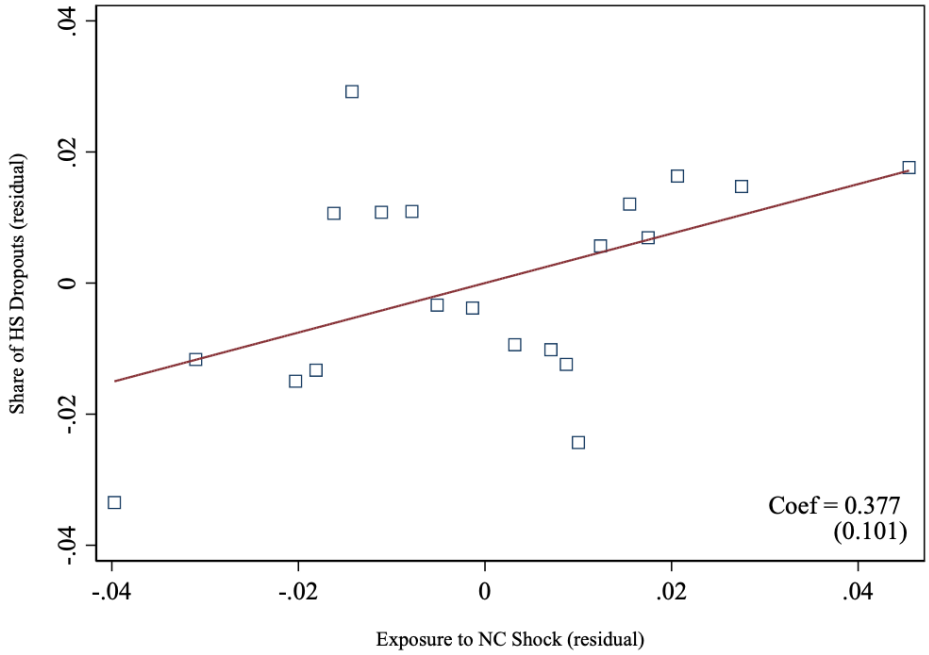


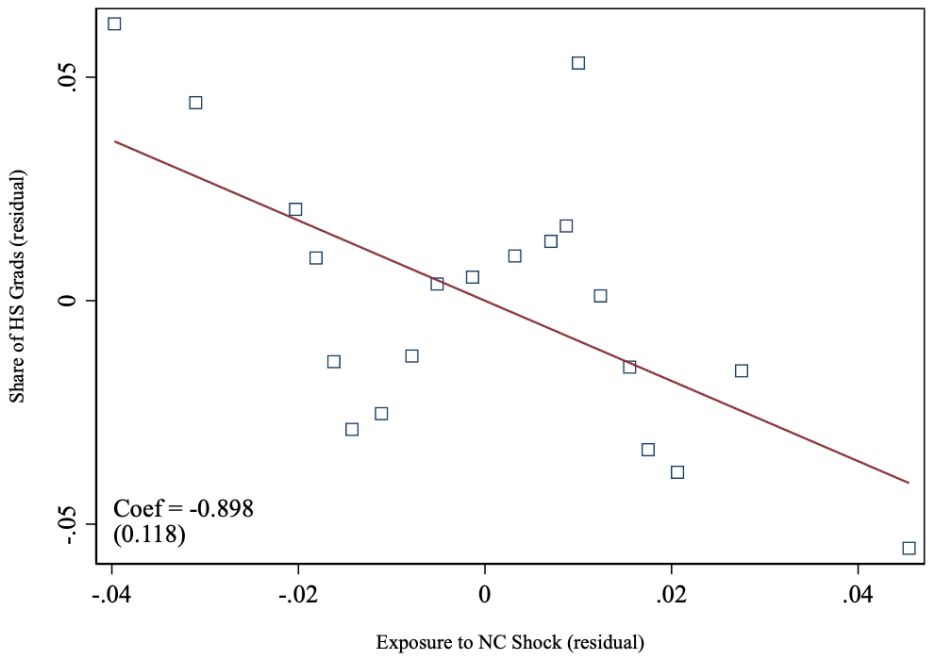
Notes: Panel A shows the logarithm of average annual capital expenditures and the logarithm of total employment in metal manufacturing from the NBER-CES Manufacturing Industry database. Panel B shows the share of metal manufacturing workers in four education categories from the CPS Annual Social and Economic Supplement. Metal industries are partitioned into two groups according to the intensity of the CNC shock they experience. Low-shock industries are motor vehicles and fabricated metals; electronics, and general industrial equipment. High shock industries are farm equipment, aircraft and parts, and precision mechanisms.

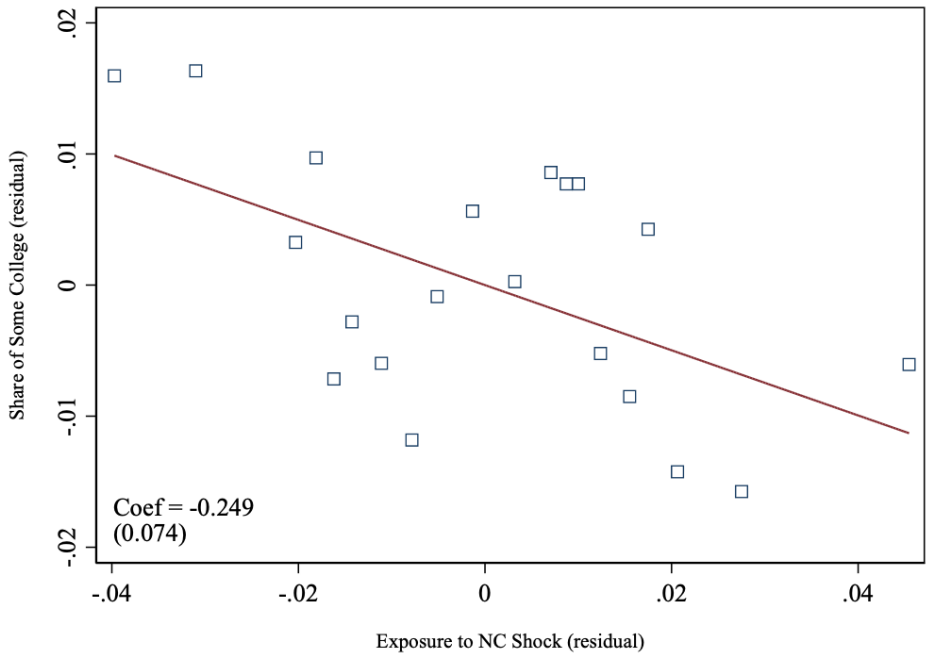
Figure 3.6: Core Economic Outcomes and CNC Shock Residualized Binscatters

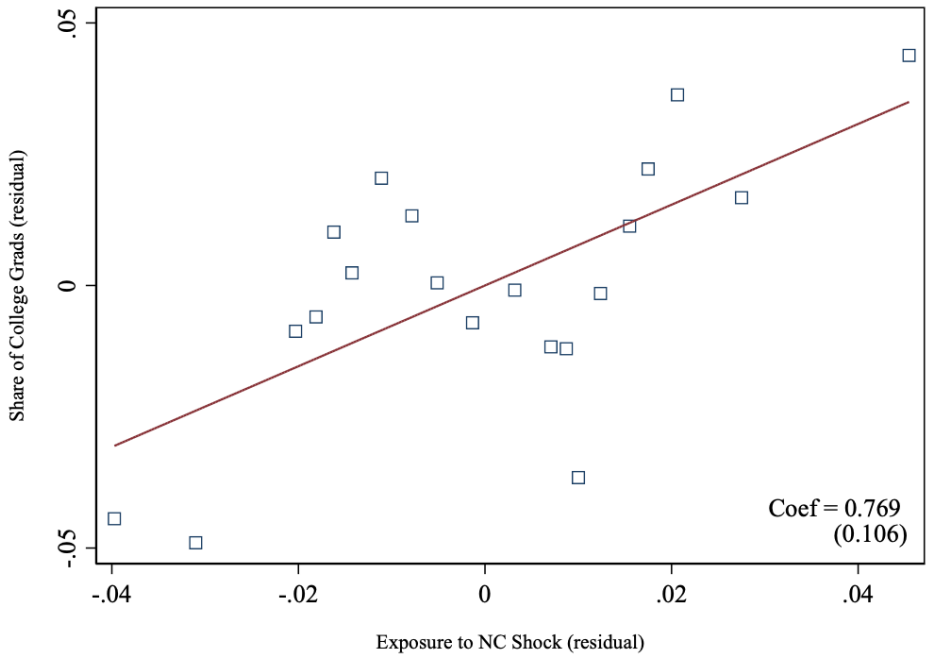





Share of HS Dropouts (residual)
.04
.02
0
-.02
-.04
-.04
-.02
0
.02
.04
Exposure to NC Shock (residual)
Coef = 0.377
(0.101)

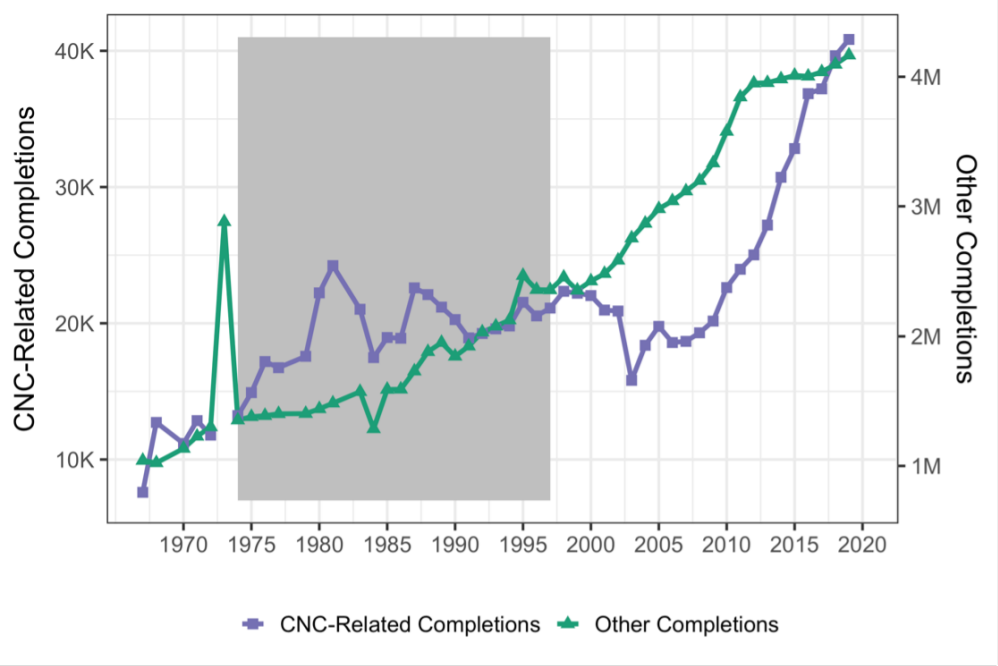

Share of HS Grads (residual)
.05
0
-.05
-.04
-.02
0
.02
.04
Exposure to NC Shock (residual)
Coef = -0.898
(0.118)





Notes: The figure shows binscatters of our CNC exposure measure and six economic outcomes residualized by year and industry fixed effects. Outcomes are the logarithms of average annual capital expenditures and total employment in metal manufacturing from the NBER-CES Manufacturing Industry database and the share of metal manufacturing workers in four education categories from the CPS Annual Social and Economic Supplement.

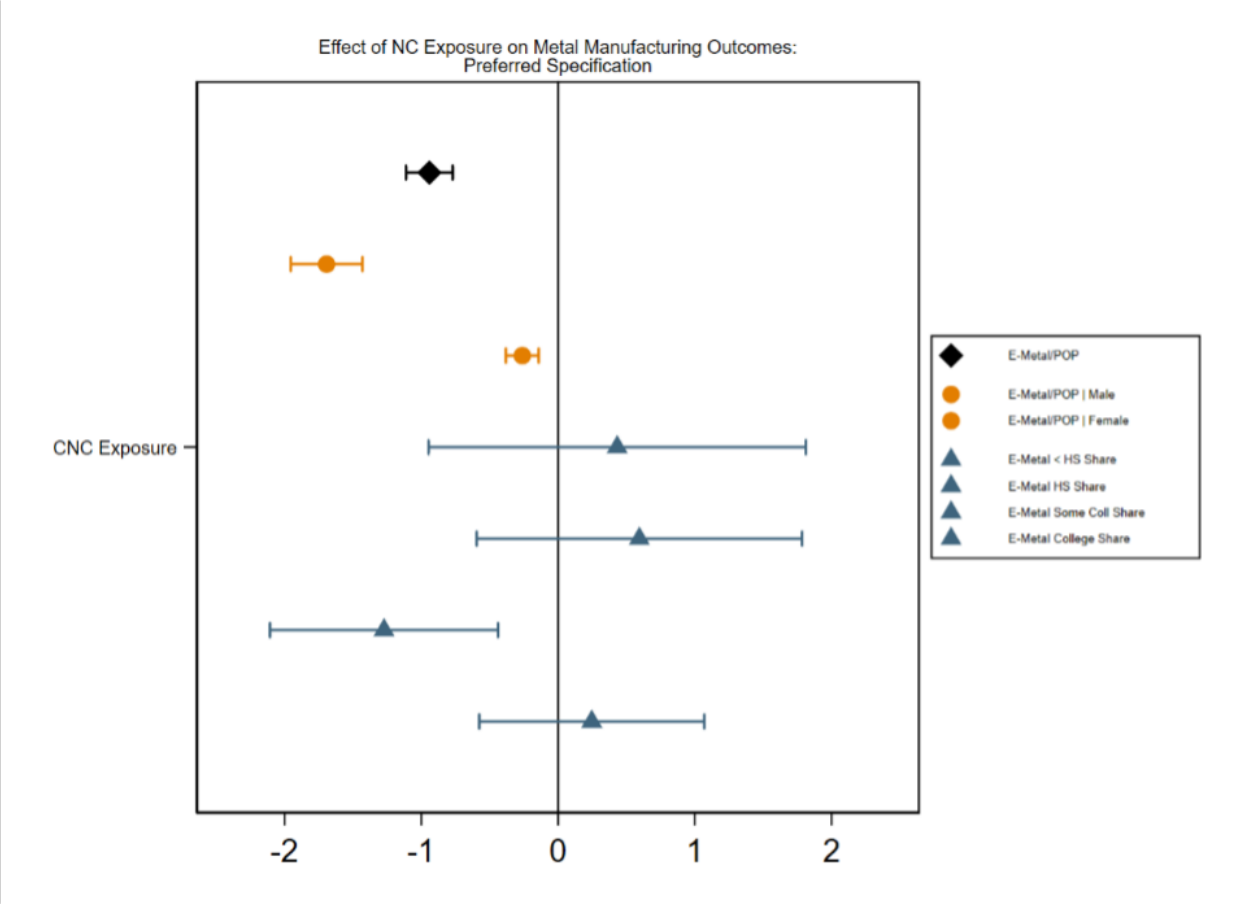
Figure 3.7: CNC-Related Degree and Program Completions



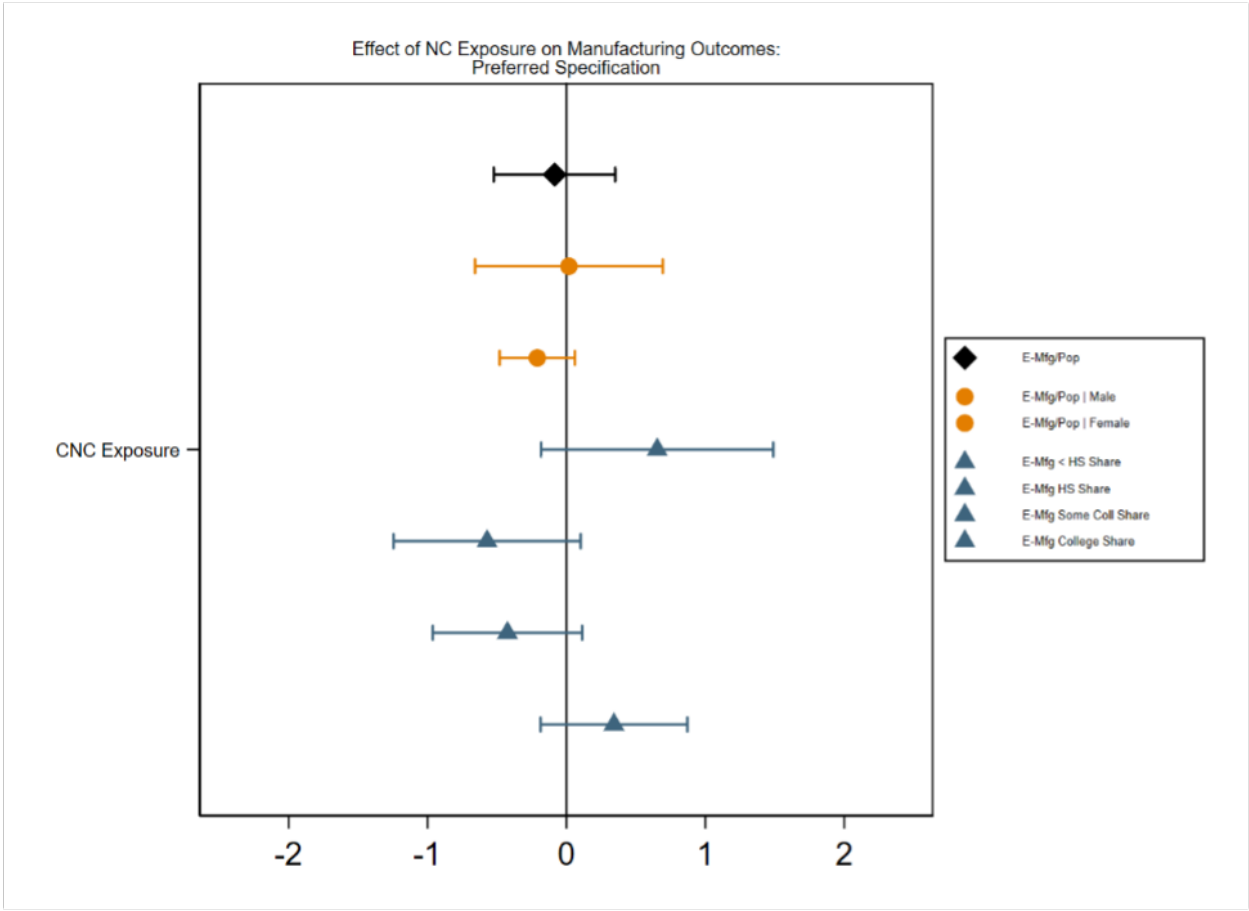
Notes: The figure shows the number of completed degrees and program in US higher educational institutions. Degrees and programs are categorized by whether their subject matter is related to CNC. Data come from the HEGIS and IPEDS databases. The gray box shows the period 1974-1997 during which the CNC shock was most intense.

Figure 3.8: CNC Exposure and Local Labor Markets

Panel A: Metal Manufacturing



Panel B: All Manufacturing



Notes: Figure shows estimates of the effect of the CNC shock on labor market outcomes at the commuting zone level based on the 1970, 1980, and 1990 Censuses. Panel A shows estimates for the metal manufacturing industries and panel B shows estimates for the entire manufacturing sector.

Table 3.1: The effect of CNC exposure on industry-level economic outcomes, 1968-2007

Row 1: Outcomes from the Census of Manufactures							
	log(Capital Exp.)	log(Value Added)	Labor Share	log(Total Emp)	log(Prod. Emp)	log(Nonprod. Emp)	
	(1)	(2)	(3)	(4)	(5)	(6)	
CNC Exposure	1.3576*	0.8916***	-0.1616**	-0.6383	-1.2019***	-0.5114	
	(0.7433)	(0.2743)	(0.0771)	(0.4100)	(0.3815)	(0.4556)	
Emp. controls	X	X					
Dep. var mean	7.8139	10.6721	0.4144	6.5796	6.1623	5.4373	
Row 2: Outcomes from the Current Population Survey							
	Share of HS Dropouts	Share of HS Grads	Share of Some Coll.	Share of Coll. Grads	Abstract	Routine	Manual
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
CNC Exposure	0.3721***	-0.8647***	-0.2174***	0.7100***	1.6772***	-0.6668	0.8396***
	(0.1042)	(0.1197)	(0.0713)	(0.1017)	(0.5619)	(0.4306)	(0.1924)
Dep. var mean	0.1742	0.3943	0.2161	0.2154	3.4032	5.3800	1.0973

Table 3.2: The effect of CNC exposure on employment by union status, sample 1973-2007

Data source:	CPS May Supplement	
	log(Non-union)	log(Union)
	(1)	(2)
CNC Exposure	-1.6135**	0.0138
	(0.6298)	(0.8401)
Industry FE	X	X
Year FE	X	X
Dep. var mean	6.3553	5.1461
N	231	231

Note: Outcome variables are log of non-union and union workforce in each metal manufacturing industry x year computed from the CPS May Supplement sample (1973-2007). All specifications include industry and year-fixed effects. Standard errors are robust.
*** = significant at 1%, ** = significant at 5%, * = significant at 10%

Table 3.3: The effect of CNC exposure on the two- and four-year college enrollment, 1976-2007

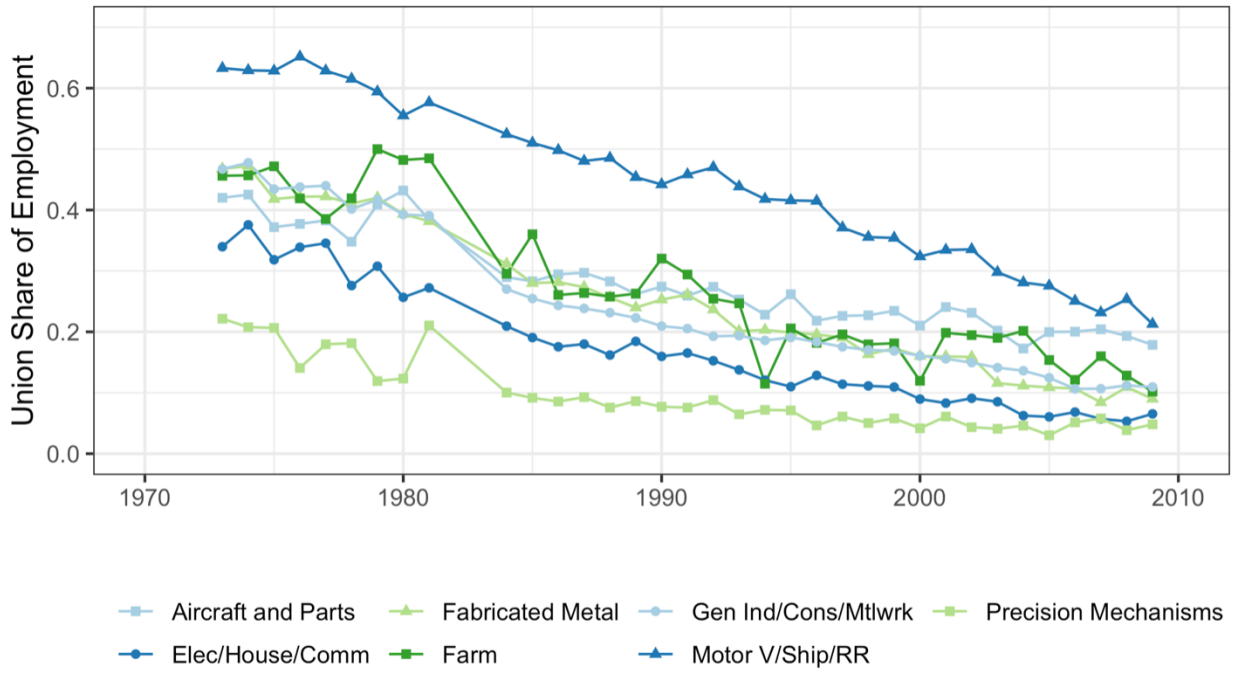
	All enrollment	2-yr, full-time	2-yr, part-time	4-yr, full-time	4-yr, part-time
	(1)	(2)	(3)	(4)	(5)
CNC Exposure	0.2502**	0.0439	0.0554	0.1267***	-0.0494
	(0.0970)	(0.0296)	(0.0468)	(0.0436)	(0.0556)
Industry FE	X	X	X	X	X
Year FE	X	X	X	X	X
Dep. var mean	0.0539	0.0047	0.0167	0.0120	0.0170
N	224	224	224	224	224

Note: Outcome variables are the annual share of two-year or four-year college enrollment among each metal manufacturing industry workforce constructed from the CPS October Supplement sample (1976-2007). We restrict our sample to prime-age men who do not hold bachelor's degree and whose current or last employment were in metal manufacturing industry. All specifications include industry and year fixed effects. Standard errors are robust. *** = significant at 1%, ** = significant at 5%, * = significant at 10%

Appendices

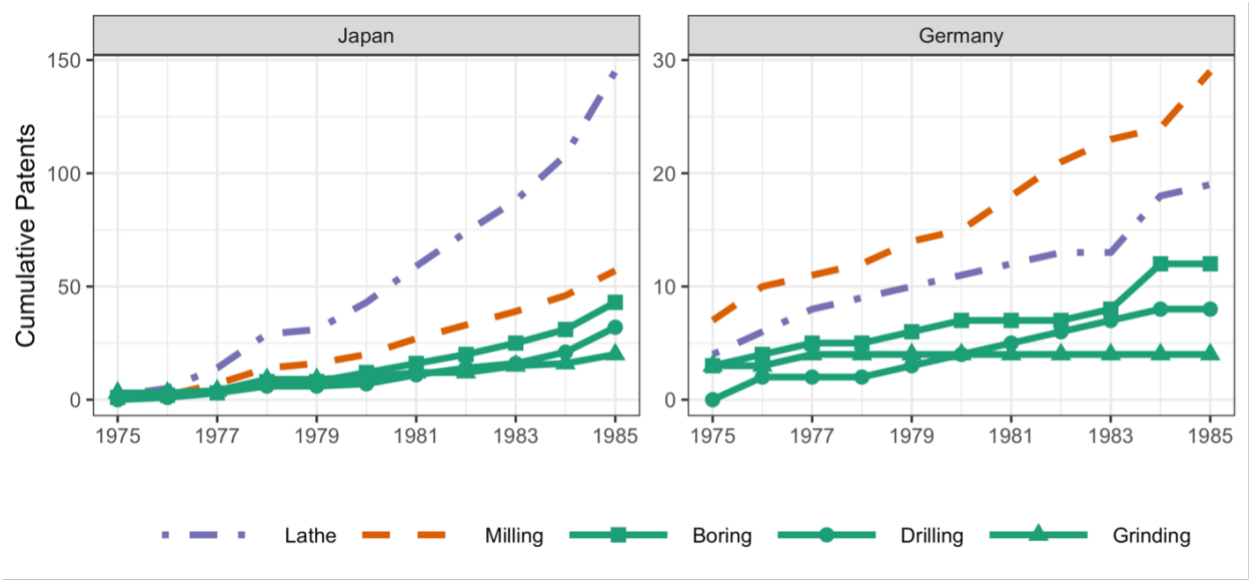
A Appendix Figures and Tables

Figure 3A.1: Union Share by Seven Metal Manufacturing Industries



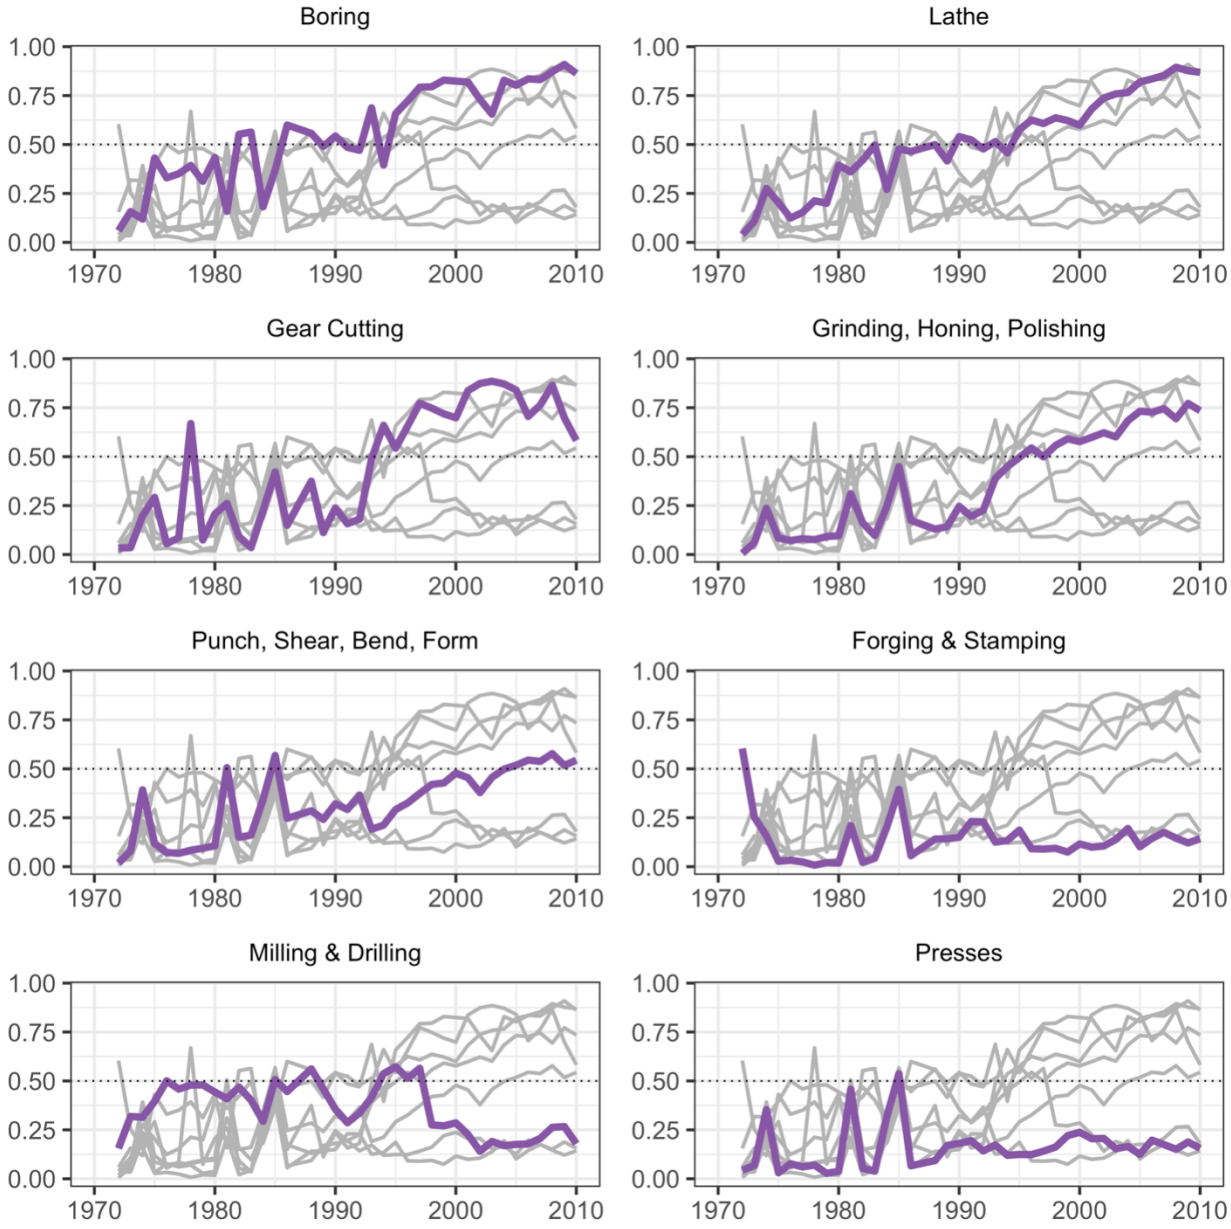
Notes: Average share of workers in each of the seven metal manufacturing industries belonging to a union. The share is computed from the CPS May Supplement (1973-1983) and CPS Outgoing Rotation Groups (1984-2009) sample.

Figure 3A.2: CNC-Related Patent Counts by Tool and Country



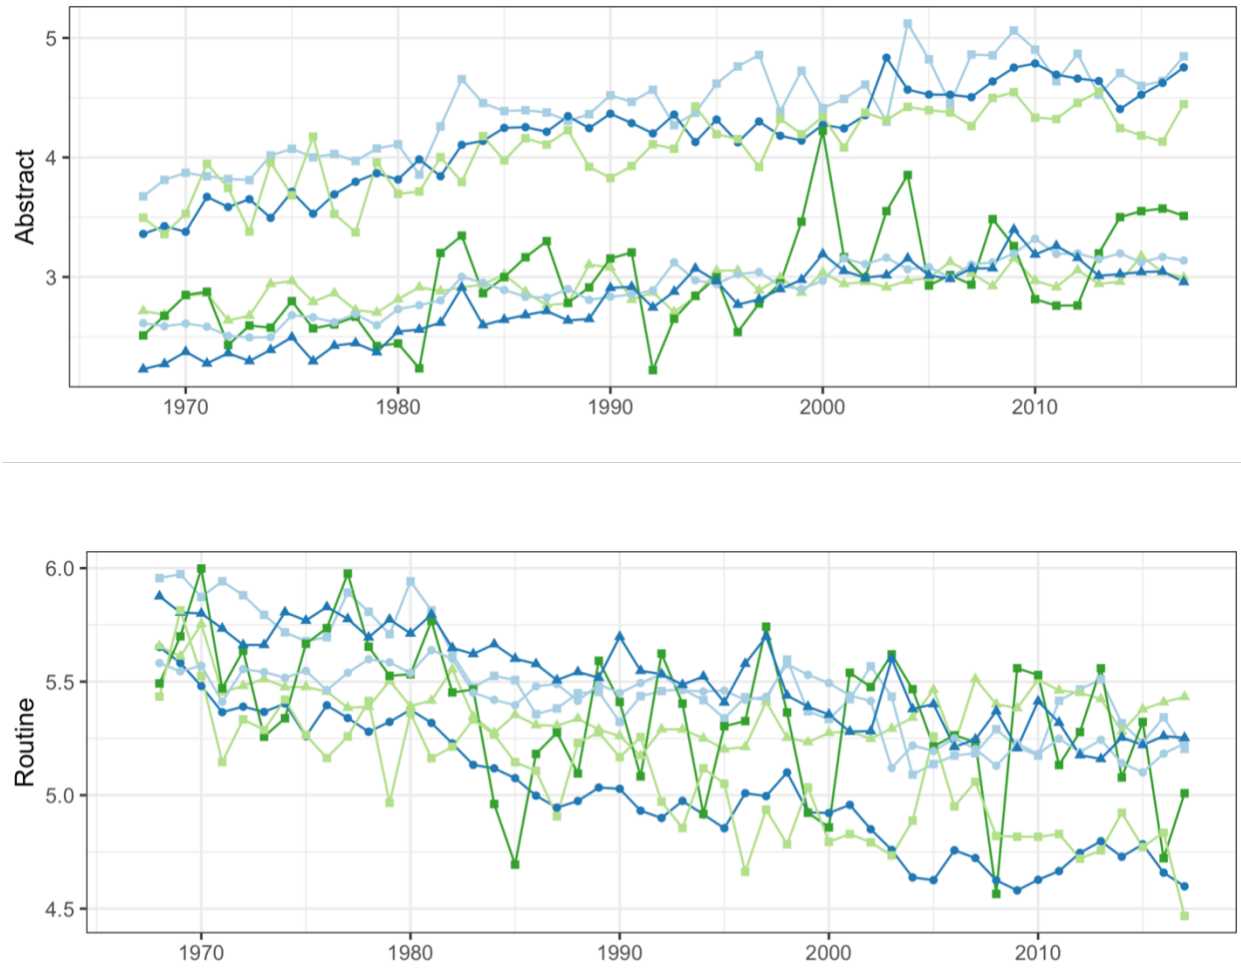
Notes: We began with a list of all US utility patents created for Kogan et al (2017) to scrape the digitized text from Google Patent. It is common for innovations created outside the United States to be patented in the US. Data from Lai et al (2011) provides the country of the grantee for each patent. We link Cooperative Patent Classification codes to the US patents. The patent subclass G05B contains classifications related to CNC. In particular, the subgroup G05B 19/18 covers "numerical control [NC], i.e., automatically operating machines, in particular machine tools..." We searched the patent text in our dataset to identify NC-related patents that are associated with five tool types: lathes, milling machines, drilling machines, boring machines, and grinding machines. We used the functions of the machines as the search terms. Our final figure examines patenting by Germany and Japan between 1975 and 1985. For those countries and period, there were 2,467 patents in G05B. The five tool types were mentioned in 52% of the patents.

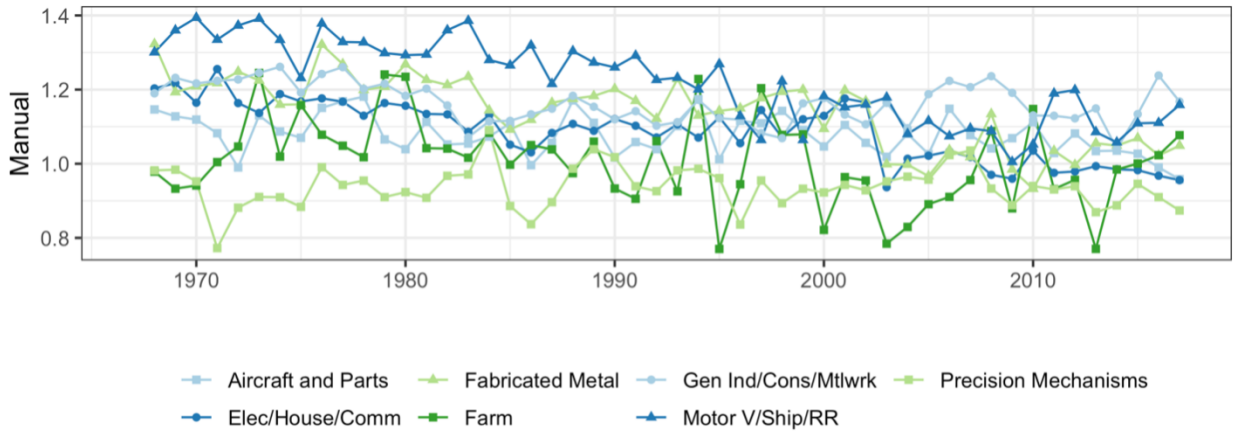
Figure 3A.3: CNC Share by Machine Tool Type for Italy



Notes: This figure presents the annual CNC shares of exports by tool type for Italy. See Figure 1 in the paper for Japan and Germany. The data come from volumes of the Economic Handbook of the Machine Tool Industry as described in Section 3.3.

Figure 3A.4: Average Task Scores of Occupations by Seven Metal Manufacturing Industries

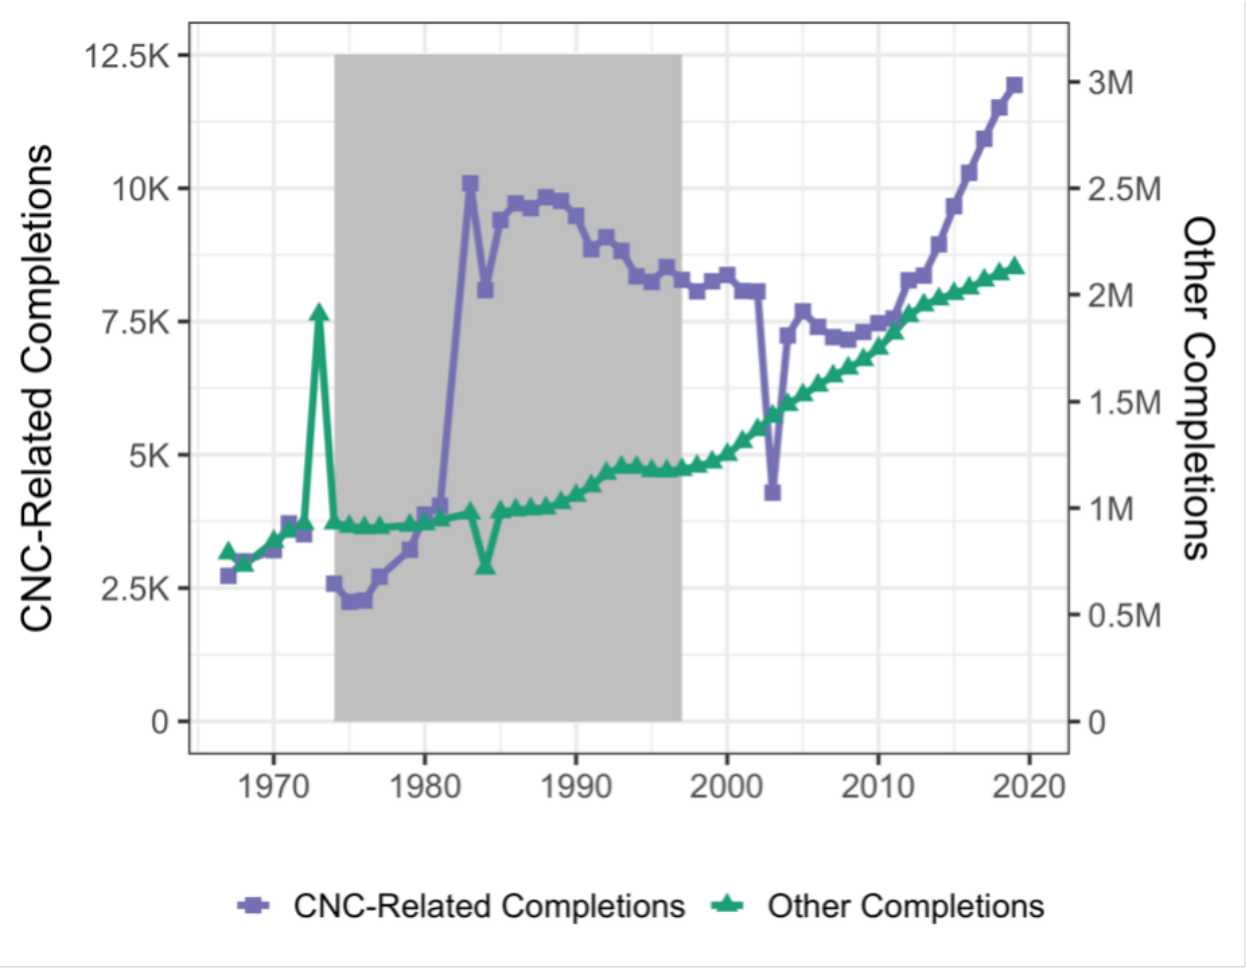




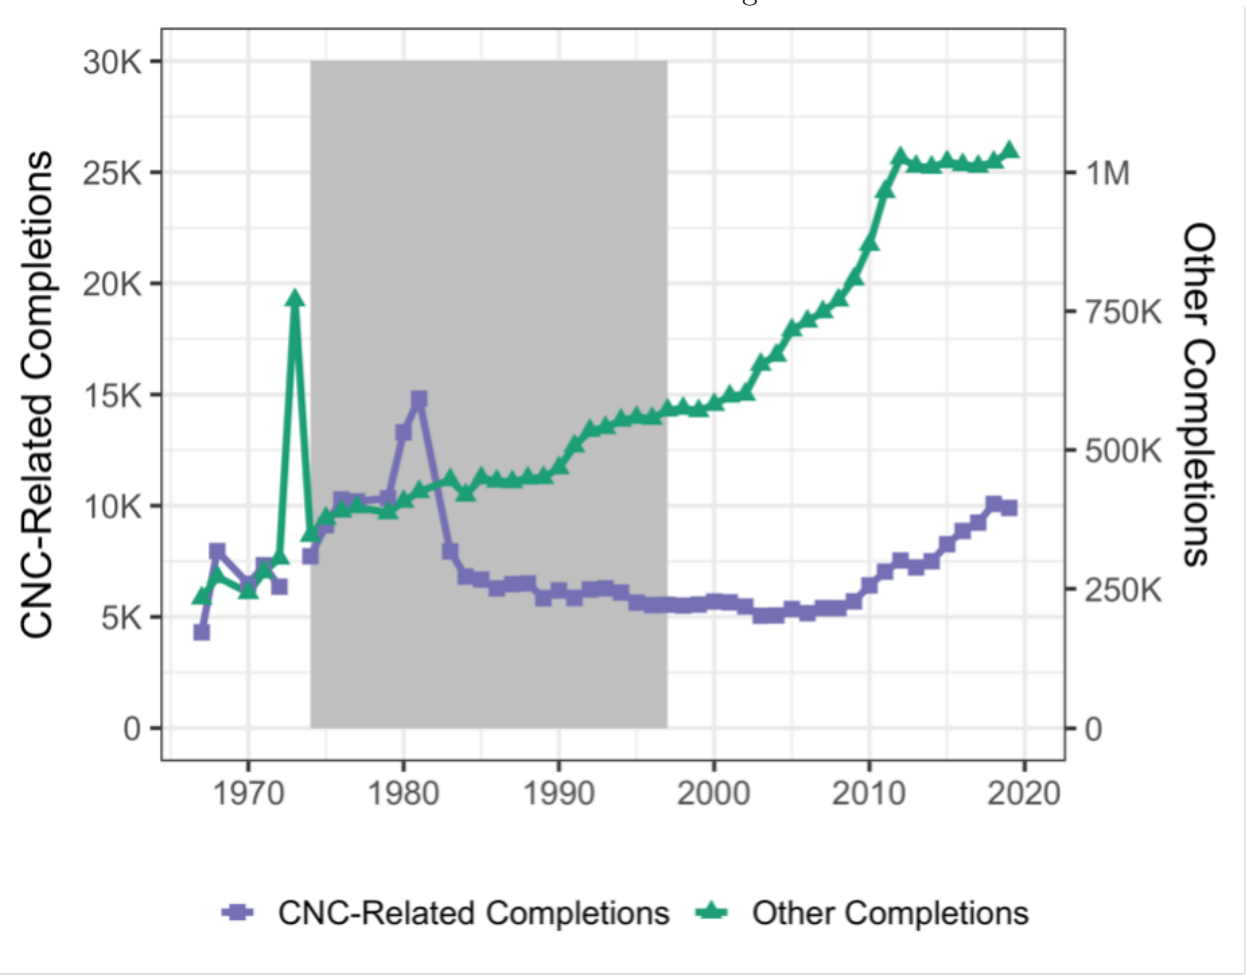
Notes: Task score for occupation are from Autor, Levy, and Murnane (2003). Workers in metal manufacturing in the CPS Annual Social and Economic Supplement are assigned the task scores corresponding to their occupations. We then compute averages for each of the seven metal manufacturing industries by year.

Figure 3A.5: CNC-Related Degree and Program Completions by Degree Type

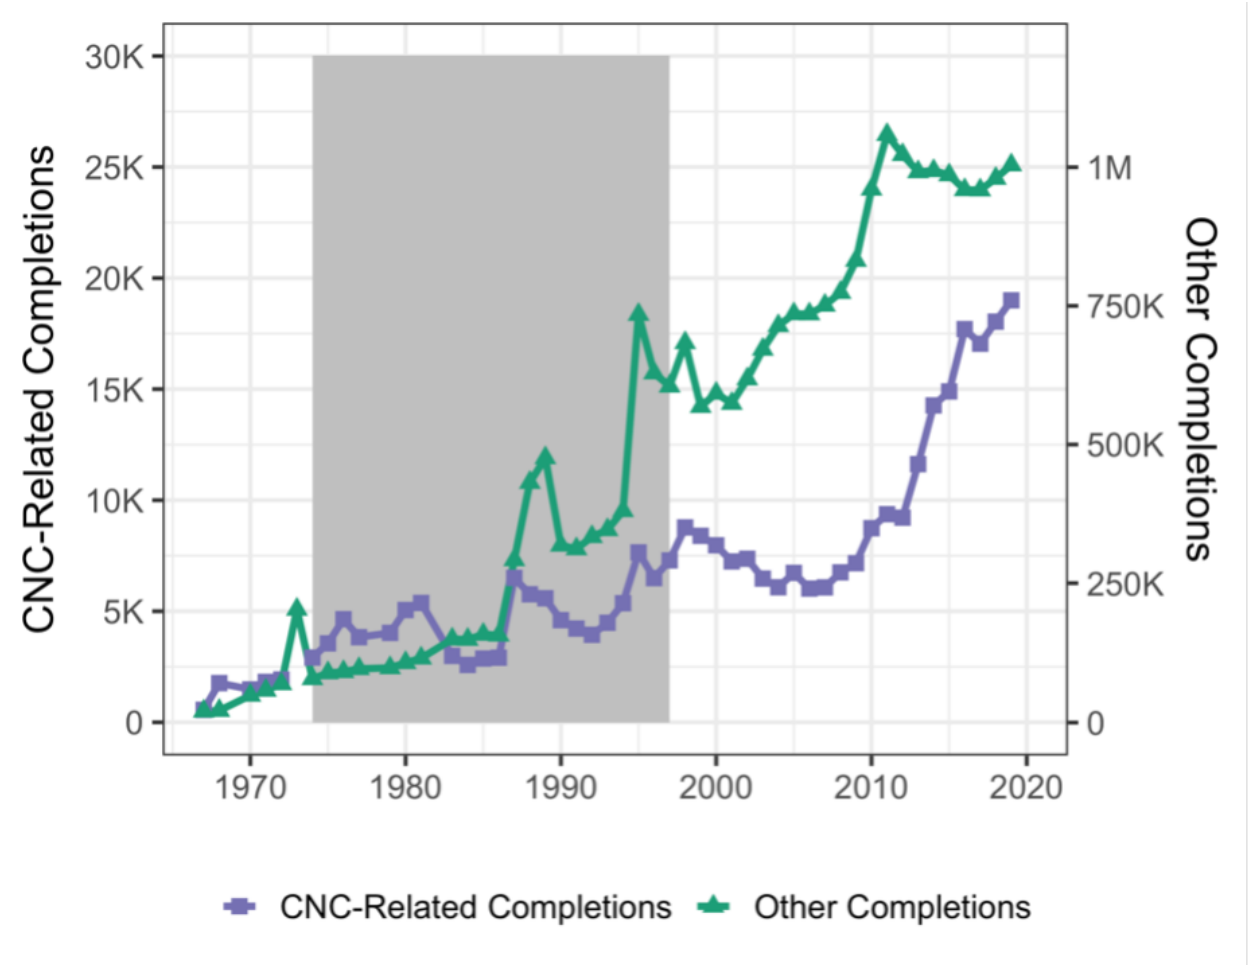
Panel A: Bachelor's Degrees



Panel B: Associate Degrees

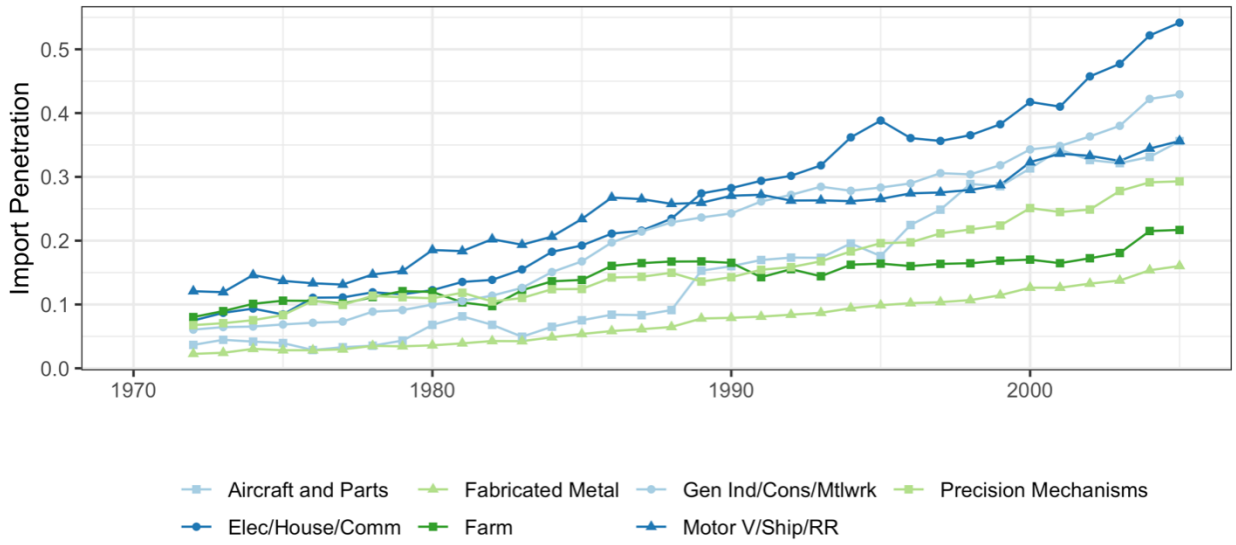


Panel C: Certificate Programs



Notes: The figure shows the number of completed degrees and program in US higher educational institutions by type of degree or program. Degrees and programs are categorized by whether their subject matter is related to CNC. Data come from the HEGIS and IPEDS databases. The gray box shows the period 1974-1997 during which the CNC shock was most intense.

Figure 3A.6: Import Penetration by Seven Metal Manufacturing Industries



Notes: Import penetration computed using trade data from Schott (2008) and the definitions found in Campbell and Lusher (2019). Import penetration is the share of imports in US domestic sales in each industry.

Bibliography

Abramitzky, Ran, Leah Boustan, Katherine Eriksson, Santiago Pérez, and Myera Rashid, "Census Linking Project: Version 2.0," *[Dataset]*, 2022.

Acemoglu, Daron and Pascual Restrepo, "Automation and new tasks: How technology displaces and reinstates labor," *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.

_ **and** _ , "Robots and jobs: Evidence from US labor markets," *Journal of Political Economy*, 2020, *128* (6), 2188–2244.

Adachi, Daisuke, Daiji Kawaguchi, and Yukiko Saito, "Robots and Employment: Evidence from Japan, 1978-2017," *Research Institute of Economy, Trade and Industry (RIETI)*, 2020, *Discussion papers 20051.*

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-share designs: Theory and inference," *The Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel, "What Are the Labor and Product Market Effects of Automation?: New Evidence from France," Sciences Po publications info:hdl:2441/170cd4sul89, Sciences Po July 2020.

Alder, Simeon, David Lagakos, and Lee Ohanian, "Competitive Pressure and the Decline of the Rust Belt: A Macroeconomic Analysis," Working Paper 20538, National Bureau of Economic Research October 2014.

Antràs, Pol, Teresa C. Fort, and Felix Tintelnot, "The Margins of Global Sourcing: Theory and Evidence from US Firms," *American Economic Review*, September 2017, *107* (9), 2514–64.

Atack, Jeremy, Robert A. Margo, and Paul W. Rhode, ""Automation" of Manufacturing in the Late Nineteenth Century: The Hand and Machine Labor Study," *Journal of Economic Perspectives*, May 2019, *33* (2), 51–70.

Atkin, David, "Endogenous Skill Acquisition and Export Manufacturing in Mexico," *American Economic Review*, August 2016, *106* (8), 2046–85.

Autor, David H and Mark G Duggan, "The rise in the disability rolls and the decline in unemployment," *The Quarterly Journal of Economics*, 2003, pp. 157–205.

Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, October 2013, *103* (6), 2121–68.

_ **, Frank Levy, and Richard J. Murnane**, "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

_ **, Lawrence F. Katz, and Alan B. Krueger**, "Computing Inequality: Have Computers Changed the Labor Market?," *The Quarterly Journal of Economics*, 1998, *113* (4), 1169–1213.

Barnard, John, *Walter Reuther and the Rise of the Auto Workers*, New York: Little, Brown, and Co., 1983.

Barro, Robert J. and Xavier Sala-i-Martin, "Convergence," *Journal of Political Economy*, 1992, *100* (2), 223–251.

Bartel, Ann, Casey Ichniowski, and Kathryn Shaw, "How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills," *The Quarterly Journal of Economics*, 2007, *122* (4), 1721–1758.

Bartelsman, Eric J and Wayne Gray, "The NBER Manufacturing Productivity Database," Working Paper 205, National Bureau of Economic Research October 1996.

Batistich, Mary Kate and Timothy N. Bond, "Stalled Racial Progress and Japanese Trade in the 1970s and 1980s," *The Review of Economic Studies*, 2022.

Berkes, Enrico, Ezra Karger, and Peter Nencka, "The census place project: A method for geolocating unstructured place names," *Explorations in Economic History*, 2022, p. 101477.

Bessen, James E., Maarten Goos, Anna Salomons, and Wiljan Van den Berge, "Automatic Reaction - What Happens to Workers at Firms that Automate?," Boston Univ. School of Law, Law and Economics Research Paper, Boston Univ. School of Law 2019.

Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge, "Firm-Level Automation: Evidence from the Netherlands," *AEA Papers and Proceedings*, May 2020, *110*, 389–93.

Black, Dan, Terra McKinnish, and Seth Sanders, "The Economic Impact Of The Coal Boom And Bust," *The Economic Journal*, 2005, *115* (503), 449–476.

Blanchard, Olivier and Lawrence Katz, "Regional Evolutions," *Brookings Papers on Economic Activity*, 1992, *23* (1), 1–76.

Blanchard, Olivier Jean, Lawrence F Katz, Robert E Hall, and Barry Eichengreen, "Regional evolutions," *Brookings papers on economic activity*, 1992, *1992* (1), 1–75.

Boone, Christopher D. A. and Laurence Wilse-Samson, "Structural Change and Internal Labor Migration: Evidence from the Great Depression," *The Review of Economics and Statistics*, 10 2021, pp. 1–54.

Borusyak, Kirill, Rafael Dix-Carneiro, and Brian Kovak, "Understanding migration responses to local shocks," 2022.

Bound, John and Harry J. Holzer, "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s," *Journal of Labor Economics*, 2000, *18* (1), 20–54.

Brown, Michael K., "Bargaining for Social Rights: Unions and the Reemergence of Welfare Capitalism, 1945-1952," *Political Science Quarterly*, 1997, *112* (4), 645–674.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 11 2014, *82* (1), 1–44.

Card, David, "The causal effect of education on earnings," in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, 1 ed., Vol. 3, Part A, Elsevier, 1999, chapter 30, pp. 1801–1863.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo, "Housing Booms and Busts, Labor Market Opportunities, and College Attendance," *American Economic Review*, October 2018, *108* (10), 2947–94.

Che, Yi, Yi Lu, Justin R. Pierce, Peter K. Schott, and Zhigang Tao, "Did trade liberalization with China influence US elections?," *Journal of International Economics*, 2022, *139*, 103652.

Chiron, Claire, "Influences of Quotas, Tariffs and Bilateral Trade agreement on Post 2005 Apparel Trade," Technical Report, Harvard Center for Apparel and Textile Research 2004.

Choi, Jiwon, Ilyana Kuziemko, Ebonya L Washington, and Gavin Wright, "Local Economic and Political Effects of Trade Deals: Evidence from NAFTA," Working Paper 29525, National Bureau of Economic Research November 2021.

Coen-Pirani, Daniele, "GEOGRAPHIC MOBILITY AND REDISTRIBUTION," *International Economic Review*, 2021, *62* (3), 921–952.

Colantone, Italo and Piero Stanig, "Global competition and Brexit," *American political science review*, 2018, *112* (2), 201–218.

Collis, David J., "The Machine Tool Industry and Industrial Policy, 1955-82," in Spence, Hazard, ed., *Trade Policies for International Competitiveness*, Pensacola, FL: Ballinger Publishing Co., 1988.

Dao, Mai, Davide Furceri, and Prakash Loungani, "Regional Labor Market Adjustment in the United States: Trend and Cycle," *The Review of Economics and Statistics*, 05 2017, *99* (2), 243–257.

Daron, Lelarge Claire Acemoglu and Pascual Restrepo, "Competing with robots: Firm-level evidence from France," *AEA Papers and Proceedings*, 2020, *110*, 383–88.

Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner, "The adjustment of labor markets to robots," *Journal of the European Economic Association*, 2021, *19* (6), 3104–3153.

Devine, Warren D, "From shafts to wires: Historical perspective on electrification," *The Journal of Economic History*, 1983, *43* (2), 347–372.

Diamond, Rebecca, "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, March 2016, *106* (3), 479–524.

Dippel, Christian, Robert Gold, and Stephan Heblich, "Globalization and its (dis-) content: Trade shocks and voting behavior," Technical Report, National Bureau of Economic Research 2015.

Dix-Carneiro, Rafael and Brian K. Kovak, "Trade Liberalization and Regional Dynamics," *American Economic Review*, October 2017, *107* (10), 2908–46.

Doms, Mark, Timothy Dunne, and Kenneth Troske, "Workers, Wages, and Technology," *The Quarterly Journal of Economics*, 1997, *112* (1), 253–290.

Dorn, David, Gordon Hanson, Kaveh Majlesi et al., "Importing political polarization? The electoral consequences of rising trade exposure," *American Economic Review*, 2020, *110* (10), 3139–83.

Eckert, Fabian, Teresa C. Fort, Peter K. Schott, and Natalie J. Yang, "Imputing missing values in the US Census Bureau's County Business Patterns.," 2021. NBER Working Paper 26632.

English, Beth Anne, "A common thread: Labor, politics, and capital mobility in the Massachusetts textile industry, 1880-1934," Dissertations, Theses, and Masters Projects 1539623415, The College of William and Mary 2003.

Eriksson, Katherine, Katheryn N. Russ, Jay C. Shambaugh, and Minfei Xu, "Trade shocks and the shifting landscape of U.S. manufacturing," *Journal of International Money and Finance*, 2021, *111*, 102254.

Farber, Henry S, Daniel Herbst, Ilyana Kuziemko, and Suresh Naidu, "Unions and Inequality over the Twentieth Century: New Evidence from Survey Data*," *The Quarterly Journal of Economics*, 04 2021, *136* (3), 1325–1385.

Feenstra, Robert C, *Offshoring in the global economy: microeconomic structure and macroeconomic implications*, MIT Press, 2010.

Feenstra, Robert C. and Gordon H. Hanson, "Aggregation Bias in the Factor Content of Trade: Evidence from U.S. Manufacturing," *American Economic Review*, May 2000, *90* (2), 155–160.

Feenstra, Robert, John Romalis, and Peter K. Schott, "U.S. Imports, Exports and Tariff Data, 1989-2001," NBER Working Paper 9387. Available at http://www.johnromalis.com/publications/ 2002.

Feigenbaum, James, Alexander Hertel-Fernandez, and Vanessa Williamson, "From the bargaining table to the ballot box: Political effects of right to work laws," Technical Report, National Bureau of Economic Research 2018.

_ **and Daniel P. Gross**, "Automation and the Future of Young Workers: Evidence from Telephone Operation in the Early 20th Century," Technical Report w28061, NBER Working Paper 2020.

Feigenbaum, James J. and Hui Ren Tan, "The Return to Education in the Mid-Twentieth Century: Evidence from Twins," *The Journal of Economic History*, 2020, *80* (4), 1101–1142.

Fishback, Price V., William C. Horrace, and Shawn Kantor, "Did New Deal Grant Programs Stimulate Local Economies? A Study of Federal Grants and Retail Sales during the Great Depression," *The Journal of Economic History*, 2005, *65* (1), 36–71.

Fiszbein, Martin, Jeanne Lafortune, Ethan G. Lewis, and José Tessada, "Powering Up Productivity: The Effects of Electrification on U.S. Manufacturing," Technical Report w28076, NBER Working Paper 2020.

Fort, Teresa C., Justin R. Pierce, and Peter K. Schott, "New Perspectives on the Decline of US Manufacturing Employment," *Journal of Economic Perspectives*, May 2018, *32* (2), 47–72.

Frank, Thomas, *What's the Matter with Kansas?*, Metropolitan Books, 2004.

Gaggl, Paul, Rowena Gray, Ioana Marinescu, and Miguel Morine, "Does electricity drive structural transformation? Evidence from the United States," *Labour Economics*, 2021, *68*.

Gethin, Amory, Clara Martinez-Toledano, and Thomas Piketty, "Brahmin left versus merchant right: Changing political cleavages in 21 Western Democracies, 1948–2020," *The Quarterly Journal of Economics*, 2022, *137* (1), 1–48.

Goldin, Claudia and Lawrence F. Katz, "The origins of technology-skill complementarity," *The Quarterly Journal of Economics*, 1998, *113* (3), 693–732.

_ and _ , *The Race Between Education and Technology.*, Belknap Press for Harvard University Press, 2008.

_ and Lawrence F Katz, *The race between education and technology*, harvard university press, 2008.

Graetz, Georg and Guy Michaels, "Robots at work," *Review of Economics and Statistics*, 2018, *100* (5), 753–768.

Gray, Rowena, "Taking technology to task: The skill content of technological change in early twentieth century United States," *Explorations in Economic History*, 2013, *50* (3), 351–367.

Greenland, Andrew, John Lopresti, and Peter McHenry, "Import competition and internal migration," *Review of Economics and Statistics*, 2019, *101* (1), 44–59.

_ , _ , and _ , "Import competition and internal migration," *Review of Economics and Statistics*, 2019, *101* (1), 44–59.

Griffith-Jones, Stephany, "Causes and lessons of the Mexican Peso crisis," in "Global Capital Flows," Springer, 1998, pp. 100–136.

Hakobyan, Shushanik and John McLaren, "Looking for local labor market effects of NAFTA," *Review of Economics and Statistics*, 2016, *98* (4), 728–741.

_ and _ , "Replication Data for: Looking for Local Labor Market Effects of NAFTA," Review of Economics and Statistics [publisher], Harvard Dataverse [distributor] https://doi.org/10.7910/DVN/MJAPQ9 2016.

Hershbein, Brad and Bryan Stuart, "The Evolution of Local Labor Markets after Recessions," Working Paper No. 22-16, FRB of Philadelphia November 2022.

Hobor, George, "Surviving the Era of Deindustrialization: The New Economic Geography of the Urban Rust Belt," *Journal of Urban Affairs*, 2013, *35* (4), 417–434.

Holt, L.T.C., *A Short History of Machine Tools*, Cambridge, MA: MIT Press, 1966.

Hounshell, David A., *From the American System to Mass Production: The Development of Manufacturing Technology in the United States*, Baltimore, MD: Johns Hopkins University Press, 1984.

Humlum, Anders, "Robot adoption and labor market dynamics," 2019. Working Paper, Princeton University.

Hyman, Benjamin G, "Can Displaced Labor Be Retrained? Evidence from Quasi-Random Assignment to Trade Adjustment Assistance," 2018.

Itohisa, Matasao, "Fanuc's Competitive Advantage and the Revolt of Machine Tool Builders: A Look at Fanuc's User and Its User's Customer Relationships from a Historical Perspective (1950s–1980s)," 2010. Manufacturing Management Research Center Discussion Paper No. 534.

Jaworski, Taylor, ""You're in the Army Now:" The Impact of World War II on Women's Education, Work, and Family," *The Journal of Economic History*, 2014, *74* (1), 169–195.

Johnson, Chalmers, *MITI and the Japanese Miracle: The Growth of Industrial Policy, 1925–1975*, Redwood City, CA: Stanford University Press, 1982.

Katz, Lawrence F. and Robert A. Margo, "Technical change and the relative demand for skilled labor: The United States in historical perspective," in "Human Capital in History: The American Record," University of Chicago Press, 2014, pp. 15–57.

Keefe, Jeffrey H, "Numerically controlled machine tools and worker skills," *ILR Review*, 1991, *44* (3), 503–519.

Khandelwal, Amit K., Peter K. Schott, and Shang-Jin Wei, "Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters," *American Economic Review*, October 2013, *103* (6), 2169–95.

Killingsworth, Mark and James Heckman, "Female Labor Supply: A Survey," in O. Ashenfelter and R. Layard, eds., *Handbook of Labor Economics*, Vol. 1, Elsevier Science Publishers BV, 1986, chapter 2, pp. 103–204.

Klinkner, Philip A, *Midterm: the elections of 1994 in context*, Routledge, 2019.

Koch, Michael, Ilya Manuylov, and Marcel Smolka, "Robots and firms," *Economic Journal*, 2021, *131* (638), 2553–2584.

Koistinen, David, "Dealing with Deindustrialization: Economics, Politics, and Policy During the Decline of the New England Textile Industry, 1920–1960," *The Journal of Economic History*, 2000, *60* (2), 501–504.

_ , "The Causes of Deindustrialization: The Migration of the Cotton Textile Industry from New England to the South," *Enterprise & Society*, 2002, *3* (3), 482–520.

_ , "Reform Politics in Hard Times: Battles Over Labor Legislation During the Decline of Traditional Manufacturing in Massachusetts, 1922-1928," *Historical Journal of Massachusetts*, 2007, *35* (1).

Kornacki, Steve, *The red and the blue: The 1990s and the birth of political tribalism*, HarperCollins, 2018.

Krueger, Alan B, Andreas Mueller, Steven J Davis, and Ayşegül Şahin, "Job search, emotional well-being, and job finding in a period of mass unemployment: Evidence from high frequency longitudinal data [with comments and discussion]," *Brookings Papers on Economic Activity*, 2011, pp. 1–81.

Krueger, Anne O., *Trade Creation and Trade Diversion Under NAFTA* number w7429, National Bureau of Economic Research Cambridge, MA, 1999.

Krugman, Paul R, "Technology, trade and factor prices," *Journal of international Economics*, 2000, *50* (1), 51–71.

Lemieux, Thomas, "The "Mincer Equation" Thirty Years After Schooling, Experience, and Earnings," in Shoshana Grossbard, ed., *Jacob Mincer A Pioneer of Modern Labor Economics*, Springer Books, Springer, November 2006, chapter 11, pp. 127–145.

Levy, Jonathan, *Ages of American Capitalism: A History of the United States*, New York: Random House, 2021.

Lewis, Joshua and Edson Severnini, "Short- and long-run impacts of rural electrification: Evidence from the historical rollout of the U.S. power grid," *Journal of Development Economics*, 2020, *143.*

Lleras-Muney, Adriana, "The Relationship Between Education and Adult Mortality in the United States," *The Review of Economic Studies*, 01 2005, *72* (1), 189–221.

López-Córdova, J. Ernesto, "NAFTA and Manuacturing Productivity in Mexico," *Economía*, Fall 2003, *4* (1), 55–98.

Malgouyres, Clément, "Trade shocks and far-right voting: Evidence from French presidential elections," *Robert Schuman Centre for Advanced Studies Research paper no. RSCAS*, 2017, *21.*

Massachusetts Department of Education, "The Development of Education in Massachusetts, 1630-1930," Technical Report, Commonwealth of Massachusetts 1930.

Michaels, Guy, Ferdinand Rauch, and Stephen J. Redding, " Urbanization and Structural Transformation *," *The Quarterly Journal of Economics*, 03 2012, *127* (2), 535–586.

Milkman, Ruth, *Farewell to the Factory: Auto Workers in the Late Twentieth Century*, Berkeley, CA: University of California Press, 1997.

Mincer, Jacob, *Schooling, Experience and Earnings*, Columbia University Press: New York, 1974.

Minchin, Timothy, *Empty Mills: The Fight Against Imports and the Decline of the U.S. Textile Industry*, Rowman & Littlefield Publishers, 2012.

Minchin, Timothy J, *Empty Mills: The Fight Against Imports and the Decline of the US Textile Industry*, Rowman & Littlefield Publishers, 2012.

_ , *Labor under fire: A history of the AFL-CIO since 1979*, UNC Press Books, 2017.

Mindell, David A., *Between Human and Machine: Feedback, Control, and Computing before Cybernetics*, Baltimore, MD: Johns Hopkins University Press, 2004.

Monte, Ferdinando, Stephen J Redding, and Esteban Rossi-Hansberg, "Commuting, migration, and local employment elasticities," *American Economic Review*, 2018, *108* (12), 3855–90.

Moretti, Enrico, *The new geography of jobs*, Houghton Mifflin Harcourt, 2012.

Noble, David F., *Forces of Production: A Social History of Industrial Automation*, New York, NY: Oxford University Press, 1986.

Notowidigdo, Matthew J., "The Incidence of Local Labor Demand Shocks," *Journal of Labor Economics*, 2020, *38* (3), 687–725.

Ohanian, Lee E., "Competition and the Decline of the Rust Belt," Economic Policy Paper 14-6, Federal Reserve Bank of Minneapolis December 2014.

Pierce, Justin R. and Peter K. Schott, "Trade Liberalization and Mortality: Evidence from US Counties," *American Economic Review: Insights*, March 2020, *2* (1), 47–64.

_ **and** _ , "Trade Liberalization and Mortality: Evidence from US Counties," *American Economic Review: Insights*, March 2020, *2* (1), 47–64.

Piketty, Thomas, *Capital and ideology*, Harvard University Press, 2020.

Raff, Daniel M. G., Timothy F. Bresnahan, Changkeun Lee, and Margaret Levenstein, "United States Census of Manufactures, Cotton Goods Industry, 1929-1935," 2015.

Rendiero, Joao, "How the Japanese Came to Dominate the Machine Tool Industry," *Long Range Planning*, 1985, *18* (3), 62–67.

Reuther, Walter P., "Policies for Automation: A Labor Viewpoint," *The Annals of the American Academy of Political and Social Science*, 1962, *340*, 100–109.

Rodrik, Dani, *Has Globalization Gone Too Far*, Peterson Institute for International Economics, 1997.

_ , *Has Globalization Gone Too Far?*, Washington, DC: Institute for International Economics, 1997.

Romalis, John, "NAFTA's and CUSFTA's Impact on International Trade," *The Review of Economics and Statistics*, 2007, *89* (3), 416–435.

Rosenbloom, Joshua L, "The Challenges of Economic Maturity: New England, 1880 - 1940," Working Paper 113, National Bureau of Economic Research February 1999.

Sarathy, Ravi, "The Interplay of Industrial Policy and International Strategy: Japan's Machine Tool Industry," *California Management Review*, 1989, *31* (3), 132–160.

Scully, Charles J., "THE WOOL INDUSTRY IN NEW ENGLAND," *Review of Social Economy*, 1951, *9* (2), 137–145.

Stein, Judith, *Pivotal decade: How the United States traded factories for finance in the seventies*, Yale University Press, 2010.

Stuart, Bryan A., "The Long-Run Effects of Recessions on Education and Income," *American Economic Journal: Applied Economics*, January 2022, *14* (1), 42–74.

U.S. Information Agency, "U.S. Trade Liberalization Initiatives in the Western Hemisphere," *Economic Perspectives*, 1998, *3* (2), 25–26.

U.S. International Trade Commission, "Potential Impact on the U.S. Economy and Selected Industries of the North American Free-Trade Agreement: Report to the Committee on Ways and Means of the United States House of Representatives and the Committee on Finance of the United States Senate on Investigation No. 332-337 Under Section 332 of the Tariff Act of 1930," Technical Report 2596, SITC PUBLICATION 1993.

Weisman, Steven R, *The Great Tax Wars: Lincoln, Teddy Roosevelt, Wilson How the Income Tax Transformed America*, Simon and Schuster, 2004.

Wieandt, Axel, "Innovation and the Creation, Development and Destruction of Markets in the World Machine Tool Industry," *Small Business Economics*, 1994, *6*, 421–437.

Woodbury, Robert S., *Studies in the History of Machine Tools*, Cambridge, MA: M.I.T. Press, 1972.

Yagan, Danny, "Employment Hysteresis from the Great Recession," *Journal of Political Economy*, 2019, *127* (5), 2505–2558.

_ , "Employment hysteresis from the Great Recession," *Journal of Political Economy*, 2019, *127* (5), 2505–2558.