

# Essays on Employment and Human Capital

by

Sean Yixiang Wang

B.A., Washington University in St. Louis (2014)

Submitted to the Department of Economics  
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2022

© Sean Yixiang Wang, MMXXII. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Author .....

Department of Economics

May 13, 2022

Certified by .....

Daron Acemoglu

Institute Professor

Thesis Supervisor

Certified by .....

David Autor

Ford Professor of Economics

Thesis Supervisor

Accepted by .....

Abhijit Bannerjee

Ford International Professor of Economics

Chair, Department Committee on Graduate Theses



# Essays on Employment and Human Capital

by

Sean Yixiang Wang

Submitted to the Department of Economics  
on May 13, 2022, in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy in Economics

## Abstract

This thesis examines how economic forces shape the nature of employment and the development of human capital. Each of the three chapters in the thesis brings economic theory and causal inference to administrative data to better understand the mechanisms that ultimately determine people's livelihoods. Collectively, the chapters emphasize how imperfect markets and institutions have the powerful potential to either reduce or exacerbate existing inequalities.

The first chapter identifies the effects of firms on the career advancement of blue-collar workers and interprets these effects through the mechanism of employer learning. I use administrative data on the universe of Brazilian formal employment to study vertical promotions from production jobs to supervisory jobs, which are an important source of wage growth for most young workers. By comparing workers around job-to-job transitions, I show that differences in average firm promotion rates reflect persistent differences in the effects of firms on workers. Workers who move to a high promotion firm become substantially more likely than other job movers to be promoted, but they are even more likely to leave formal employment altogether. Correspondingly, their average long-term wage gains are negligible. I explain these effects using a model where firms differ in the rate they learn about the abilities of employed workers. High learning firms improve the efficiency of matching between workers and jobs, but these firms also exacerbate the adverse selection of unemployed workers and increase occupational wage inequality. By quantifying the parameters of the model using my estimated effects, I show that skill misallocation remains high and ex-post market power for employers can be large.

The second chapter, written jointly with Samuel Young, studies the effect of private-sector unionization on establishment employment and survival. Specifically,

we analyze National Labor Relations Board (NLRB) union elections from 1981 to 2005 using administrative Census data on the universe of establishments in the U.S. Our research design combines difference-in-differences and regression discontinuity extrapolation methods to estimate treatment effects including elections that win by larger margins of support. We show that unionization decreases an establishment's employment and likelihood of survival. We hypothesize that two reasons for these effects are firms' ability to avoid dealing with new unions and managers' opposition to unions. We test this hypothesis for unionization in manufacturing, the largest sector where we find substantial negative effects. There, the negative effects are significantly larger for elections at multi-establishment firms, especially those with no other unionized establishments. We provide direct evidence suggesting that some of these differences are driven by multi-establishment firms shifting employment from newly unionized establishments to other establishments. Finally, we use the length of delays during the election process as a proxy for managers' opposition to the union and find substantially larger effects of successful elections with longer delays. Taken together, our results are consistent with firms' union avoidance tactics playing a role in explaining the overall negative effects of unionization.

The third chapter directly estimates a theoretically motivated measure of schools' competitive pressures using centralized assignment data from a large urban school district's deferred-acceptance mechanism. I find that competitive pressure within the district is dispersed, and most of the variation in competition is unexplained by concentration. While there is substantial pressure to attract more students to some schools, these competitive incentives do not induce schools to raise their school effectiveness on academic achievement. Instead, schools respond by shifting discretionary expenditures from administration to instruction.

**JEL Classification:** I21, J51, M51

Thesis Supervisor: Daron Acemoglu  
Title: Institute Professor

Thesis Supervisor: David Autor  
Title: Ford Professor of Economics

## Acknowledgments

This thesis represents the culmination of a long and rewarding journey. I would not be here today without the mentorship and support of many people. I will use this limited space to acknowledge a few people in particular.

First, I am grateful to my advisors, Daron Acemoglu, David Autor, and John Van Reenen for their generosity, wisdom, and support. They never questioned what I would be capable of accomplishing, and they were patient and unwavering every step along the way. Thanks to them, I will forever remember to pursue important questions, act with integrity and humanity, and be relentlessly diligent.

I am also grateful to my classmates and colleagues, who have been patient teachers, supportive peers, and caring friends over the past years in Cambridge. I was drawn to the program for its tight-knit and collaborative community. Jonathan Cohen, Martina Uccioli, and Samuel Young easily confirmed that MIT was a wise choice. I also regularly benefited from the kindness, perspective, and community of many others, including Martin Aragonese Torrego, Mayara Felix, David Hughes, Clemence Idoux, Jeremy Majerovitz, Matthew Ridley, Parinitha Sastry, Garima Sharma, and Michael Wong.

Graduate school is also the endpoint of a twenty-four-year educational marathon. Countless teachers and mentors helped lay the path that led me to MIT. Liu Laoshi (柳老师), Lesley Avery, Marlys Davidson, Kathy Currie, Brent Eldridge, and Dorothy Haas are exceptional teachers who taught me to learn for myself. Marcus Berliant, George-Levi Gayle, Limor Golan, and John Nachbar fostered my interest in economics and encouraged me to pursue graduate studies. John Beshears, James Choi, David Laibson, and Brigitte Madrian were pivotal mentors, role models, and a key reason that I am at MIT.

Finally, I can never recognize enough the love and support from my family. My partner, Carolyn Sharzer, is my closest friend and my limitless source of strength. My parents, Wang Jianqing (王建卿) and Wen Xiaoping (温小萍), are responsible for who I am today. I am deeply rooted in their unyielding perseverance, optimism, and joy.

I dedicate this thesis to my late grandfathers, Wang Xi (王玺) and Wen Xucai (温续才). They grew up in the most difficult times and were given nothing. But they exemplify the fact that wealth and formal schooling are not prerequisites for living an exceptional life. This thesis is one of their countless legacies.

THIS PAGE INTENTIONALLY LEFT BLANK

# Contents

<b>1</b>	<b>What Is the Price for Opportunity? The Effects of Employer Learning on Worker Promotions and Turnover</b>	<b>17</b>
1.1	Introduction . . . . .	17
1.2	Model of promotions and exit under employer learning . . . . .	23
1.2.1	Basic setup . . . . .	24
1.2.2	Partial equilibrium and direct effects . . . . .	25
1.3	Setting and data . . . . .	29
1.3.1	Background on the Brazilian labor market . . . . .	29
1.3.2	Data and sample restrictions . . . . .	30
1.3.3	Measuring promotions . . . . .	31
1.4	Identifying the effects of firms . . . . .	33
1.4.1	Defining high promotion firms . . . . .	34
1.4.2	Baseline research design . . . . .	35
1.4.3	Extensions incorporating additional variation . . . . .	38
1.5	Direct effects of high promotion firms . . . . .	41
1.5.1	Baseline results . . . . .	41
1.5.2	Robustness of direct effects . . . . .	44
1.6	Additional evidence for the employer learning mechanism . . . . .	46
1.6.1	Promotion effects reflect positive worker outcomes . . . . .	46
1.6.2	Exits from formal labor force are involuntary . . . . .	47
1.6.3	Promotion and exit effects are stronger for high potential workers	48
1.6.4	Survey of firm labor practices is consistent with estimated effects	49
1.7	Equilibrium Effects of Employer Learning . . . . .	50
1.7.1	Endogenizing vacancy creation . . . . .	51
1.7.2	Testing for equilibrium effects on occupational wage structure	53
1.8	Structural quantification . . . . .	54
1.8.1	Model identification and estimation . . . . .	54

1.8.2	Parameter estimates . . . . .	56
1.9	Conclusion . . . . .	57
<b>2</b>	<b>Unionization, Employer Opposition, and Establishment Closure</b>	<b>79</b>
2.1	Introduction . . . . .	79
2.2	Unionization through NLRB Elections . . . . .	86
2.3	Election, Contract, and Establishment Data . . . . .	90
2.4	Empirical Strategy and Identifying Assumptions . . . . .	93
2.5	Empirical Results: Overall Employment and Survival Effects . . . . .	99
2.6	Testing for Manager Opposition and Union Avoidance . . . . .	109
2.7	Relation to Literature and Implications . . . . .	119
2.8	Conclusion . . . . .	121
2.9	Figures . . . . .	123
2.10	Tables . . . . .	125
<b>3</b>	<b>Competition and School Quality: Evidence from Centralized Assignment</b>	<b>141</b>
3.1	Introduction . . . . .	141
3.2	Setting and data . . . . .	146
3.3	Model of school competition . . . . .	147
3.4	Empirical model of the market . . . . .	149
3.4.1	Demand estimation . . . . .	150
3.4.2	School supply and assignment mechanism . . . . .	152
3.4.3	Validation . . . . .	153
3.5	Estimates of competitive pressure . . . . .	154
3.6	Relationship between competition and school outcomes . . . . .	156
3.6.1	Cross sectional variation . . . . .	156
3.6.2	Variation from centralized assignment . . . . .	158
3.7	Conclusion . . . . .	159
<b>A</b>	<b>Appendix for Chapter 1</b>	<b>173</b>
A.1	Additional Tables and Figures . . . . .	173
A.2	Proofs of Propositions . . . . .	190
A.2.1	Partial Equilibrium Model . . . . .	190
A.2.2	General Equilibrium Model . . . . .	193
A.3	Additional Empirical Details . . . . .	196
A.3.1	Data Construction . . . . .	196
A.3.2	Calculating Average Treatment Effects . . . . .	200



A.3.3	Structural Quantification . . . . .	201
<b>B</b>	<b>Appendix for Chapter 2</b>	<b>203</b>
B.1	Appendix Figures . . . . .	203
B.2	Appendix Tables . . . . .	209
B.3	Data and Matching Details Appendix . . . . .	215
<b>C</b>	<b>Appendix for Chapter 3</b>	<b>227</b>
C.1	Additional Tables and Figures . . . . .	227

THIS PAGE INTENTIONALLY LEFT BLANK

# List of Figures

1-1	Example of job output functions . . . . .	59
1-2	Model Timeline . . . . .	59
1-3	Prevalence of Promotions for Production Workers . . . . .	60
1-4	Timeline of empirical approach . . . . .	60
1-5	Test for movers balance . . . . .	61
1-6	Impact of high promotion firms on worker promotions . . . . .	62
1-7	Impact of high promotion firms on formal labor market attachment . . . . .	63
1-8	Impact of high promotion firms on log earnings for employed workers . . . . .	64
1-9	Robustness of main estimates to controls . . . . .	65
1-10	Comparison of main effects between identification strategies . . . . .	66
1-11	Differential impact of high promotion firms by worker job levels . . . . .	67
1-12	Impact of high promotion firms on each type of exit . . . . .	68
1-13	Comparison of main effects by worker potential . . . . .	69
1-14	Correlation in firm practices . . . . .	70
2-1	Testable Implications of Parallel Trends Identifying Assumption . . . . .	123
2-2	Characteristics of Close Elections that Motivate Including Larger Margin-of-Support Elections . . . . .	124
2-3	Employment and Survival Estimates, 20-80 % Vote-Share Elections, All Industries . . . . .	127
2-4	Employment and Survival Estimates, 20-80 % Vote-Share Elections, Manufacturing . . . . .	128
2-5	Nonparametric Vote-Share Heterogeneity Estimates, Manufacturing . . . . .	129
2-6	Nonparametric Vote-Share Heterogeneity Estimates, All Industries . . . . .	130
2-7	Single- Versus Multi-Establishment Firm Heterogeneity . . . . .	131
2-8	Employment Effects of Successful Elections on Firms' Other Establishments . . . . .	132
2-9	Unionized versus Non-Unionized Firm Heterogeneity . . . . .	133
2-10	Election Delay Heterogeneity . . . . .	134

2-11	Establishment-Level Total Factor Productivity Heterogeneity . . . . .	135
3-1	Model fit: substitution patterns . . . . .	161
3-2	Model fit: offer shares . . . . .	162
3-3	Distribution of competitive pressure . . . . .	163
3-4	Competitive pressure by sector . . . . .	164
3-5	Competitive pressure by grade . . . . .	165
3-6	Event study coefficients for school effectiveness . . . . .	166
A-1	Prevalence of Occupations with Clear Promotion Tracks . . . . .	173
A-2	Distribution of firm promotion propensities . . . . .	174
A-3	Age profile of supervisors and promotions . . . . .	174
A-4	Visual intuition for IV first stage . . . . .	175
A-5	Impact of high promotion firms on a worker ever being promoted . .	175
A-6	Impact of high promotion firms on worker promotions (by cohort) . .	176
A-7	Impact of high promotion firms on formal labor market attachment (by cohort) . . . . .	176
A-8	Impact of high promotion firms on log earnings for employed workers (by cohort) . . . . .	177
A-9	Impact of high promotion firms on worker promotions (continuous measure) . . . . .	177
A-10	Impact of high promotion firms on formal labor market attachment (continuous measure) . . . . .	178
A-11	Impact of high promotion firms on log earnings for employed workers (continuous measure) . . . . .	178
A-12	Impact of high promotion firms on worker promotions (mass layoffs) .	179
A-13	Impact of high promotion firms on formal labor market attachment (mass layoffs) . . . . .	179
A-14	Impact of high promotion firms on log earnings for employed workers (mass layoffs) . . . . .	180
A-15	Impact of high promotion firms on worker promotions (local hiring IV)	180
A-16	Impact of high promotion firms on formal labor market attachment (local hiring IV) . . . . .	181
A-17	Impact of high promotion firms on log earnings for employed workers (local hiring IV) . . . . .	181
A-18	Comparison of main effects between alternative firm types . . . . .	182
A-19	Robustness of model estimates to alternate assumptions . . . . .	183

B-1	Number of Unique Case Numbers Across Datasets versus NLRB Annual Reports . . . . .	203
B-2	Election Vote-Share Histogram, 50 + Vote Elections . . . . .	204
B-3	Log Employment and Payroll Estimates, 20-80 % Vote-Share Elections	205
B-4	DHS Employment Estimates, 20-80 % Vote-Share Elections, 10 Yr Pre- and Post-Periods . . . . .	206
B-5	Election Win Rates and Challenged Vote Rates by Delay Time . . . .	207
B-6	Establishment-Level Total Factor Productivity Heterogeneity, Multi-Establishment Firms . . . . .	208
C-1	Number of schools ranked in application . . . . .	227
C-2	Substitution patterns (no heterogeneity) . . . . .	228
C-4	Match replication rate . . . . .	229
C-5	Event study coefficients (other specifications) . . . . .	230

THIS PAGE INTENTIONALLY LEFT BLANK

# List of Tables

1.1	Example of occupational group with observable line of progression . .	71
1.2	Estimates of promotion wage premia . . . . .	71
1.3	Role of promotions in lifecycle wage profile . . . . .	72
1.4	Differences between high and low promotion firms . . . . .	73
1.5	Analysis sample summary statistics . . . . .	74
1.6	Estimates of regional wage premia . . . . .	75
1.7	Summary of parameters and moments for identification . . . . .	76
1.8	Baseline model estimates . . . . .	77
2.1	Winning versus Losing Election Establishment Summary Statistics . .	125
2.2	Pre-Election Employment Growth Trends by Vote Share, 20-80 % Elections . . . . .	126
2.3	Post-Election Outcome Trends by Vote Share, 20-80 % Vote-Share Elections . . . . .	136
2.4	Employment and Survival Estimates by Industry, 20-80 % Vote-Share Elections . . . . .	137
2.5	Effects of Successful Elections on Firms' Other Establishments . . . .	138
2.6	Election Delay Heterogeneity, Continuous Delay Time Specification .	139
3.1	Summary statistics . . . . .	167
3.2	Demand estimates . . . . .	168
3.3	Relationship between competitive pressure and concentration . . . . .	169
3.4	Cross sectional effects of competitive pressure on school outcomes . .	170
3.6	Pooled differences in differences estimates . . . . .	171
A.1	Correlation between measures of firm promotions . . . . .	184
A.2	Pooled estimates on the impact of high promotion firms on other worker outcomes . . . . .	185
A.3	Worker potential subsample summary statistics . . . . .	186

A.4	World Management Survey questions and scoring criteria . . . . .	187
A.5	Estimates of regional differences in labor market attachment . . . . .	188
A.6	Sensitivity of model parameters to empirical moments ( $\Lambda$ ) . . . . .	189
B.1	Union Election Matched Sample Construction . . . . .	209
B.2	Post-Election Outcome Trends by Vote Share, 20-80 % Vote-Share Elections, Employment and Industry Ctrls. . . . .	210
B.3	Employment and Survival Bargaining Unit Share Interaction, 20-80 % Vote-Share Elections . . . . .	211
B.4	Manufacturing versus Services Employment and Survival Estimates, Robustness Checks . . . . .	212
B.5	Single- Versus Multi-Establishment Firm Heterogeneity, Robustness Checks . . . . .	213
B.6	Unionized versus Non-Unionized Firm Heterogeneity, Robustness Checks	214



# Chapter 1

## What Is the Price for Opportunity? The Effects of Employer Learning on Worker Promotions and Turnover

### 1.1 Introduction

Research has consistently shown that where you work matters. Across a variety of countries and time periods, the same worker can expect to earn considerably different wages at different employers (Card et al., 2018). Employers, however, influence many aspects of the employment relationship besides the *level* of wages, including whether a worker is given the opportunity to advance to more complex and higher paying jobs. In Brazil, the focus of this paper, the average probability that a production worker is promoted to supervisor is nearly zero at most firms, but approximately two percent per year at the top quartile of firms.

Evidence from my data confirms that promotions are a key channel for worker wage growth and skill accumulation. Direct promotions from production worker to

---

I am grateful to Daron Acemoglu, David Autor, and John Van Reenen for their enduring guidance and support. I also thank Jonathan Cohen, Viola Corradini, Mayara Felix, Peter Ganong, Robert Gibbons, Jonathan Gruber, Ahmet Gulet, Simon Jäger, Danielle Li, Michael Piore, Charles Rafkin, Tobias Salz, Garima Sharma, Carolyn Sharzer, Carolyn Stein, Pedro Bessone Tepedino, Joonas Tuhkuri, Martina Uccioli, Michael Wong, Samuel Young, and participants at the MIT Labor Lunch and Organizational Economics Lunch for helpful comments and discussions. Access to Brazil's RAIS database is governed by the Data Use Agreement between MIT and Brazil's Ministry of Labor. I thank David Atkin and Mayara Felix for procuring MIT's access to the database, and Mayara Felix for de-identifying, harmonizing, and translating the RAIS datasets pursuant to MIT COUHES guidelines. All unmodeled errors are my own. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302 and the Washington Center for Equitable Growth.

supervisor are accompanied by persistent earnings increases equivalent to the returns to two years of schooling. These promotions also explain five to ten percent of the lifecycle wage profile for young, blue-collar workers. Understanding the employer's role in creating promotion opportunities, especially for low skill workers, is thus critical to understanding the contributions of firms to economic mobility and inequality.

This paper uses administrative data on the universe of formal employment in Brazil to identify and interpret the effects of firms on the career advancement of blue-collar workers. I show that the unique structure of the occupational data in Brazil allows me to observe direct, vertical promotions for 70% of the formal labor market. By using flexible panel-based identification strategies, I find that the most upwardly mobile quartile of firms have persistent effects on the positive outcome of worker promotion as well as the negative outcome of formal labor force exit. I argue that these new facts are most consistent with the interpretation that firms systematically differ in their rate of learning about worker ability. I also show theoretically and empirically that this mechanism has meaningful implications for the equilibrium wage structure and the implied degrees of skill misallocation and employer ex-post market power.

The direct measurement of worker promotions using the Brazilian administrative linked employer-employee data is a key basis of this paper. The data are unique in that they clearly distinguish within occupational groups between workers who are focused on production and advanced workers who have supervisory tasks. As a result, I can directly observe when workers are promoted from production jobs to directly related and more advanced supervisor jobs. These promotions are applicable for the majority of the Brazilian workforce, including almost all blue-collar workers.

Employers are initially uncertain about worker ability, and labor markets are frictional. I establish a conceptual framework for the paper by incorporating these two features into a stylized model of job assignment and highlighting a key testable prediction about the overall effects of employers: under employer learning, firms that promote more often also fire more often. The model combines a standard learning and job assignment problem (as in Waldman, 1984; Gibbons and Waldman, 1999) with a frictional labor market where asymmetric information between employers results in the adverse selection of job movers (as in Greenwald, 1986; Acemoglu and Pischke, 1998). In the model, output is complementary in a worker's unidimensional ability and their job's complexity. It is efficient to assign high ability workers to complex jobs, keep workers of unknown ability in simple jobs, and fire low ability workers. I assume that employers are more likely to learn about a worker's ability when the worker is assigned to the complex job, so high learning employers are employers that are more

willing or able to try out workers of unknown ability in complex jobs. Crucially, high learning employers are more likely to promote and fire workers because they are more active in acquiring information about the abilities of their workers.

I test the predictions of the employer learning model by estimating the effects of high promotion firms on long-term worker outcomes, and my results robustly confirm the model's predictions. High promotion firms increase the promotion probability for workers who join the firm by an additional 1.2 percentage points relative to an average baseline promotion rate of 1.7 percentage points for similar workers who join other firms. The effects decay only slightly over time and are persistent for at least seven years after moving, including across subsequent moves to other firms. On the other hand, high promotion firms also reduce the long-run formal sector employment rate for their workers by an additional 1.9 percentage points, and the effects are similarly persistent, so workers are even more likely to leave formal employment than to become promoted as a result of joining a high promotion firm. Finally, although workers who are promoted experience persistent increases in earnings, the average earnings for workers who moved to high promotion firms are comparable to average earnings for workers who moved to other firms. Taken together, my results suggest that most workers do not benefit on net from moving to a high promotion firm.

There are two key threats to identifying the effects of employers on workers: the composition of workers may differ across firms, and firms may affect workers through transient firm shocks in addition to systematic practices. I use a two-step estimation strategy to address both concerns. I first define high promotion firms as firms in the top quartile of (composition-adjusted) promotion rates for blue-collar workers using the first four years of my data. I then use the subsequent years of my data to estimate the effects of high promotion firms by comparing workers who move to high promotion firms to workers who move to other, low promotion firms. My core identification assumption is that, conditional on moving, the identity of the worker's destination firm is uncorrelated with idiosyncratic changes in the worker's labor market outcomes. This identifying assumption is weaker than the standard assumptions for estimating firm wage effects primarily because I focus on estimating long-term effects, which include any effects stemming from workers' subsequent mobility. As a result, I do not require that all mobility decisions are as good as random; instead, I assume that the job movers are comparable to each other at the time of the move, which is supported by several falsification exercises.

I also validate my baseline research design using two additional sources of aggregate variation that shift workers' job choices. The first approach follows Gibbons and Katz

(1992) by using mass layoff events to purge potential biases arising from a worker's potential selection into moving employers. The second approach follows Oreopoulos et al. (2012) in spirit by instrumenting for the worker's destination firm with the local hiring share from all high promotion firms (excluding the worker's destination firm) to purge potential biases arising from the worker sorting based on idiosyncratic shocks or unobserved trends. My estimates remain similar when using these aggregate variation for identification, which supports the interpretation that the estimated effects are driven by the effects of firms rather than by worker sorting.

I then use a series of additional empirical exercises to show that the estimated effects are consistent with employer learning and inconsistent with several alternative explanations. Differences in promotions between firms appear to reflect real differences in job assignments and earnings rather than simply differences in the willingness to label otherwise identical workers as supervisors. Similarly, worker exits from formal employment are explained by negative outcomes like layoffs or firings rather than voluntary quits. The negative employment outcomes are also concentrated on workers who are likely to be promoted, which is consistent with learning and inconsistent with greater volatility in firm labor demand. Finally, corroborating survey evidence from managers across a variety of countries supports the interpretation that the effects reflect systematically different firm responses to general economic mechanisms rather than any Brazil-specific institutional feature.

After establishing that the direct employer effects are consistent with employer learning, I then explore the effects of high learning employers on *all* workers in their market in addition to direct effects on their employed workers. I theoretically clarify this indirect mechanism by endogenizing the labor market parameters of my model through an initial vacancy creation stage, and I empirically test the resulting predictions about the overall wage structure. High learning firms are more likely to fire low ability workers, which exacerbates the adverse selection of workers who change jobs. As a result, the secondary market is more pessimistic about the expected quality of incoming workers and makes lower wage offers. In equilibrium, wage differentials between promoted and non-promoted workers rise since non-promoted workers are more exposed to the softer secondary market competition. These additional predictions about occupational wage inequality are also robustly supported by the data. Moving from the 10th to the 90th percentile of municipalities (in terms of the local employment share of high promotion firms) increases the local wage premium for supervisors by 39% of its average value.

The model also establishes a framework for quantifying the economic magnitude

of the learning mechanism using my estimated employer effects. Although I make strong assumptions on wage setting to close the model in an analytically tractable manner, the basic components of the model that govern job assignment and turnover can be mapped to the estimated effects under a wide range of realistic wage setting mechanisms, including bargaining and binding wage floors. The key restriction for model identification is that workers in the analysis sample do not systematically differ based on their destination firms. This restriction is stronger than the identification assumption used to estimate the employer effects, but it is also consistent with that identification assumption and supported by the data.

Quantifying the model shows that the rate of employer learning is low on average, including at high promotion firms. As a result, skill misallocation is high – 87% of workers who would be suitable for supervisory occupations end up in elementary occupations or outside of formal employment. Nevertheless, the rate of employer learning is sizable relative to idiosyncratic turnover, so the scope for employer ex-post market power is also substantial. I estimate that job movers are only 71% as likely to be of high ability as the general population, which suppresses outside options for workers and amplifies the value of private information held by employers.

## **Related Literature**

This paper primarily contributes to the literatures on firm heterogeneity in labor market outcomes, worker dynamics within the firm, and firm promotion policies. To my knowledge, it is one of a few papers to focus on occupational outcomes using economy-wide administrative data, and it is the first of these papers to document the firm-level link between promotion opportunity and separation risk. It is also part of a series of papers that combines quasi-experimental estimates of firm heterogeneity with a theoretical framework for quantifying their economic implications, and it is the first of these papers to focus on the implications of employer learning.

A well-established literature has documented that differences between firms pass through to workers. A series of papers use linked employer-employee data and an exogenous movers research design to estimate the dispersion in the level of firm wage premia (for example, Abowd et al., 1999; Card et al., 2013, 2016; Song et al., 2019), including in the Brazilian context (Alvarez et al., 2018; Gerard et al., 2018). Related papers using worker data have also found that firms experiencing shocks tend to share a portion of the shocks with workers (Kline et al., 2019; Lamadon et al., 2019). There are relatively fewer papers focusing on firm differences in long-run effects, but the papers that do so tend to focus on interpreting long-term wage changes through search,

human capital accumulation, or a mixture of the two (Bagger et al., 2014; Herkenhoff et al., 2018; Gregory, 2019; Jarosch et al., 2019; Arellano-Bover, 2020; Taber and Vejlin, 2020; Addario et al., 2021).<sup>1</sup> Within this literature, Jarosch (2015) shows that firms differ in layoff risk that erodes workers' long-term earnings, and Arellano-Bover and Saltiel (2020) shows that firms differ in wage growth that may be correlated with occupational growth. My paper supports the idea that these two patterns reflect firms' causal effects rather than worker sorting. Moreover, I argue that the two outcomes are linked through the mechanism of employer learning.<sup>2</sup>

A similarly rich literature has analyzed the firm's role in changing a worker's skill mix. One prominent explanation for a worker's rising task complexity over time is that firms provide either direct or indirect training (Becker, 1964; Mincer, 1974; Acemoglu and Pischke, 1998; Jovanovic and Nyarko, 1997; Lazear, 2009). An alternative explanation is that information about a worker's productivity is revealed over time, so tenure profiles reflect selection rather than investment (Jovanovic, 1979; Waldman, 1984; O'Flaherty and Siow, 1995; Farber and Gibbons, 1996; Altonji and Pierret, 2001; Golan, 2005; Lange, 2007). Of course, the two explanations are not mutually exclusive and may interact (Gibbons and Waldman, 1999; Autor, 2001; Kahn and Lange, 2014; Pastorino, 2019). This paper shows that learning is relevant even for older production workers and can rationalize the observable patterns in job assignments and turnover. However, my data are not suited to directly study training, and I do not reject the role of either direct training or indirect training (which may be a conduit for learning or promotions). In fact, I find evidence in support of asymmetric information as a meaningful source of labor market frictions, which would imply that firms also benefit from paying for general training.

Finally, there is substantial interest in describing and interpreting firms' promotions decisions. Case studies have documented that firms commonly draw higher level workers from their pool of lower level workers rather than from external sources

---

<sup>1</sup>Note, though, that interpreting wage growth as human capital typically requires making strong assumptions about wage setting, since many mechanisms like job search and dynamic contracting would also generate wage growth without any changes in worker productivity. For example, workers who continue to search for new jobs while employed may receive wage increases whenever they receive a competing outside offer (Postel-Vinay and Robin, 2002; Caldwell and Harmon, 2019). Models of optimal dynamic contracting often feature increasing wages over a worker's tenure even if the worker's productivity is constant or decreasing (Lazear, 1979; Burdett and Coles, 2003).

<sup>2</sup>Another portion of the literature seeks to answer whether worker *utility* is equalized across firms and concludes that non-wage characteristics are dispersed across firms (Sorkin, 2018; Maestas et al., 2018). My paper supports those conclusions by showing that career advancement potential is a form of indirect compensation that differs across firms. I also raise the possibility that workers may be uncertain about employers' non-wage qualities at the time of hire.

(Doeringer and Piore, 1971; Baker et al., 1994). Recent studies on white-collar workers also show that, although firms attempt to target workers for promotion, their decisions are not always efficient or fair (Benson et al., 2019; Cullen and Perez-Truglia, 2021). Particularly related is Friedrich (2020), which uses administrative Danish data to show that more productive firms are more likely to use internal labor markets for filling top and middle managerial positions. My conclusions are consistent with Friedrich (2020) – actively promoting firms improve matching efficiency at the cost of higher adverse selection and wage inequality. I additionally provide quasi-experimental evidence supporting the causal effects of firms, and I show that these mechanisms are also relevant for the majority of the blue-collar workforce and have long-term consequences on workers’ labor force attachment.

The rest of the paper proceeds as follows. Section 1.2 establishes a conceptual framework by constructing a model of job assignment and deriving predictions about employers’ effects. Section 1.3 describes the Brazilian institutional setting and the administrative data. Section 1.4 discusses the identification assumptions for estimating employers’ effects on workers. Section 1.5 reports estimates of the effects of high promotion firms. Section 1.6 rules out potential alternative mechanisms using a series of additional checks. Section 1.7 then extends the baseline model to characterize the additional equilibrium effects of high learning firms. Section 1.8 discusses the structural quantification approach and reports the model’s estimates. Section 1.9 concludes.

## **1.2 Model of promotions and exit under employer learning**

To fix ideas about the possible effects of high promotion firms, I begin by characterizing a stylized model of job assignment. I incorporate three critical features that are realistic for my setting. First, firms learn about the quality of their employed workers over time. Second, information about worker quality is asymmetric between firms. Third, labor market matching is frictional and random. The resulting model generates a key prediction – firms that are more likely to promote workers are also more likely to fire workers – which I test in Section 1.5 by identifying the effects of firms on worker promotions and exit. The model in this section also serves as a basis for exploring the equilibrium effects of employer learning on the regional wage structure in Section 1.7 and for quantitatively interpreting the estimated employer effects as structural

parameters in Section 1.8.

### 1.2.1 Basic setup

I follow the framework of Waldman (1984); Gibbons and Waldman (1999) in modeling vertical job ladders as the optimal matching between workers' unidimensional ability and jobs' returns to ability. I assume that workers are one of two types: high ability ( $\theta = \theta_H$ ) with probability  $\alpha$ , and low ability ( $\theta = \theta_L$ ) with probability  $1 - \alpha$ . There are two possible job assignments  $j \in \{1, 2\}$  that are supermodular in worker ability and crossing:

$$f_2(\theta_L) < f_1(\theta_L) < 0 < f_1(\theta_H) < f_2(\theta_H).$$

Figure 1-1 shows an example of the expected output from each of the two occupations as a function of the probability that the worker is of high ability. The complex job ( $f_2$ ) has higher returns to ability but lower output for low ability workers. As a result, expected output is maximized in the high complexity job ( $j = 2$ ) if the worker is likely to be of high ability and in the low complexity job ( $j = 1$ ) if the worker is likely to be of low ability. Moreover, I assume that  $f_j(\theta_L) < 0$  for  $j \in \{1, 2\}$ , so workers who are revealed to be of low ability are not productive in either job, which introduces a motive for firms to fire workers. Finally, I assume that

$$\alpha f_1(\theta_H) + (1 - \alpha) f_1(\theta_L) > \alpha f_2(\theta_H) + (1 - \alpha) f_2(\theta_L),$$

so it is output maximizing to assign unknown workers to the low complexity occupation and the promotions problem is not trivial.

The timing of worker and firm actions and the realization of events are summarized in Figure 1-2. Workers are initially employed at either a high learning firm ( $f = H$ ) with probability  $\rho$  or a low learning firm ( $f = L$ ) with probability  $1 - \rho$ . At the start of employment, each firm  $f$  randomly assigns each worker to a trial in the complex job ( $f_2$ ) with probability  $Q_f$  or a trial in the simple job ( $f_1$ ) with probability  $1 - Q_f$ . The firm then observes the worker's output from the trial and updates its beliefs about the worker's ability. The complex job always reveals the worker's ability and the simple job never reveals the worker's ability. So, a high learning firm is more likely to learn about the ability of each worker because it is more likely to assign them to the more complex job (i.e.,  $Q_H > Q_L$ ).<sup>3</sup>

---

<sup>3</sup>An equivalent formulation, which generates the same predictions, is to assume that the firm simply learns about the ability of each worker with some probability  $Q_f$  before making job and wage offers. A high learning firm is then a firm that learns about the worker's ability more often (so



Based on its updated beliefs, the firm then decides either to fire the worker or to offer them a wage and job assignment. A fraction  $\delta$  of workers exogenously separates from their employers. The remaining workers choose to either accept the firm’s offer or leave the firm. All workers who separate from their initial employers encounter the secondary market with probability  $g$ .

Firms in the secondary market compete by making wage offers to workers. These firms have access to the same set of production technologies as incumbent firms, but they do not observe the ability of any worker, so there is asymmetric information. If a worker was offered a promotion at their previous employer, there is a probability  $\kappa$  that they can convince the secondary market of this fact. Firms do not observe any other information about incoming workers, so they do not know whether the worker was fired by their previous employer or they voluntarily quit.

At the end of the period, each firm employing a worker of type  $\theta$  in job  $j$  with wage  $w$  receives  $\pi = f_j(\theta) - w$ , while the worker receives  $w$ . Workers and firms are risk-neutral, and their outside options are normalized to 0.

### 1.2.2 Partial equilibrium and direct effects

The partial equilibrium of the model takes the share of high learning firms ( $\rho$ ) and the secondary market contact rate ( $g$ ) as given and assumes that workers and firms maximize their expected wages and profits. Since the game is one of imperfect information, I use the Perfect Bayesian equilibrium as my solution concept. The Perfect Bayesian equilibrium in this model is a set of worker and firm strategies such that the following conditions are satisfied:

1. Workers make turnover decisions that maximize expected wages given the incumbent firm’s wage and job assignment and the expected secondary market wage offers
2. Incumbent firms make wage and job offers that maximize expected profits given the worker’s turnover decision and the expected secondary market wage offers
3. Firms in the secondary market make wage offers that maximize expected profits given their beliefs about the expected ability of each worker in the secondary market

---

$Q_H > Q_L$ ). I use the “trial promotion” formulation because it is more consistent with the dynamics of my estimated promotion effects, as discussed in Section 1.5. It is also equivalent to assume instead that both jobs are partially informative as long as the complex job is more informative. I make the more stark assumption here to simplify notation.

4. Firms in the secondary market have rational beliefs about the ability of each worker given the turnover decisions of workers and the wage and job assignment decisions of incumbent firms

I consider the primary and economically interesting equilibrium in which an incumbent firm promotes workers who are revealed to be of high ability, fires workers who are revealed to be of low ability, and retains workers whose abilities remain unknown. Working backward, the secondary market is perfectly competitive, so firms offer the expected output for each worker and make zero profits in equilibrium. Any worker that successfully convinces the secondary market that they were previously promoted must have been revealed to be of high ability. Meanwhile, other workers in the secondary market are a combination of high ability workers (who exogenously separated and were not recognized as high ability by the secondary market) or low ability workers (who either exogenously separated or were fired). The probability that an unknown secondary market worker is of high ability is then

$$\alpha' = \frac{\alpha\delta(1 - \bar{Q}\kappa)}{\alpha\delta(1 - \bar{Q}\kappa) + (1 - \alpha)(\delta + (1 - \delta)\bar{Q})}, \quad (1.1)$$

where  $\bar{Q} = \rho Q_H + (1 - \rho) Q_L$  is the average rate of learning in the economy. These workers are adversely selected ( $\alpha' < \alpha$ ) both because high ability workers are more likely to enter the secondary market as promoted workers and because low ability workers are also more likely to be fired. Secondary market wages  $w_j^S$  for each job  $j$  are

$$w_1^S = \underbrace{\alpha' f_1(\theta_H) + (1 - \alpha') f_1(\theta_L)}_{\equiv E[f_1(\theta)|\alpha']} \quad (1.2)$$

$$w_2^S = f_2(\theta_H)$$

where I assume that  $\alpha'$  is sufficiently high so that the adversely selected workers are still expected to be productive:

$$E[f_1(\theta)|\alpha'] > 0. \quad (1.3)$$

An incumbent firm makes take-it-or-leave-it offers to workers, so it is sufficient to offer the expected outside option for the worker to accept. I assume that job take-up decisions are made before the worker makes contact with the secondary market (i.e., there is no on-the-job search), so secondary market wages are discounted by the probability that the worker encounters the secondary market ( $g$ ). Promoted workers run the additional risk that they may not successfully convince the secondary market

that they were previously promoted (i.e.,  $\kappa < 1$ ). So, optimal incumbent wage offers  $w_j^I$  for each job  $j$  are

$$\begin{aligned} w_1^I &= g w_1^S \\ w_2^I &= g [\kappa w_2^S + (1 - \kappa) w_1^S]. \end{aligned} \tag{1.4}$$

Notice that even in the case when  $g = 1$ , so there are no re-employment frictions, the incumbent employer still earns positive profits from each unpromoted worker due to adverse selection in the secondary market, since  $w_1^S = E[f_1(\theta) | \alpha'] < E[f_1(\theta) | \alpha]$ . Meanwhile,  $w_2^S$  reflects the true productivity of promoted workers, so the imperfect transmission of information about promoted workers ( $\kappa < 1$ ) ensures that an incumbent firm also retain some informational rents from promoted workers.

For the conjectured job assignments to be optimal for an incumbent firm, the following conditions need to hold

$$\begin{aligned} E[f_1(\theta) | \alpha] - w_1^I &\geq 0 \\ f_2(\theta_H) - w_2^I &\geq f_1(\theta_H) - w_1^I. \end{aligned}$$

The first condition ensures that the firm will find it profitable to retain workers whose abilities remain unknown by offering them the low complexity job with wage  $w_1^I$ .<sup>4</sup> This condition is implied by my assumption that employment in the secondary market is viable for the adversely selected workers (Equation 1.3). The second condition is an incentive compatibility condition that ensures the firm will find it more profitable to promote high ability workers and pay the higher wage  $w_2^I$  than to keep them in the low complexity occupation and pay the lower wage  $w_1^I$ .

The conditions for the existence and uniqueness of the equilibrium and the characterization of the equilibrium strategies are summarized in the following Proposition 1, and a detailed proof of the proposition is in Appendix A.2.

**Proposition 1.** *If information about job assignments on the secondary market is sufficiently weak (so  $\kappa$  is sufficiently small), then a unique Perfect Bayesian equilibrium exists. In this equilibrium, (i) job assignments for workers are efficient (given firms' information about workers) (ii) all turnover is involuntary (iii) wages are given by Equations 1.2 and 1.4.*

To see why the degree of asymmetric information ( $\kappa$ ) is key to satisfying the conditions for existence, note that a low  $\kappa$  relaxes the key constraints for both the

---

<sup>4</sup>The firm will never find it optimal to offer those workers the high complexity job since under these assumptions, their expected output is lower but the required wages to retain the worker are higher.

secondary market and the incumbent firm. A low  $\kappa$  ensures that promoted high ability workers may nevertheless enter the secondary market in the same pool as fired low ability workers, which helps offset the degree of adverse selection in the market. Meanwhile, a low  $\kappa$  also softens wage competition for promoted workers at the incumbent firm, since workers are then less likely to maintain their higher position if they go to the secondary market. This reduces the wage increase necessary to retain a worker upon promotion.

The partial equilibrium setup is sufficient for comparing the outcomes of otherwise identical workers who were initially matched to a high learning firm ( $Q_f = Q_H$ ) as opposed to a low learning firm ( $Q_f = Q_L$ ). Since workers may be of either high ability or low ability, being matched to a high learning firm introduces greater opportunities for promotion, but also greater risk of becoming unemployed whenever the probability of re-employment upon separation ( $g$ ) is less than 1. Proposition 2 formalizes this comparison, and the accompanying proof is also in Appendix A.2.

**Proposition 2.** *In the equilibrium described in Proposition 1, workers initially employed at high learning firms are (i) more likely to be promoted and (ii) more likely to become unemployed than workers initially employed at low learning firms.*

Proposition 2 makes the key predictions about the direct effects of firms that I will empirically test in Section 1.5. Specifically, the testable predictions are that in the presence of systematic differences in employer learning:

1. Some firms will be systematically more likely to promote workers
2. The firms that are more likely to promote workers are also more likely to fire workers

The model assumes that labor market matching is random, so the workers initially working at high learning firms are comparable to the workers initially working at low learning firms. However, the workers that exit firms are not comparable between high and low learning firms due to firings. In the presence of this selection, simply comparing the cross sectional differences between workers at different firms would overstate the causal effect of high learning firms on worker promotions and understate the causal effect on worker exits. Correspondingly, I focus on testing the model's predictions by estimating the effect of firms on newly hired workers, and I ensure that the effects are not driven by differences in worker composition.

## 1.3 Setting and data

The empirical setting is the Brazilian formal employment sector between 2003 and 2015, which is particularly well suited to examining the effects of employers on workers' careers for three reasons. First, Brazil's uniquely detailed administrative data allow researchers to follow approximately 70% of the formally employed workforce along direct lines of progression in their occupation groups. Second, although Brazilian labor market institutions have rigid components, employers generally have the flexibility to assign workers to different jobs. Finally, the rate of formal higher education in Brazil is low compared to high income countries, so employers can be expected to play a larger role in human capital accumulation or signaling.

### 1.3.1 Background on the Brazilian labor market

Brazil's labor market environment and institutions have been the subject of extensive research. Instead of trying to write a complete account of the Brazilian labor market, I focus on two aspects that are particularly important for the interpretation of my empirical strategy and results. I also briefly summarize several other details that provide additional context.

First, Brazil, like many other developing countries, has a sizable informal sector of the labor market. In Brazil, informal jobs do not have a signed "work card" (Carteira de Trabalho) and consequently are not subject to taxes and labor market regulations. Estimates of the size of the informal sector can vary since informal jobs are by definition missing from official registers. Recent estimates range as high as approximately 50% in metropolitan areas to as low as approximately 20% for all prime age workers; these estimates are comparable to informality rates in other developing economies. Firms in the formal labor market are generally more productive, pay higher wages, employ a greater share of educated workers, and are required to provide legally mandated employment protections and unemployment insurance for their workers.<sup>5</sup> In my data, I can observe whether a worker leaves the formal labor market, but not whether the worker enters the informal labor market versus unemployment. I generally interpret these exits as negative employment outcomes for the worker (and provide supporting evidence for this in Section 1.6.2), but it is important to note that I do not equate leaving the formal labor market as necessarily reflecting unemployment.

Second, although the Brazilian labor code mandates some employment protections

---

<sup>5</sup>See, for example, the summary in Perry et al. (2007) and estimates in Gerard et al. (2018); Haanwinckel and Soares (2020); Dix-Carneiro et al. (2021).

for workers in the formal sector, firms have the latitude to determine both the size and composition of their workforce. Legal protections mandating job security for workers were largely eliminated for employment contracts after 1966, and firms are allowed to dismiss workers without cause as long as they provide notice and severance penalties. The length of notice is typically 30 days (and can be as high as 90 days for workers with long periods of service), and the severance payment is up to 4% of the worker's total earnings while employed at the firm.<sup>6</sup> Firms are exempt from severance payments for voluntary quits or dismissals with cause (for more details, see Gonzaga et al. 2003; OECD-IDB 2014). By the OECD summary index of employment protections for workers with regular contracts, Brazil is scored as more flexible than countries like Denmark, the United Kingdom, and New Zealand. A key aspect of my interpretation is that some firms are more likely to layoff or fire workers after learning new information about the worker, and the legal flexibility afforded to employers is consistent with this interpretation.

Several other factors about the labor market in Brazil provide helpful context but are less central to the interpretation of my results. Labor unions are prominent in Brazil and bargain at the sectoral and the firm level. Firm-level collective bargaining agreements tend to cover all workers at an establishment (rather than varying by occupations or union membership) due to the Brazilian practice of universal coverage. The Brazilian labor code also explicitly prohibits nominal wage reductions, except those that are negotiated through collective bargaining (Lagos, 2019). The rate of tertiary education is generally low but has been experiencing rapid growth from 10% for 25-34 year olds in 2007 to 17% in 2017 (OECD, 2019). Finally, the macroeconomic environment around my period of study is generally stable. Brazil's period of hyperinflation ended with the introduction of the real in 1994, and inflation during my period of study ranged between 4-9%. While the period that I study straddles the Great Recession, its effects were muted in Brazil; instead, the country experienced a larger and more persistent recession starting in 2014.

### 1.3.2 Data and sample restrictions

My data on worker outcomes and firm characteristics come from the Brazilian Ministry of Labor's *Relação Anual de Informações Sociais* (RAIS), a worker-level dataset containing reporting data on all formal employment contracts. The data are likely to be fairly complete and high quality since the government mandates reporting to RAIS

---

<sup>6</sup>In the case of dismissals without cause, 80% of the penalty is paid out to workers and 20% is added to the state unemployment insurance fund

for all formal sector employers and penalizes late or missing filings. The dataset has also been used in several recent studies about the Brazilian labor market (including Menezes-Filho et al., 2008; Dix-Carneiro and Kovak, 2017; Alvarez et al., 2018; Gerard et al., 2018; Dix-Carneiro et al., 2021).

I observe key information about each formal employment contract, including its occupation, duration, contracted hours, contracted salary, and average monthly earnings. Furthermore, I observe detailed demographic information about the worker, including their gender, age, educational attainment, and race, as well as basic information about the employer’s industry and location. The records are linked over time by longitudinal identifiers, so I can follow workers over time across different firms. The RAIS separately records each employment contract, and a worker may hold several jobs over a year due to either job transitions or multiple part-time jobs. I construct an annual panel of worker employment histories from the collection of all contracts by selecting the long-term employment contract with the highest average earnings for each worker.<sup>7</sup> Additional details about the data construction are in Appendix A.3.1.

I consider outcomes between 2003, the first year when worker data are reported under the revised Brazilian occupational codes, and 2015, the last year of my data. For most worker-level analyses (unless otherwise specified), I restrict my sample to prime age workers who are between the ages of 25 and 50, since they have a high attachment to the formal labor force and are less likely to experience formal schooling or retirement. I discuss any additional sample restrictions for the implementation of my identification approaches in the corresponding parts of Section 1.4.

### 1.3.3 Measuring promotions

A distinctive feature of the Brazilian data is that I can observe direct lines of progression from worker to supervisor for most blue-collar occupational groups. Worker occupations are reported under the *Classificação Brasileira de Ocupações* (CBO) system, which is similar to the *International Standard Classification of Occupations* (ISCO) in that it groups jobs into a hierarchical structure based primarily on the type and complexity of the tasks involved. However, the Brazilian system, especially after its 2002 revision (CBO-02), is particularly suitable for the measurement of worker promotions for two reasons. First, it organizes occupations into occupational groups by natural lines of progression in addition to task content, so occupations that are often related for

---

<sup>7</sup>Defined as a contract that covered at least 6 months out of the year and entailed at least 20 contracted hours per week.

workers but do not share narrowly-defined tasks (like loom operators and fabric dyers) are explicitly grouped. Second, it consistently distinguishes between supervisors, who are advanced production workers with additional managerial responsibilities, and more elementary line workers within each occupational group. As a result of these two features, I can interpret changes from production occupations to supervisor occupations within the same general occupational group as reflecting a direct, vertical promotion.

Table 1.1 shows an example of an occupational group with a directly observable line of progression. The first two digits of the occupation code indicate the main occupational group, a “0” in the third digit indicates the sub-group of supervisory occupations, whereas other codes in the third digit indicate other sub-groups that contain production occupations but do not necessarily have a clear hierarchical structure. This basic structure applies to nearly all of the non-professional occupational groups in the Brazilian data. Figure A-1 shows that around 70% of workers in the formal sector belong to a 2-digit occupational group with a clear supervisor-worker line of progression and that this share is fairly stable over the period that I study.<sup>8</sup>

In this paper, I define a promotion as moving from any production occupation belonging to a supervisor-track occupational group to a supervisor occupation. Figure 1-3 shows that promotions are reasonably common. Furthermore, around half of promotions are purely vertical moves within the same occupational group. Promotion rates are fairly stable and free of secular trends over my sample, although they did begin to taper around 2014 when Brazil experienced a recession.

Two descriptive facts support the argument that my measure of promotions captures an important source of worker growth. First, promotions are valuable. Table 1.2 shows the estimated wage premium for supervisors relative to other production workers. The average cross-sectional wage premium is 63 (s.e. 1.8) log points. Controlling for worker characteristics accounts for 39% of this difference and controlling for unobserved differences through worker fixed effects accounts for another 34%, which implies that supervisors are positively selected. Even when I focus on within-worker changes in earnings, promotions are accompanied by wage increases of 17 (s.e. 0.35) log points, which is equivalent to the returns to an additional 1.7 years of education in Brazil.<sup>9</sup>

---

<sup>8</sup>A reclassification of the CBO-02 system occurred in 2008 and slightly increased the share of workers that belong to a relevant occupational group, but it does not substantively affect my approach or results.

<sup>9</sup>Controlling for firm-wage premia by subtracting the estimated AKM firm wage fixed effect from worker earnings increases the estimated promotion premium from 17 log points to 20 log points. These estimates are generally larger than the average wage increases accompanying promotions in the literature, but smaller than the large jumps in pay at the top of corporate hierarchies (see, for



Second, promotions consistently explain part of the lifecycle wage profile for young production workers. Table 1.3 compares the estimated lifecycle wage profile for workers between the ages 25 and 35 before and after including the promotion measure as an explanatory variable. When controlling for only basic worker characteristics, an extra year in age increases earnings by 1.73 (s.e. 0.03) log points. Adding promotions as an explanatory variable decreases the age coefficient to 1.63 (s.e. 0.03) log points, which implies that promotions explain approximately 5% of the age profile. The estimated age coefficient falls as I add additional controls for occupational characteristics and firm wage premia, but the estimated promotion wage premium stays relatively stable. As a result, the share of the age profile that is explained by promotions increases to 11%.<sup>10</sup> Although this exercise is descriptive, the results confirm that promotions capture a meaningful change that is distinct from other factors like schooling and employer upgrading.

## 1.4 Identifying the effects of firms

I focus on identifying the effects of *existing firms* on worker promotions and job turnover, so the ideal experiment is to randomly assign workers to various firms while taking as a given that those firms are a bundle of underlying practices. As a result, my identification strategy centers around changes in a worker's firm assignment, rather than changes in a firm's practices or changes in a worker's promotion likelihood. The baseline approach compares changes in worker outcomes among job-movers by the type of the destination firm, and two extensions relax the key identification assumptions by using aggregate shifters of workers' job choices.

My approach is similar to the two-step grouped fixed effects estimator of Bonhomme et al. (2019). I first separate firms into distinct groups using observational data, and I then estimate the effects of each group of firms on workers while allowing for flexible dynamics. However, my identification assumption is weaker than that of Bonhomme et al. (2019) since my estimand of interest is the long-term effect of differences between firms.<sup>11</sup>

---

example, Murphy, 1985; Baker et al., 1994; McCue, 1996; Blau and Devaro, 2007).

<sup>10</sup>My estimate is generally comparable to the estimate from McCue (1996) using self-reported promotions from the U.S. PSID. For comparison, Topel and Ward (1992) estimate that a third of wage growth over the first 10 years of employment in the U.S. is accounted for by wage gains at job changes, which would include promotions, and Bagger et al. (2014) estimate that human capital growth accounts for 20-25% of the life-cycle wage profile for low to medium educated workers in Denmark.

<sup>11</sup>As a concrete example, suppose firm wages are static and vary only by level. In this setting,

Generalizing the movers research design into a two-step event study framework has three additional benefits. First, I can explicitly separate the timing between the classification of firms and the estimation of employer effects to ensure that the effects reflect permanent differences between firms rather than different realizations of transitory firm shocks. Second, I can follow standard event study methodology in conducting falsification tests of my identifying assumptions through pre-trend and balance tests. Finally, it is straightforward to use flexible estimators that estimate the relevant average treatment effect by combining cohort-level treatment effects with explicit weights.

### 1.4.1 Defining high promotion firms

As a first step, I classify firms into two groups using firms' composition-adjusted promotion rates between 2004 and 2006. For each year between 2004 and 2006, I estimate the worker-level regression

$$\text{Promoted}_{it} = \beta_t X_{i,t-1} + \gamma_{ot} + \eta_{jt} + \varepsilon_{it}, \quad (1.5)$$

where  $\text{Promoted}_{it} = 1$  if a worker was promoted from production worker to supervisor between years  $t - 1$  and  $t$ .  $X_{i,t-1}$  and  $\gamma_{ot}$  adjust for differential promotion rates by the worker's observable characteristics (a quadratic in age interacted with gender and indicators for education, race, and state) and the worker's occupational group, respectively.<sup>12</sup>  $\eta_{jt}$  is the firm's residualized promotion rate in year  $t$ . I average over all three years of estimates for each firm to calculate the firm's average promotion propensity  $\eta_j = E[\eta_{jt}]$ . Firms with fewer than 10 promotion-track workers in at least one of those three years and firms who did not exist in all three years are unlikely to generate precise estimates of  $\eta_j$  or survive through to the movers analysis, so I drop them from the classification sample. I similarly exclude public sector firms, which may have different organizational structures and internal incentives.

Figure A-2 plots the distribution of the average promotion propensity  $\eta_j$  for the

---

firms may have dynamic effects on worker earnings if they disproportionately lead their workers to move to high wage firms. Estimators that seek to identify firms' wage policies would need to make assumptions about workers' mobility decisions to net out the contribution of subsequent high wage firms from the dynamic effects of prior firms. On the other hand, I seek to estimate the overall effect of the prior firm, including any effects that arise from workers subsequently moving to high wage firms.

<sup>12</sup>To facilitate interpretation, I restrict the sample to workers who are in the same occupational group in both years. However, the estimates are highly correlated if I include workers who switch occupations or restrict the sample further to include only job stayers. For more details, see Table A.1 and Appendix A.3.1.

remaining sample of firms, winsorized at the 5th and 95th percentiles. The distribution is highly skewed, with a small tail of firms that are highly active in promoting workers and a majority of firms that promote few (if any) workers. Some of the dispersion is mechanical since residualizing promotion rates by worker characteristics can generate dispersion in the firms' estimated fixed effects even when the firms' promotion rates were uniformly zero. Splitting firms into two groups based on  $\eta_j$  then ensures that my estimates of differences between firms are not primarily driven by small differences in residualized rates between firms that did not promote any workers.

For the rest of the paper, I define high promotion firms as firms with an average promotion propensity ( $\eta_j$ ) that is in the top quartile of all firms in this classification sample, and low promotion firms as all other firms in the sample. Table 1.4 shows the differences in firm characteristics between high and low promotion firms between 2003 and 2006. High promotion firms are 49 log points larger on average, so they actually constitute 34% of employment during this period. They are also higher paying and have a greater share of supervisors in the relevant occupational groups.<sup>13</sup> High promotion firms are responsible for almost the entirety of promotions within this sample of firms, which supports pooling the remaining firms together as one group. Finally, promotions may reflect in part the realization of positive firm shocks – high promotion firms are both faster growing and have higher wage growth for incumbent workers. This correlation between promotions and firm growth rates motivates testing for persistent effects using workers who later join these firms.

## 1.4.2 Baseline research design

### Estimating equation and identifying assumption

My baseline research design compares the change in outcomes between workers who move to high promotion firms and workers who move to low promotion firms. The primary identification assumption is that for a worker who is making an employer-to-employer transition, the type of the worker's destination firm is uncorrelated with idiosyncratic changes in the worker's labor market outcomes. This identification assumption is implied by standard exogenous mobility assumptions that are used in the literature for estimating firm wage effects, but it is strictly weaker. Specifically, I also allow for workers to select into moving based on idiosyncratic worker-level shocks and for firms to have persistent effects on workers after they leave.

---

<sup>13</sup>There may be concerns that supervisor shares are mechanically higher at high promotion firms since these firms were defined by having promoted workers. I also consider the prior year's supervisor share as a check, and the results are slightly smaller but similar.

To fix ideas, consider workers who make an employer-to-employer move in a single cohort year  $c$ . Let  $H_i = 1$  if worker  $i$  moved to a high promotion firm and  $H_i = 0$  if they moved to a low promotion firm. My estimation equation for the dynamic effects of the high promotion firm on worker outcome  $y_{it}$  is

$$y_{it} = \sum_{\tau} \beta_{\tau} (H_i \times I_t^{\tau}) + \alpha_i + \pi X_{it} + \theta_t B_{i,c-1} + \gamma_{sot} + \varepsilon_{it}, \quad (1.6)$$

where  $I_t^{\tau}$  are event-time indicators that equals 1 when  $t - c = \tau$  and 0 otherwise (with  $\tau = -1$  as the omitted time period).<sup>14</sup>  $\alpha_i$  is a worker fixed effect,  $X_{it}$  is a vector of time-varying worker covariates, and  $B_{i,c-1}$  is a vector of baseline worker covariates from year  $c - 1$  that may affect worker outcomes in other periods through the time varying coefficients  $\theta_t$ . Finally, I saturate the regression with  $\gamma_{sot}$ , a set of state-by-baseline-occupational-group-by-time fixed effects that control flexibly for aggregate trends in market conditions or skill prices.<sup>15</sup>

Note that by defining the estimand of interest  $\beta_{\tau}$  as the effect of moving *to* a high promotion firm after  $\tau$  years, rather than the effect of staying *at* the high promotion firm for  $\tau$  years, I do not distinguish between workers who remain at the high promotion firm from workers who subsequently leave. So, my estimate includes the direct effects of the firm on its stayers as well as the persistent effects of the firm on workers' mobility decisions and subsequent outcomes at other firms. I make this choice for two reasons. First, some models of human capital accumulation predict that workers would realize most of the gains after leaving the employer responsible for providing the human capital, and this estimand ensures that I appropriately capture these channels.<sup>16</sup> Second, I do not need to make any assumptions about the worker's mobility decisions after the year of the initial move, which itself may be a result of the worker's destination firm.

The key assumption for the identification of  $\beta_{\tau}$  is that

$$E[\varepsilon_{it} H_i] = 0.$$

---

<sup>14</sup>In pooled specifications, I collapse event time into two periods – the near term ( $\tau \in [0, 2]$ ) and the long-term ( $\tau > 2$ ).

<sup>15</sup>In the baseline specification, the time varying covariates  $X_{it}$  are quadratic trends in age that vary by gender and indicators for educational attainment, while the baseline covariate  $B_{i,c-1}$  is the origin firm's estimated AKM firm effect. These controls explain any differences in baseline characteristics of the high versus low promotion firm movers, as shown in the balance test discussed later in this section.

<sup>16</sup>For example, employers that provide training may demand compensating differentials through lower wages (Becker, 1964), or employers may specialize in jobs that are “stepping stones” for more complex jobs at other firms (Jovanovic and Nyarko, 1997).

In words, I assume that changes in worker outcomes that are unexplained by my controls are uncorrelated with the type of the worker’s destination firm. Economically, this assumption would be satisfied if high and low promotion firms are observationally indistinguishable to workers at the point of job take-up. This assumption would also be satisfied if firms observationally differ, but workers (with the same covariates) draw offers at random from a common distribution and make the same take-up decisions upon receiving an offer. On the other hand, this assumption would be violated if workers who are more likely to experience a positive shock are also more likely to receive or accept an offer from a high promotion firm.

Although my identifying assumption is about the counterfactual *changes* in worker outcomes and not directly testable, a related falsification test suggested by Kahn-Lang and Lang (2020) is whether workers moving to high versus low promotion firms are different in *levels*. Figure 1-5 plots the difference in baseline covariates between high and low promotion firm movers from the regression

$$X_{i,c-1} = \beta H_i + \theta M_{i,c-1} + \varepsilon_i,$$

where  $X_{i,c-1}$  are worker characteristics from the baseline year,  $H_i$  is the indicator for whether the worker moved to a high promotion firm, and  $M_{i,c-1}$  are any controls from the baseline year. The raw differences in baseline means between the two groups show that movers to high promotion firms have slightly higher baseline earnings, are more likely to be male, and have fewer years of education. However, most of the differences are explained by the state and occupational group of the worker. Any remaining differences in baseline earnings between the two groups are explained by the fact that movers from high promotion firms are more likely to come from high wage firms (as proxied by the estimated AKM firm fixed effect), and remaining differences in education and gender are quantitatively small.

## Estimation details

I estimate average employer effects by stacking the five cohorts of workers who make an employer-to-employer transition to a high or low promotion firm between 2008 and 2012.<sup>17</sup> Figure 1-4 shows a timeline of my approach. Since my data end in 2015, I observe at least three years of worker outcomes following the move, but any effects for longer-term outcomes are identified by the earlier cohorts. A recent

---

<sup>17</sup>In the cases when the same worker belongs to multiple mover cohorts, I consider the worker’s earliest move. However, the results are similar when I include all moves as separate events.

literature on event studies with staggered treatment timing has cautioned that simply pooling all cohorts and estimating treatment effects by OLS may yield unintuitive and potentially negatively weighted means of the cohort-specific treatment effects (e.g., de Chaisemartin and D’Haultfoeuille, 2020; Goodman-Bacon, 2021; Sun and Abraham, 2021). I avoid this issue by allowing all coefficients in Equation 1.6 to vary arbitrarily by cohort and then averaging the cohort-specific treatment effects explicitly with uniform weights as  $\beta_\tau = E[\beta_{c\tau}]$ .<sup>18</sup> This estimation approach also clarifies that the identification of  $\beta_\tau$  comes solely from aggregating comparisons between workers that are in the same cohort. For more details, see Appendix Section A.3.2.

I make several sample restrictions on the set of movers to ensure that I capture the effect of interest. First, I restrict the estimation sample to workers who were continuously employed for the three years before the move to facilitate the assessment of pre-trends and to ensure that my sample consists of workers with a high degree of attachment to the formal labor force. I additionally restrict the sample to workers who are between the ages of 25 and 50 at the time that they switch firms. This ensures that my sample includes ages where promotions are most relevant, but worker outcomes are unlikely to be driven by external shocks like schooling and retirement.<sup>19</sup> Finally, I focus on workers for whom promotions would be relevant – production workers working in an occupational group with a directly observable supervisor track in the year before the move. Some employer-to-employer transitions in the data appear to reflect firm organizations or spinoffs, so I use a worker-flows approach to eliminate employer changes that appear to be spurious.<sup>20</sup> The remaining transitions reflect worker-level variation and do not require higher clustering under the framework of Abadie et al. (2017), but I cluster standard errors conservatively at the destination firm level to allow for arbitrary correlations in outcomes between workers moving to the same firm and over time.

### 1.4.3 Extensions incorporating additional variation

The exogenous job movers assumption that is also sufficient for my identification assumption makes strong restrictions on worker mobility. Although the restrictions

---

<sup>18</sup>This approach is algebraically identical to estimating Equation 1.6 separately for each cohort. Estimating the coefficients jointly in this approach ensures that the standard errors are correct when clustering across cohorts.

<sup>19</sup>Figure A-3 plots the empirical distribution of supervisors and promotions. Approximately 70% of all promotions are between the ages of 25 and 50.

<sup>20</sup>Specifically, I follow the mass layoffs literature (e.g., Schmieder et al., 2020) by dropping any origin firms where at least 20% of exiters went to the same destination firm.

are generally consistent with the data, researchers have also pointed out that worker sorting can violate the identifying assumptions but generate similar empirical patterns (for example, see Eeckhout and Kircher, 2011). To address these concerns, I extend the baseline design to incorporate two sources of aggregate variation that relax the two key components of the assumption – that workers separate from their jobs for exogenous reasons, and that conditional on separating, the type of firm the worker moves to is exogenous. In the framework of Abaluck et al. (2021), I test the balance and fallback conditions, respectively. It is worth noting that my approach is not specific to promotions or the Brazilian data, and it may be generally useful as tests for the validity of movers research designs.

To ensure that my effects are not driven by workers’ differential selection into moving employers, I follow Gibbons and Katz (1992) and use mass layoffs as an exogenous shock that separates workers from their current employers. I follow the methodology from Schmieder et al. (2020) to identify mass layoff events as firm events where employment dropped by at least 30% year over year and where no more than 20% of separated employees were re-employed at the same firm. Around 127,000 workers from 16,000 firms in my baseline sample were separated as a result of a mass layoff, which is 12% of the baseline sample. Column 2 of Table 1.5 summarizes the characteristics of the laid-off workers. Relative to the baseline sample, they are slightly older, have fewer years of education, and lower earnings, but they are more likely to move to high promotion firms.

I test for whether selection into moving firms is driving my results by estimating the baseline Equation 1.6 on the subsample of exogenously separated workers. It is worth emphasizing, however, that I compare laid-off workers who move to high promotion firms to laid-off workers who move to low promotion firms. So, I use mass layoffs as a particularly powerful instrument that shifts all affected workers into moving, but I assume that the effects of the mass-layoff itself are homogeneous and absorbed by my time-varying controls. This is in contrast to the literature on estimating the effects of mass layoffs by comparing laid-off workers to workers at other, non-layoff firms.<sup>21</sup>

To ensure that my effects are not driven by workers differentially sorting into high promotion firms based on idiosyncratic shocks, I consider aggregate regional variation in the types of employers that are hiring workers (which is similar in spirit to Oreopoulos et al., 2012). I construct  $z_m$  for each cohort as the jack-knife share

---

<sup>21</sup>Of course, a concern with this approach is that mass layoffs themselves differentially affect workers. This concern is mitigated in my setting by the fact that my comparison is within the set of mass layoff workers, and also by the fact that any worker who is included in the employer-to-employer mover sample would have been unemployed for less than 12 months following the mass layoff.

of hiring at the worker’s municipality that is from high promotion firms rather than low promotion firms, excluding any hires from the worker’s destination firm.<sup>22</sup> I then instrument  $H_i$  with  $z_m$  by estimating the following just-identified system of equations using two stage least squares:

$$\begin{aligned} H_i \times I_t^\tau &= \phi_\tau(z_m \times I_t^\tau) + \tilde{\pi}X_{it} + \tilde{\alpha}_i + \tilde{\theta}_t B_{i,c-1} + \tilde{\gamma}_{sot} + \eta_{it} \\ y_{it} &= \sum_{\tau} \beta_\tau(H_i \times I_t^\tau) + \pi X_{it} + \alpha_i + \theta_t B_{i,c-1} + \gamma_{sot} + \varepsilon_{it}. \end{aligned} \tag{1.7}$$

The second stage equation is identical to the baseline Equation 1.6, including the use of worker fixed effects and flexible controls for skill prices. Since variation is at the municipality level, I differ from the other specifications by clustering standard errors by the worker’s destination municipality.

Figure A-4 shows graphical intuition for my instrument by plotting the average  $H_i$ , the probability that a worker moves to a high promotion firm, against binned values of  $z_m$ , the jack-knifed municipal hiring share. I control for observable differences in worker characteristics (a quadratic in age interacted with gender along with indicators for the worker’s education, race, and state), the firm wage premia of the worker’s origin firm, and state-by-occupation fixed effects.  $z_m$  has clear predictive power and the conditional expectation function is approximately linear. As an additional check, I estimate the cross-sectional regression

$$H_i = \phi z_m + \tilde{\pi}X_{ic} + \tilde{\theta}B_{i,c-1} + \tilde{\gamma}_{so} + \nu_i \tag{1.8}$$

separately for each cohort. The minimum F-statistic for the instrument across all five cohorts is 627, and the average F-statistic is 1161.

The key identifying assumption for the instrumental variables approach is:

$$E[\varepsilon_{it}z_m] = 0.$$

In words, I assume that changes in worker outcomes that are unexplained by my time-varying controls are uncorrelated with the worker’s municipality. This municipality-level assumption substantially relaxes the worker-level assumption in Section 1.4.2

---

<sup>22</sup>The workers in my analysis sample contribute to the total number of new hires, which introduces a functional dependence between  $H_i$  and the hiring share  $z_m$ . Although this problem is minor in my setting since municipalities in the sample are generally large, I exclude any hiring from the worker’s destination firm from  $z_m$  to avoid the reflection problem that would otherwise arise (Manski, 1993). Technically, the jack-knife procedure introduces some slight variation in  $z_m$  by the worker’s destination firm, but I slightly abuse notation to make the source of identification clear.



by allowing workers to sort arbitrarily across firms within municipalities, including in anticipation of worker-level changes. However, this assumption also imposes the additional exclusion restriction that  $z_m$  is uncorrelated with any other unobserved factors that may shift worker outcomes. This concern is mitigated by the rich state-by-occupational-group-by-time fixed effects and testable pre-trends. Moreover, I focus on assessing whether some firms are more likely to generate positive outcomes (promotions) as well as negative outcomes (exits). To the extent to which regions with more hiring by high promotion firms experience different local labor market shocks, these additional shocks are unlikely to increase both worker promotions and formal labor force exits. Nevertheless, potential violations of the exclusion restriction are an important caveat to interpreting the IV results.

## 1.5 Direct effects of high promotion firms

My estimates match the model’s predictions on the direct effects of employer learning: high promotion firms have persistently positive effects on worker promotions and negative effects on formal labor market attachment. Consequently, these firms have only negligible long-term effects on average worker earnings. I begin by reporting the estimates from the baseline research design described in Section 1.4.2, which are my preferred estimates since they most directly reflect the effects for the typical worker making a job-to-job transition. To address any remaining concerns that my results may be driven by unobservable differences between workers, I next discuss the results of a series of robustness checks and alternate identification strategies. I show that including rich time-varying trends by additional worker characteristics or removing most of the controls from the baseline specification has little effect on my estimates. I also show that alternative identification strategies that incorporate quasi-experimental shifters of workers’ firm choices, as discussed in Section 1.4.3, produce estimates that are quantitatively and qualitatively similar to my baseline results.

### 1.5.1 Baseline results

My first result is that differences in firm promotion rates reflect persistent differences in firms’ causal effects on workers. Figure 1-6 plots the estimates of  $\beta_\tau$  from estimating Equation 1.6 on the outcome of whether a worker is working as a supervisor  $\tau$  years after moving.<sup>23</sup> Workers are 1.19 (s.e. .09) percentage points more likely to be working

---

<sup>23</sup>For this outcome, I include all workers who leave the formal labor market, and I assume these workers are not working as a supervisor. The effects could be considered a lower bound if high

as a supervisor within the first two years after moving to a high promotion firm rather than a low promotion firm. The effects are large relative to the average promotion rate of 1.68 percentage points for movers to low promotion firms over the same period, as well as the overall annual promotion rate of 0.779 percentage points for production workers. Moreover, the effects are persistent. Workers are still 0.936 (s.e. 0.08) percentage points more likely to be working as a supervisor more than two years after initially moving, even though 77% of workers are no longer working at the destination firm by then.<sup>24</sup>

This result rules out two alternative explanations for the role of high promotion firms. First, promotions are not due to the realization of transient firm shocks. Firms are classified as high or low promotion firms based on their promotion rates for workers between 2004 and 2006, whereas workers in the analysis sample joined the firms between 2008 and 2012. The fact that firms that are active in promotions from the early data also have clear effects on workers who join the firm several years afterward indicates that these effects are driven by persistent differences between firms.<sup>25</sup> Second, promotions at high promotion firms cannot be determined by seniority rules alone. My identification strategy relies upon comparing workers around employer transitions, so all workers in the sample would be at the bottom of the seniority ladder. The fact that I detect large effects on promotions immediately after the worker moves firms indicates that firms are willing to consider relatively new workers for more senior positions.

My second result is that high promotion firms are also more likely to lead workers to leave formal employment altogether. Figure 1-7 plots the baseline event study specification for the outcome of whether the worker is working for *any* firm in the formal employment sector  $\tau$  years after moving. Workers are 1.93 (s.e. 0.25) percentage points less likely to be found in formal employment more than two years after moving to a high promotion firm rather than a low promotion firm. Although the effects are smaller in relative terms given that 21.4% of workers who move to low promotion

---

promotion firm movers who leave the formal labor market are also more likely to work as a supervisor in the informal labor market.

<sup>24</sup>Note, however, that these effects do *not* necessarily imply that the long term promotion effects are driven by a single group of workers who were immediately promoted. Figure A-5 plots the event study coefficients from the same estimating equation on the outcome of whether a worker was *ever* promoted by  $\tau$  years after moving. Promotions within a year of moving the firm account only for half of the cumulative number of promotions. So, the relative stability of the effects on promotions at  $\tau$  masks the fact that workers are entering and exiting the supervisor role at roughly similar rates in the later years.

<sup>25</sup>Moreover, Figure A-6 shows that these effects are not substantially different between the early and later cohorts.

firms also leave formal employment over the same period, the differential effect of high promotion firms on exit rates is twice as large as their effect on promotion rates. So, a production worker who joins a high promotion firm is substantially more likely to leave formal employment altogether than to become a supervisor.

The key prediction from Proposition 2 is that workers joining high learning firms are both more likely to be promoted and more likely to become unemployed. The results on employers' effects in this section confirm this prediction. In other words, *high promotion* firms seem to be *high learning* firms. Moreover, the details of the empirical results support the model's assumptions. I assume in the model that the matching of workers to firms is random. Focusing on a sample where this assumption is plausible (based on the evidence in Figure 1-5) and additionally isolating the effects of the high promotion firms help ensure that the connection between the model and the empirical support is tight. In addition, I model learning as the result of uninformed "trial promotions" that generate information about worker ability. The fact that promotion effects appear immediately after workers move firms and then partially decay over time is consistent with this formulation.

Finally, high promotion firms have only mixed effects on workers' long-term earnings. Figure 1-8 plots the event study estimates for the worker's average monthly earnings conditional on remaining in formal employment. Workers who move to a high promotion firm earn 1.26 (s.e. 0.51) log points more within the first two years of moving, which is consistent with higher overall wage policies at those firms. However, the short term wage gains dissipate quickly, and those workers earn no more than workers who initially moved to low promotion firms three years after the initial move. Given that earnings are defined only for workers who remain in formal employment, the estimates can be considered an upper bound on the overall effects on earnings, since workers who move to high promotion firms are also more likely to leave formal employment afterward.<sup>26</sup> The event study on earnings is also likely to be a particularly sharp test of my identifying assumptions. If workers who move to high promotion firms are more likely to have systematically higher productivity growth, then they are likely to have differentially higher earnings growth in the three years before moving; my results show that this is not the case.

---

<sup>26</sup>Table A.2 provides a lower bound on earnings effects by imputing earnings outcomes for workers who leave the formal labor force under the assumption that all workers who leave are unemployed. To implement this, I assume that counterfactual earnings for workers who leave the formal labor market would have grown at the average rate for their baseline occupation-by-state cell, but are discounted by the ratio between the value of non-work time and the market wage (estimated as 0.58 by Mas and Pallais, 2019).

## 1.5.2 Robustness of direct effects

Several checks support the interpretation that my estimates reflect high promotion firms' causal effects rather than differences in worker sorting or market level shocks. The estimated effects remain quantitatively and qualitatively similar when I use a simplified specification with only worker and year fixed effects or a rich specification with time trends that absorb additional differences between workers at baseline. The estimates are also comparable when I use shifters of worker separations or worker destinations, as discussed in Section 1.4.3, although these estimates are generally less precise. Finally, the effects do not depend on the details of how I classify firms, but the effects of firms with high promotion rates are distinct from the effects of firms with high wage growth.

Figure 1-9 compares the pooled estimates of the main outcomes across three different specifications of controls. Compositional differences between high and low promotion firm movers are small in magnitude and become largely indistinguishable from zero conditional on the worker's baseline geography, firm wage premia, and occupation, which motivated controlling for flexible trends at the state-by-occupation level and by baseline firm wage premia. However, dropping these additional controls only slightly decrease the estimated effects on workers' formal labor market attachment and earnings. Meanwhile, adding additional flexible trends for each gender-by-education cell, the worker characteristics for which the difference between movers was statistically significant, as well as flexible trends by baseline worker earnings does not meaningfully change the estimated effects. The stability of estimates across these controls is reassuring that my identification assumption appears valid.

Meanwhile, Figure 1-10 compares the pooled estimates of the main outcomes across the three sources of variation. The mass layoff estimates are from estimating Equation 1.6 on the subsample of movers whose baseline firm was experiencing a separation shock that is plausibly exogenous to the worker. The estimates are less precisely estimated than the baseline due to the smaller sample but are otherwise nearly identical. The estimated earnings effects are slightly higher than the baseline estimates, although the two are not statistically distinguishable.<sup>27</sup> Meanwhile, the local hiring estimates are from estimating the instrumental variables Equation 1.7. Variation is at the municipality level, and the standard errors are correspondingly larger, but the estimates remain comparable to my baseline results. The hiring instrument is strong in the full specification – the F-statistic for the joint significance

---

<sup>27</sup>Figures A-12, A-13, and A-14 replicate the main event study figures for this subsample. The outcomes are free of pre-trends and the dynamics are similar to that of the main specification.

of the excluded instruments in the pooled IV first stage regression is 119, so I calculate confidence intervals based on the standard asymptotically valid critical values (as suggested by Lee et al., 2021).<sup>28</sup> The magnitudes of the IV estimates on formal labor market attachment and earnings gains are both larger, and the IV estimate on promotions is again nearly identical to my baseline estimates.<sup>29</sup>

Finally, I find that the results do not depend on the details of how I classify firms based on promotion rates, but the use of actual promotions data is critical. Figures A-9, A-10, and A-11 replicate the event study estimates of the baseline Equation 1.6 using the continuous measure of the firm’s promotion propensity  $\eta_j$  (as defined in Section 1.4.1 and winsorized at the 5th and 95th percentiles) instead of the discrete high versus low classification. The qualitative patterns are nearly identical when using this continuous measure, and the magnitudes of the estimates are consistent with the overall difference in average promotion rates between high and low promotion firms. Another possible concern is that the classification process includes both external and internal promotions, so the estimated effects are not necessarily representative of firms that focus on developing their workers. I test for this explanation by replicating my baseline empirical strategy on an alternate definition of high and low promotion firms that uses information from internal promotions alone. Figure A-18 compares the pooled estimates from this alternative classification, and the effects are again similar.<sup>30</sup>

To highlight the importance of direct occupational data, I also benchmark the effects of high promotion firms against the effects of high *wage growth* firms (defined as firms in the top quartile of wage growth for incumbent workers between 2004-2006) in Figure A-18. Unsurprisingly, firms with high wage growth do increase the earnings of their workers, but they have almost no effect on workers’ career progressions and they have similarly negative effects on the probability that workers remain in formal employment. This suggests that the effects of high wage growth firms may primarily be driven by other channels of wage determination like bargaining, on-the-job search, or seniority pay.

---

<sup>28</sup>The F-statistic for the earnings outcome is 105. It is slightly different since the sample is conditional on remaining in formal employment.

<sup>29</sup>Figures A-15, A-16, and A-17 replicate the main event study figures for the IV specification. The effects are noisier (especially for the promotion and exit outcomes), but the dynamics again remain similar. There is some evidence that workers in regions with more high promotion employers are more likely to be working as a supervisor in the years before the move, but this difference is less than a third of the immediate promotions effects. Importantly, those workers also do not appear to be on differential earnings trajectories in the years before the move.

<sup>30</sup>The similarity of these two effects is unsurprising since most of the promotions that determined a firm’s promotion propensity were internal promotions. For more details on classification, see Appendix A.3.1.

## 1.6 Additional evidence for the employer learning mechanism

The analysis in Section 1.5 shows that the effects of high promotion firms on worker promotions, labor force exit, and earnings are robust to alternate specifications and unlikely to be driven by worker sorting. Furthermore, given the setup of the empirical strategy, the effects are unlikely to be explained by transient shocks or seniority rules at high promotion firms. In this section, I use a variety of additional empirical exercises to show that the effects of high promotion firms are consistent with employer learning and inconsistent with several other alternative explanations.

### 1.6.1 Promotion effects reflect positive worker outcomes

I empirically examine whether differences in promotions reflect differences in real outcomes by comparing the promotion wage premium for workers who move to high promotion firms versus low promotion firms. To do so, I extend Equation 1.6 and estimate

$$y_{it} = \beta_{\tau}^{HP} (H_i \times P_{i,c+1} \times I_t^r) + \beta_{\tau}^H (H_i \times (1 - P_{i,c+1}) \times I_t^r) + \beta_{\tau}^P (P_{i,c+1} \times I_t^r) + \pi X_{it} + \alpha_i + \theta_t B_{i,c-1} + \gamma_{sot} + \varepsilon_{it} \quad (1.9)$$

where  $P_{i,c+1} = 1$  if worker  $i$  is working as a supervisor one year after moving firms and  $P_{i,c+1} = 0$  otherwise. These estimates are descriptive rather than causal since promotions are outcomes from moving to high promotion firms. But the exercise is useful for testing for the patterns of selection that would result if high promotion firms are simply more likely to classify workers as supervisors.

Workers who are promoted following a move to a high promotion firm experience persistent increases in earnings that are at least as large as the promotion wage-premia at low promotion firms. Figure 1-11 plots the estimates of  $(\beta_{\tau}^{HP}, \beta_{\tau}^H, \beta_{\tau}^P)$ , the relative wage changes for promoted high promotion firm movers, unpromoted high promotion firm movers, and promoted low promotion firm movers, respectively. The omitted reference group is the unpromoted low promotion firm movers. Promotions are meaningful for these workers – high promotion firm movers who are promoted within a year of moving receive wage increases of 15.6 (s.e. 0.97) log points, while corresponding low promotion firm movers receive wage increases of 11.9 (s.e. 0.66) log points. These effects decay over time, but remain economically and statistically significant even seven years after the move. Meanwhile, high promotion firm movers

who were not promoted within a year of moving initially earn slightly more than corresponding low promotion firm movers, but their earnings are nearly identical more than two years after the move.

These results are inconsistent with the alternative explanation that the promotion effects reflect only surface-level firm differences in the use or reporting of job titles.<sup>31</sup> Under this explanation, both promoted and unpromoted workers at high promotion firms would be more negatively selected than the same groups at low promotion firms. As a concrete example, suppose the distribution of workers' wage growth  $\Delta w$  is  $F(\Delta w)$  for both high and low promotion firms, and a firm  $f$  labels a worker as a supervisor when  $\Delta w$  exceeds some cutoff  $c_f$ . A high promotion firm  $H$  is a firm with a lower cutoff for promoting workers (i.e.,  $c_H < c_L$ ), so it follows that

$$E[\Delta w | \Delta w > c_H] < E[\Delta w | \Delta w > c_L]$$

$$E[\Delta w | \Delta w \leq c_H] < E[\Delta w | \Delta w \leq c_L].$$

My results from Figure 1-11 are the opposite, which suggests that differences in promotions reflect real differences in job assignments (and corresponding wage structures) rather than differences in the willingness to report similar jobs as supervisory jobs.

### 1.6.2 Exits from formal labor force are involuntary

The main results in Section 4 show robust evidence that high promotion firms increase the likelihood that workers exit from the formal labor force. Although employer-to-unemployment transitions are a common proxy for involuntary turnover (for example, in Moscarini and Postel-Vinay, 2018; Sorokin, 2018), the large informal sector in Brazil may complicate their interpretation in my setting. I present two additional pieces of evidence that bolster the interpretation that increased exits reflect the higher use of separations by high promotion firms rather than alternate explanations like different outside options or voluntary quits.

The first result is that the higher formal labor force exit effects are consistent with higher firm separations at high promotion firms. Table A.2 presents the pooled effects of high promotion firms on the probability that the worker is still working at the destination firm. Workers are 4.47 (s.e. 0.82) percentage points more likely to

---

<sup>31</sup>Although the use of administrative data mitigates some of these concerns, worker occupations are self reported by firms, and there may be additional incentives to classify workers differently due to collective bargaining agreements or the use of job titles as status (Baron and Bielby, 1986; Lagos, 2019).

leave their destination firm more than two years after moving to a high promotion firm instead of a low promotion firm, which is more than double the effect on formal labor force exit. However, firm exit is a less direct measure of negative employment outcomes than formal labor force exit, so I use it as a supporting fact rather than the main outcome. Nevertheless, it is reassuring that the effects on formal sector exit can be rationalized by effects of separations from the destination firm.

In addition, I directly test for whether high promotion firms affect worker-initiated turnover and find a precise zero effect. A feature of RAIS is that it records the reason why an employment contract is terminated, including whether the termination was initiated by the firm or by the worker. This distinction matters in Brazil because the employer pays higher separation penalties for employer terminations without cause (including layoffs). I classify employer terminations with or without cause (i.e., firings or layoffs) as firm-initiated separations, and I classify voluntary quits as worker-initiated separations. I observe reasons for separation preceding around half of spells outside of formal employment, of which 77.6% are firm-initiated and 14.5% are worker-initiated.<sup>32</sup> Figure 1-12 reports the estimated effect of moving to a high promotion firm on the likelihood of each type of exit. Firm-initiated exits account for virtually all of these additional separations, whereas high promotion firms' effects on worker-initiated exits are precisely zero.

### 1.6.3 Promotion and exit effects are stronger for high potential workers

To test whether promotions mediate both positive and negative job outcomes, I compare the effects of high promotion firms for workers who are ex-ante more likely versus less likely to be promoted. Abadie et al. (2018) cautions that even using prior characteristics to predict subsequent outcomes can introduce bias due to overfitting, so I split my analysis sample into a 25% hold-out sample and a 75% estimation sample. I first estimate the equation

$$\text{Promoted}_{i,c+1} = \theta Z_{i,c-1} + \varepsilon_i \tag{1.10}$$

on the 25% hold-out sample of movers. I then rank workers in the 75% estimation sample by their predicted promotion potential  $(\hat{\theta} Z_{i,c-1})$ , and I separately estimate the baseline event study Equation 1.6 for the top and bottom tercile of workers by

---

<sup>32</sup>The remaining reasons for separations are primarily contract expirations, transfers, and retirements.



promotion potential. Table A.3 reports summary statistics for the high and low potential worker subsamples.<sup>33</sup> As expected, high potential workers have more years of education and higher baseline earnings than low potential workers. They are also more likely to work as supervisors and less likely to exit.

Figure 1-13 compares the main effects of high promotion firms on worker promotions, formal labor force exit, and average earnings for the baseline, high promotion potential, and low promotion potential samples. The fact that the effects on worker promotions are strongest for the high potential workers is reassuring but unsurprising. On the other hand, the fact that the effects on formal labor force exit are also larger for these high potential workers contradicts the alternative explanation that high promotion firms simply have more volatile labor demand. For example, firms with “dual labor markets” may offer some workers both job security and promotion opportunities and other workers precarious jobs that are eliminated during demand downturns. My results instead suggest that promotions and exits are related outcomes that are two sides of the same coin. These results lend further support to employer learning as the key mechanism that rationalizes the employer effects.<sup>34</sup>

#### 1.6.4 Survey of firm labor practices is consistent with estimated effects

Finally, I validate my worker-level evidence using structured interview data from an internationally comparable survey of manufacturing plants’ human capital management practices.<sup>35</sup> I view these data as supporting evidence that my results about the worker-level effects on promotions and exits are consistent with managers’ actions. Furthermore, the international comparability of the survey allows me to assess the extent to which the patterns are specific to Brazil

Two questions from the World Management Survey connect to the worker-level

---

<sup>33</sup>I define worker potential terciles prior to making the additional sample restrictions discussed in Section 1.4.2 to ensure that the definitions would be stable across alternate sample restrictions. This introduces some differences between the size the two subsamples but does not affect the interpretation of the effects.

<sup>34</sup>In the model, the effects on high potential workers can be rationalized by workers for whom trial promotions are more relevant (so  $Q$  is higher). On the other hand, it would be difficult to explain this pattern without learning (e.g., differences in the types of workers at different firms) since high potential workers are generally *less* likely to leave formal employment.

<sup>35</sup>There are general caveats to interpreting survey responses since managers’ reported promotions practices may reflect over-optimism or worker sorting rather than the firm’s true causal effects. Some of these concerns are mitigated by the survey I use, which comes from a series of structured interviews between MBA students and plant managers that are designed to minimize incentives to misreport firm practices. For more details about the survey, see Bloom et al. (2014).

outcomes that I study. Table A.4 describes each question from the questionnaire and its corresponding scoring criteria. The question on promotions from the questionnaire is primarily about the *reason* for promotion rather than the *frequency* of promotions. However, in my setting, a firm that primarily promotes workers based on tenure should have negligible estimated effects on promotions since my sample focuses on workers who move firms and consequently are at the bottom of the seniority ladder. Meanwhile, the question on firings directly connects to my results on formal labor force exits given the discussion in Section 1.6.2.

Figure 1-14 plots the average firings score by firms' promotions score for three groups of countries. First, I find that firms that are more likely to develop and promote high performing workers are also more likely to fire low performing workers. This is consistent with my baseline empirical results, as well as the general finding in the management and organization literatures that optimal personnel policies tend to be a suite of complementary practices (Milgrom and Roberts, 1992; Ichniowski et al., 1997; Bloom et al., 2019; Benson et al., 2019; Cornwell et al., 2021).<sup>36</sup> Second, the correlations between firms' promotions practices and firing practices are similar between Brazil, OECD countries, and non-OECD countries.<sup>37</sup> These three groups differ substantially in their labor market institutions, industrial compositions, and the allocation of production factors. The similarity of the relationship across all three settings further supports the interpretation that my results are driven by a general economic mechanism rather than any features that are specific to Brazil.

## 1.7 Equilibrium Effects of Employer Learning

In addition to direct effects on the labor market outcomes of their employed workers, high learning firms also have equilibrium effects on the wage structure for all workers. To fully explore these implications, I endogenize the employment share of high learning firms and the secondary market contact rate from the partial equilibrium model in Section 1.2. I follow Lise and Robin (2017) in assuming that firms set the number of vacancies based on expected profits, and that the total cost of vacancies is convex.

---

<sup>36</sup>The correlation between these two practices is also notable since analyses of internal labor markets usually argue that investment is more likely when turnover is low (for example, see Doeringer and Piore, 1971; Prendergast, 1993).

<sup>37</sup>The set of OECD countries in the sample are Australia, Canada, Chile, Columbia, France, Germany, Great Britain, Greece, Italy, Japan, New Zealand, Poland, Portugal, Ireland, Spain, Sweden, United Kingdom, and the United States. The set of non-OECD countries in the sample are Argentina, China, Ethiopia, Ghana, India, Kenya, Mexico, Mozambique, Myanmar, Nicaragua, Nigeria, Singapore, Tanzania, Turkey, Vietnam, and Zambia.

If the vacancy supply function is sufficiently convex (e.g., if it is log-convex), then high learning employers post a proportionately larger share of vacancies in more productive areas. A higher share of high learning employers increases the adverse selection of job movers, which decreases market wages for promoted and especially unpromoted workers. As a result, occupational wage inequality increases. I show that these equilibrium predictions also hold robustly in the data – municipalities with more high learning employers have higher wage differentials between promoted and unpromoted workers even after controlling for a wide range of worker-, job-, and municipality-level characteristics.

### 1.7.1 Endogenizing vacancy creation

In addition to the baseline setup, suppose that locations differ in a multiplicative productivity term,  $\psi$ , so output in region  $r$  for a worker of type  $\theta_i$  in occupation  $o$  is

$$f_{roi} = \psi_r f_o(\theta_i).$$

There is an equal mass of atomistic firms, each choosing a density of vacancies  $v$  at a total cost of  $c(v)$  at the start of the period. The matching technology is constant returns to scale and does not distinguish between the type of firm offering the vacancy or the type of worker. Initial job matches are formed through matching workers with vacancies. Similarly, contact with the secondary market is established by matching initially employed workers with initially unfilled vacancies, so the number of workers who will make contact with the secondary market can be written as

$$m^S = M(m^I, v - m^I),$$

where  $m^I = M(l, v)$  is the mass of initially employed workers,  $l$  is the total mass of workers,  $v$  is the total mass of vacancies, and  $M$  is the matching function. I assume that, conditional on arriving at the secondary market, firms still compete for workers as in Section 1.2.2.

Implicit in this setup are two simplifying assumptions to ensure the model remains tractable. First, I rule out any subsequently vacated positions (e.g., from a firm that fired its initially matched worker) from also joining the secondary market so that employment shares always correspond to vacancy shares. Second, I rule out any initially unmatched workers joining the secondary market so the initial level of vacancies will not affect the equilibrium degree of adverse selection.

The Perfect Bayesian equilibrium in this extended model is a set of worker and firm strategies such that each firm sets the number of vacancies that maximizes expected profits given the vacancy setting decisions of other firms, and such that the conditions in the partial equilibrium model in Section 1.2.2 are also satisfied. Proposition 3 compares the different outcomes that arise in equilibrium between regions that differ only in general productivity.

**Proposition 3.** *Suppose (i) vacancy creation as a function of expected profits is log-convex and (ii) the elasticity of re-employment with respect to the number of high learning employer vacancies exceeds the elasticity of expected output for unknown secondary market workers. Then, in more productive regions: (i) high learning employers post a greater share of vacancies, (ii) offered wages for incumbent workers are higher (iii) wage differentials between promoted and unpromoted workers are larger.*

The full proof for Proposition 3 is in Appendix A.2 and relies on two assumptions. The first assumption ensures that vacancy creation is sufficiently responsive so that increases in the difference in expected profits between high and low learning employers will increase the share of total vacancies created by high learning employers. The second assumption ensures that the net effect of an additional vacancy from high learning firms is positive for workers' expected wages.

The need for the second assumption highlights the negative equilibrium effects of high learning employers on wages. From Equation 1.4, offered wages at incumbent firms match expected offers in the secondary market. Holding the share of high learning employers equal, increasing the total number of vacancies would raise expected outside options by increasing the likelihood that workers encounter the secondary market. On the other hand, holding the total number of vacancies equal, increasing the share of high learning employers would exacerbate the adverse selection of workers entering the secondary market. Secondary market wages for unknown workers would correspondingly fall, which lowers expected outside options for promoted and especially unpromoted workers. The second assumption in Proposition 3 ensures that the positive wage effects of the vacancy dominate the negative wage effects of high learning employers for unpromoted workers. However, wage inequality would increase in either case since promoted workers are more insulated from the beliefs of the secondary market due to the informativeness of their job assignments.<sup>38</sup>

---

<sup>38</sup>Although I motivate differences in high learning employer shares as arising from vacancy creation in response to productivity differences, neither of those components are crucial for this mechanism.

## 1.7.2 Testing for equilibrium effects on occupational wage structure

Proposition 3 also yields testable implications about the relationship between a region’s high learning employer share and its market-level wages. As an empirical analog, I define

$$\bar{H}_m = \frac{\sum_{Q_f=Q_H} L_f}{\sum_{Q_f=Q_H} L_f + \sum_{Q_f=Q_L} L_f}$$

as the (employment-weighted) share of high promotion firms in a municipality, and I estimate the regression

$$y_{imt} = \beta_1 \bar{H}_m + \beta_2 [\bar{H}_m \times P_{it}] + \theta_1 X_{it} + \theta_2 (Z_m, Z_m \times P_{it}) + \sigma_{op} + \psi_f + \gamma_{sot} + \epsilon_{it} \quad (1.11)$$

on a longitudinal 5% sample of all promotion-track workers between the ages of 25 and 50 in Brazil.  $P_{it}$  is an indicator for whether the worker is working as a supervisor,  $X_{it}$  controls for worker characteristics (a quadratic in age interacted with gender along with indicators for the worker’s education, race, and state), and  $Z_m$  controls for region size. Finally,  $\sigma_{op}$  is a full set of occupational-group-by-supervisor fixed effects,  $\psi_f$  are firm fixed effects, and  $\gamma_{sot}$  are state-by-occupational-group trends.

Table 1.6 reports the estimates of  $(\beta_1, \beta_2)$ , the correlation between regional wages and the high promotion employment share, as I progressively add the controls in Equation 1.11. Both  $\beta_1$  and  $\beta_2$  are positive and substantial in the basic specification that controls only for flexible trends in wages at the state-by-occupation level. Adding additional controls generally has little effect on the estimates despite their substantial explanatory power on the variance in wages, which suggests that the estimates of  $(\beta_1, \beta_2)$  are not driven by observable differences between the characteristics of jobs, workers, or municipalities. The one exception is that the estimate of  $\beta_1$ , the overall regional wage premium, becomes small and insignificant upon the inclusion of firm fixed effects, but the regional supervisor wage premia remain substantial. Under my preferred specification that controls flexibly for worker characteristics and different wage structures within occupations (column 3), moving from the 10th to the 90th percentile in the population-weighted municipal high promotion firm share (from 17.3% to 50.8%) increases regional wages for production workers by 4.69 log points and the region wage premium for supervisors by an additional 6.57 log points, which is 39% of the average supervisor wage premium from Table 1.2.

---

The equilibrium channel where high learning employers suppress wages and increase wage inequality would exist whenever wage determination follows the structure from the partial equilibrium model.

## 1.8 Structural quantification

The model in Section 1.2 and its extension in Section 1.7 are both highly stylized to clarify the mechanism of employer learning and characterize its equilibrium implications. However, the model’s basic components governing learning and job assignment also have empirical content, and they can be mapped to the data under minimal assumptions about wage setting or the labor market parameters. In this section, I show how those key parameters can be identified using the employer effects estimates under internally consistent model restrictions. The resulting estimates allow me to quantify the degree of employer learning by high and low promotion firms as well as the implied degree of skill misallocation and adverse selection in my setting.

### 1.8.1 Model identification and estimation

The partial equilibrium model in Section 1.2 can be separated into two parts. The first half of the model characterizes the incumbent employer’s learning and job assignment decision while taking market wages as given. The second half of the model specifies the equilibrium wage from the secondary market. This separability between the two parts is crucial for transparently mapping my estimates to the learning and job assignment problem while remaining agnostic about the exact form of wage setting. The key condition is that employers must find the efficient job assignments to be incentive compatible. In other words, employers must find it optimal to promote workers of known high ability, fire workers of known low ability, and retain workers of unknown ability. This incentive compatibility condition would be satisfied under a variety of wage setting institutions that would be relevant in Brazil, including rent sharing, wage bargaining, wage floors, or firm-wide wage schedules. As a result, predictions from the learning and job assignment problems are particularly likely to be robust to departures from the stylized setting of the model, and I focus on quantifying their relevant parameters.

To further bolster the transparency and robustness of the quantification exercise, I use a methods of moments approach for estimation. The approach allows me to be explicit about the exact moments from the data that identify the parameters and the relationship between the two (Andrews et al., 2020). I can also choose the moments to reflect the model’s assumptions about timing and heterogeneity (similar to the approach from Lamadon et al., 2019). Table 1.7 summarizes the parameters of interest and the moments used to identify the parameters.

I first calibrate two moments “outside” of the model. I use the probability that

a promoted worker remains a supervisor upon moving employers to calibrate the asymmetric information parameter  $\kappa$ . I also use the average job destruction rate estimated in Dix-Carneiro et al. (2021), aggregated to the average length of the outcome period in the sample, to calibrate the exogenous separation rate  $\delta$ . Since there is greater uncertainty in these moments, I systematically vary the calibrated values to ensure that my conclusions are not sensitive to the choice of these values.

I jointly identify the remaining parameters governing job assignment and turnover  $(\bar{Q}_H, \bar{Q}_L, \alpha, g)$  by matching the sample means and treatment effects from the movers analysis using classical minimum distance. The key model restriction that allows me to identify the model is that I assume movers to high versus low promotion firms differ only in the rate of learning at their destination firms. This restriction ensures that the model analogs for the means and treatment effects estimates for promotions and formal labor force exit are determined by the four parameters alone.<sup>39</sup> In theory, I can calibrate the model using observational data from all workers in my sample. However, the comparability of baseline characteristics and the absence of systematic pre-trends in earnings between high promotion firm and low promotion firm movers (as discussed in Sections 1.4.2 and 1.5) suggest that the model restriction is most plausible in the movers analysis. I use the sample means and treatment effects from at least two years after workers initially move firms to allow enough time for learning and turnover to take place.

It is straightforward to implement the quantification approach since the model-implied analogs of the moments can be computed in closed form. The estimate  $\hat{\theta} = (\bar{Q}_H, \bar{Q}_L, \alpha, g)$  is defined as

$$\hat{\theta} = \arg \min_{\theta} (\hat{\pi} - h(\theta))' W (\hat{\pi} - h(\theta)), \quad (1.12)$$

where  $\hat{\pi}$  is the vector of empirical moments,  $h(\theta)$  is the vector of corresponding model outcomes (given  $\theta$ ), and  $W$  is a positive definite weighting matrix. Since  $\hat{\pi}$  comes directly from sample calculations, I calculate its full cluster-robust variance-covariance matrix  $\hat{\Omega}$  by stacking the estimation equations and clustering the stacked equation by the worker's destination firm. I use the inverse of the variance-covariance matrix as my weighting matrix (so  $W = \hat{\Omega}^{-1}$ ), although I can generally match the empirical

---

<sup>39</sup>As an example, suppose that both sample means and treatment effects on promotions are large, but sample means and treatment effects on labor force exit are small. The moments on promotions would imply that the relative difference in learning between high and low employers is large, overall learning is sufficiently high, and the share of high ability workers in the population is large. The moments on employment would then distinguish between a high overall rate of learning versus a large share of high ability workers and pinpoint the implied re-employment rate.

moments exactly so the choice of the weighting matrix affects only the efficiency of my estimator. Asymptotically valid standard errors for  $\hat{\theta}$  also follow from  $\hat{\Omega}$  using the sandwich estimator from Newey and McFadden (1994).<sup>40</sup> For more details, see Appendix A.3.3.

## 1.8.2 Parameter estimates

Table 1.8 summarizes the parameter estimates under my baseline calibration, which have intuitive interpretations given the simple structure of the model. I estimate from  $\alpha$  that 21.4% of workers in the sample are of high ability. Given that the average share of workers who are in supervisory roles at least two years after initially moving firms is 2.8%, my estimate implies that a high ability worker has an 86.9% chance of being mismatched to a production job or outside formal employment.<sup>41</sup> Part of the reason for the misallocation is because turnover is reasonably high, and information about worker quality is likely to be lost when workers change firms. Under the estimated parameters, a worker who is promoted by their incumbent employer has only a 75.3% chance of remaining in the supervisor position after potential separation shocks and secondary market matching are realized.

But the primary reason for the high rate of mismatch is that the estimated rate of learning for even high promotion employers is 20.8%, so most workers' abilities remain unknown to employers. Although learning is generally low in this setting, it is still substantial compared to the overall rate of worker turnover. As a result, I estimate that adverse selection can be a meaningful source of ex-post market power – workers who change employers are only 70.8% as likely to be of high ability as the overall population of workers in this sample.

Finally, the structural model also gives me a framework to quantitatively assess the overall degree to which high and low promotion firms differ. I estimate that high promotion firms, on average, are 5.8 percentage points more likely than low promotion firms to learn about the ability of their workers. This difference is larger than the treatment effects on either promotions or formal labor force exit, as well as the naive

---

<sup>40</sup>This approach does ignore the uncertainty due to the calibrated parameters, as discussed by Cocci and Plagborg-Møller (2019). I assess the sensitivity of  $\theta$  to calibrated values by explicitly assuming alternative calibrations to account for model uncertainty in addition to the (known) sampling uncertainty of the sample means.

<sup>41</sup>A stronger interpretation of the results would be that under perfect information, 21% of the sample would be working as supervisors. However, I do not model the full organizational structure of the firm, and diminishing returns would imply that the actual counterfactual share of supervisory workers would be lower. I view these calculations as more indicative of the degree of imperfect information that a high ability worker faces in the current market.



sum of the two effects. Intuitively, the estimated long-run effects on promotions and exit are attenuated by worker turnover and re-employment. The structural model provides a principled approach to adjust for this attenuation using the labor market parameters, which allows me to map differences in long-run outcomes to differences in firm behavior.

Finally, Figure A-19 shows how the main conclusions change when I vary the two calibrated parameters by 50% to 150% of the assumed baseline values. The assumed value for the asymmetric information parameter  $\kappa$  has a quantitatively small effect on my estimates. On the other hand, the assumed rate of exogenous separations  $\delta$  is important for inferring whether a job mover is a low ability worker who was previously fired or a high ability worker who exogenously separated. Moreover, at 50% of the baseline assumed value for  $\delta$ , my model can no longer exactly fit the sample means and treatment effect estimates, so those estimates should be interpreted with caution. For all other assumed values of  $\delta$ , my estimates for skill misallocation and adverse selection remain similar. Since the baseline assumed value of  $\delta$  is likely a lower bound, the robustness of my conclusions to higher values of  $\delta$  is particularly reassuring.<sup>42</sup>

## 1.9 Conclusion

I show that employers do influence workers' subsequent careers. Production workers who join high promotion firms are more likely to eventually work as supervisors, but they are also more likely to leave formal employment altogether, and they do not earn more on average. I argue that these results are most consistent with the explanation that employers' information about worker ability at the time of hiring is imperfect, and that employers vary systematically in the degree to which they learn about their workers' abilities. Both treatment effects estimates and a structural quantification show that employer learning is key to rationalizing the effects of employers on worker promotion and exit.

More generally, this paper highlights the role of information, particularly asymmetric information, in determining worker outcomes. Firms acting on information revealed after hiring introduce worker-level risk in employment that is separate from firm-level or market-level risks. This additional risk has implications for the design of employment

---

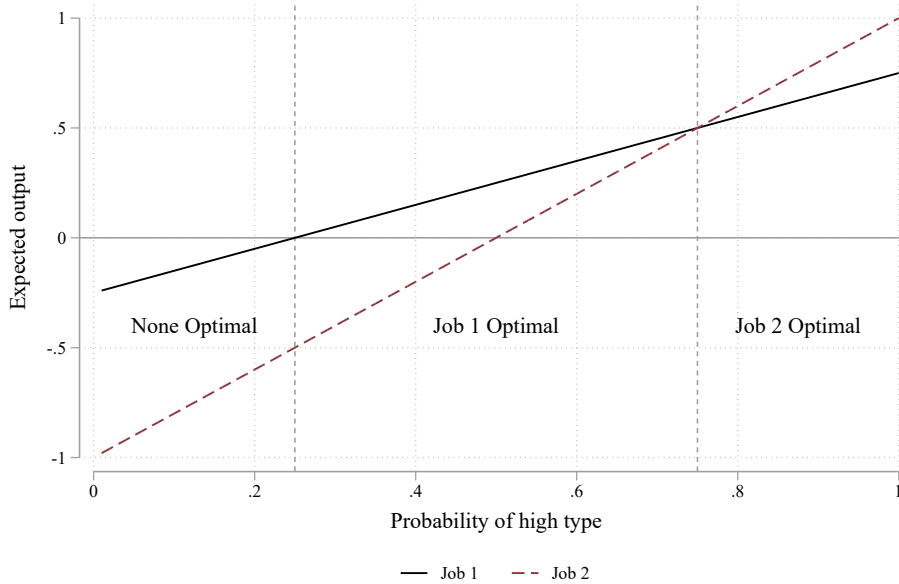
<sup>42</sup>The external estimate for  $\delta$  is the estimated *firm* destruction rate from Dix-Carneiro et al. (2021), whereas the parameter in the model is the exogenous *worker* separation rate. If workers have a positive probability of exogenously separating without the job itself disappearing (e.g., if a worker leaves due to geographic reasons or idiosyncratic preference shocks), then the firm destruction rate would be an underestimate of the exogenous separation rate.

protections and unemployment insurance. Adverse selection in the pool of job movers also suppresses workers' outside options, implying that asymmetric information can be a substantial source of ex-post market power for employers. Finally, my model provides an example where the improvement in information about workers exacerbates labor market frictions when that information is privately held. This highlights that the overall effects of more sophisticated employment practices may differ from the direct effects.

The current analysis also points to several fruitful avenues for further research. Adding information on firm accounts can help quantify the rents generated by employer learning as well as the other implications of organizational design. Similarly, adding survey evidence on worker outcomes in the informal sector would help inform the overall welfare trade-offs of high promotion firms. In addition, unions are prominent institutions in Brazil, and it would be useful to assess the degree to which collective bargaining contracts mediate firm practices. Finally, there may be gender and racial differences in the overall effects of firms in my setting; these questions are outside of the scope of this paper, but they are important for understanding who is benefiting from good opportunities and who is getting shaken off the career ladder.

# Tables and Figures

Figure 1-1: Example of job output functions



Note:: The figure shows an example of two job output functions that generate a motive for promotions. Workers who are believed to be of high ability have higher expected output in the high risk occupation (job 2), whereas workers who are believed to be of low ability have higher expected outcomes in the low risk occupation (job 1). Both occupations provide negative output when workers are sufficiently likely to be of low ability, so the outside option is preferred.

Figure 1-2: Model Timeline

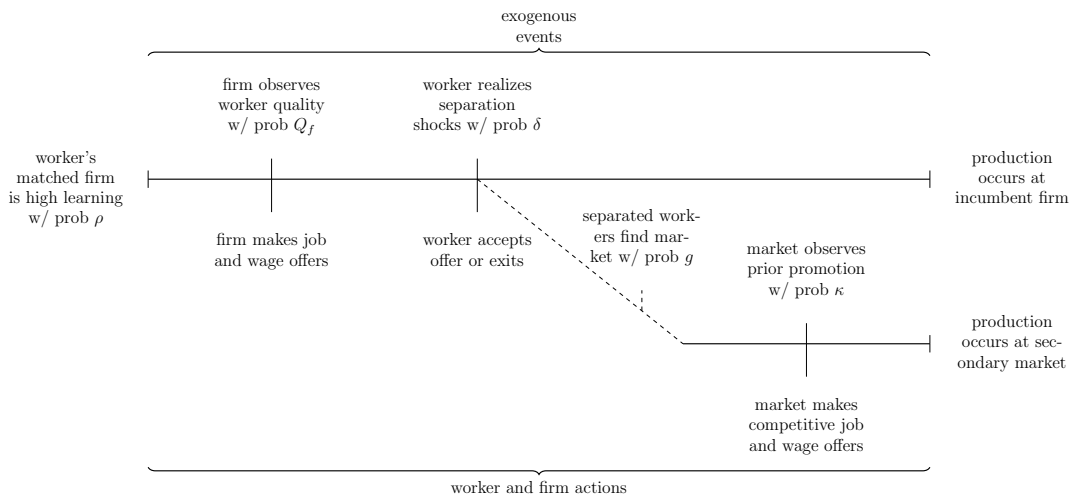




Figure 1-3: Prevalence of Promotions for Production Workers



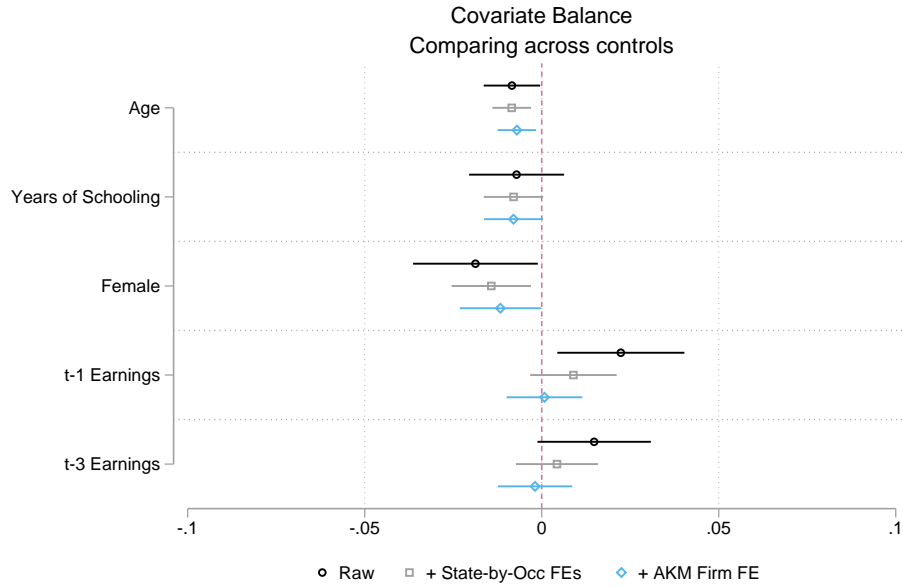
Note:: The figure plots the share of workers each year that advanced from an elementary production occupation in the previous year to a supervisor occupation in the current year. The “within occupation” series restricts the definition of promotions to purely vertical moves that are within the same occupational group. The sample is all workers in the Brazilian formal sector between the ages of 25 and 50. For more details, see Section 1.3.3.

Figure 1-4: Timeline of empirical approach




Note:: The figure describes the years of data used in each section of the empirical approach. The first four years of data (2003-2006) are reserved for estimating firms’ promotion propensities (see Section 1.4.1). The analysis sample consists of workers who moved firms between 2008 and 2012 and have three years of pre-move data (see Section 1.4.2). My data end in 2015, so the length of post-move outcomes ranges from 4 years (for the 2012 cohort) to 8 years (for the 2008 cohort).

Figure 1-5: Test for movers balance



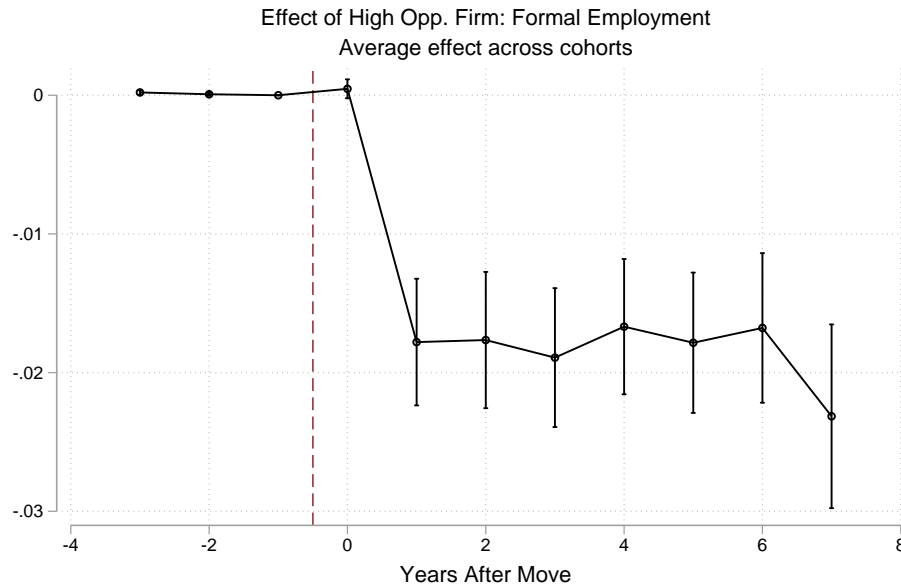
Note:: The figure plots the difference in average worker characteristics between those that moved to high promotion firms and those that moved to low promotion firms. Age and years of schooling are both scaled by their overall means, so the coefficients are interpreted as relative differences. The first set of estimates is the raw difference across all worker cohorts and only controls for possible differences in composition across worker cohorts. The second set of estimates adds state-by-occupation-by-cohort fixed effects. The final set of estimates also controls for the estimated firm wage premium at the worker's origin firm. Standard errors are clustered by the worker's destination firm, and lines indicate the 95% confidence interval.

Figure 1-6: Impact of high promotion firms on worker promotions



Note:: The figure plots event study coefficients from estimating Equation 1.6 on the baseline sample of all promotion track workers between the ages of 25 and 50 that made an employer-to-employer transition to a high or low promotion firm. The coefficients are the estimated yearly differences in the outcomes of workers who moved to a high promotion firm relative to workers who moved to a low promotion firm, averaged across all five cohorts of movers. The outcome is whether the worker is working as a supervisor and includes workers who have left formal employment (assumed to not be working as supervisors). Standard errors are clustered by the worker's destination firm, and lines indicate the 95% confidence interval.

Figure 1-7: Impact of high promotion firms on formal labor market attachment



Note:: The figure plots event study coefficients from estimating Equation 1.6 on the baseline sample of all promotion track workers between the ages of 25 and 50 that made an employer-to-employer transition to a high or low promotion firm. The coefficients are the estimated yearly differences in the outcomes of workers who moved to a high promotion firm relative to workers who moved to a low promotion firm, averaged across all five cohorts of movers. The outcome is whether a worker is working in formal employment. Standard errors are clustered by the worker's destination firm, and lines indicate the 95% confidence interval.

Figure 1-8: Impact of high promotion firms on log earnings for employed workers




Note:: The figure plots event study coefficients from estimating Equation 1.6 on the baseline sample of all promotion track workers between the ages of 25 and 50 that made an employer-to-employer transition to a high or low promotion firm. The coefficients are the estimated yearly differences in the outcomes of workers who moved to a high promotion firm relative to workers who moved to a low promotion firm, averaged across all five cohorts of movers. The outcome is log earnings and is only defined for workers in formal employment. Standard errors are clustered by the worker's destination firm, and lines indicate the 95% confidence interval.




Figure 1-9: Robustness of main estimates to controls

(a) Effects within 2 years of move



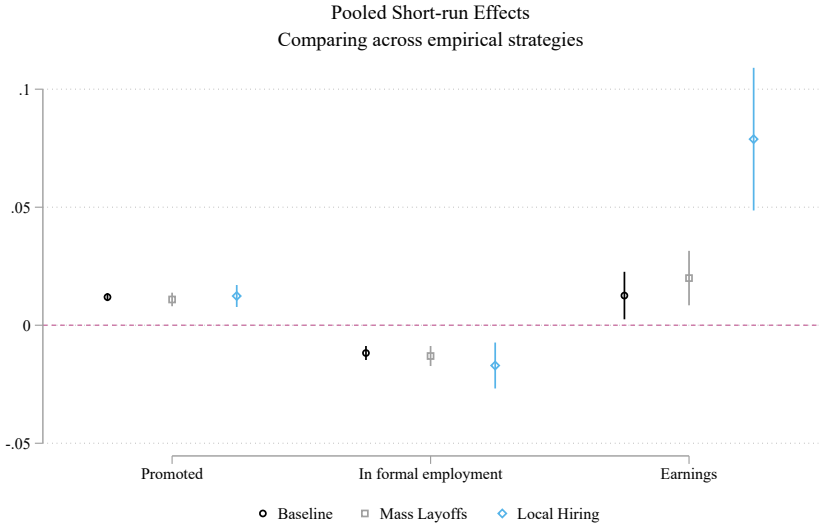
(b) Effects more than 2 years after move



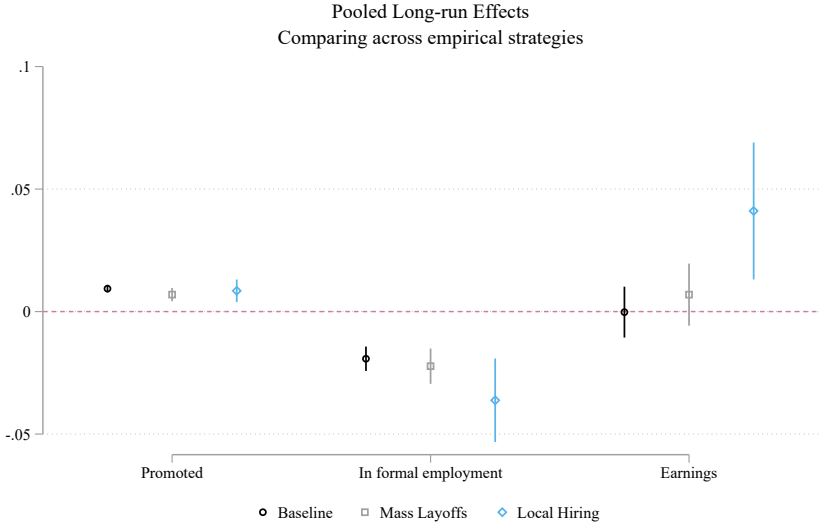
Note:: The figures compare pooled event study coefficients from estimating variations of Equation 1.6 on the baseline sample. The top panel plots the pooled effects within two years of moving, and the bottom panel plots the pooled effects more than two years after the move. The outcomes are whether the worker is working in a supervisor position, whether the worker is working in any job in formal employment, and the worker’s log earnings conditional on being in formal employment. Basic controls only include age, gender, education controls, and two-way worker and year fixed effects. Extended controls include all of the baseline controls, and additionally include flexible trends based on workers’ baseline earnings as well as time-varying changes to the gender and educational wage structure. Standard errors are clustered by the worker’s destination firm, and lines indicate the 95% confidence interval.

Figure 1-10: Comparison of main effects between identification strategies

(a) Effects within 2 years of move

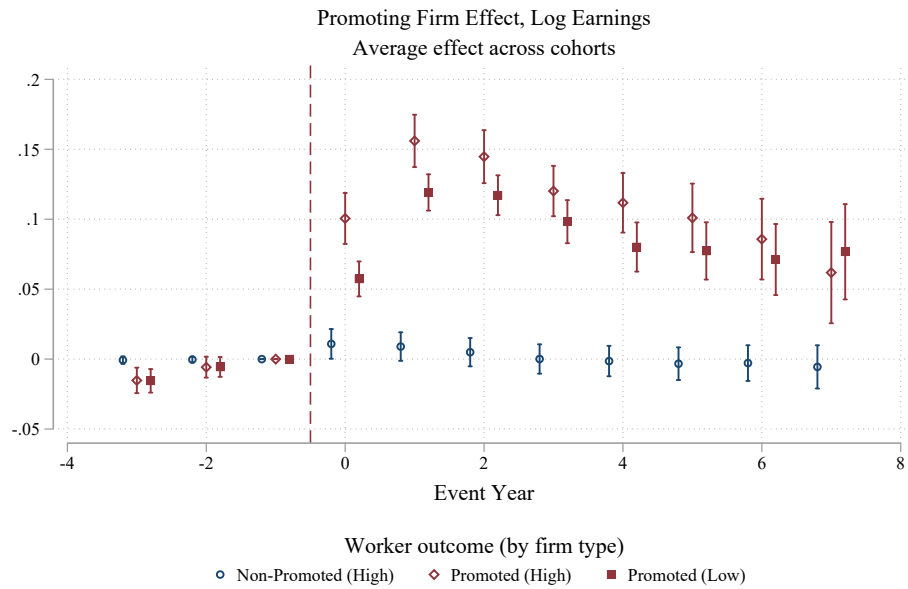


(b) Effects more than 2 years after move



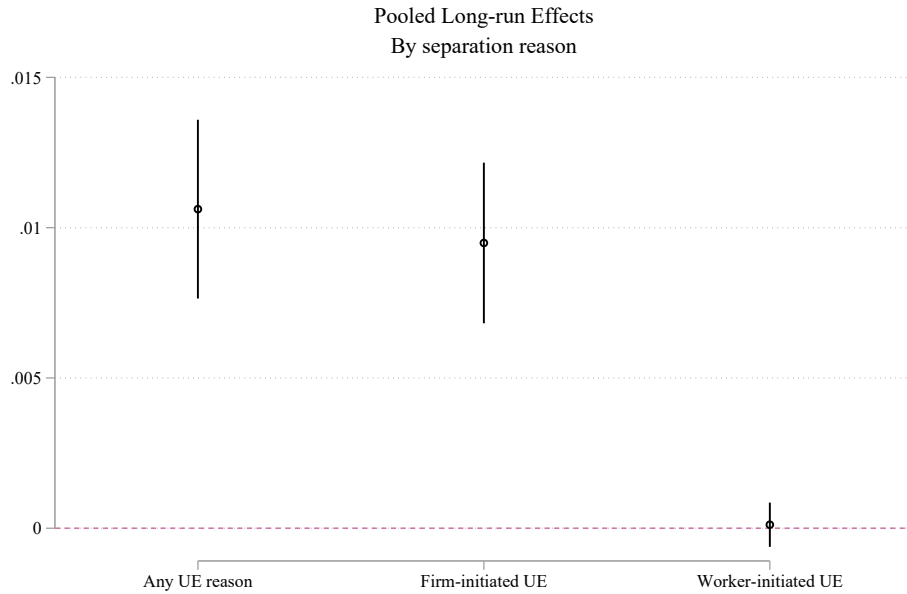
Note:: The figures compare pooled event study estimates from the baseline research design in Section 1.4.2 to estimates from the extensions in Section 1.4.3 that incorporate additional sources of variation. The top panel plots the pooled effects within two years of moving, and the bottom panel plots the pooled effects more than two years after the move. The outcomes are whether the worker is working in a supervisor position, whether the worker is working in any job in formal employment, and the worker’s log earnings conditional on being in formal employment. The mass layoff estimates restrict the sample to workers whose origin firms experienced mass layoff events at the time of the workers’ moves. The local hiring estimates are the instrumental variables estimates from using the jack-knifed municipal hiring share as an instrument for the worker’s destination firm. Standard errors are clustered by the worker’s destination firm for the mass layoff estimates and by the destination municipality for the local hiring estimates, and lines indicate the 95% confidence interval.

Figure 1-11: Differential impact of high promotion firms by worker job levels



Note:: The figure plots event study coefficients from estimating Equation 1.9 on the baseline sample. Workers are split into four groups based on the type of their destination firm as well as whether they were promoted within a year of moving. The plotted groups are high promotion firm movers who were not promoted, high promotion firm movers who were promoted, and low promotion firm movers who were promoted, respectively. The reference group is low promotion firm movers who were not promoted. The outcome is log earnings conditional on being in formal employment. Standard errors are clustered by the worker's destination firm, and lines indicate the 95% confidence interval.

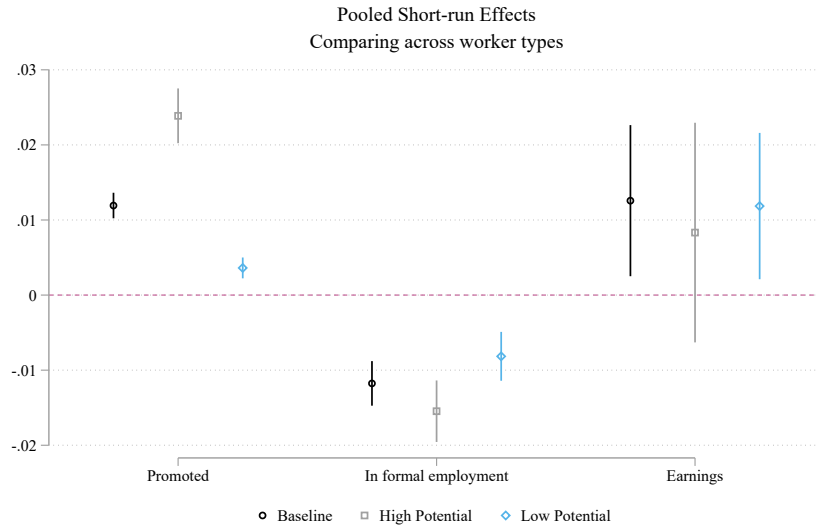
Figure 1-12: Impact of high promotion firms on each type of exit




Note:: The figure decomposes the baseline estimate of the effects of high promotion firms on formal labor market attachment by estimating Equation 1.6 on each type of worker separation. The coefficients are the estimated difference in the outcomes of workers who moved to a high promotion firm relative to workers who moved to a low promotion firm, averaged across all five cohorts of movers. The first outcome is a spell outside formal employment following any recorded reason for separating from their previous employer. The second outcome is a spell outside formal employment following an employer-initiated separation. The third outcome is a spell outside formal employment following a worker-initiated separation. Standard errors are clustered by the worker's destination firm, and lines indicate the 95% confidence interval.

Figure 1-13: Comparison of main effects by worker potential

(a) Effects within 2 years of move



(b) Effects more than 2 years after move



Note: The figures compare pooled event study coefficients from separately estimating Equation 1.6 on high and low potential workers from the baseline sample. The top panel plots the pooled effects within two years of moving, and the bottom panel plots the pooled effects more than two years after the move. The outcomes are whether the worker is working in a supervisor position, whether the worker is working in any job in formal employment, and the worker's log earnings conditional on being in formal employment. Worker potential is defined as the worker's predicted likelihood of being promoted within a year after moving firms based on their characteristics before moving firms (estimated from Equation 1.10 using a holdout sample of workers). High potential workers are workers in the top tercile of worker potential, and low potential workers are workers in the bottom tercile. Standard errors are clustered by the worker's destination firm, and lines indicate the 95% confidence interval.

Figure 1-14: Correlation in firm practices



Note:: The figure plots the correlation between the plant's willingness to promote high performing workers and the plant's willingness to fire low performing workers, as scored from structured interviews from the World Management Survey. Scores range from 1 to 5 on each question, with 5 as reflecting the most active worker management practices and 1 as reflecting the least active. For more details, see Section 1.6.4.

Table 1.1: Example of occupational group with observable line of progression

Occupation Code	Title
7601XX	<i>Supervisors of the textile industry</i>
7610XX	Multipurpose workers in the textile industry
7611XX	Workers in the classification and washing of fibers
7612XX	Operators of spinning machines
7613XX	Operators of looms
7614XX	Workers in finishing, dyeing, and stamping
7618XX	Inspectors and reviewers of textile production

Table 1.2: Estimates of promotion wage premia

	Outcome: Log Earnings				
	(1)	(2)	(3)	(4)	(5)
Supervisor	0.628*** (0.0180)	0.520*** (0.0185)	0.384*** (0.0191)	0.170*** (0.00347)	0.200*** (0.00309)
Number of observations	11581373	11581278	11581208	11092695	5616458
Number of workers	2473852	2473850	2473846	1985349	1071791
Adjusted $R^2$	0.234	0.407	0.520	0.890	0.874
Controls:					
State-occupation-year FE		Y	Y	Y	Y
Worker controls			Y	Y	Y
Worker FE				Y	Y
Firm wage premia					Y

Note:: The table reports estimates for the average supervisor wage premium in the Brazilian formal employment sector. The sample is a 5% sample of all workers between the ages of 25 and 50 that work in occupational groups with a well-defined supervisor role. The first column reports the bivariate wage premium and controls only for year fixed effects. Worker controls are a quadratic in age interacted with gender along with indicators for the worker's education, race, and state. Firm wage premia controls net out the estimated AKM firm fixed effect from worker earnings. Standard errors are two-way clustered at the worker and firm level. Stars indicate the level of significance: \* 5%, \*\* 1%, and \*\*\* 0.1%.

Table 1.3: Role of promotions in lifecycle wage profile

	Outcome: Log Earnings					
	(1)	(2)	(3)	(4)	(5)	(6)
Age	0.0173*** (0.000294)	0.0163*** (0.000321)	0.0115*** (0.000185)	0.0105*** (0.000195)	0.0133*** (0.000238)	0.0118*** (0.000258)
Supervisor		0.412*** (0.0130)		0.393*** (0.0112)		0.434*** (0.0135)
Share of age coefficient explained		0.0547		0.0880		0.111
Number of obs.	6056674	6056674	6056525	6056525	3217367	3217367
Number of workers	1697830	1697830	1697814	1697814	978304	978304
Adjusted $R^2$	0.474	0.489	0.538	0.551	0.483	0.504
Controls:						
Worker controls	Y	Y	Y	Y	Y	Y
Occupation controls			Y	Y	Y	Y
Firm wage premia					Y	Y

Note: The table reports estimates of the average lifecycle wage profile with and without controlling for the average supervisor wage premium in the Brazilian formal employment sector. The sample is a 5% sample of all workers between the ages of 25 and 35 that work in occupational groups with a well-defined supervisor role. The share of age coefficient explained is the relative decrease in the age coefficient after adding an indicator for whether the worker is working as supervisor. Worker controls are indicators for gender, worker education, and race as well as state-by-year fixed effects. Occupational controls are the worker's occupational tenure and state-by-occupational-group-by-year fixed effects. Firm wage premia controls net out the estimated AKM firm fixed effect from worker earnings. Standard errors are two-way clustered at the worker and firm level. Stars indicate the level of significance: \* 5%, \*\* 1%, and \*\*\* 0.1%.



Table 1.4: Differences between high and low promotion firms

Outcome:	Difference	Standard errors	Low prom. firm mean	Adj. $R^2$
Promotions	0.0172***	(0.000152)	0.000159	0.354
Exit Formal Employment	0.0165***	(0.000744)	0.165	0.00740
Log Employment	0.490***	(0.00934)	3.584	0.0415
Age	-0.310***	(0.0186)	35.27	0.00376
Log Earnings	0.0900***	(0.00417)	6.462	0.00677
AKM Firm FE	0.0542***	(0.00346)	-	0.00443
Supervisor Share	0.0392***	(0.000612)	0.0302	0.0789
Log Earnings Growth	0.00990***	(0.000424)	-	0.00869
Log Employment Growth	0.0338***	(0.00162)	0.0355	0.00746
Number of firms			69417	

Note:: The table reports the difference between high and low promotion firms in average firm characteristics between 2004-2006, as well as the means for low promotion firms in each relevant category. The adjusted  $R^2$  reports the share of firm-level variation in each outcome that is explained by the firm's high versus low promotion status. Stars indicate the level of significance: \* 5%, \*\* 1%, and \*\*\* 0.1%.

Table 1.5: Analysis sample summary statistics

	Baseline	Mass layoffs
Number of workers	1100590	127058
Number of origin firms	162573	16103
Number of destination firms	49785	21170
Worker characteristics before move:		
Age	32.94 [6.761]	33.89 [7.007]
Female	0.262 [0.439]	0.241 [0.428]
Years of schooling	10.50 [2.941]	9.866 [3.033]
Monthly earnings (2010 Reals)	1242.9 [900.4]	1186.4 [831.7]
Share to high promotion firm	0.417	0.442
Worker outcomes > 2 years following move:		
In formal employment	0.777	0.764
At destination firm	0.305	0.278
In supervisor occupation	0.0280	0.0284

Note:: The table reports summary statistics about the job-to-job movers in the baseline sample and the mass-layoffs subsample. Pre-move worker characteristics refer to the snapshot of worker data from the year before the job-to-job transition ( $t = b - 1$ ). Outcomes more than two years following the move are averaged over all relevant years where the data are available. All statistics on characteristics and outcomes are in means, and standard deviations for continuous measures are in brackets.

Table 1.6: Estimates of regional wage premia

	Outcome: Log Earnings				
	(1)	(2)	(3)	(4)	(5)
High Firm Share ( $\beta_1$ )	0.132*** (0.0162)	0.140*** (0.0151)	0.140*** (0.0152)	0.0397** (0.0134)	0.00850 (0.00544)
High Firm Share $\times$ Super. ( $\beta_2$ )	0.244*** (0.0423)	0.237*** (0.0405)	0.196*** (0.0334)	0.106*** (0.0287)	0.159*** (0.0306)
Number of observations	11569458	11569388	11569388	11569388	11187050
Number of municipalities	5500	5500	5500	5500	5499
Adjusted $R^2$	0.408	0.521	0.523	0.539	0.771
Controls:					
State-occupation-year FE	Y	Y	Y	Y	Y
Worker controls		Y	Y	Y	Y
Occupation-super.-year FEs			Y	Y	Y
Municipality controls				Y	
Firm FEs					Y

Note:: The table reports estimates for municipal differences in wages and supervisor wage premia in the Brazilian formal employment sector (estimated using Equation 1.11). The sample is a 5% sample of all workers between the ages of 25 and 50 that work in occupational groups with a well-defined supervisor track in municipalities with at least 100 workers. The estimands of interest are the wage premium for all workers in municipalities with a high share of high promotion employers ( $\beta_1$ ) as well as the differential wage premium for promoted workers in municipalities with a high share of high promotion employers ( $\beta_2$ ). Worker controls are a quadratic in age interacted with gender along with indicators for the worker's education, race, and state. Municipality controls are controls for log employment and the overall hiring share of either high or low promotion firms that enter both linearly and as interactions with supervisor status. Standard errors are clustered by municipality. Stars indicate the level of significance: \* 5%, \*\* 1%, and \*\*\* 0.1%.

Table 1.7: Summary of parameters and moments for identification

	Description	Moments	Assumptions
Calibrated parameters:			
$\kappa$	Degree of asymmetric information	Share of job changers that remain supervisors	Calibrated from movers data
$\delta$	Prob. of exogenous separation	Estimated job destruction rate	From Dix-Carneiro et al. (2021)
Jointly estimated parameters:			
$\bar{Q}_H$	Rate of learning for high promotion firms	Average share promoted	Constrained efficient promotions + firings
$\bar{Q}_L$	Rate of learning for low promotion firms	Average share in formal labor market attachment	Random labor market matching
$\alpha$	Share of high ability workers	Effect of high prom. firms on promotions	
$g$	Prob. of re-employment upon separation	Effect of high prom. firms on formal labor market attachment	

Note:: The table summarizes the moments and assumptions required to identify the parameters in the partial equilibrium model of employer learning and job assignment. Two parameters - the degree of asymmetric information  $\kappa$  and the probability of exogenous separation  $\delta$  are directly calibrated. The remaining parameters are jointly identified from the means and treatment effects from the baseline movers analysis. The identifying assumptions and model restrictions are discussed in more detail in Section 1.8.1.

Table 1.8: Baseline model estimates

Definition	Description	Estimate	S.E.
Parameters:			
$\kappa$	Degree of asymmetric information	0.300	-
$\delta$	Prob. of exogenous separation	0.282	-
$\bar{Q}_H$	Rate of learning for high promotion firms	0.208	(0.029)
$\bar{Q}_L$	Rate of learning for low promotion firms	0.150	(0.022)
$\alpha$	Share of high ability workers	0.214	(0.031)
$g$	Prob. of re-employment upon separation	0.414	(0.028)
Implied measures:			
$(\alpha - \alpha') / \alpha$	Degree of adverse selection in job movers	0.292	(0.039)
$(\alpha - E[P]) / \alpha$	Share of high ability workers not in supervisor role	0.869	(0.019)

Note:: The table reports the estimates from quantifying the partial equilibrium model of employer learning and job assignment. The quantification approach is described in Section 1.8.1. The definition of high and low promotion firms match the classification used for the estimation of treatment effects, as described in Section 1.4.1, and the resulting estimates quantify the difference in economic behavior between the two groups. The degree of adverse selection in job movers is the relative decrease in the likelihood that a job mover is of high ability as compared to the general population. Meanwhile, the share of high ability workers that are not working as supervisors reflects skill misallocation. Standard errors are calculated from the joint variance-covariance matrix of empirical moments, which are clustered by the worker's destination firm.

THIS PAGE INTENTIONALLY LEFT BLANK

# Chapter 2

## Unionization, Employer Opposition, and Establishment Closure

*Written jointly with Samuel Young*

### 2.1 Introduction

Union elections in the U.S. are extremely contentious. Managers frequently threaten to close establishments if they unionize, and survey evidence suggests that some follow through on these threats (Bronfenbrenner, 1996).<sup>1</sup> The conventional economic explanation for establishments closing after unionization is that unions' demands for wage increases or other workplace changes make the businesses unprofitable. This explanation, however, is not supported by existing research, which finds little evidence of successful union elections leading to higher wages or lower productivity.<sup>2</sup>

---

We are grateful to Daron Acemoglu, David Autor, and Simon Jäger for guidance and advice throughout this project. We thank Josh Angrist, Jon Cohen, David Hughes, Sylvia Klosin, Tom Kochan, Felix Koenig, Mike Piore, Frank Schilbach, Garima Sharma, Martina Uccioli, John Van Reenen, Michael Wong, and Josef Zweimüller and seminar participants at the U.S. Census Bureau for helpful comments. This paper benefited greatly from Henry Hyatt and Kirk White's data expertise. We thank Stephanie Bailey, Jim Davis, and Nathan Ramsey for their assistance with the data access and the disclosure process. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. (1745302). All errors are our own. *Disclaimer: Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2389 (CBDRB-FY22-P2389-R9311 and CBDRB-FY22-P2389-R9358).*

<sup>1</sup>Establishments are distinct locations where employees work. Firms are groups of establishments under the same ownership. Union elections are generally at the establishment level. Thus, a firm could have unionized and non-unionized establishments.

<sup>2</sup>Frandsen (2021); DiNardo and Lee (2004); Freeman and Kleiner (1990b) do not find wage increases following successful union elections. Additionally, Dube et al. (2016); Sojourner et al.

Consider two examples of how employers responded to unions that suggest alternative reasons why unionization may lead to establishment closure. First, during a 2017 campaign to unionize the news website Gothamist, the owner wrote to employees, “as long as it’s my money that’s paying for everything, I intend to be the one making the decisions.” One week after the workers voted to unionize, the owner shut down the business (Wamsley, 2017). This example suggests that some establishment closures could be driven by managers’ unwillingness to operate alongside unions or their general dislike of working with unions. Second, consider Boeing’s production of 787 airplanes. In 2011, Boeing shifted production of some 787s away from a unionized plant in Washington state to a non-union plant in South Carolina. In 2021, it produced all 787s in South Carolina (Cameron, 2020). According to a Boeing executive, the motivation “was not the wages we’re paying today. It was that we cannot afford to have a work stoppage, you know, every three years” (Greenhouse, 2011). This example illustrates how some firms can avoid working with unionized workers without eliminating their production. Additionally, the example suggests that conflictual labor relations between unions and firms may also lead to establishment closures. Both examples illustrate how unionization may lead to establishment closure even without direct wage or productivity effects.

This paper assesses whether these examples generalize by analyzing the effect of unionization on establishment employment and survival. We then examine whether firms’ ability to avoid working with new unions and managers’ general opposition to unions help explain the employment and survival effects. Our setting is around 27,000 U.S. private-sector union certification elections through the NLRB from 1981-2005. We link these elections to administrative Census data on establishment employment and survival from the Longitudinal Business Database (LBD) and total factor productivity from the Annual Survey of Manufactures (ASM). We supplement these data with union contract data from the Federal Mediation and Conciliation Service (FMCS).

We analyze these elections using a novel research design that extends standard difference-in-differences techniques with falsification tests from the regression discontinuity extrapolation literature. This strategy allows us to estimate treatment effects that include elections that win by large margins of support. Using this design, we find that unionization decreases establishments’ employment, primarily by lowering their likelihood of survival. We estimate a five-year effect on establishment survival of four

---

(2015) find productivity *increases* following unionization, and the overall unions and productivity literature generally finds zero or *positive* effects (Freeman and Medoff, 1984). As exceptions, Knepper (2020) finds non-wage benefit increases and LaLonde et al. (1996) find productivity declines following unionization.



percentage points (pct. pts.) relative to a survival rate of 82 % for establishments where the union lost. The overall employment declines are also bigger for larger margin-of-victory elections. Additionally, we document substantial effect heterogeneity across three broad industry groups: manufacturing, services, and other industries.<sup>3</sup> In the service sector, where the majority of recent union organizing has occurred, we find much smaller and sometimes insignificant effects of unionization. Alternatively, the overall employment and survival declines are driven by large effects in manufacturing and “other industries.” For example, the ten-year effect on survival for manufacturing elections is eight pct. pts.

Motivated by these overall effects of unionization, we test whether firms’ ability to avoid working with new unions and managers’ opposition to unions help explain the effects. For this analysis, we focus on manufacturing elections because we have better data to test the specific parts of our hypothesis, and it is the largest sector with substantial negative employment effects.<sup>4</sup>

The first part of our hypothesis is that some firms can avoid working with new unions by shifting production from unionized establishments to other establishments.<sup>5</sup> To test this, we first estimate whether the effects of unionization are larger at establishments part of multi-establishment (or multi-unit, MU) firms than single-establishment (SU) firms. We find significantly larger employment and survival decreases from elections at MU firms. For example, the ten-year effects on survival are twelve pct. pts. versus three pct. pts. at MU and SU firms, respectively. This heterogeneity is consistent with MU firms avoiding working with new unions by shifting production to other establishments. In contrast, SU firms need to either work with the union or shut down entirely.

Next, we more directly test for production shifting after successful elections. Specifically, following successful versus unsuccessful elections at MU firms, we compare the employment growth at the firms’ *other* establishments. Overall, we do not find

---

<sup>3</sup>Examples of service-sector elections include hospitals, nursing homes, grocery stores, and janitors. The “other” industry elections include transportation, warehouse, and construction elections.

<sup>4</sup>The manufacturing data are better for two reasons. First, in manufacturing, the union avoidance method we test is shifting employment away from newly unionized establishments, which we can measure with the establishment and firm linkages. Second, we have high-quality establishment-level productivity measures that we use to test effect heterogeneity by baseline productivity.

<sup>5</sup>While employment shifting is the most prominent union avoidance tactic in manufacturing (see, e.g., Bluestone and Harrison (1982); Verma (1985); Kochan et al. (1986a)), other tactics may be more prominent in other industries. For example, Hatton (2014) document replacing unionized workers with independent contracts in several industries, and Evans and Lewis (1989) document construction firms opening separate non-union firms to avoid hiring unionized workers. We hypothesize that all of these tactics may explain the employment effects of unionization, but we only have the data to test for employment shifting.

different post-election employment growth between the firms' *other* manufacturing establishments. Yet, these establishments may produce different products than the election establishment, making production shifting difficult. To test for production shifting between establishments that produce similar products, we focus on establishments in the same three-digit NAICS industry as the election establishment. With this restriction, we find significantly higher employment growth for the other establishments at firms with successful elections. However, these effects are insignificant five years after the election, at which point the firm may have shifted production to *new* establishments. Both pieces of evidence support firms avoiding unions through production shifting as one explanation for the impact of unionization on employment and survival.

The second part of our hypothesis is that the effects of unionization are greater when management is more opposed to the union. To test this, we estimate treatment effect heterogeneity using two proxies for managers' opposition. First, we estimate effects separately for MU firms with and without any *other* unionized establishments. Survey evidence indicates that less unionized firms would more "vigorously resist dealing with unions," and managers' anti-union philosophies were often a key motivation for this opposition (Freedman, 1979; Foulkes, 1980). Additionally, similar to Selten (1978)'s "chain store paradox," non-unionized firms may aggressively resist the first unionization campaign to deter future attempts, even if not economically profitable when considering the attempt in isolation. Consistent with this evidence, we find significantly larger long-run employment and survival declines from successful elections at non-unionized firms than (partially) unionized firms.

Our second proxy for managers' opposition to the union is delays during the election process. Strategies that delay elections are a key way that managers attempt to influence elections and consequently a proxy for their opposition. For example, in "Confessions of a Union Buster," Levitt and Conrow (1993) write that the National Labor Relations Act "presents endless possibilities for delays, roadblocks, and maneuvers that can undermine a union's efforts" and that delay "steals momentum from a union-organizing drive." We define election delay as the time between the date the union filed for the election and the election date. We estimate separate treatment effects for elections with shorter and longer delays, and find significantly larger employment and survival decreases following longer delay elections. For example, the ten-year survival effect for MU elections in the top tercile of election delay times is 20 pct. pts. versus 7 pct. pts. for the bottom tercile.

Finally, we test for effect heterogeneity between establishments with different baseline productivity, which is implied by the conventional explanations for why

unionization leads to establishment closures. Specifically, theoretical and empirical work from other contexts identifies productivity as a key determinant of establishment closure decisions. This suggests that wage increases or productivity declines from unionization should lead to larger survival effects for lower-productivity establishments.<sup>6</sup> However, we do not find significant differences in the survival effects of unionization between establishments with different baseline total-factor productivity, measured using establishment-level input and output data from the ASM and Census of Manufacturers. Thus, this evidence is more consistent with alternative explanations for why unionization leads to establishment closures (e.g., our union avoidance hypothesis) than the conventional wage and productivity explanations.

There are several potential interpretations of our evidence on managers' opposition to unions and their use of union avoidance tactics. One interpretation is that the opposition and avoidance are motivated by a dislike of working with unions unrelated to their direct costs. This interpretation helps resolve the puzzles discussed earlier. First, existing research has not found that recent union elections have raised wages or negatively affected establishment productivity. Second, it is hard to rationalize unions making demands that lead to establishment closures, as this would directly harm their members (Friedman, 1951). Under the conventional explanations for closures, both points are difficult to reconcile with our large survival estimates. Yet, they are consistent with closures being driven by managers' idiosyncratic dislike of working with unions.<sup>7</sup> This interpretation is also consistent with our finding of no treatment effect heterogeneity by baseline establishment productivity. On the other hand, we cannot rule out that our proxies for manager opposition simply measure where unions would have been the costliest. Supporting this interpretation, there is research suggesting direct costs of recent union elections (e.g., Lee and Mas (2012)'s evidence of equity declines following successful elections). Overall, while our results do not measure the direct costs of unionization, they suggest that the employment and survival effects of unionization may be excessive relative to these costs.

We next summarize our econometric methodology and multiple falsification tests of our identifying assumption. We start by implementing a difference-in-differences design that compares outcomes before and after union elections at establishments where the

---

<sup>6</sup>See, for example, the theoretical and empirical literature on the reallocation effects of minimum wages (Berger et al., 2021; Dustmann et al., 2020) and the relationship between productivity and establishment exit (Foster et al., 2008).

<sup>7</sup>As further support for this interpretation, survey evidence has not found that the firms most opposed to unions are also where unions are likely the costliest (Freedman, 1979; Bronfenbrenner, 2001).

union won versus lost. Our identifying assumption is that outcomes at establishments with different election vote shares but the same baseline characteristics would have followed parallel trends had no election occurred. To support this assumption, we first show that only conditioning on baseline employment and industry yields similar pre-election employment and payroll growth rates between establishments with winning and losing elections. This similarity holds when we add much richer baseline covariates and for up to ten years before the union elections. Additionally, we show that our treatment effects are increasing in the share of workers covered by the union and not driven by firm-level trends, which are consistent with our estimates being driven by unionization.

To further support our design, we assess several additional testable implications of our identifying assumption that are possible since we observe election vote shares. These checks extend tests from the regression discontinuity extrapolation literature to panel-data settings (Angrist and Rokkanen, 2015; Bennett, 2020). First, we show that the similarity in *pre-election* employment growth rates holds across the entire vote-share distribution in our sample. In other words, we test for trends in pre-election growth rates by vote share, which is a stronger test of “pre-trends.” Second, we show that establishments’ *post-election* employment growth and survival were similar between *losing elections* with different vote shares. If our treatment effects were driven by contemporaneous shocks correlated with vote shares, we would also expect different post-election outcomes between losing elections with different vote shares. Overall, these tests show that our identifying assumption holds for several sets of observations where we observe untreated potential outcomes.

Our empirical strategy combines features of regression discontinuity (RD) and panel data methods that have previously been used to analyze union elections.<sup>8</sup> Although RD methods are appealing due to their internal validity, there are some disadvantages in this setting. First, there is manipulation around the 50 % threshold that leads to *pre-election* discontinuities in establishment characteristics (Frandsen, 2017). Second, the effects of close union elections may be different than elections with larger margins of support. To address these issues, our paper expands the bandwidth to include all 20-80 % vote-share elections and uses the panel dimension to account for selection into winning versus losing elections. The wider bandwidth also gives us more power to estimate heterogeneous treatment effects. Relative to the other panel-data analyses,

---

<sup>8</sup>See DiNardo and Lee (2004); Sojourner et al. (2015); Knepper (2020); Bradley et al. (2017) for RD analyses and Freeman and Kleiner (1990b); LaLonde et al. (1996); Lee and Mas (2012); Dube et al. (2016); Goncalves (2021) for panel data analyses. Frandsen (2021) also combines these methods by implementing a regression discontinuity design on first differenced outcomes.

we better exploit observing the vote shares to implement tests of our identifying assumption.<sup>9</sup> These tests could also be implemented in other DiD analyses where the “forcing variable” is observed (e.g., Ganong and Noel (2020); Harju et al. (2021)).

Our overall employment and survival estimates contribute to the literature on the effects of unionization in the U.S. Due to the different empirical strategies, our estimates complement Frandsen (2021)’s RD estimates of short-run decreases in establishment employment for close union elections and his suggestive evidence of negative survival effects.<sup>10</sup> Our estimates also generalize other research finding employment declines following successful union elections in specific sectors (Sojourner et al., 2015; LaLonde et al., 1996).<sup>11</sup> As a consequence, our results contrast with the null effects of close union elections in DiNardo and Lee (2004) and other research finding no relationship between unions and business survival (Freeman and Kleiner, 1999). These differences are potentially due to our use of higher-quality establishment survival data. Finally, our evidence of larger employment declines from successful elections with larger vote shares mirrors Lee and Mas (2012)’s finding of larger stock market declines from larger vote-share elections.

Our evidence supporting the manager opposition and union avoidance hypothesis is novel relative to economic research on union elections but supports other research about firms’ responses to unionization. For example, Bronfenbrenner (2000, 2001) report similar results from a survey of union organizers in elections during the 1990s. She finds survival effects of twelve pct. pts. following successful elections. She also finds that establishment closing threats were more common at the types of elections where we find larger survival effects (e.g., in manufacturing and at MU firms). In addition, our evidence that managers more opposed to the union are more likely to shut down establishments after successful elections adds to the literature on anti-union firms’ broader union avoidance tactics (Freeman and Kleiner, 1990a; Kleiner, 2001; Flanagan, 2007). In particular, this result complements Ferguson (2008)’s finding that successful elections with unfair labor practice charges (another proxy for employers’ opposition) are less likely to reach a first contract. Finally, our production shifting

---

<sup>9</sup>Lee and Mas (2012) and Frandsen (2021) present pre-trends and post-election outcomes *across the vote-share distribution* but do not use these estimates as formal tests of their identifying assumptions.

<sup>10</sup>His survival estimates are differences in survival probabilities around the 50 % threshold and he writes that “a causal interpretation of the differences in survival probability should be made with caution” due to manipulation around the threshold.

<sup>11</sup>LaLonde et al. (1996) analyze the employment and output effects of manufacturing union elections from 1977-1989 using a difference-in-differences design. Yet, their analysis differs from ours on several dimensions. First, they do not analyze the effect on establishment closures, which makes interpreting the results conditional on survival difficult. Second, due to their smaller sample size, their pre-trend estimates are often imprecise, making it difficult to assess the parallel trends assumption.

evidence is consistent with firms becoming less unionized during this time through investing in and opening non-union establishments (Kochan et al., 1986a; Verma, 1985).

The rest of the paper is structured as follows. Section 2.2 describes the institutional details of NLRB union elections. Section 2.3 describes the election and Census data. Section 2.4 discusses our empirical strategy and tests of our identifying assumption. Section 2.5 presents our estimates of the overall effects of unionization on employment and survival. Section 2.6 provides multiple tests of our manager opposition and union avoidance hypotheses. Finally, Sections 2.7 and 2.8 discuss our results.

## 2.2 Unionization through NLRB Elections

The National Labor Relations Act (NLRA) guarantees most workers in the U.S. the right to collective bargaining and action. Under the NLRA, when a union represents a group of workers, their employer is required to bargain with the union over the conditions of employment.<sup>12</sup> This bargaining generally occurs at the establishment level (Traxler, 1994). During negotiations, the union may go on strike or the employer may “lockout” workers to pressure the other party. The NLRA also created the National Labor Relations Board (NLRB), a quasi-judicial agency that administers union elections and enforces unfair labor practice violations. Much of the current U.S. policy discussion around organized labor focuses on increasing representation at non-unionized establishments.<sup>13</sup> Our results speak directly to the potential consequences of these efforts to increase unionization.

The primary way for private-sector workers to gain union representation is through a secret-ballot NLRB election. The organizing drive is initiated by workers at the establishment, either on their own initiative or prompted by outreach from a union. The first step is getting cards indicating union support signed by workers in the proposed “bargaining unit” (i.e., the workers the union would represent). The bargaining unit generally only contains workers at a single establishment. It can range from workers in a single occupation (e.g., delivery truck drivers) to all non-managerial workers. After gathering signatures from at least 30 % of the bargaining unit, the union files

---

<sup>12</sup>The goal of these negotiations is a contract. Contracts commonly specify wage and non-wage compensation for each job title, grievance procedures for disputes, policies for implementing layoffs, and promotion policies (Slichter et al., 1960).

<sup>13</sup>For example, the currently debated *Protecting the Right to Organize (PRO) Act of 2021* would limit employers’ ability to campaign against union elections and increase penalties for unfair labor practices during elections.

an election petition with the NLRB. The NLRB then confirms that the cards show sufficient support for the union, resolves any disagreements over the composition of the bargaining unit, and schedules the election. After the petition is filed, employers frequently attempt to delay the election to reduce union support (e.g., contesting the composition of the bargaining unit) (Levitt and Conrow, 1993).

Before the election, the union and employer often actively campaign for and against union representation. Union organizers and pro-union workers can campaign by (1) speaking with other workers at work or during “house calls,” (2) publicly showing solidarity among union supporters (e.g., rallies or wearing pro-union attire), or (3) enlisting the support of community groups (Bronfenbrenner and Juravich, 1998). Employers also have many campaign tools at their disposal, including one-on-one meetings with supervisors and “captive audience meetings,” where employees are required to attend. Employers also frequently hire “union avoidance” consultants and law firms (Logan, 2002). Finally, although there are legal restrictions on firing pro-union workers and threatening to close establishments, these tactics still occur (Weiler, 1983; Schmitt and Zipperer, 2009). If a majority of workers vote for the union, the union is certified by the NLRB to represent the bargaining unit. After the union is certified, the employer is required to bargain “in good faith” with the union. But the parties are not required to reach an agreement.<sup>14</sup> If a contract is not reached one year after certification, employees can vote out the union by holding a *decertification* election.

NLRB elections are the primary method for private-sector workers at an establishment to gain union representation. However, there are two reasons why some unionization occurs without an election. First, the NLRA does not cover all workers (General Accounting Office, 2002).<sup>15</sup> Second, workers covered by the NLRA can gain union representation without an election through voluntary “card check” recognition. However, card check is much less common than elections.<sup>16</sup>

---

<sup>14</sup>In a review, CRS (2013) find that 56-85 % of successful elections result in first contracts during the period we consider.

<sup>15</sup>Some workers lack collective bargaining rights (e.g., some small business employees, independent contractors, domestic workers, and “agricultural laborers”). Other workers have collective bargaining rights but are not covered by the NLRA. For example, airline and railroad employees’ collective bargaining rights are covered by the *Railway Labor Act*. Similarly, public-sector workers’ bargaining rights are covered by various federal, state, and local statutes.

<sup>16</sup>Schmitt and Zipperer (2009) estimate that from 1998-2003, 60 % of workers were organized through NLRB elections but assume that before then 90 % of organizing occurred through elections.

## Selection into Union Elections and the Determinants of Winning Elections

Since our empirical design compares winning and losing elections with similar baseline characteristics, it is helpful to review the literature on selection into holding and winning elections. This literature motivates which baseline characteristics we condition on and our additions tests of whether election vote shares are related to remaining unobservable shocks. For selection *into* elections, Dinlersoz et al. (2017) find that elections are more likely at larger, more productive, and younger establishments. We account for this selection by only comparing establishments that hold elections.

For election outcomes, workers, employers, and other factors could all influence whether the union wins. For our empirical strategy, a concern is that vote shares may be related to future establishment productivity changes. For example, workers who expect their establishment to become more productive and have more rents to share may be more likely to vote for a union. This would generate a positive bias between vote shares and establishment growth. Alternatively, firms that expect to become more productive may campaign harder against unions, leading to a negative bias.

Research on election outcomes finds that these factors all play some role. The most consistent finding is higher union win rates for smaller bargaining units (Heneman and Sandver, 1983; Farber, 2001). Win rates also vary substantially across industries (Bronfenbrenner, 2002). In the 2000s, the win rate in manufacturing was around 40 % versus 60 % for services. These factors motivate our first specification that just conditions on establishments' baseline industry and employment. In terms of the influence of employer versus union campaigns, Bronfenbrenner (1997) finds that "union tactic variables explain more of the variance in election outcomes than any other group," including employer tactics or characteristics. Yet, other research finds that the strength of firms' anti-union campaigns is associated with lower win rates (Freeman and Medoff, 1984). To address the concern that firms' anti-union campaigns lead to a negative bias between vote shares and establishment growth, we implement multiple tests of how vote shares are related to firm productivity shocks. Additionally, past research has found that winning versus losing elections have similar pre-election productivity trends. For example, Dube et al. (2016) find similar productivity pre-trends for nursing home elections, and Lee and Mas (2012) find similar stock-market trends, which is a stronger test since it incorporates expectations of future productivity growth.



## Motivation for Estimating the Effects of Larger Margin-of-Support Elections

An advantage of our empirical strategy is that it does not rely on comparing only elections that barely won or lost. One motivation for this is evidence of non-random sorting of elections just around the 50 % threshold (i.e., “vote-share manipulation”). Figure 2-2 Panel A plots the vote-share distribution for the elections in our sample and shows manipulation around the 50 % threshold previously documented by Frandsen (2017) (e.g., a missing mass of elections that barely win).<sup>17</sup> Frandsen (2021) also documents that this manipulation leads to large differences in observable establishment characteristics across the threshold (e.g., 13-22 % differences in employment).

Another motivation for our empirical strategy is that the treatment effect of unionization may depend on the election vote share. For example, Lee and Mas (2012) find that the negative stock market effects of unionization are larger for higher margin-of-victory elections. One potential reason for this heterogeneity is that close union elections are often followed by lengthy delays before bargaining begins (e.g., debates about challenged votes). Figure 2-2 Panel B shows this by plotting the average number of days between the election date and the case closing date (e.g., when the union is officially certified). The figure shows a striking increase in this delay time for close elections (e.g., the median (mean) for elections that barely win is around 118 (223) days versus only 11 (57) days for 60 % vote-share elections). Since delays can dampen the unions’ bargaining power, this evidence suggests that the effects of close elections may be different than higher vote-share elections.

Second, for close elections, firms may delay the bargaining process anticipating a future decertification election. Figure 2-2 Panel C provides evidence that close elections are more likely to be decertified by plotting the probability of each certification election experiencing a *decertification* election in the five years following the original election. It shows that more than twelve percent of very close winning elections experience a future decertification election compared to less than five percent of larger margin-of-victory elections. This suggests that higher margin-of-victory elections may be more likely to reach first contracts, leading to more changes at the establishment. A final reason why the treatment effect of unionization may vary by vote share is that unions that

---

<sup>17</sup>It is difficult to see manipulation in this figure because of the *discrete* running variable and since our sample includes elections small numbers of votes. Consequently, we plot elections with exactly 50 % of votes separately to make the manipulation easier to see. Frandsen (2017) finds evidence of manipulation using formal tests that accommodate discrete running variables. Additionally, Figure B-2 plots vote-share density for elections with more than 50 votes, where it is clearer to see manipulation.

win with more support may be able to more credibly threaten to strike. Figure 2-2 Panel D supports this by showing that within manufacturing, where strikes were more common, the probability of a post-election works stoppage increases in the election vote share (see Appendix B.3 for details). Overall, these results show that several proxies for the unions' bargaining power increase in the election vote share, suggesting that the effects of unionization may also differ along this margin.

## 2.3 Election, Contract, and Establishment Data

For our analysis, we combine union election and contract data with administrative establishment-level data from the U.S. Census Bureau. These data are uniquely suited to study union elections. First, the data contain the universe of *establishments*, the level at which most elections are held. Analysis of more aggregated data would include establishments not directly affected by the elections and attenuate the effects of unionization. Second, the Census constructs high-quality longitudinal establishment linkages that allow us to separate real establishment exit from spurious exit due to administrative reasons or ownership changes (Haltiwanger et al., 2013). These links are important for our analysis because survival is a key outcome of interest. Finally, the rich establishment covariates allow us to compare similar winning and losing elections (e.g., same size, age, and industry).

**NLRB Union Election Data** We combine data from multiple sources to construct a comprehensive dataset of union elections from 1962 to 2018. Specifically, we use datasets assembled by Henry Farber, J.P. Ferguson, and Thomas Holmes and public data from the NLRB.<sup>18</sup> The data contain election vote counts that we use to define treatment. Additionally, they include employers' names and addresses that we use to match elections to Census data. Finally, the data include the election petition filing date, the actual election date, and the closing date. We define our treatment time based on the filing date of each election because this is the earliest date we observe for each election. We also use these dates to define the time between filing the election petition and holding the election, a proxy for managements' opposition to the elections described further in Section 2.6.

---

<sup>18</sup>For duplicates across datasets, we pick one observation for each NLRB case number (see Appendix B.3 for details). Appendix Figure B-1 shows that this yields a similar number of cases each year to the number of cases from the NLRB's annual reports.

**FMCS Contract Notice Data** To measure whether an establishment is covered by any collective bargaining agreement, we use contract notice data from the Federal Mediation and Conciliation Service (FMCS) from 1984-2019. We combine data from Thomas Holmes and the FMCS. The data include both notices of *initial contracts* (i.e., first-contract negotiation after an election) and *contract renegotiation or reopening* for existing contracts. These “notices of bargaining” are provided to the FMCS so it can be ready to provide mediation. Although filing is legally incentivized, underreporting is possible. These data also include names and addresses for matching. We use these data to measure whether an election establishment has any other workers covered by a collective bargaining agreement and whether the election establishment’s firm has any other unionized establishments.

**Employment, Payroll, and Survival Data from the LBD** Our primary source of establishment-level outcomes is the Longitudinal Business Database (LBD). It contains annual employment and payroll for the universe of non-farm, private sector establishments from 1976-2015 (Jarmin and Miranda, 2002). Our employment measure is the total number of employees in March of each year. The payroll measure is employees’ total “wages, tips, and other compensation” over the entire year. Consequently, we would expect larger effects on “event-time 0” payroll than employment. The data also contain high-quality longitudinal establishment IDs that identify the same establishments over time, even across ownership changes. We use these IDs to define establishment survival. Specifically, we define survival based on the last year the establishment has non-zero employment. Finally, we use the Fort and Klimek (2016) 2012 NAICS codes to classify each establishment into consistent industries across the entire time period.

We address potential biases from how the Census defines employment at establishments part of *multi-establishment (MU) firms* by focusing on longer-run outcomes. In particular, although the LBD is at the establishment level, some of the annual employment and payroll data are received at higher levels of aggregation. These aggregate measures are initially allocated proportionately across establishments based on their previous employment. Consequently, if a unionized establishment at an MU firm shrinks, some of this decrease in employment may be initially allocated to the firm’s other establishments, creating a short-run underestimate of the effect of unionization. To avoid these allocation issues biasing our results, we focus on longer-run outcomes (e.g., five to ten years after the elections) since the Census receives establishment-level employment measures at least every five years (see Appendix B.3

for details).

**Sample Selection and Matching Elections to Census Establishments** Before matching the election data to the Census data, we impose sample restrictions to focus on elections likely to shift an establishment’s union status. Appendix Table B.1 shows how these restrictions affect the number of elections and eligible voters in our final sample. First, we restrict the sample to elections held between 1981-2005. Since the LBD starts in 1976 and ends in 2015, this gives us a five-year pre-period and ten-year post-period for all elections. Second, we drop non-representation election cases (e.g., decertification elections). Third, we drop contested elections, which are elections with multiple unions on the ballot. These elections often involve incumbent unions (e.g., “union raids”) and consequently may not be associated with changes in union representation (Sandver and Ready, 1998). Fourth, we drop elections with fewer than six workers in the bargaining unit to ensure that the election could lead to a non-trivial increase in union representation.

After these sample restrictions, we implement a name and address matching procedure to link each election to a unique establishment in the LBD (our strategy is similar to Kline et al. (2019)). We match each election to the universe of LBD establishments by calculating a weighted average of the *Soft TF-IDF* distance between employer names and the geographic distance between geocoded addresses. We match each election to the Census establishment with the highest match score above a minimum threshold. This procedure yields a match for 70 % of elections.<sup>19</sup> We also apply the same procedure for each FMCS contract notice. See Appendix B.3 for details on our matching algorithm.

We further restrict the election sample based on the requirements of our empirical strategy. For each establishment, we only keep the first election. As discussed in Section 2.4, this means that our estimates should be interpreted as the effects of winning the *first* union election at an establishment. Next, we drop elections at establishments less than three years old. Since a key test of the identifying assumption is that the outcomes for winning and losing elections evolved similarly before the election, we do not want to include observations where we cannot evaluate this for at least three time periods. Finally, to keep our sample the same across model specifications, we require that each observation have non-zero payroll and employment one year before the election. These restrictions result in an overall sample of approximately 27,000

---

<sup>19</sup>Although matching introduces measurement error in our binary treatment variable, such measurement error should bias us against finding effects of unionization (see e.g., Card (1996) for measurement error in *individual-level* union status).

elections (see Appendix Table B.1).

Finally, for much of our analysis, we restrict the sample to 20-80 % vote-share elections. Appendix Table B.1 shows that this decreases our sample to 19,000 elections. The motivation for this restriction is that some of the tests of our identifying assumption discussed in Section 2.4 fail for the extreme vote-share elections. To assuage concerns that this choice of bandwidth drives our results, we show that our main results are robust to instead including a 30-70 % bandwidth.<sup>20</sup>

Table 2.1 presents summary statistics for winning and losing union elections in our sample. The estimates confirm the patterns of selection into winning elections described in Section 2.2. In particular, we find that winning elections are at establishments that are, on average, smaller, less likely to be part of multi-establishment firms, and more likely to already have another unionized bargaining unit. The differences, however, are less striking for workers' average wages or establishment age.

## 2.4 Empirical Strategy and Identifying Assumptions

Our research design combines standard difference-in-differences (DiD) techniques with tests of our main identifying assumption from the regression discontinuity extrapolation literature. Our identifying assumption is a conditional parallel trends assumption between elections with different vote shares. Since we observe vote shares that determine treatment assignment, we can assess several testable implications of this assumption that are not possible in a standard DiD setting.

**Potential Outcomes** To fix ideas, consider establishments,  $i$ , that held an election in one year,  $E_i$  (e.g., all elections in 1995). We refer to these elections as *cohort*  $E_i$ . Treatment at time  $t$ ,  $D_{it}$ , is defined as both holding an election and the union receiving a vote share,  $V_i$ , of more than 50 %<sup>21</sup>

$$D_{it} = \mathbb{1}[V_i > .5 \ \& \ t \geq E_i]. \quad (2.1)$$

---

<sup>20</sup>Specifically, our vote-share heterogeneity estimates in Figure 2-6 show that the overall estimates are not driven by the 20-30 or 70-80 % elections. Additionally, Tables B.4, B.5, and B.6 present the heterogeneity estimates with a 30-70 % bandwidth and show that the results are qualitatively the same although sometimes less precise than with the wider bandwidth.

<sup>21</sup>This definition assumes that treatment is absorbing (i.e.,  $D_{it} = 1 \Rightarrow D_{it'} = 1 \ \forall \ t' > t$ ). This assumption ignores that workers may lose union representation through a *decertification* election. Additionally, after losing an election, unions may hold another election. Since we only include the first election at each establishment, we interpret treatment as the dynamic effects of winning a *first union election* which does not correspond one-to-one with union representation or having a contract.

An establishment's non-unionized potential outcome is  $Y_{it}^0$ . Its unionized potential outcome is  $Y_{it}^E(V)$  which depends on its cohort  $E$  and election vote share  $V$ . This allows for dynamic treatment effects and heterogeneous treatment effects by vote share, respectively. We assume *no anticipation* before the year of the election (i.e.,  $Y_{it}^E(V) = Y_{it}^0$  for all  $t < E_i$ ). Observed outcomes are thus<sup>22</sup>

$$Y_{it} = Y_{it}^0 + D_{it} (Y_{it}^{E_i}(V_i) - Y_{it}^0). \quad (2.2)$$

Our estimand of interest is the treatment effect  $n$  years after a successful election with vote share  $V$

$$\delta_n(V) = \text{E} [Y_{it}^{E_i}(V_i) - Y_{it}^0 | V_i = V \ \& \ t - E_i = n]. \quad (2.3)$$

**DiD Specifications** For a single cohort, we can estimate the following specification

$$Y_{it} = \gamma_i + \alpha_t + \sum_n \delta_n \cdot \mathbb{1}[t - E_i = n] \times \mathbb{1}[V_i > .5] + X_i' \beta_n + \varepsilon_{it} \quad (2.4)$$

where  $\gamma_i$  are establishment fixed effects (FEs) and  $\alpha_t$  are year FEs.<sup>23</sup> The coefficients of interest,  $\delta_n$ , capture the average, dynamic treatment effects of a successful union election.  $X_i$  are baseline, one year before the election, establishment characteristics whose coefficients vary with event time  $n$  (i.e., flexible trends by baseline establishment size).

**Identifying Assumption** Our identifying assumption conditional parallel trends *by vote share*. Specifically, we assume that outcomes at establishments with different election vote shares but the same baseline characteristics would have followed parallel trends had no election occurred

$$\text{E} [Y_{it}^0 - Y_{it-1}^0 | X_i, V_i] = \text{E} [Y_{it}^0 - Y_{it-1}^0 | X_i]. \quad (2.5)$$

There are several things to note about this assumption. First, it does not restrict selection *into* union elections (e.g., organizers targeting productive establishments)

---

<sup>22</sup>Here, we assume that losing elections have no causal effect. This assumption is stronger than what we make in our empirical approach since we cannot disentangle the effect of losing an election from the selection into holding an election. We make this assumption for simplicity, but we could also index losing election potential outcomes by cohort to relax the assumption.

<sup>23</sup>We exclude establishment FEs for outcomes that are identical for all establishments in the baseline year,  $t - E_i = -1$  (e.g., establishment survival and DHS growth rates). We include them for log outcomes. See the outcome discussion for details.

or selection *on gains* based on the effects of unionization (e.g., workers only voting for effective unions). Second, the assumption is stronger than the standard DiD assumption because it requires parallel trends by vote share instead of only, on average, between the treated and control observations. Yet, this stronger assumption yields a richer set of testable implications discussed next. Third, the assumption imposes a functional form restriction on potential outcomes (Kahn-Lang and Lang, 2020; Roth and Sant’Anna, 2021), which we discuss for each specific outcome later. Finally, as discussed in Section 2.2, vote shares are influenced by workers, employers, and other factors that could lead to violations of this assumption. This possibility motivates our conditioning on particular baseline  $X_i$ s and assessing multiple testable implications of this assumption to provide reassurance that such selection is not biasing our results.

Our empirical strategy also addresses the concern that vote-share manipulation around the 50 % threshold could violate assumption 2.5 because elections just around the threshold are only a small share of our overall sample. For example, our vote-share heterogeneity estimates show that excluding elections right around the 50 % threshold would not qualitatively change our results.

**Testable Implications of the Identifying Assumption** Our identifying assumption yields several testable implications. The intuition for these tests is that we observe  $Y_{it}^0$  for many observations and can test whether equation 2.5 holds for different subsets of these observations.

The first testable implication is that, if equation 2.5 holds, there should be conditional parallel trends in **pre-election** outcomes across **all vote shares**

$$E[Y_{it} - Y_{it-1}|X_i, V_i] = E[Y_{it} - Y_{it-1}|X_i] \text{ for all } t < E_i. \quad (2.6)$$

This test nests the standard DiD pre-trends test between all winning versus losing elections. Moreover, we can test for similar pre-trends between finer vote-share groups. For example, we can estimate whether establishments where the union won by different margins of victory grew at different rates before the election by comparing pre-trend estimates for 50-60 % versus 60-70 % elections. This test mirrors the tests proposed by Angrist and Rokkanen (2015) and Bennett (2020) for regression discontinuity identification away from the threshold. They argue that conditional mean independence of potential outcomes and the running variable for a given bandwidth around the RD threshold is strong support for being able to estimate treatment effects within that bandwidth. One reason that we only include 20-80 % vote-share

elections in our preferred specification is that for some outcomes, we find violations of equation 2.6 for extreme parts of the vote-share distribution.

The second testable implication is that there should be conditional parallel trends in **post-election outcomes** between **losing elections** with different vote shares

$$E[Y_{it} - Y_{it-1}|X_i, V_i] = E[Y_{it} - Y_{it-1}|X_i] \text{ for all } t \geq E_i \text{ \& } V_i \leq .5. \quad (2.7)$$

To implement this test, we can estimate whether post-election outcomes are different between losing elections with different vote shares (e.g., compare conditional post-election survival rates for 30-40 % versus 40-50 % elections). This test gives us one way to address the concern that election vote shares are correlated with future productivity shocks. If this were the case, we would also expect these shocks to cause differences between the outcomes at losing elections with different vote shares.<sup>24</sup>

Figure 2-1 illustrates our identifying assumption and these testable implications. It plots average outcomes two years before the election,  $Y_{i,-2}$  and  $Y_{i,-1}$ , and one year afterward,  $Y_{i,1}$ , by vote share. Testing parallel *pre-trends* by vote share corresponds to comparing the distance between  $Y_{i,-2}$  and  $Y_{i,-1}$ . Likewise, testing parallel *post-trends for losing elections* corresponds to comparing the distance between  $Y_{i,-1}$  and  $Y_{i,1}$  for losing elections.

**Estimating Effects for Multiple Cohorts** Our sample includes all election cohorts from 1981-2005. To estimate the effect across all cohorts, we pool these elections and estimate

$$Y_{it} = \gamma_i + \alpha_{t,E_i} + \sum_n \delta_n \cdot \mathbb{1}[t - E_i = n] \times \mathbb{1}[V_i > .5] + X_i' \beta_{n,E_i} + \varepsilon_{it}. \quad (2.8)$$

This specification is the same as the single-cohort specification in equation 2.4, except that the year FEs and baseline controls can now vary by cohort (i.e.,  $\alpha_{t,E_i}$  and  $\beta_{n,E_i}$  have  $E_i$  subscripts). The motivation for this flexibility is that with cohort-specific controls, our estimates are the same as estimating  $\delta_n$  cohort-by-cohort except we use regression weights to aggregate the estimates. Consequently, there are two differences between our setting and the standard “staggered adoption” DiD setting. First, we avoid the potential negative weight issues that arise from heterogeneous, cohort-specific

---

<sup>24</sup>This test also allows us to evaluate one version of the “union threat” hypothesis. In particular, it allows us to test whether losing a union election by a small margin of victory affects an establishment differently than losing by a larger margin. This test, however, would not capture across-the-board union threat effects that don’t vary by vote shares.



treatment effects (Sun and Abraham, 2020; Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020).<sup>25</sup> Second, we only need to assume that our identifying assumption in equation 2.5 holds *within each cohort*.<sup>26</sup> Both differences are because our estimates come from comparing winning and losing elections *within the same cohort* rather than across cohorts which might lead to negative weights or alternative parallel trend assumptions. Finally, we cluster standard errors at the *firm* level.<sup>27</sup>

**Establishment-Level Controls** To account for observable determinants of election outcomes, we control for progressively richer establishment-level characteristics. All controls are from one year before the election and interacted with event time. The event-time interaction allows for flexible pre- and post-election trends by baseline characteristics (e.g., differential employment growth rates for large versus small establishments). Our first *industry and employment controls* specification includes baseline employment and three-digit NAICS industry-by-year controls.<sup>28</sup> The motivation for starting with these covariates is that they are among the strongest predictors of union election victory (see Section 2.2), and they are key determinants of establishment growth and survival dynamics (Dunne et al., 1989; Haltiwanger et al., 2013). Next, we add other characteristics in the LBD (baseline payroll, establishment age, and single/multi-establishment status) and an indicator for whether we observe a previous FMCS contract at the establishment (i.e., another bargaining unit already unionized at the establishment).<sup>29</sup> We refer to this specification as the *pooled controls* specification. Finally, we interact all controls from the previous specification with cohort (i.e., year of election). This is our preferred *flexible controls* specification. The cohort interactions result in the within-cohort identification assumption discussed previously. We show,

---

<sup>25</sup>We test for negative weights on each cohort treatment effect using Sun and Abraham (2020)’s `eventstudyweights` package.

<sup>26</sup>With multiple cohorts, our identifying assumption is  $E[Y_{it}^0 - Y_{it-1}^0 | X_i, E_i, V_i] = E[Y_{it}^0 - Y_{it-1}^0 | X_i, E_i]$ . Thus, we do not require that selection into elections in the 1980s is the same as selection into elections in the 2000s.

<sup>27</sup>This accounts for serially correlated establishment-level outcomes across time and across elections at different *establishments* within the same *firm*. Our regression weighting to aggregate the  $\delta_n$  estimates easily accommodates this level of clustering.

<sup>28</sup>Our baseline specification interacts industry by year and event time because some of our outcomes are cumulative measures (e.g., the DHS growth rates and survival). Thus, just industry-by-year FEs would capture industry growth rates over different time horizons. For all continuous variables, we flexibly parameterize their functional form with *decile* fixed effects.

<sup>29</sup>The motivation for including the previous contract control is that union elections are more successful when other workers at the same establishment are already unionized (Bronfenbrenner, 2002). The selection into such elections may also differ from the selection into elections for an establishment’s first bargaining unit. When we pool all industries together, we interact controls in this specification with three coarse industry groups (e.g., manufacturing, services, and “other”). This keeps them at the same level of granularity for our overall and manufacturing estimates.

however, that our main results are robust to pooling controls across cohorts or only including the employment and industry controls.

**Establishment-Level Outcomes** The first outcome we consider is the Davis, Haltwanger and Schuh (1996) (DHS) symmetric growth rate for employment and payroll

$$G_{it} = 2 \times \frac{Y_{i,t} - Y_{i,t=E_i-1}}{Y_{i,t} + Y_{i,t=E_i-1}}. \quad (2.9)$$

This growth rate is a second-order approximation of the log difference from time  $t$  to one year before the union election,  $E_i - 1$ . Yet, it accommodates establishment exit as  $G_{it}$  equals  $-2$  for establishments that do not exist (i.e., have zero employment).<sup>30</sup> Consequently, a  $-0.2$  value of  $G_{it}$  could represent either an approximately 20 % decline in intensive margin employment with no survival effects or a 10 percentage point decrease in the likelihood of survival. Since the growth rate accommodates exit, we can simultaneously evaluate pre-trends and interpret treatment effects, even if unionization affects establishment survival which could lead to a selected group of survivors. For this reason, the DHS growth rate is commonly used to analyze firm growth dynamics.<sup>31</sup>

To estimate the effect of unionization on extensive margin employment growth, we include establishment survival as an outcome (an indicator for whether the establishment exists at time  $t$ ). We can compare the survival effect to the DHS growth rate effect to answer how much of the DHS growth rate effect is *mechanically* due to exit (e.g.,  $G_{it} = -0.2$  could be completely explained by a 10 pct. pt. decrease in survival). However, the residual, the part of  $G_{it}$  unexplained by exit, could be either intensive-margin employment changes or selective exit based on employment growth rates.

Finally, we define the outcome as log employment or log payroll. A challenge with interpreting the effects on log outcomes is that treatment effects on establishment survival can bias comparisons of potentially selected *survivors*. The pre-trends for these log outcomes, however, are a useful complement to the DHS growth rate pre-

---

<sup>30</sup>Conventionally, the growth rate is defined annually (e.g., from  $t - 1$  to  $t$ ) but we define it over longer time-horizons to measure cumulative changes. Additionally, since our sample restrictions impose non-zero employment at  $t = E_i - 1$ ,  $G_{it}$  is never equal to 2 which it usually equals for entrants. Establishments that do not exist at time  $t$  before the election have  $G_{it} = -2$ .

<sup>31</sup>See Haltiwanger et al. (2013); Chodorow-Reich (2014) for general use and Arnold (2019); Davis et al. (2014) for DiD contexts.

trends.<sup>32</sup> For interpreting the treatment effects on log outcomes, we provide two ways of partially alleviating the selective survival concern. First, all specifications with log outcomes include establishment FEs that account for *level* differences between the surviving and exiting establishments.<sup>33</sup> Second, for some results, the timing of the log outcome versus survival effects suggests intensive margin effects (e.g., large effects on log outcomes before any substantial survival effects). Yet, we still recommend interpreting the treatment effects for log outcomes with caution since we cannot completely eliminate potential bias from selective survival.

For our outcomes, we make related parallel trends functional form assumptions. For log outcomes, we assume that log employment and payroll would have (conditionally) evolved in parallel, which we view as a reasonable restriction in this setting.<sup>34</sup> Additionally, we can test whether the restriction holds in the pre-period (i.e., equation 2.6). For establishment survival, we assume that the survival probabilities between elections with different vote shares would have (conditionally) been equal had no election occurred at the establishments. We cannot test whether this assumption holds in the pre-period since all establishments exist at event-time zero. However, we can test whether this functional form assumption holds between the losing elections with different vote shares (i.e., equation 2.7). For the DHS growth rate, the outcome is approximately a linear combination of log employment changes and survival probabilities, so the two functional form assumptions we have already made imply parallel trends in the DHS growth rate.<sup>35</sup>

## 2.5 Empirical Results: Overall Employment and Survival Effects

In this section, we estimate the effects of successful union elections on establishment employment and survival. We first analyze the differences in employment growth rates

---

<sup>32</sup>The DHS pre-trends combine intensive and extensive margin employment changes. However, in specifications where we control for baseline establishment age, the DHS pre-trends will closely approximate pre-trends for log outcomes.

<sup>33</sup>For DHS growth rates and survival, we do not include establishment FEs. For DHS growth rates we capture the time-invariant component by differencing relative to  $t = E_i - 1$ . For survival, it is unclear what time-invariant characteristic FEs would capture.

<sup>34</sup>For example, consider two firms with the same Cobb Douglas production function parameters but different baseline TFP and/or input and output prices. In response to the same demand shock (e.g., the same proportional change in the price of output), their log payroll and log employment would both evolve in parallel while their levels would diverge.

<sup>35</sup>Specifically, we assume  $E[\Delta \ln Y_{it}^0 | X_i, V_i] = E[\Delta \ln Y_{it}^0 | X_i]$  and  $E[\mathbb{1}[Y_{i,t}^0 = 0] | X_i, V_i] = E[\mathbb{1}[Y_{i,t}^0 = 0] | X_i]$  which imply  $E[G_{it}^0 | X_i, V_i] \approx E[G_{it}^0 | X_i]$ .

between establishments with winning and losing elections. Next, we implement several tests of our parallel trends identifying assumption described in Section 2.4. Finally, since we later focus on manufacturing, we present our estimates and falsification checks separately for all industries and just for elections in manufacturing.

## Overall Employment and Survival Estimates

We start by estimating establishment employment growth for successful versus unsuccessful elections. Figure 2-3 plots the  $\delta_n$  coefficients from estimating the “pooled cohort” specification in equation 2.8 for elections with 20-80 % vote shares. Panel A. plots the estimates for DHS employment growth relative to one year before the election. Panel B. includes log employment as the outcome. Both panels include estimates with no controls (i.e., only year by cohort FEs), the industry and employment controls, and the flexible control specification described in Section 2.4.

The estimates in Panels A. and B. show that establishments with successful elections had similar conditional pre-election growth rates to establishments with unsuccessful elections but experienced large relative employment decreases following the election. The first, “no control” estimates, however, show that, without any controls, establishments where the union won had relatively slower pre-election employment growth rates than establishments where the union lost.<sup>36</sup> However, the next “industry + emp ctrls.” estimates show that just conditioning on baseline employment and industry yields similar pre-election growth rates for DHS and log employment. As discussed in Section 2.4, we start with these controls because they are prominent predictors of election outcomes and establishment employment growth. Starting one year after the election, this specification also shows decreased employment for establishments with successful union elections. The effects stabilize approximately three years after the election. Finally, the results from the third “flexible control” specification show that our pre- and post-election employment growth estimates are very similar adding the richer and more flexible establishment-level controls. In the next section, we show that our other “vote-share heterogeneity” tests of our identification assumption also yield similar estimates with just the industry and employment versus flexible control specifications.

To help interpret the magnitude and timing of the employment effects, Panel C. in

---

<sup>36</sup>Without any controls, the DHS growth rates and log employment pre-trends measure somewhat different growth rates. The DHS employment growth rates combine intensive and extensive margin changes, while log employment only captures intensive margin changes. The measures are approximately the same in the control specifications that include establishment age.

Figure 2-3 additionally plots payroll and establishment survival estimates. Specifically, it includes estimates of DHS employment and payroll growth and establishment survival with the flexible control specification. We find that establishment payroll initially declines faster than employment. This difference could be due to either compositional shifts to low-wage workers or differences in the timing of the payroll versus employment measures described in Section 2.3. Five years after a successful union election, the cumulative DHS employment and payroll growth rates are -0.13 and -0.14 lower, respectively, than establishments with unsuccessful elections (consistent with a 14 % decrease in payroll or a seven pct. pt. decrease in survival likelihood). Appendix Figure B-3 presents estimates from the same specification for *log* employment and payroll. These estimates allow us to reject five-year, pre-election growth rate differences of more than 3.5 % for employment and 1.8 % for payroll. Unlike the DHS measures, we find larger five-year log payroll than employment declines. Although this evidence would be consistent with long-run compositional changes, we recommend interpreting it cautiously given potential biases from selective exit.

The survival estimates in Panel C. of Figure 2-3 indicate that most of the decrease in DHS employment and payroll growth rates is from a lower likelihood of establishment survival. To decompose what share of the DHS effects is from survival effects, we plot the survival estimates on a separate y-axis scaled to be one-half the DHS growth rate axis. Comparing the exit and DHS coefficients illustrates how much of the DHS effect can be mechanically explained by the survival effect (see Section 2.4). The estimates show that five years after an election, establishments with successful elections are four pct. pts. less likely to survive, and this effect increases slightly to five pct. pts. after ten years. Consequently, about two-thirds of the -0.13 five-year DHS employment growth rate estimate is mechanically due to decreased establishment survival. Finally, the relatively slower timing of the survival versus employment effects is consistent with an increased legal risk of immediately closing an establishment following an election. For example, Munger et al. (1988) describe how a short time between an election and establishment closure could be used as evidence that the closure is an unfair labor practice due to its “intent to chill unionism” across an entire firm.

Given our later focus on manufacturing, Figure 2-4 presents the same estimates including only manufacturing elections. For these elections, we find similar pre-election employment growth rates (i.e., a lack of pre-trends) even without baseline industry and employment controls. For example, Panel B. shows that, without any controls, we can rule out five-year employment growth rate differences of more than five percent. One explanation for the lack of detectable pre-trends without controls is that by only

comparing elections in manufacturing, we may account for sector and employment differences that the controls capture when we include all industries. Additionally, for manufacturing elections, the magnitude of the treatment effects is larger than the effects for all industries (e.g., the five-year DHS employment estimates are -0.17 versus -0.13, and the five-year survival effects are -0.05 versus -0.04, respectively). We show later that this difference is because the effects of unionization in the service sector are much smaller.

## Vote-Share Heterogeneity Tests of Identifying Assumption

We next provide further evidence that our results are driven by unionization by assessing several testable implications of our identifying assumption. Additionally, we estimate treatment effect heterogeneity by the unions' margin of support. To implement these tests, we first present visual evidence of how treatment effects and pre-election trends vary across the vote-share distribution and then implement parametric tests of linear trends in establishment outcomes by election vote shares.

**Nonparametric Vote-Share Heterogeneity** To estimate pre-trends and treatment effects for different parts of the vote-share distribution, we estimate the following modified version of our main DiD specification

$$Y_{it} = \alpha_{t,E_i} + \sum_g \sum_n \delta_{g,n} \cdot \mathbb{1}[t - E_i = n] \times \mathbb{1}[V_i \in \mathcal{V}^g] + X_i' \beta_{n,E_i} + \varepsilon_{it} \quad (2.10)$$

where  $\mathcal{V}^g$  are exhaustive subsets of the vote-share distribution.<sup>37</sup> We partition the vote-share distribution into eight groups (0 – 20 %, 20 – 30 %, 30 – 40 %, 40 – 50 %, 50 – 60 %, 60 – 70 %, 70 – 80 %, and 80 – 100 %). We omit the 20 – 30 % group, so the estimates for each group are relative to outcomes for 20 – 30 % vote-share elections. This specification allows us to assess the two testable implications of our identifying assumption described in Section 2.4. First, we test whether establishments' *pre-election* outcomes are similar across the vote-share distribution by comparing  $\delta_{g,n}$  estimates for  $n < 0$  (i.e., testing equation 2.6). Second, we test whether post-election outcomes differ between losing elections with different vote shares by comparing  $\delta_{g,n}$  estimates for  $n > 0$  and  $V_i \leq .5$  (i.e., testing equation 2.7). For presenting these estimates, we start with manufacturing where the results closely support our identifying assumption, making the setup easier to explain. We then turn to estimates for all industries where

---

<sup>37</sup>We omit the establishment FEs here because we only estimate this specification for DHS growth rates and establishment survival where we never include establishment FEs.

there is evidence of rejections of these tests for some outcomes. We find, however, that the violations are driven by elections that lost by exactly 50 %, where we would expect such differences.

Figure 2-5 presents results from estimating equation 2.10 for all manufacturing elections. The estimates include our *flexible controls* specification (see the following parametric vote-share heterogeneity analysis for robustness to alternative controls). Panel A. includes pre-period and treatment-effect estimates for each vote-share group with DHS employment growth as the outcome. First, the five-, three-, and two-year pre-trend estimates are similar across almost the entire vote-share distribution relative to 20-30 % elections (the one exception is 0-20 % elections which we exclude from our main analysis). These results support our identifying assumption by showing that the similarity between pre-election employment growth rates holds between much finer vote-share groups. Second, the figure shows that none of the five- and ten-year treatment effect estimates for losing elections are significantly different than the estimates for 20-30 % elections. These results provide reassurance against the concern that our main estimates are driven by future productivity shocks correlated with vote shares. In that case, we would also expect these shocks to cause different outcomes for losing elections with different vote shares. Finally, the five- and ten-year treatment effect estimates for winning elections increase in the union vote share but are not statistically different (e.g., -0.18 versus -0.28 ten-year estimates for 50-60 % and 70-80 % elections, respectively).

Figure 2-5 Panel B. plots the same estimates with establishment survival as the outcome variable. Although we cannot test for pre-trends in establishment exit rates, we can test our parallel trends assumption for establishment survival by estimating whether losing elections' with different vote shares had different post-election survival rates. Reassuringly, the survival rates for all losing election vote-share groups are not statistically different than the survival rate for 20-30 % elections. For winning elections, however, the figure shows that the long-run effects on survival increase in the union vote share, although the differences are not statistically different across groups.

Figure 2-6 presents the vote-share heterogeneity estimates for elections in all industries. Panel A. shows that for our main sample of 20-80 % vote-share elections, we find very similar pre-election employment growth rates for elections with different vote shares. For 0-20 and 80-100 % vote-share elections, however, we find evidence of different pre- and post-election growth rates, which is one motivation for excluding these elections from our main analysis. For post-election outcomes, we find similar

DHS employment growth rates between 20-30 % and 30-40 % elections but find somewhat slower employment growth for 40-50 % vote-share elections. The ten-year estimate for 40-50 % elections is also significantly different from zero at the 10 % level. However, these negative estimates are driven by elections where the union received exactly 50 % of votes, and there are multiple reasons to expect differences between 50 % vote-share elections and elections where the union lost by slightly larger margins.<sup>38</sup> To see this, when we estimate the 40-50 % effects excluding elections with a 50 % vote share, the five-year estimate is -0.015 (SE 0.025) and the ten-year estimate is -0.032 (SE 0.028). Both estimates are much smaller than the treatment-effect estimates for the neighboring group of winning elections with 50-60 % of votes (-0.11 and -0.16 at the two time horizons). Furthermore, Panel B. of Figure 2-6 shows that there is no evidence of differential survival rates between 20 – 30, 30 – 40, and 40 – 50 % losing elections.

**Parametric Vote-Share Heterogeneity** To complement the previous *nonparametric* analysis, we estimate a series of *parametric* vote-share heterogeneity tests. Specifically, we test for linear trends in pre- or post-election outcomes by vote share. There are two motivations for this extension. First, these tests may have more power. Second, they provide a parsimonious way to assess robustness to different sets of controls. Specifically, we show that our estimates from these tests are qualitatively the same with just the *employment and industry controls* and the *flexible control* specification.

We first test for a linear trend in pre-election employment growth rates across the vote-share distribution. Second, we test for linear trends in post-election outcomes separately for winning and losing elections. To implement these tests, we estimate a modified version of the specification in equation 2.8. Specifically, instead of only interacting event-time with the winning indicator (e.g.,  $\mathbb{1}[V_i > .5]$ ), we include the

---

<sup>38</sup>There are multiple potential reasons for outcome differences at establishments with 50 % vote-share elections. First, due to the discreteness of total votes, elections with exactly 50 % vote shares have a small number of total votes cast (see the “integer problem” in DiNardo and Lee (2004)). For example, based on the NLRB data, the median (mean) number of voters in 50 % vote-share elections is 12 (22) compared to 50 (96) voters in elections with vote shares in the [45, 50) range. Although our employment controls capture corresponding establishment size differences, they do not capture potential differences in the bargaining unit size to employment shares. Second, the manipulation around the 50 % threshold is largely due to challenges to single votes which disproportionately affects elections with 50 % vote shares Frandsen (2017). Finally, experiencing an election where the union loses by just one vote may have a different effect on employers than losing by larger margins of support.



following interactions with event-time<sup>39</sup>

$$\underbrace{\mathbb{1}[t - E_i = n]}_{\text{Event Time}} \times \begin{cases} \rho \cdot V_i & \text{if } n < 0 \\ \eta \cdot \mathbb{1}[V_i > .5] + \theta \cdot V_i + \tau \cdot [V_i - .5]^+ & \text{if } n \geq 0 \end{cases} \quad (2.11)$$

For the pre-trend, vote-share heterogeneity test (i.e.,  $n < 0$ ), we do not include an interaction with treatment, so pre-period “treatment effects” would be captured in the  $\rho$  estimates. For the post-election outcome tests, we include an interaction with treatment so  $\eta$  estimates the treatment effect for close winning elections (e.g., a linear RD estimate). Consequently, the  $\theta$  and  $\tau$  coefficients estimate slope differences that do not include the change in outcomes right around the 50 % threshold.

Table 2.2 includes estimates of pre-election growth rate trends by vote share,  $\rho$ , for one to five years before the election. The estimates are for the main 20-80 % vote-share sample and are presented separately for all industry and manufacturing elections. We present estimates with the *employment and industry* and *flexible control* specification. Across all estimates, we never find significant pre-election growth rate trends.<sup>40</sup> These estimates complement the nonparametric evidence in Figures 2-5 and 2-6 by showing that the lack of pre-trends across the vote-share distribution holds formally testing for linear trends and only including more limited controls.

Table 2.3 presents the estimates testing for post-election outcome vote-share trends. We present estimates of separate slopes for losing elections (i.e.,  $\theta$ ) and winning elections (i.e.,  $\theta + \tau$ ). This table includes our preferred *flexible control* specification but Appendix Table B.2 shows qualitatively similar results only including the *employment and industry* controls. Motivated by the potential issues with 50 % vote-share elections (see footnote 38), we also present the estimates with and without excluding 50 % elections from the estimates.

The results for all industries in Table 2.3 Panel A. indicate significant negative

---

<sup>39</sup> $[V_i - .5]^+$  is equal to  $[V_i - .5] \times \mathbb{1}[V_i > .5]$ . Since we only estimate this specification for elections with 20-80 % vote shares, we actually shift the vote-share variables to all start at zero (e.g., subtracting 0.2 from the  $V_i$  variables and 0.3 from the winning vote-share variable). This ensures that the vote-share coefficients only capture slope and not level differences.

<sup>40</sup>To assess the magnitude of the estimates, the largest positive point estimate is 0.05. A reasonable benchmark is what the estimates imply for the differences between 20-30 and 70 -80 % elections presented in Figure 2-6. Since the midpoints between those bins are 0.5 apart, the 0.05 coefficient implies a small difference in pre-election employment growth rates of around 2.5 % between 20-30 and 70 -80 % elections. For all industries, the confidence intervals also allow us to reject large trends in pre-election employment growth rates (e.g., with our flexible controls we can rule out five-year growth rate differences between the previous groups of more than around four percent). For manufacturing, however, the 95 % confidence intervals on some of the estimates would include relatively large pre-election growth rate differences.

trends in DHS employment growth rates by vote share for both losing and winning elections. However, mirroring the nonparametric analysis, when we exclude the 50 % elections, we do not detect significant trends for losing elections. However, we find significant negative trends by vote share for winning elections, consistent with increasing treatment effects for larger margin-of-support elections. For example, we estimate a vote-share trend of -0.066 (SE 0.122) for losing elections and -0.389 (SE 0.149) for winning elections for five-year DHS employment growth rates. For establishment survival, we never find significant trends for winning or losing elections. For some specifications, the losing election trends are actually positive, further supporting our overall survival estimates not being driven by negative productivity shocks correlated with election vote shares.

The manufacturing estimates in Table 2.3 Panel B. are similar to the estimates for all industries pooled together. Without excluding the 50 % elections, we find negative although insignificant DHS employment trends for losing elections. However, dropping the 50 % elections results in smaller trends for losing elections and large although insignificant vote-share heterogeneity estimates for winning elections (e.g., five-year DHS trend estimates of -0.072 (SE 0.199) for losing elections and -0.406 (SE 0.299) for winning elections). The manufacturing survival estimates are also never significant for winning or losing elections and, at times, positive for losing elections.

Overall, these estimates in Table 2.3 also confirm the nonparametric post-election estimates in Figures 2-5 and 2-6. First, they show that, excluding the 50 % elections, the lack of a trend in post-election DHS employment growth by vote share holds testing for linear trends and only including the employment and industry controls. For establishment survival, we also cannot detect trends with and without excluding the 50 % vote-share elections. We note, however, that the 95 % confidence intervals for some of these estimates include relatively large post-election growth rate differences for losing elections. Additionally, the estimates for winning elections provide a formal test of treatment effect heterogeneity by vote share. Specifically, for the overall DHS employment growth rate estimates, we find significant vote-share heterogeneity.<sup>41</sup> For our establishment survival effects, however, we do not find significant evidence of vote-share heterogeneity.

**Employment and Survival Effect Robustness** We next present two additional checks that further validate our overall estimates of the negative impacts of unionization

---

<sup>41</sup>For manufacturing, the estimates are only significant at the 10 % level although, Appendix Table B.2 presents more significant manufacturing vote-share heterogeneity estimates only including the industry and employment controls.

on establishment survival and employment. First, we assess whether our estimated effects increase in the size of the bargaining unit (Lee and Mas (2012) conduct a similar test). The motivation is that the relative share of unionized workers should mediate many direct effects of unionization. However, potential violations of our identifying assumption may not be mediated by the share of unionized workers (e.g., workers voting based on their expectations of future company performance or managers' competence). Appendix Table B.3 presents the coefficient estimates from interacting the three-, five-, and ten-year treatment indicators with the share of each establishment's total employment included in the bargaining unit (see Appendix B.3 for details). It shows that the three and five-year treatment effects are significantly increasing in the bargaining unit share for both outcomes. These estimates confirm that the effect seems mediated by the share of workers gaining union representation. The interactions, however, are no longer significantly different than zero at the ten-year horizon. One explanation for the lack of persistence is that the relative size of the bargaining unit versus establishment employment could change substantially over time.

Second, Appendix Figure B-4 plots DHS employment growth rate estimates with *ten-year* pre- and post-periods.<sup>42</sup> First, it shows no evidence of large pre-trends in employment growth rates up to ten years before elections in manufacturing or for all industries pooled together. Although we find significant pre-period estimates six, seven, and eight years before the election in all industries, the estimates are economically small (e.g., approximately 1.7 to 2.0 percent differences). Moreover, the ten-year pre-period estimate is insignificant, and its confidence interval allows us to rule out employment growth differences of more than approximately 3.2 percent. Second, the figure shows that the post-election effects are relatively stable starting three years after the election. For manufacturing, however, there is a slight increase in the effect from years five to ten.

## Industry-Specific Employment and Survival Estimates

We next separately estimate the effects for different industries and show that the overall effects are driven by non-service-sector elections. There are multiple reasons to expect heterogeneity across industries. First, the quality of labor relations may

---

<sup>42</sup>Note, since our data start in 1976 and our elections start in 1981, the -10 to -6 estimates are from an unbalanced panel (e.g., the -6 point estimate is based on a different number of election years than the -5 estimate). This is one motivation for why we focus on the -5- to 10-year estimates with the balanced panel for the main analysis.

differ across sectors (e.g., the higher post-election strike propensity for manufacturing elections in Figure 2-2 suggests more adversarial relations). Second, firms in different industries may differ in how easily they can “avoid unionization.” For example, mobile, multi-establishment manufacturing firms may avoid working with new unions by shifting production to other establishments. However, this tactic may be difficult in non-tradable industries (e.g., hospitals) or tradable industries with ties to their local area (e.g., hotels).

To estimate this heterogeneity, we classify our elections into three broad industry groups: manufacturing, services, and a residual “other” group.<sup>43</sup> Weighted by the number of eligible voters, 70 % of our service-sector elections are for healthcare (e.g., hospitals and nursing homes), security, restaurants, grocery stores, universities, and print media establishments. The other category includes agriculture, construction, mining, transportation and warehousing, utilities, and wholesale trade.

To estimate the industry-specific heterogeneity, we use the following specification for a categorical heterogeneity variable  $H_i$  (e.g., the three industry groups)<sup>44</sup>

$$Y_{it} = \alpha_{t,E_i} + \sum_h \sum_n \delta_{h,n} \cdot \mathbb{1}[t - E_i = n] \times \mathbb{1}[V_i > .5] \times \mathbb{1}[H_i = h] + X_i' \beta_{n,E_i} + \varepsilon_{it}. \quad (2.12)$$

The  $\delta_{h,n}$  coefficients now estimate the dynamic effects of successful union elections for elections with  $H_i = h$ . We also estimate all subsequent heterogeneity in Section 2.6 using equation 2.12.

Table 2.4 presents the DHS employment growth and survival effects estimated

---

<sup>43</sup>These classifications are based on the Fort and Klimek (2016) 2012 NAICS sectors of the Census establishment we match to each election. We define manufacturing as NAICS sectors 31-33 and services as NAICS 51-81 and NAICS 44-45. Our classification of services differs slightly from other measures. For example, compared to the sampling frame for the Census’s Service Annual Survey, we include retail trade in the services group and exclude utilities and transportation and warehousing. Bronfenbrenner (2002) also excludes utilities and transportation and warehousing from service-sector unions. The motivation for these changes is that we want to capture a notion of “service-sector unionization.” Retail workers (e.g., grocery store workers) are commonly referred to as part of service-sector unionization. As evidence of this, the "OUR Walmart" campaign was frequently described as attempting to unionize service workers (Brown, 2011). Alternatively, most elections in utilities involved electrical workers (more similar to crafts unions in the building trades), and many elections in transportation and warehousing involved drivers.

<sup>44</sup>This specification has two advantages relative to restricting the sample for each value of  $H_i$ . First, we can pool the control coefficients across heterogeneity groups and use all the data to estimate their coefficients. For all heterogeneity estimates, we also add the specific heterogeneity group as an additional control in  $X_i$  so that we account for any differential trends by the specific heterogeneity groups. Second, it allows us to easily conduct Wald tests of equality across the different heterogeneity groups.

separately for each industry group. First, there is limited evidence of employment growth pre-trends for any of the groups. The only marginally significant pre-period estimate is for the service sector, where we find the smallest main effects. Second, the overall employment and survival decreases are driven by similarly sized effects for elections in manufacturing and the other sector. For elections in the service sector, the effects of unionization are substantially smaller. For example, the five-year DHS employment growth estimates for manufacturing and services are -0.174 (SE 0.029) and -0.057 (SE 0.024), respectively. Moreover, the ten-year survival estimate for the service sector is not significantly different than zero, and the confidence interval allows us to reject effects of more than four pct. pts.

Appendix Table B.4 shows that the smaller effects of unionization in the service sector are robust to alternative sets of controls and sample selection criteria. Specifically, it presents the point estimate and standard error of the difference between the manufacturing and service-sector coefficients over each time horizon. The effects in manufacturing remain significantly larger when we (1) pool controls across cohorts, (2) restrict the sample to 30-70 % vote-share elections, and (3) restrict the sample to elections where the size of the bargaining unit was at least 25 % of total establishment employment. The results with the last restriction show that the smaller effects of unionization in the service sector are not because service-sector elections are more likely to only include a relatively small share of the establishments' overall employment.

## 2.6 Testing for Manager Opposition and Union Avoidance

After documenting the large overall impacts of successful union elections on establishment employment and survival, we test whether some of this effect is due to managers' dislike of working with unions or firms' ability to easily avoid working with new unions. For this analysis, a sector-specific analysis is the most appropriate because the tactics that employers can use to avoid unions may differ across sectors. In manufacturing, a common union avoidance tactic for multi-establishment firms during this time was shifting production away from unionized establishments to non-unionized establishments (Bluestone and Harrison, 1982; Verma, 1985; Kochan et al., 1986a). However, in construction, one of the industries in the "other" industry group where we also find negative effects, most firms are single-establishment firms, so they cannot shift production across establishments (Butani et al., 2005).<sup>45</sup> So, the same test might

---

<sup>45</sup>Alternatively, in construction, there is evidence of firms avoiding unions by going "dual shop" and opening a new non-unionized shop that does previous work of the unionized shop (Evans and Lewis,

not capture union avoidance across sectors. Consequently, we focus on elections in manufacturing for three reasons. First, manufacturing is the largest sector where we find negative effects.<sup>46</sup> Second, as discussed above, we can use our data’s high-quality establishment and firm linkages to test for union avoidance via production shifting. Finally, we have detailed measures of establishment-level productivity in manufacturing that we use for this analysis. In this section, we refer to manufacturing “establishments” and “plants” interchangeably.

## Multi- versus Single-Establishment Manufacturing Firms

The first part of our hypothesis is that firms avoid working with unions by shifting production away from newly unionized plants to other plants. Since this shifting is only possible for firms with multiple plants, we start by estimating whether the effects of union elections are larger at establishments part of multi-establishment (or multi-unit, MU) firms versus single-establishment (SU) firms. Specifically, we define “an election at an MU firm” based on whether the establishment’s firm had at least one other establishment under its control one year before the election.

Figure 2-7 plots the estimates for the DHS employment growth and survival effects for elections at SU versus MU firms. The left panel plots the cumulative DHS employment growth rates for five years before and three, five, and ten years after the election. Below each x-axis label, we include the p-value of the difference between the SU and MU estimates. Reassuringly, there is no evidence of differential pre-election employment growth rates for either group. After the election, we find significantly larger employment declines for elections at MUs at the three- and ten-year horizon. The estimates for SUs are, however, still negative and significant. For the establishment survival estimates in the right panel, the differences are even more striking. For all post-election time horizons, the effects are significantly larger for MUs, and none of the estimates for SUs are significantly different than zero. For example, the ten-year survival estimates are  $-0.122$  (SE 0.021) versus  $-0.029$  (SE 0.029) for MUs and SUs, respectively.

Appendix Table B.5 shows the robustness of these estimates to (1) including

---

1989). Although we may see these new establishment openings in our data by linking establishments across Census *firmids*, owners may try to disguise the common ownership of these establishments to avoid potential labor-law issues with going “dual shop” (Milne, 1985). In other sectors, there is also evidence of employers using temporary workers to replace newly unionized workers (Hatton, 2014).

<sup>46</sup>For all sectors in our manufacturing and other industry groups where we find large negative effects, manufacturing makes up 54 % of elections compared to 18 % for transportation and warehousing (the next largest sector). Weighting by the number of eligible voters, manufacturing comprises 68.8 % of voters due to its relatively large bargaining units.

controls pooled across cohorts and (2) only using 30-70 % vote-share elections to estimate effects. It presents the difference and standard errors between the SU and MU estimates. The estimates are very similar with only pooled controls. For the 30-70 % bandwidth, we still estimate substantially larger survival effects for MU firms (e.g., six pct. pts. at the ten-year horizon), but the larger standard errors only lead to a significant difference for survival at the five-year horizon.

We interpret these results as showing that the effects of unionization on establishment survival in manufacturing are driven by plant closings at MU firms. For the overall employment declines, the effects are also significantly larger at the multi-establishment firms but still significant for SUs. This evidence is consistent with MU firms responding to unionization by shifting production across plants which we investigate more directly next. As an alternative explanation, MU firms may have a greater incentive to react strongly to unionization due to concerns about unionization spreading to their other establishments. We also investigate this later by focusing on entirely non-unionized MUs where this incentive may be even sharper.

## Employment Shifting after Successful Elections

Next, we directly test the hypothesis that manufacturing firms avoid working with new unions by shifting production to other plants. Specifically, we analyze whether a successful election at one of a firm's plants increases the employment and survival of the firm's other plants. While the production-shifting hypothesis predicts positive effects on other plants, other mechanisms like input-output linkages or firms' financial constraints predict negative spillovers (Boehm et al., 2019; Giroud and Mueller, 2017). However, one prediction of the production-shifting hypothesis is that the positive effects should be the largest at plants where it is easiest to produce the same products as the election plant. Consequently, we start by only considering the effects on other manufacturing plants and then restrict to plants in the same three-digit NAICS industry as the election plant.<sup>47</sup>

To construct the sample for this analysis, we start with all manufacturing elections in a specific year at MU firms. Next, we take all of the firms' *other* manufacturing plants that existed during the election year and never experienced their own union

---

<sup>47</sup>LaLonde et al. (1996) similarly analyze within-firm employment spillovers of successful union elections. They do not find any evidence of spillovers but only consider the effects on all other manufacturing plants, where we also do not find only spillovers. We only find evidence of spillovers when we focus on other plants within the same three-digit NAICS industry. Bradley et al. (2017) similarly find that firms shift R&D activity away from newly unionized establishments.

election.<sup>48</sup> We then calculate these plants' DHS employment growth rates before and after the election relative to one year before the election. Finally, we stack these observations from all cohorts together and estimate a modified version of our main DiD specification 2.8.<sup>49</sup> The two differences from our main specification are that (1) relative time and vote-share variables are defined from the election at the firms' *other* plant, and (2) we weight the regression by each plant's share of its firms' total employment.<sup>50</sup> The reason for the weighting is that the sample could include multiple plants matched to each election, and we want to weight each election equally (i.e., not give the most weight to elections at firms with the most other plants). For this specification, we two-way cluster the standard errors by firm and establishment.<sup>51</sup>

Figure 2-8 Panel A. plots the dynamic employment effects of successful elections on the firms' non-election plants. It presents estimates that include all manufacturing plants and that only include plants in the same three-digit NAICS industry as the election plant. For all manufacturing plants, there is no evidence of relatively higher employment growth at the other plants following successful elections. There are two things to note about this result. First, even if firms shifted employment away from newly unionized plants, it is not surprising that we do not find spillovers when we include *all* other plants. Specifically, many of these plants may have produced different products than the election plant, making production shifting more costly. Second, it is reassuring that we do not estimate *lower* post-election employment at the *other* plants of firms with successful elections. If our plant-level productivity shocks bias our estimates of the direct effects of unionization, we might expect some of these shocks to be firm-wide. Yet, the estimates in Figure 2-8 allow us to rule out differences in five-year DHS employment growth rates of more than -0.04 which is much smaller

---

<sup>48</sup>We exclude plants that ever experienced an election so our "spillover estimates" are not contaminated by direct effects. Yet, this conditioning could selectively bias our sample. The most plausible mechanism, however, biases us *against* finding positive spillovers. Specifically, assume that successful elections lead to *more* future elections at a firm. Since elections occur at relatively fast-growing plants and the plant needs to survive to hold a future election, we would drop faster-growing plants at firms with successful elections. This would downward bias our overall spillover estimates.

<sup>49</sup>This construction results in some establishments being in the data set multiple times if their firms experience multiple union elections. For our baseline analysis, we avoided this problem by taking the first election at each establishment. For this analysis, similar conditioning is more difficult because the Census `firm` IDs change over time, even for firms that stay in business, and establishments can switch to different `firm` IDs. This also motivates our two-way clustering by firm and establishment.

<sup>50</sup>For the denominator, we only include employment at plants in the sample so the employment weights sum to one.

<sup>51</sup>We use each establishments' `firmid` during the election year (e.g., the clustering variable is fixed over time). Yet, since an establishment can appear multiple times in the sample, establishment clustering is not nested by `firmid` clustering.



than our overall estimate for elections at MUs of -0.21.

We find significant employment growth effects when we restrict the sample to other plants that produced similar products to the election plants. The solid estimates in Panel A. present the coefficients estimated from just the firm's other plants in the same three-digit NAICS industry as the election plant. Two years after the election, we estimate growth rate differences of 0.043 (SE = .019) for plants at firms with successful versus unsuccessful elections. These effects persist three and four years after the election. However, the effect becomes insignificant five years after the election and remains insignificant ten years following the election. Additionally, Table 2.5 presents both DHS growth rate and survival estimates and indicates that some of the overall increase in employment growth is due to an increased likelihood of plant survival.

Figure 2-8 Panel B. further splits up the same-industry elections based on whether or not the *election plant* made up a large share of the firm's total employment. The motivation is that we would not expect to have enough power to detect spillovers when the election plant was only a small share of the firm's overall employment. We specifically split up elections based on whether the election plant was more than 10 % of the firm's employment in the same three-digit NAICS industry during the election year.<sup>52</sup> The estimates in Panel B. show that the overall increase in other plants' employment growth is driven by relatively large elections. It is reassuring for two reasons that the effects are driven by same-industry plants and by relatively large elections. First, these are the types of plants where we would expect to detect the most production shifting. Second, it is not clear why we would also expect potential threats to our parallel trend assumption to be more pronounced for these specific groups.<sup>53</sup>

These other plant employment growth estimates are both economically and statistically significant. For the same-industry plant estimates, the increase in DHS employment growth rate of around 0.04 is consistent with a two percentage point increase in survival probabilities. When we focus on the same-industry plants at high-employment share elections, the spillover effects are even larger between 0.07-0.09. As a benchmark, the direct three-year DHS employment growth rate effect of unionization for elections at MU manufacturing plants is -0.23. While our spillover estimates suggest that a sizeable share of the overall negative effects of unionization may be offset by employment shifting, there are several reasons that we cannot use

---

<sup>52</sup>This heterogeneity specification is estimated the same as other heterogeneity specifications (e.g., estimated jointly with pooled controls and controlling for the heterogeneity group by event time).

<sup>53</sup>For this analysis, the concern that would violate the parallel trends assumption is that the other plants at winning election firms are growing faster than other plants at losing election firms.

these estimates to calculate this share. First, our spillover estimates are average establishment-level employment changes, while we would need firm-level estimates to calculate the total share offset by reallocation.<sup>54</sup> Second, we focus on a specific subset of plants where we are most likely to detect spillovers. However, calculating the total share offset by reallocation requires the total firm-level employment changes (e.g., the estimates for all manufacturing plants where we do not find significant spillovers).

Overall, this evidence of successful union elections leading to faster employment growth at the firm's other plants is consistent with firms shifting production away from newly unionized plants. Furthermore, the higher survival probabilities for firms' other plants suggest that some of this production shifting occurs via decisions over which plants to close. Although we do not find significant long-run employment spillover estimates, this does not necessarily indicate a lack of long-run production shifting. First, given the increased variance of long-run employment growth rates, we may not have enough power to detect effects. Second, we may not be capturing all margins of production shifting that could occur over longer time horizons. For example, our analysis does not include shifting production by opening *new* plants or shifting production to plants in other countries (see e.g., Bluestone and Harrison (1982) and Bronfenbrenner (2000) for evidence of shifting production internationally following successful union elections).

## Firms' Unionization Status

The second part of our hypothesis is that the effects of unionization are greater when management is more opposed to the union. To test this, we estimate treatment effect heterogeneity based on two proxies for managers' opposition. First, we estimate effects separately for elections at MU firms with and without other unionized establishments. The motivation for this analysis is evidence that, during this time period, non-unionized firms (e.g., firms without any unionized establishments) were more opposed to unions than (partially) unionized firms. For example, Freedman (1979) and Kochan et al. (1986b) show that less unionized firms were more committed to remaining non-union and they provide accounts of managers at non-union firms "vigorously resist[ing] dealing with unions."<sup>55</sup> To test for heterogeneity by firms' unionization status, we

---

<sup>54</sup>We conduct an establishment-level analysis for two reasons. First, the longitudinal establishment linkages are higher quality than firm-level linkages (Haltiwanger et al., 2013). Second, we may have more power at the establishment level because we can include age, baseline employment, and time-varying industry controls that explain some of the employment growth variation.

<sup>55</sup>One reason unionized firms would respond less aggressively to new unionization attempts is that their other unionized workers could apply pressure on the entire firm to discourage aggressive

split up our elections at MUs based on whether we observe an FMCS contract at any of the firm’s establishments in the five years before the election.<sup>56</sup> Since the contract data start in 1984, we classify MUs as unionized versus non-unionized for elections starting in 1985 and show robustness to instead starting in 1990.

Figure 2-9 presents the DHS employment growth and survival estimates for elections at unionized versus non-unionized firms. The estimates are presented the same as the previous heterogeneity results (e.g., Figure 2-7). For overall employment growth rates, elections at non-unionized firms lead to larger employment decreases than elections at unionized firms. These differences are significant at the five- and ten-year horizons. For establishment survival, the differences are rather small and insignificant at the three- and five-year horizon. However, at the ten-year horizon, the negative survival effect is substantially larger for elections at non-unionized firms (e.g., -0.20 (SE 0.040) versus -0.09 (SE 0.027) for elections at non-unionized versus unionized firms, respectively).

Appendix Table B.6 shows that the larger effects at non-unionized firms are robust to alternative sets of controls and sample selection criteria. Specifically, it present estimates of the difference between effects at unionized versus non-unionized firms when (1) pooling controls across cohorts, (2) only classify firms into unionized versus non-unionized firms starting in 1990, and (3) only using 30-70 % vote-share elections to estimate the effects. The differences between estimates for non-unionized versus unionized firms are larger than our baseline specification when we define firms’ unionization status starting in the 1990s. For the other two specifications, the estimates are qualitatively the same as our baseline estimates.

These estimates show that the long-run negative effects of unionization are substantially larger at firms without any previously unionized establishments. This evidence is consistent with these firms being more opposed to and more rigorously resisting unionization. As one explanation for this opposition, Foulkes (1980) documents that some non-unionized firms were motivated by a philosophical opposition to unions even if they did not have previous bad experiences with unions. Alternatively, similar to Selten (1978)’s “chain store paradox”, non-unionized firms may have a strong incentive to aggressively deter the first unionization attempt to prevent unionization from spreading across the firm (even if not economically profitable when considering each establishment in isolation).<sup>57</sup> Both cases suggest that the larger effects at non-

---

responses. An anecdotal example is the failure of GM’s “southern strategy” of opening non-unionized plants in the South due to pressure from the UAW (Nelson, 1996).

<sup>56</sup>We include previous contracts at both the election establishment (i.e., a separate bargaining unit already unionized) and at all of the firm’s other establishments. See Appendix B.3 for details.

<sup>57</sup>As an extreme example, consider Walmart’s switch to pre-packaged meat across *all* stores days

unionized firms may be quite excessive relative to the direct costs of unions at these firms.

## Election Delay Time

Our second proxy for managers' opposition to the union is delay during the election process. The motivation is that managers frequently use tactics that delay the election date to try to win the election. First, delay itself can reduce support for the union. In "Confessions of a Union Buster," Levitt and Conrow (1993) write that the NLRA "presents endless possibilities for delays, roadblocks, and maneuvers that can undermine a union's efforts and frustrate would-be members" and that this delay "steals momentum from a union-organizing drive, which is greatly dependent on [...] the sense of urgency among workers." Additionally, other tactics managers employ to influence elections also delay the election (e.g., challenging the composition of the bargaining unit or filing unfair labor practice charges). Furthermore, research has found that delay is associated with lower election win rates which supports delay time being a proxy for the intensity of managers' anti-union campaigns (Roomkin and Block, 1981; Ferguson, 2008).

We start by defining election delay time and verifying that it is related to election outcomes in our sample. We define delay time as the number of days between the date the election petition was filed to the NLRB and the date the election was held (see Appendix B.3 for details). The average election delay in our sample is 62 days, and the 10th and 90th percentiles are 31 and 80 days. Appendix Figure B-5 shows that our measure of delay time is negatively associated with union election success rates and positively associated with the probability of any challenged votes in the election (another proxy for managers' anti-union campaign intensity). These relationships also hold conditioning on other election characteristics that may be correlated with delay time.

To analyze whether the negative effects of unionization differ by election delay, we start by estimating treatment effect heterogeneity separately by terciles of the within-year delay time distribution. Figure 2-10 plots the estimated effects for the first and third terciles for DHS employment growth (left panels) and establishment survival (right panels). Panel A. includes results for all elections and Panel B. just includes elections at MU firms. The p-values below the labels are from testing whether the effects for the first and third terciles are equal. Across both figures, the effects

---

after ten Walmart meat cutters at one Texas store voted to unionize in 2000 (Zimmerman, 2000).

of unionization on establishment employment and survival are larger for elections in the top tercile of the delay time distribution. Focusing on elections at MUs, the first versus third tercile estimates are significantly different for both outcomes at the three- and ten-year horizon. For example, the estimated ten-year effect on survival for the top tercile is -0.20 (SE 0.037) versus -0.071 (SE 0.036) for the bottom tercile.

We next assess the robustness of these results to instead using a continuous measure of delay time. While the within-year terciles are appealing because they only rely on within-year variation in delay time and allow for a flexible, functional form, we might have more power using the entire distribution of delay times. To implement this, we add an interaction between the event-time treatment indicators times the *log* election delay time to the specification in equation 2.8.<sup>58</sup> Table 2.6 presents the coefficient estimates on the log election delay interaction for three, five, and ten years post-election. The first two columns show that the negative effects of unionization are significantly larger for elections with longer delays across all time horizons. For the ten-year survival effect, an approximately 10 % increase in election delay is associated with a .7 pct. pt. increase in the probability of a plant closing.<sup>59</sup> Columns (3) and (4) show that the effects are robust to including the controls pooled across cohorts. Columns (5) and (6) address the concern that our election delay time measure is just capturing larger bargaining units. Specifically, these columns show that our estimates are qualitatively the same when we first residualize the log delay time on bargaining unit size deciles, although the ten-year estimates are only significant at the ten-percent level.

Our primary interpretation of these results is that the negative effects of unionization are largest at establishments where the employer initially campaigned harder against the union. This is supported by anecdotal accounts linking election delays to the intensity of firms' anti-union campaigns (Levitt and Conrow, 1993). Another interpretation is that delay may be a proxy for hostile labor relations conditions. For example, more adversarial unions and management might have more disagreement before the election that could delay the process. Overall, this heterogeneity adds to our results showing that managers' opposition to unions plays a role in the overall negative effects of successful union elections.

---

<sup>58</sup>We also control for log delay time interacted with event-time directly in the specification.

<sup>59</sup>Reassuringly, the magnitudes of the continuous-specification estimates are similar to the tercile specification estimates. The implied survival difference from the continuous specification between the 10th and 90th percentiles of the delay time difference is  $[\ln(80) - \ln(30)] \times -0.07 = -0.066$ . The ten-year survival difference between the first and third terciles is  $-0.089$ .

## Unionization and Productivity Reallocation

Finally, we examine how the negative effects of unionization vary by establishment productivity. The motivation for this analysis is that theoretical and empirical research in other contexts predicts that wage increases or productivity declines should have larger impacts at lower productivity establishments.<sup>60</sup> Since these channels are two leading “economic” reasons why unionization might cause decreased employment or exit, this research suggests that unionization may also have a larger impact on lower productivity plants. Consequently, substantial heterogeneity by establishment productivity may be more consistent with the survival effects being driven by direct effects on wages or productivity than our union avoidance hypothesis.

To measure establishment-specific TFP for our manufacturing elections, we use cost-share-based productivity measures from the Annual Survey of Manufacturers (ASM) and Census of Manufacturers (CM) calculated by Foster et al. (2016). We use within-industry TFP comparisons to address potential measurement or productivity differences across industries. Specifically, we classify each establishment into three productivity terciles based on their pre-election, within year and six-digit NAICS industry TFP ranking (see Appendix B.3 for details). Figure 2-11 plots the estimated effects for the first and third terciles of the baseline TFP distribution. We find evidence that the three- and five-year employment and survival effects are larger for lower-productivity establishments. But these differences are never significant and, at the five-year and ten-year horizon, are not economically very large (e.g., -0.066 (SE 0.023) versus -0.041 (SE 0.022) at the five-year horizon for the first and third terciles, respectively). Appendix Figure B-6 shows that these patterns hold when we separately estimate heterogeneity by baseline TFP only for MU firms. Overall, we do not interpret this evidence as supporting economically larger survival effects for less productive establishments. Thus, the evidence is more consistent with alternative explanations for why unionization leads to plant closures (e.g., our union avoidance hypothesis) than conventional explanations.

---

<sup>60</sup>For the effect of wage increases, Berger et al. (2021) show that minimum wage increases can cause relatively larger employment declines at less productive firms in oligopsonistic labor markets with firm heterogeneity. Luca and Luca (2019); Dustmann et al. (2020) provide empirical evidence of minimum wages increasing the exit rates of low-productivity and smaller firms. We note, however, that the settings are different as minimum wages are market-wide wage increases while unionization is an establishment-specific wage increase. We are not, however, aware of any models of unionization with firm heterogeneity and imperfectly competitive labor markets. For the effect of productivity declines, see (Foster et al., 2008) for evidence that plants’ level of revenue-based productivity is a key determinant of exit. One additional caveat to extending both predictions to the effects of unionization is that unions could base their bargaining demands on establishment productivity (e.g., try to extract more from higher productivity establishments).

## 2.7 Relation to Literature and Implications

In this section, we relate our estimates to the literature on unionization’s employment and survival effects. We also discuss the implications of our union avoidance results for interpreting these effects.

**Unionization, Employment, and Survival Literature** Our estimates of overall employment and survival decreases following successful union elections add to other recent research finding similar effects following elections. The most comparable results to ours are Frandsen (2021)’s regression discontinuity estimates also using the LBD. We qualitatively match his short-run employment and long-run survival declines but find somewhat smaller effects (e.g., five-year survival effects of 4 pct. pts. versus 8-10 pct. pts.). Some explanations for this are different samples or different empirical strategies and identifying assumptions.<sup>61</sup> Additionally, our smaller employment effects and insignificant survival effects for service-sector elections match Sojourner et al. (2015)’s estimates of negative employment but no survival effects of nursing home elections. However, even for relatively close elections, our estimates are inconsistent with the null effects that DiNardo and Lee (2004) find for establishment survival and employment. One potential explanation is that the LBD longitudinal linkages we use to define survival are higher quality than linkages in the telephone-book-based InfoUSA or the LRD data used in DiNardo and Lee (2004).<sup>62</sup>

**Implications of Firm Union Avoidance and Opposition** We also present evidence consistent with firms’ union avoidance tactics and their opposition to unions playing a role in these overall negative effects of unionization in manufacturing. First, we find evidence of firms avoiding new unions by shifting production away from newly unionized plants. Second, we find larger effects of successful union elections at firms more opposed to unions (based on whether the firm has other unionized establishments and on measures of delay during the election process).

One interpretation of these results is that firms’ opposition to unions and attempts to avoid unions are driven by managers’ dislike of working with unions and unrelated

---

<sup>61</sup>For example, Frandsen (2021) requires at least 20 votes cast to implement the regression discontinuity analysis while we only require more than five eligible voters. This could explain the differences in the magnitudes of our estimates since the average election in our sample may have a lower bargaining unit to overall establishment employment share. Furthermore, we find larger treatment effects for elections where the bargaining unit is a larger share of overall employment.

<sup>62</sup>See, Jarmin and Miranda (2002) and Crane and Decker (2019) for comparisons of the linkages across these datasets.

to the direct costs of unions. This interpretation is consistent with accounts suggesting that our measures of manager opposition were not based on expectations of the costs of unions. For example, Freedman (1979) finds that non-unionized firms placed the most weight on resisting new unions but that these firms were also where unions were least able to attain higher wages. For example, the threat of strikes was more limited at less unionized firms because these firms could still produce at non-union plants during a strike. Similarly, Foulkes (1980) documents that some non-unionized firms were motivated by a philosophical opposition to unions even if they did not have previous bad experiences with unions. Finally, Bronfenbrenner (2001) finds that the intensity of firms' anti-union campaigns was "unrelated to the financial condition of the employer, but rather were a function of the extreme atmosphere of anti-union animus." This interpretation of our results also helps resolve the puzzle that it has been difficult to find evidence of unionization raising wages or negatively affecting productivity (see the citations in footnote 2), but we find large effects on establishment survival. Additionally, it is puzzling that unions would make demands that push firms out of business, directly harming their workers (Friedman, 1951). The interpretation of our results that managers' idiosyncratic opposition to unions drives the survival effects of unionization provides a resolution to these puzzles.

An intermediate interpretation of our results is that while there may be some direct costs unionization, the survival and employment effects we estimate are excessive relative to these costs. For example, if firms can easily shift production away from unionized establishments, even small wage or productivity effects could lead to large survival effects. Furthermore, the larger effects we find at non-unionized firms could be due to efforts to prevent unionization from spreading to other establishments by establishing a reputation for vigorously resisting unions. One channel that could magnify such effects is evidence that successful unionization negatively affects managers' careers. For example, Clark (1980) finds increased manager turnover after successful elections and Dunlop (1994) documents some managers' expectations that unionization would hurt their career prospects.

On the other hand, we cannot rule out that our proxies for manager opposition simply reflect rational expectations of where unions would have been the costliest. This interpretation is supported by evidence suggesting direct costs of unions (e.g., stock price declines following successful union elections and a large body of literature on unions reducing firm profits (Freeman and Medoff, 1984; Lee and Mas, 2012)). However, even parts of Lee and Mas (2012)'s stock price results are difficult to reconcile with the interpretation that our results are driven by direct costs of unions.



In particular, as noted in Frandsen (2021), Lee and Mas (2012) only find equity value declines away from the 50 % threshold while we and Frandsen (2021) also find survival and employment declines for close elections. A way to reconcile these results is that the survival effects for these very close elections may not be driven by direct costs of unions that would also cause stock-market declines.

Overall, since we do not provide estimates of the magnitude of the direct costs of unions, we cannot rule out any of these interpretations. Yet, our evidence that the largest negative effects of unionization are where firms are the most opposed to unions and can avoid working with new unions suggests that the overall negative effects may not necessarily imply large direct costs.

## 2.8 Conclusion

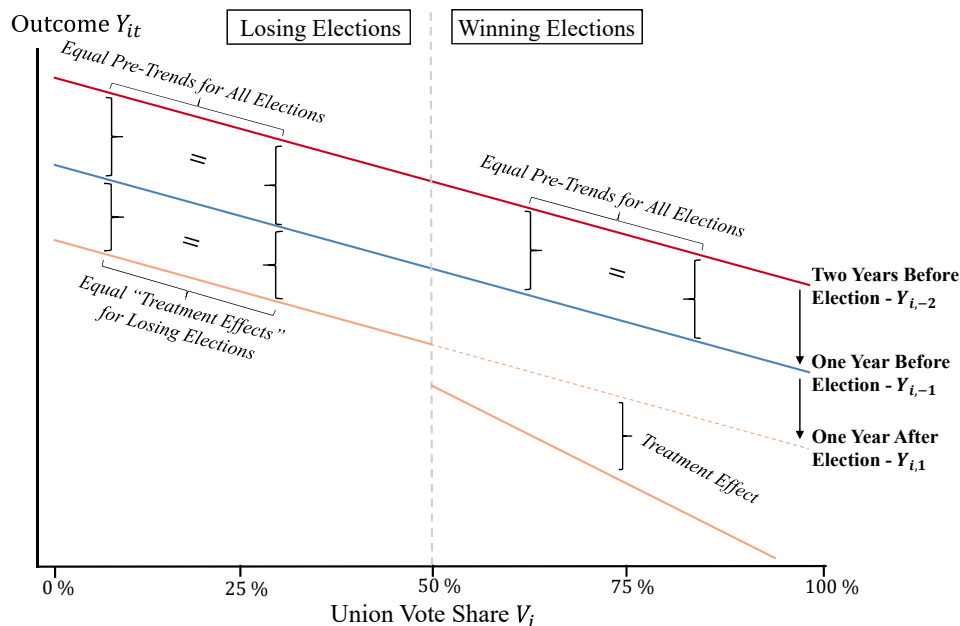
This paper revisits the effects of successful NLRB union elections on establishments' employment and survival. We first show that winning union recognition decreases establishments' employment and long-run survival. While one interpretation of these results is that unions must have large direct costs on businesses, we raise alternative explanations for these effects and explore whether our results for manufacturing elections are consistent with these explanations. First, we hypothesize that the effects of unionization could be magnified by firms' ability to avoid working with the union. We support this hypothesis by showing that the largest effects are at multi-establishment firms and by providing direct evidence of increased employment at firms' *other* establishments following successful elections. Both results are consistent with firms shifting production away from newly unionized establishments, one version of avoiding working with unions. Second, the overall negative effects may be partially driven by managers' dislike of working with unions even if unions do not have large direct costs. Supporting this, we find the largest effects at non-unionized firms and at elections with the longest delay during the election process, both proxies for a firms' opposition to unions. Overall, these results are consistent with firms' union avoidance tactics and managers' opposition to unions playing a role in explaining the overall negative effects of unionization.

Finally, our results raise many questions about the impact of unionization that suggest opportunities for further research. First, our results highlight how firms that strongly oppose unions go to great lengths to avoid working with unions. Another strategy they may pursue is raising workers' wages to discourage unionization. These "union threat effects" may be one of the main channels through which unions have

affected the U.S. wage structure (Taschereau-Dumouchel, 2020). Further research quantifying this channel would be helpful to understand unions' total effect on wages (Neumark and Wachter, 1995; Farber, 2005). Second, we present evidence of large differences in how firms in different sectors respond to unionization. These findings mirror research that finds small employment effects of minimum wages in the service sector but larger negative effects in manufacturing (Cengiz et al., 2019; Harasztsosi and Lindner, 2018). Together, these results suggest that understanding how tradable industries respond to wage increases would be useful for forecasting the effects of mandating higher wages in such industries. Specifically, the production shifting channel we explore in this paper may also explain why other labor market policies have larger effects in tradable industries. Third, our results raise the question of whether the managers' opposition to unions is driven by the incentives created under the U.S.'s collective bargaining institutions. In particular, since unionization occurs at the establishment level, a single establishment may be the only one in its labor market or in its firm that needs to deal with a union. This setup could exacerbate managers' incentives to resist unions (Estlund, 1993). Consequently, further research into whether the adverse consequences of unionization in the U.S. reflect the adversarial setup of its institutions would help inform future policies to foster more collaborative labor relations.

## 2.9 Figures

Figure 2-1: Testable Implications of Parallel Trends Identifying Assumption




*Note:* This figure illustrates our empirical strategy’s identifying assumption and its testable implications discussed in section 2.4. It plots hypothetical average establishment-level outcomes before and after union elections with different vote shares.  $Y_{i,-2}$  and  $Y_{i,-1}$  correspond to outcomes one and two years before the union election.  $Y_{i,1}$  corresponds to outcomes one year after the election. Testing parallel *pre-trends* by vote share corresponds to comparing the distance between  $Y_{i,-2}$  and  $Y_{i,-1}$ . Testing parallel *post-trends for losing elections* corresponds to comparing the distance between  $Y_{i,-2}$  and  $Y_{i,1}$  for losing elections.

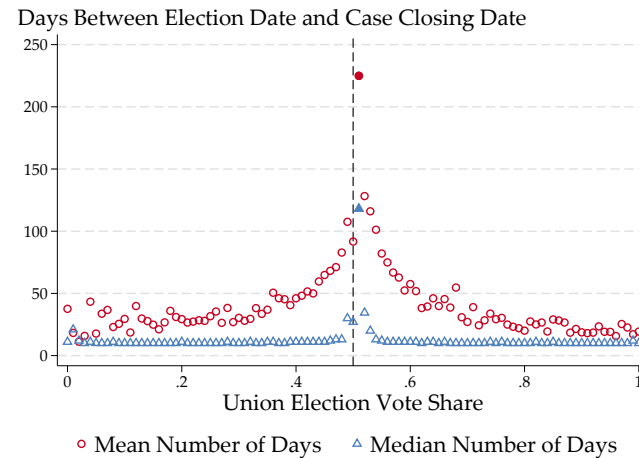
*Figure 2-2 Note:* Figure 2-2 presents four panels illustrating characteristics of close union elections. All panels are constructed using external union election data (e.g., not our final sample matched to the Census) but the sample was constructed to mirror the overall sample construction (see Appendix B.3 for details). Panel A. plots the vote-share histogram of elections included in our sample. Given the discreteness of the running variable and the fact that our sample includes elections with a small number of votes, it is difficult to detect manipulation from the vote-share density figure. Consequently, we plot elections with exactly 50 % of votes separately to make the manipulation easier to see. See Frandsen (2017) for evidence of manipulation using formal tests that accommodate discrete running variables. Panel B. plots the average and median number of days between the union election date and the date that the case closed. Panel C. plots the probability of each union election experiences a decertification election in the five years following the case closing. The decertification elections are also from our combined NRLB datasets but excluded from our main analysis. Panel D. plots the probability of each union election experiencing a works stoppage in the five years following the case closing. The works stoppage data is from the FMCS and covers works stoppages from 1984-2019. Consequently, we only plot follow-up works stoppages for elections from 1984-2005. For the decertification and works stoppage figures, we match based on exact company names and cities rather than the SoftTFIDF algorithm we use for the main analysis. The “conditional regression coefficients” are the coefficients from regressing the stoppage indicators on the vote share for winning elections including controls for deciles of the number of workers in the bargaining unit, the four-digit NAICS industry, and election state.

Figure 2-2: Characteristics of Close Elections that Motivate Including Larger Margin-of-Support Elections


Panel A. Election Vote-Share Histogram



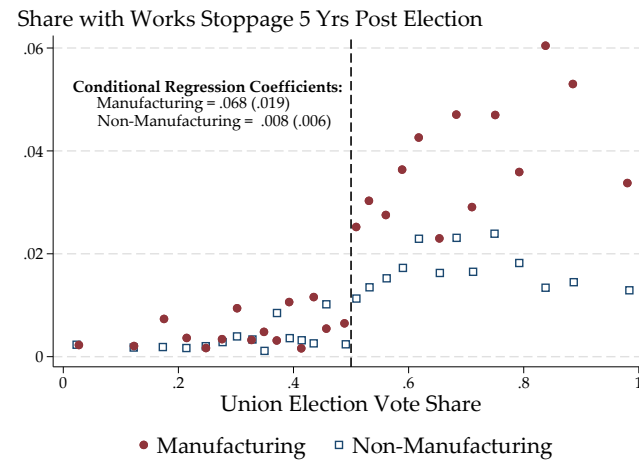
Panel B. Number of Days Between Election and Case Closing Dates



Panel C. Probability *Decertification Election* Five Years Following Election



Panel D. Probability of *Works Stoppage* Five Years Following Election



Note: See the previous page.

## 2.10 Tables

Table 2.1: Winning versus Losing Election Establishment Summary Statistics

	All Industries		Manufacturing	
	Union Loses	Union Wins	Union Loses	Union Wins
<b>Establishment Characteristics</b>				
Employees	154	137	167	148
Payroll/Worker (\$ 2019)	49,400	49,700	49,900	48,700
Establishment Age	9.65	10.0	9.82	9.58
Multi-Establishment Firm	0.512	0.476	0.525	0.464
Previous Contract at Establishment	0.090	0.147	0.088	0.152
<b>Survival Base Rates</b>				
5-Year Survival	0.818	0.765	0.847	0.779
10-Year Survival	0.667	0.610	0.702	0.608
<b>Approximate Number of Elections</b>	27,000		7,000	

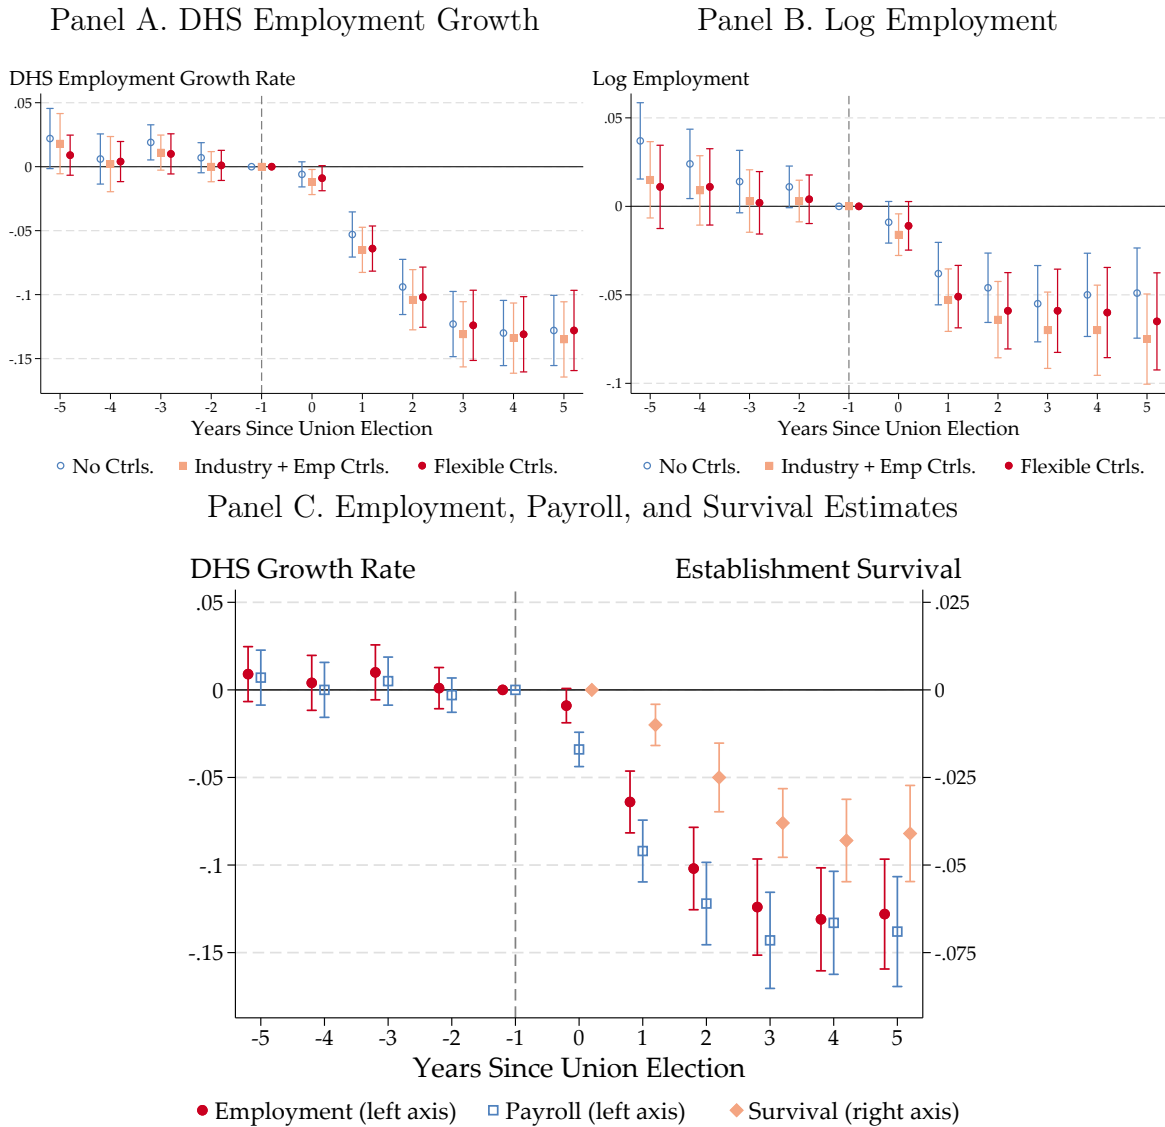
*Note:* This table presents summary statistics for all union elections included in our analysis sample with vote shares between 0-100 %. All establishment characteristics are measured one year before the union election. Since the FMCS contract data are only available starting in 1984, we only calculate the share of establishments with a previous contract using elections from 1985 onward. The five- and ten-year survival rates are the probability of surviving five and ten years after the union election, respectively. To satisfy the Census' disclosure requirements, all estimates are rounded to only include three significant digits, and sample sizes are round to the nearest 1,000.

Table 2.2: Pre-Election Employment Growth Trends by Vote Share, 20-80 % Elections

Outcome:	DHS Employment Growth Rate			
Industry Group:	All Industries		Manufacturing	
5-Year Pre Election $\times$ Vote Share	0.050 (0.037)	0.033 (0.025)	0.029 (0.069)	-0.018 (0.047)
4-Year Pre Election $\times$ Vote Share	0.018 (0.032)	0.019 (0.024)	0.026 (0.059)	0.008 (0.045)
3-Year Pre Election $\times$ Vote Share	0.028 (0.023)	0.029 (0.023)	0.022 (0.041)	0.018 (0.042)
2-Year Pre Election $\times$ Vote Share	0.006 (0.018)	0.012 (0.019)	-0.026 (0.035)	-0.018 (0.037)
Industry + Employment Ctrl.	X	X	X	X
Flexible Ctrl.		X		X
Number of Elections	19,000	19,000	6,000	6,000

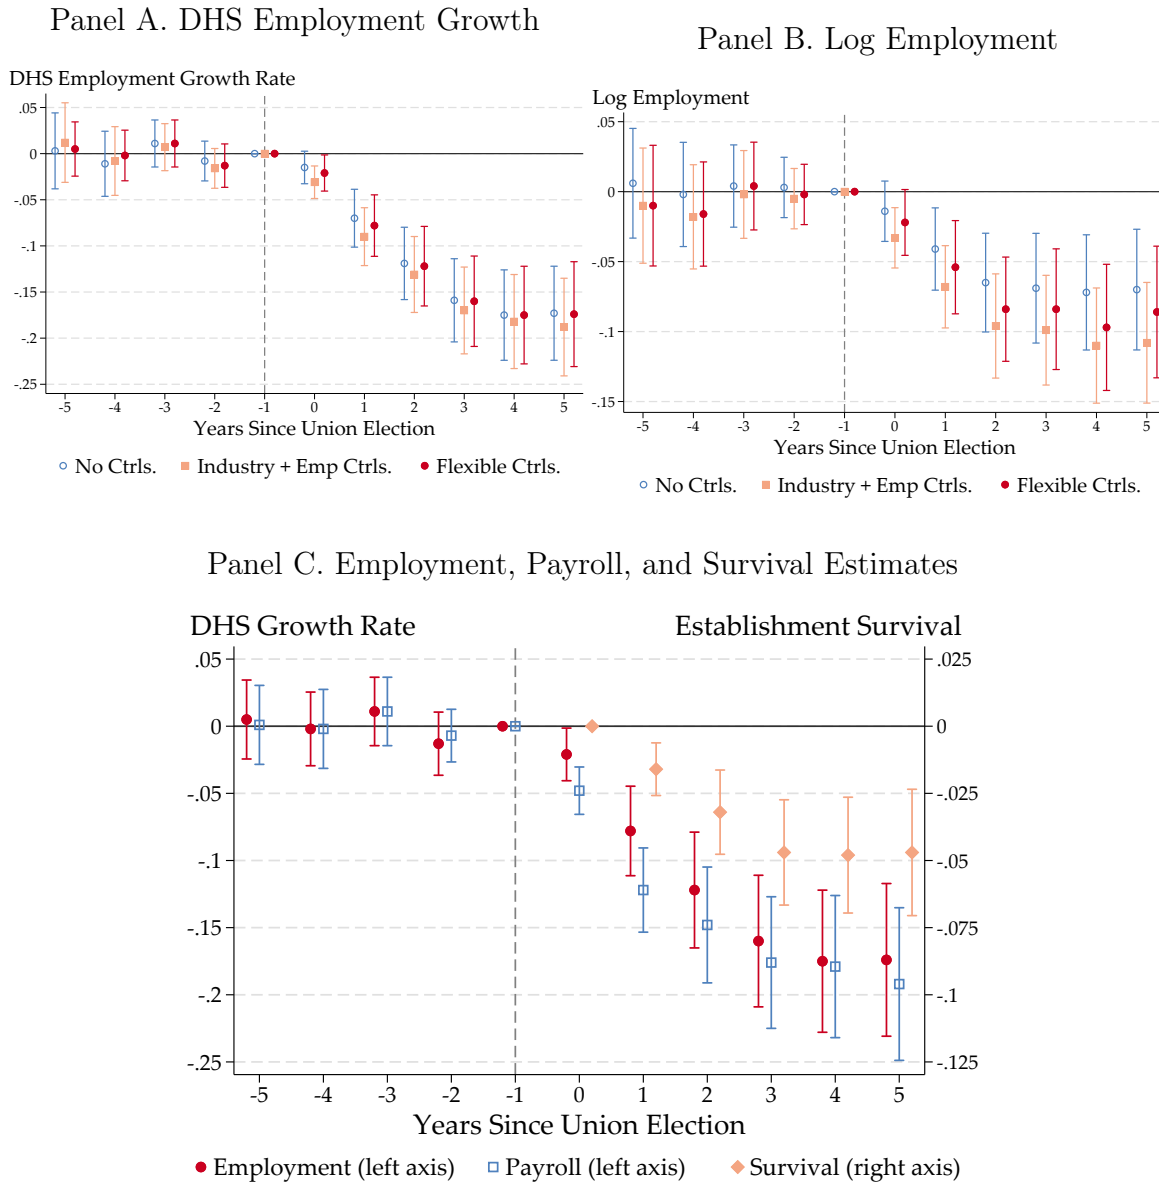
*Note:* This table presents estimates testing for linear trends by vote share in pre-election employment growth rates. Significant estimates would violate a testable implication of our parallel trends by vote share assumption (see equation 2.5). Specifically, the table reports the estimated coefficients on interactions between event-time indicators and the continuous election vote-share (i.e., the  $\rho$  coefficients from equation 2.11). A five-year coefficient of 0.03 implies that elections with 75 % of votes grew approximately 1.5 percent slower during the five years before the election than an election with 25 % of votes. The outcome for all specifications is establishment-level DHS employment growth relative to time  $-1$ . The sample includes 20-80 % vote-share elections. The first two columns include elections in all industries and the last two columns include just manufacturing elections. The odd columns include only industry and employment controls and the even columns include our flexible control specification (see Section 2.4 for details). Standard errors are clustered by establishments' firmid during the year of the election (e.g., the clustering variable is fixed over time for each establishment). \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Figure 2-3: Employment and Survival Estimates, 20-80 % Vote-Share Elections, All Industries



*Note:* This figure plots the  $\delta_n$  coefficients (i.e., the interaction between winning a union election and being  $n$  years from the election) from estimating specification 2.8 for all union elections with 20-80 % vote shares inclusive. The sample includes observations -10 to 10 years before and after each union election but we only plot the -5 to 5 coefficients. The outcome variable for Panel A. is establishment-level DHS employment growth relative to time  $-1$ . The outcome variable for Panel B. is establishment-level log employment. The outcome variables for Panel C. are DHS employment and payroll growth rates and an indicator for whether the establishment exists at time  $t$ . For Panel C., the survival y-axis is scaled to be one-half the DHS growth rate axis. Consequently, comparing the exit and DHS coefficients illustrates how much of the effect on the DHS growth rate can be mechanically explained by the exit effect. Panels A. and B. include estimates with no controls, just industry and employment controls, and the flexible control specification (see Section 2.4 for details). Panel C. includes estimates from the flexible control specification. The log outcome estimates in Panel B. include establishment fixed effects but these are not included in Panel A. or Panel C. Standard errors are clustered by establishments' firmid during the year of the election (e.g., the clustering variable is fixed over time for each establishment).

Figure 2-4: Employment and Survival Estimates, 20-80 % Vote-Share Elections, Manufacturing




*Note:* These estimates are identical to Figure 2-3 except that they are only estimated for manufacturing elections.




Figure 2-5: Nonparametric Vote-Share Heterogeneity Estimates, Manufacturing

Panel B. DHS Employment Growth Rate




Panel B. Establishment Survival




*Note:* This figure plots the  $\delta_{g,n}$  coefficients from estimating the vote-share heterogeneity specification 2.10 with the vote-share distribution partitioned into eight groups indicated on the x-axis. We omit the 20-30 % election group so the other estimates are relative to that group. The sample includes all manufacturing elections. We include observations -10 to 10 years before and after each union election but we only plot a subset of coefficients. The outcome variable for Panel A. is establishment-level DHS employment growth relative to event time  $-1$ . The outcome variable for Panel B. is an indicator for establishment survival. The estimates include the flexible control specification (see Section 2.4 for details). Standard errors are clustered by establishments' firmid during the year of the election (e.g., the clustering variable is fixed over time for each establishment).

Figure 2-6: Nonparametric Vote-Share Heterogeneity Estimates, All Industries

Panel A. DHS Employment Growth Rate

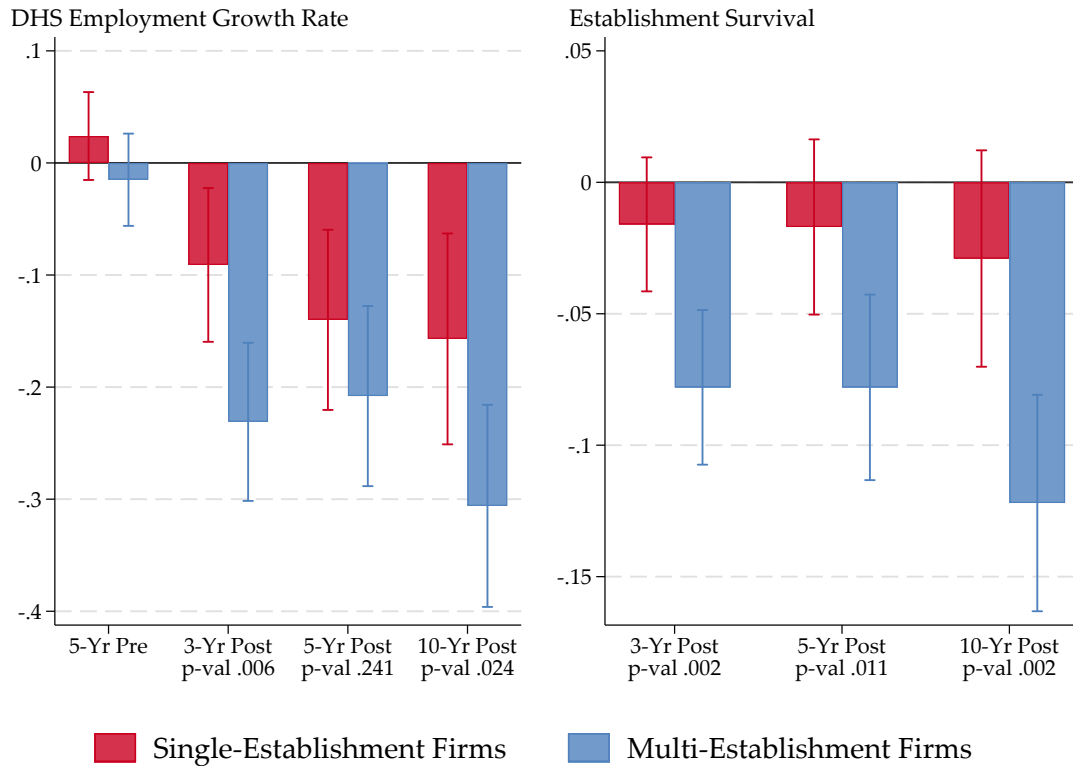


Panel A. Establishment Survival



*Note:* This figure is identical to Figure 2-5 except it includes elections across all industries. The alternative estimates listed in the text box in Panel A. are the 40-50 % estimates excluding elections with exactly 50 % of votes (rather than restrict the sample, we include a separate category for 50 % vote elections).


Figure 2-7: Single- Versus Multi-Establishment Firm Heterogeneity




*Note:* This figure plots the  $\delta_{h,n}$  coefficients from estimating our heterogeneity specification in equation 2.12 for elections at single- versus multi-establishment firms. An election at a multi-establishment firm is defined based on whether the establishment's firm has any other establishments one year before the election. The sample includes all manufacturing union elections with 20-80 % vote shares inclusive. It includes observations -10 to 10 years before and after each union election but we only plot a subset of these coefficients. The outcome variable for the left panel is DHS employment growth rates relative to time  $-1$  (see Section 2.4 for their definition). The outcome variable for the right panel is an indicator for establishment survival. The estimates include the flexible control specification (see Section 2.4 for details). The control coefficients are pooled across the heterogeneity groups. See Appendix Table B.5 for robustness to alternative controls specifications. Standard errors are clustered by establishments' firmid during the year of the election (e.g., the clustering variable is fixed over time for each establishment).

Figure 2-8: Employment Effects of Successful Elections on Firms' Other Establishments

Panel A. All Plants and Three-Digit NAICS Plants

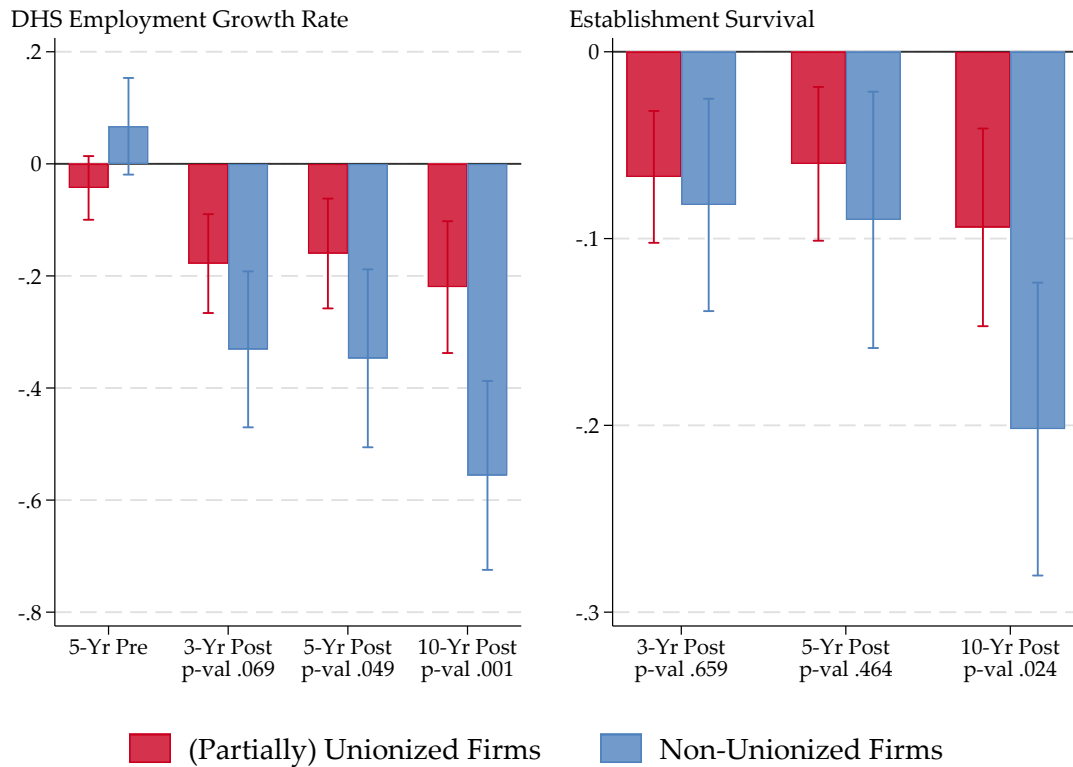


Panel B. Estimates by Election's Employment Share



*Note:* This figure plots the  $\delta_n$  coefficients from estimating specification 2.8. The sample is manufacturing plants at multi-establishment firms where *another* plant experienced a union election. See Section 2.6 for details about the sample construction. The relative time and vote-share variables are defined from the election at the firm's *other* establishment. We weight the regression by the observation's share of total firm-level employment across all plants included in the sample the year of the election. The outcomes in both panels are establishment-level DHS employment growth rates relative to one year before the union election. The estimates include the flexible control specification (see Section 2.4 for details) except we do not include a control for establishments SU/MU status (all plants are part of MUs) or for establishments' previous contract status. Since we match plants based on the election year, the industry is also from the year of election. The "All Manufacturing Estabs" estimates in the left panel include all manufacturing establishments with at least two employees during the year of the election. The "Within-NAICS 3 Estabs" estimates restrict the sample to plants that are in the same 3-digit NAICS industry as the election plant. The right panel includes 3-digit NAICS industry matches but separately estimates the effects by whether or not the election establishment comprised more than 10 % of the firm's employment in the same three-digit NAICS industry during the year of election. The estimates in Panel B. are from the same specification with the controls pooled across both groups and the treatment indicators interacted with the two employment share groups. In this panel, we also directly control for the effect of the two employment share groups interacted with event time.


Figure 2-9: Unionized versus Non-Unionized Firm Heterogeneity




*Note:* This figure plots similar heterogeneity results as Figure 2-7 except that the heterogeneity is for elections at multi-establishment firms with at least one unionized establishment versus firms without any unionized establishments. See Appendix B.3 for how we define firms' unionization status. The controls additionally directly include these heterogeneity groups interacted with cohort and event time.

Figure 2-10: Election Delay Heterogeneity

Panel A. All Elections




Panel B. Elections at Multi-Establishment Firms



*Note:* These figures plots the  $\delta_{h,n}$  coefficients from estimating our heterogeneity specification in equation 2.12 for elections in different terciles of the *election delay* distribution. These terciles are defined within each year based on the number of days between the election petition filing date and the election date (see Section B.3 for details). We plot the coefficients for the first and third terciles but estimate the effects for all three. The sample includes all manufacturing union elections with 20-80 % vote shares inclusive. It includes observations -10 to 10 years before and after each union election but we only plot a subset of these coefficients. The outcome variable for the left panel is DHS employment growth rates relative to time  $-1$  (see Section 2.4 for their definition). The outcome variable for the right panel is an indicator for establishment survival. The estimates include the flexible control specification (see Section 2.4 for details). Standard errors are clustered by establishments' firmid during the year of the election (e.g., the clustering variable is fixed over time for each establishment). Panel A. defines the election delay terciles across all elections. For Panel B. the election delay terciles are only defined for elections at multi-establishment manufacturing firms. Consequently, we estimate but do not report separate coefficients for elections at single-establishment firms.

Figure 2-11: Establishment-Level Total Factor Productivity Heterogeneity



*Note:* This figure plots the  $\delta_{h,n}$  coefficients from estimating our heterogeneity specification in equation 2.12 for elections in different terciles of baseline TFP distribution. These terciles are defined based on plants' pre-election cost-share-based productivity measures from the Annual Survey of Manufacturers (ASM) calculated by Foster et al. (2016). The TFP terciles are defined based on within-year and within six-digit NAICS productivity rankings. See Appendix B.3 for details. We plot the coefficients for the first and third terciles but estimate effects for all three terciles and a fourth group of plants without TFP defined. The sample includes all manufacturing union elections with 20-80 % vote shares inclusive. It includes observations -10 to 10 years before and after each union election but we only plot a subset of these coefficients. The outcome variable for the left panel is DHS employment growth rates relative to time  $-1$  (see Section 2.4 for their definition). The outcome variable for the right panel is an indicator for establishment survival. The estimates include the flexible control specification (see Section 2.4 for details). The controls additionally include these heterogeneity groups interacted with cohort and event time. Standard errors are clustered by establishments' firmid during the year of the election (e.g., the clustering variable is fixed over time for each establishment).

Table 2.3: Post-Election Outcome Trends by Vote Share, 20-80 % Vote-Share Elections

Outcome:	DHS Emp Growth Rate		Establishment Survival	
<b>Panel A: All Industries</b>				
<i>Event-Time</i> × 0-50 % Vote Share				
3-Year Post Election	-0.216** (0.095)	-0.085 (0.103)	0.013 (0.036)	0.036 (0.040)
5-Year Post Election	-0.220** (0.110)	-0.066 (0.122)	-0.031 (0.045)	0.004 (0.049)
10-Year Post Election	-0.332*** (0.125)	-0.193 (0.140)	-0.080 (0.054)	-0.038 (0.060)
<i>Event-Time</i> × 50-100 % Vote Share				
3-Year Post Election	-0.280** (0.131)	-0.286** (0.131)	-0.028 (0.052)	-0.029 (0.052)
5-Year Post Election	-0.381** (0.149)	-0.389*** (0.149)	-0.052 (0.063)	-0.053 (0.063)
10-Year Post Election	-0.271* (0.164)	-0.278* (0.164)	-0.071 (0.073)	-0.073 (0.073)
<b>Panel B: Manufacturing</b>				
<i>Event-Time</i> × 0-50 % Vote Share				
3-Year Post Election	-0.236 (0.159)	-0.145 (0.170)	0.017 (0.061)	0.035 (0.065)
5-Year Post Election	-0.216 (0.187)	-0.072 (0.199)	-0.023 (0.076)	0.025 (0.081)
10-Year Post Election	-0.425* (0.226)	-0.210 (0.241)	-0.151 (0.097)	-0.049 (0.104)
<i>Event-Time</i> × 50-100 % Vote Share				
3-Year Post Election	-0.462* (0.266)	-0.470* (0.266)	-0.009 (0.104)	-0.010 (0.104)
5-Year Post Election	-0.394 (0.299)	-0.406 (0.299)	0.008 (0.126)	0.004 (0.126)
10-Year Post Election	-0.559* (0.336)	-0.578* (0.336)	-0.162 (0.150)	-0.171 (0.150)
Exclude 50 % Elections		X		X
Industry + Employment Ctrls.	X	X	X	X
Flexible Ctrls.	X	X	X	X

*Note:* This table presents estimates testing for linear trends by vote share in post-election outcomes. We test for trends separately across winning versus losing elections. The *Event-Time* × 0-50 rows present estimates of the  $\theta$  coefficients from equation 2.11 and capture linear trends in post-election outcomes for losing elections. The *Event-Time* × 50-100 rows present estimates of  $\theta + \tau$  and capture linear trends in post-election outcomes for winning elections. Since the specification separately includes an interaction with a winning election indicator, these slope estimates are in excess of any treatment effect right around the 50 % threshold. The outcome for the first two columns is establishment-level DHS employment growth relative to time  $-1$ . The outcome for the last two columns is an indicator of whether the establishment exists at time  $t$ . All specifications include our flexible control specification (see Section 2.4 for details). See Appendix Table B.2 for the same results with alternative included controls. The columns that "Exclude 50 % Elections" include an interaction between having a vote share of exactly 50 % and event time. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .



Table 2.4: Employment and Survival Estimates by Industry, 20-80 % Vote-Share Elections

Industry Group:	Manufacturing		Services		Other	
	DHS Emp	Survival	DHS Emp	Survival	DHS Emp	Survival
5-Year Pre Election	0.005 (0.015)		0.010 (0.012)		0.011 (0.016)	
2-Year Pre Election	-0.013 (0.012)		0.017* (0.009)		-0.009 (0.012)	
5-Year Post Election	-0.174*** (0.029)	-0.047*** (0.012)	-0.057** (0.024)	-0.026*** (0.010)	-0.192*** (0.030)	-0.058*** (0.013)
10-Year Post Election	-0.231*** (0.033)	-0.075*** (0.015)	-0.059** (0.027)	-0.017 (0.012)	-0.229*** (0.033)	-0.083*** (0.015)
Industry + Employment Ctrls.	X	X	X	X	X	X
Flexible Ctrls.	X	X	X	X	X	X
Industry Group Number of Elections	6,000	6,000	8,000	8,000	5,000	5,000
Industry Group Share of Elections	0.302	0.302	0.414	0.414	0.284	0.284

*Note:* This figure plots the  $\delta_{n,n}$  coefficients from estimating our heterogeneity specification in equation 2.12 for elections in three different broad industry groups. Manufacturing is defined as NAICS sectors 31-33, services are defined as NAICS 51-81 and retail trade (NAICS 44-45), and other is the remaining industries. Elections are classified into industries based on their Fort and Klimek (2016) NAICS 2012 codes. Otherwise, the sample, controls, and standard errors are the same as in Figure 2-3. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 2.5: Effects of Successful Elections on Firms' Other Establishments

Outcome:	DHS Employment	Survival
1-Year Post Election	0.017 (0.016)	0.006 (0.006)
2-Year Post Election	0.043** (0.019)	0.012 (0.008)
3-Year Post Election	0.044** (0.021)	0.022** (0.009)
4-Year Post Election	0.048* (0.025)	0.015 (0.010)
5-Year Post Election	0.034 (0.025)	0.023** (0.011)
Industry + Employment Ctrls.	X	X
Flexible Ctrls.	X	X

*Note:* This table presents the DHS employment growth rate and survival estimates estimated as described for Figure 2-8. The DHS employment growth rate estimates exactly match the DHS employment estimates presented in that table.

Table 2.6: Election Delay Heterogeneity, Continuous Delay Time Specification

Treatment:	Log Delay Time				Residualized Log Delay	
Outcome:	DHS Emp	Survival	DHS Emp	Survival	DHS Emp	Survival
3-Year Post Election	-0.124** (0.058)	-0.057** (0.023)	-0.117** (0.055)	-0.045** (0.022)	-0.120** (0.059)	-0.055** (0.023)
5-Year Post Election	-0.121* (0.063)	-0.064** (0.026)	-0.104* (0.060)	-0.051** (0.025)	-0.113* (0.065)	-0.060** (0.026)
10-Year Post Election	-0.147** (0.073)	-0.071** (0.033)	-0.152** (0.073)	-0.074** (0.032)	-0.132* (0.074)	-0.061* (0.033)
Industry + Employment Ctrls.	X	X	X	X	X	X
Pooled Ctrls.			X	X		
Flexible Ctrls.	X	X			X	X
Observations	6,000	6,000	6,000	6,000	6,000	6,000

*Note:* This table presents coefficient estimates from a modified version of specification 2.8. Specifically, we interact the treatment by event time indicators with the continuous *log delay time*. See Appendix B.3 for details on how we calculate the delay time. The table reports the coefficients on these interactions at various time horizons. Thus, a survival coefficient of -0.05 means that the effect of successful unionization on survival is 0.5 pct. pts. higher for elections with a 10 % longer delay time. The first four columns use the raw number of days between petition filing and election dates to define the log delay time. For the last two columns, we first regress log delay time on within-year deciles of the election bargaining unit size and use the residuals from this regression as the interaction. The sample includes all elections at manufacturing establishments -10 to 10 years before and after each union election but we only include a subset of these coefficients. The even columns include the DHS employment growth rate relative to time  $-1$  as the outcome variable (see Section 2.4 for their definition). The odd columns include an indicator for whether the establishment exists at time  $t$  as the outcome. Standard errors are clustered by establishments' firmid during the year of the election (e.g., the clustering variable is fixed over time for each establishment). \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

THIS PAGE INTENTIONALLY LEFT BLANK

# Chapter 3

## Competition and School Quality: Evidence from Centralized Assignment

### 3.1 Introduction

There has been a surge in market-based reforms to U.S. public school districts over the past two decades. While proposals like school vouchers or charter schools expansions are generally mired in controversy, many large public school districts across the U.S. have introduced sophisticated mechanisms that facilitate school choice within the existing portfolio of schools in the district.<sup>1</sup> These mechanisms are appealing to policymakers since they are an inexpensive way to increase competition between schools and boost student achievement. However, skeptics argue that competition between public schools is inherently limited due to capacity constraints, residential segregation, and informational frictions.

Existing evidence on the effects of competition on school quality is mixed in the public-school context. For example, Hoxby (2000) and Card et al. (2010) show that increasing the number of school districts in an area causally raises the area's average test scores, but it's not clear whether the effects are driven by school responses to competition or other channels like student sorting.<sup>2</sup> Recent papers using student-level

---

I thank Daron Acemoglu, Nikhil Agarwal, Josh Angrist, David Autor, Jonathan Cohen, Clemence Idoux, Parag Pathak, Carolyn Stein, John Van Reenen, and Samuel Young for helpful comments and discussions. I am especially grateful to Eryn Heying, Maggie Ji, Jim Shen, and Anna Vallee for their data and administrative support, and to Jim Carpenter, Laurie Premer, Rob Schaller and Monica Villarreal for their institutional knowledge and expertise. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302.

<sup>1</sup>Whitehurst (2017) estimates that the share of large school districts in the U.S. that allow school choice nearly doubled from 29% in 2000 to 56% in 2016.

<sup>2</sup>A large literature in IO has emphasized that a higher concentration of firms does not necessarily

data show that families do not appear to choose schools based on school effectiveness (for example, Abdulkadiroglu et al., 2017b; Walters, 2018), casting doubt on the competitive incentives for raising school effectiveness.

This project revisits the magnitude and effects of competition in public schools to provide two contributions — a direct estimate of the competitive pressure that schools face due to intradistrict school choice and novel evidence on the relationship between competition and school effectiveness. I show that competitive pressure in my sample is generally low, but can be substantial at some schools, including ones in the traditional sector. The competitive pressure is negatively correlated with conventional measures of concentration, but the majority of the variation in competitive pressure is unexplained by concentration. I find that competition alone does not lead to increases in school effectiveness, but it does induce schools to devote discretionary spending towards instruction rather than administration.

The setting of the project is a large urban school district in the U.S. between 2009 and 2017. This setting offers several key advantages for measuring the level and effects of competition between public schools. First, the district has a large and diverse range of school options, and they give individual schools unusually high autonomy in determining the personnel and course offerings at the school. Second, the district introduced a centralized student assignment mechanism during my sample period; the mechanism generates detailed data on student preferences that allow me to directly quantify competition in the district. Third, I observe student-level panels in test scores as well as school-level expenditures, which allow me to characterize school inputs and outputs at an unusually high level of detail. Finally, I can leverage the introduction of coordinated choice as a source of external variation in the competitive environment to test for the effects of competition on school outcomes.

I begin with a standard model to show that schools' competitive enrollment incentives can be summarized by one conceptual measure: the change in school enrollment in response to a change in its quality. Like past literature that modeled public school competition, I assume that schools choose a level of quality to balance the tradeoff between enrollment and the cost of quality provision. However, I deviate from past work by allowing for capacity constraints to bind, which is relevant for many schools in my setting. The model formalizes a simple intuition for when competition matters — schools face higher incentives to provide quality on the margin when quality improvements yield a greater relative increase in student enrollment.

A key focus of the project is to explicitly estimate the level of competitive pressure

---

imply less competition between firms (Bresnahan, 1989).

that schools face, so I show how the conceptual measure of competitive pressure can be mapped to empirical objects in my setting. A full empirical model of the market requires estimates of student demand, school capacities, and the market assignment mechanism. To quantify student demand for schools, I estimate an exploded logit discrete choice model using students' submitted rank order lists while allowing for rich interactions between observable student characteristics and school characteristics. Meanwhile, I directly observe proxies for schools' capacities and the assignment mechanism used to determine student offers. The resulting empirical model appears to capture the main elements of the assignment system, closely replicating observed patterns in rank choices and outcomes for a holdout sample of students.

I find that the estimated competitive incentives for enrollment within the district are generally low, but can be substantial at some schools. Capacity constraints bind at 32% of the schools in the district. The median school expects enrollment to increase by 8.3% in response to an improvement equivalent to the utility value of reducing the commuting distance by 1km. Traditional schools and high schools tend to face the lowest degrees of competition, while elementary schools and alternative model schools (like magnet and charter schools) generally face substantial competitive incentives at the margin. Competitive pressure is lower, on average, in schools facing high market concentration, but over 95% of the variation in competitive pressure remains after conditioning on concentration.

With the estimates of competitive pressure in hand, I then use two complementary sources of variation to test for the effects of competition on school outcomes. I draw similar conclusions from each approach. The first empirical strategy exploits cross-sectional variation in competitive pressure across schools. The second empirical strategy studies within-school changes following the introduction of centralized assignment. The conclusions from the two empirical strategies are remarkably similar despite the different sources of variation used in the two approaches. Schools appear to respond to competition by shifting expenditures from administration to instruction but not by increasing school value-added on test scores.

In the first approach, I compare observably similar schools that face different degrees of competitive pressure and find little difference in their school effectivenesses but large differences in their discretionary expenditures. I estimate a reasonably precise and approximately zero correlation between competitive pressure and a school's value-added on state standardized exams. The value-added measures are estimated from student-level panel data and control for the differences in student composition by construction (Chetty et al., 2014). As a result, it is unsurprising that the relationship between

competition and value-added does not change meaningfully after adding controls for student composition. On the other hand, I find that schools facing high competitive pressures devote a greater share of their budgets towards instructional purposes. These estimates are also robust to controls for student composition and school type, bolstering the interpretation that the results reflect the effects of competition as opposed to other channels.

To address any remaining concerns that cross-sectional estimates are confounded by unobserved differences between schools, I compare changes in school outcomes following the introduction of the centralized assignment mechanism for schools facing different levels of competitive pressure. This differences-in-differences strategy allows me to control flexibly for omitted variables by including school fixed effects and flexible time trends by school type. My results mirror the conclusions from the cross sectional analysis. Although the assignment mechanism drastically lowered the costs of applying to different schools within the district, the greater degree of choice did not differentially increase value-added at competitive schools (relative to uncompetitive schools) over the five years after the introduction of centralized assignment. On the other hand, competitive schools were more likely to increase instructional spending following the introduction of the assignment mechanism. Moreover, both sets of results are robust to controlling for changes in student composition and do not appear to be driven by subsequent policy changes.

This paper contributes to the large literature on the downstream effects of school competition. Economists have long emphasized competition as a potentially important channel for inducing school-level improvements (Friedman, 1962; Hoxby, 2003). But empirical evidence on this channel has generally been indirect. Existing research often focuses on the effects of changes in the menu of schooling choices on school outcomes. In the school voucher context, examples include Hsieh and Urquiola (2006); Figlio and Hart (2014); Muralidharan and Sundararaman (2015), as well as the survey in Epple et al. (2017). In the charter school context, examples include Ridley and Terrier (2018); Gilraine et al. (2021). Studies using more general, market-level variation in district boundaries include Hoxby (2000); Rothstein (2007); Card et al. (2010). These studies tend to find mixed to positive effects of increasing choice on school outcomes, but as Urquiola (2016) points out, it is difficult to disentangle the competition channel from other spillover effects without explicitly modeling the competitive framework.

A more recent literature has taken a more IO approach to specify the market structure for schools and identify the competition channel. Neilson (2021) and Allende (2019) develop full equilibrium models of competition between private schools in Chile



and Peru to determine schools' pricing and quality provision decisions. Bayer and McMillan (2005) model competition between public elementary schools in the San Francisco Bay Area to derive a theoretically motivated competition measure and assess its relationship to supply-side characteristics. Campos and Kearns (2022) estimate demand for schools using rank-order-preference data from the centralized assignment mechanism at the Los Angeles Unified School District. They show that introducing school choice at the neighborhood level increased school effectiveness, especially at schools facing more intense competition. My paper is closest to Campos and Kearns (2022) in the setting and approach, although I differ by explicitly considering enrollment-based fiscal concerns and potentially binding capacity constraints in determining competitive incentives. I also study a broader, market-level adoption in centralized choice rather than the more localized rollout in Campos and Kearns (2022), which may explain in part the differences in our findings.

The paper also relates to the literature on estimating student preferences over schools. Recent papers highlight that student demand for school effectiveness appears to be low, which would limit the potential for school choice to improve student outcomes (Hastings et al., 2008; Abdulkadiroglu et al., 2017b; Walters, 2018). However, without variation in school quality, it is difficult to infer whether the observed lack of responsiveness to school effectiveness is due to true preferences over school effectiveness or unobserved factors that are correlated with school effectiveness. My paper complements these results by providing an alternative supply-side test — if students value school effectiveness, then we should expect to see schools provide more school effectiveness when there is greater competitive pressure to do so. I also provide a methodological contribution to the growing literature on flexibly estimating models of student demand from centralized school assignment data (Hastings et al., 2008; Abdulkadiroglu et al., 2017a; Agarwal and Somaini, 2019). I show that by adding information on school capacities and the assignment algorithm, it is also straightforward to use the demand estimates to characterize the school-side enrollment incentives.<sup>3</sup>

The rest of the paper proceeds as follows. Section 2 first describes the school district and the student- and school-level data in more detail. Section 3 starts by specifying

---

<sup>3</sup>More broadly, by characterizing school-side incentives as reflecting enrollment-based fiscal motives, the paper relates to the literature on soft budget constraints and the financial incentives facing non-profit and government owned institutions (Duggan, 2000; Glaeser and Shleifer, 2001; Kornai et al., 2003; Dafny, 2005). By consider potential responses in school characteristics rather than prices, the paper also relates to the growing IO literature on endogenous product attributes (Fan, 2013; Sweeting, 2013; Guajardo et al., 2016; Prince and Simon, 2017; Ito and Sallee, 2018).

a simple model where schools compete for student enrollment and derives the key statistic capturing competitive pressure. Section 4 then describes how I empirically implement the model using data from the district. Section 5 presents estimates of competitive pressure using the empirical model of the market. Section 6 assesses the relationship between competitive pressure and school effectiveness. Finally, Section 7 concludes.

## 3.2 Setting and data

The setting of the project is a large urban school district in the U.S. The district's size and a high degree of school-level autonomy make it particularly suitable for studying competition between schools. Focusing on a single district allows me to observe detailed data on student preferences, academic performance, and school budgets, including from before the introduction of the centralized assignment system. This level of granularity is crucial for quantifying competition and its effects.

Like most other large urban school districts in the U.S., the district primarily enrolls lower-income, minority students. Table 3.1 presents summary statistics on the characteristics of students and schools in the district over the entire sample period (in the first column) as well as during the introduction of centralized assignment (in the second column). 71% of the students in the district qualify for either free or reduced lunch based on their family incomes. The district is majority-minority — 57% of students in the district are Hispanic, and 14% of students are black. Approximately 23% of students in the district are English language learners.

The district has drawn distinction for implementing a balanced approach to school choice that emphasizes school autonomy and experimentation within the district. In addition to traditional and magnet schools, which constitute approximately half of the schools in the district, the district also contains a growing number of charter schools and “innovation schools” that have additional autonomy over programming, budget, and staffing decisions. Even in the traditional sector, school administrators have direct authority in determining the size and composition of their staff, and most funding is allocated to schools on a per-student basis.

In Fall 2011, the district introduced a centralized assignment mechanism to match students to schools. Previously, the process through which students applied to non-neighborhood schools was uncoordinated and often undocumented. Centralizing admissions not only made the choice process transparent to all students, but the process also required students to actively participate in the choice mechanism and

rank any schools they would prefer to their neighborhood schools. Student rankings are then combined with school-side priorities and a random tiebreaking lottery number through the student-proposing deferred acceptance algorithm, which matches each student to their most preferred stable match (Gale and Shapley, 1962).

I bring together several datasets for this project. The student-level data are provided by the school district from their administrative records and contain longitudinal unique identifiers that can be linked across datasets. The data cover the universe of students in the district, including those attending charter schools, but do not include students at private schools or in other districts. The data also vary in the time periods they cover due to changes in the data storage practices in the district.<sup>4</sup> I observe students’ ranked choices and offers from the centralized assignment mechanism from the introduction of the system in 2012 to 2017. I also observe students’ enrollment and demographic information from 2010 to 2017, and students’ test scores on state exams from 2009 to 2017. Finally, I digitize school-level expenditures data between 2012 and 2014 from published budget books and match them to the administrative data. However, schools from the charter sector are not required to report their budgets in the budget books, so I do not observe expenditures from that sector.

### 3.3 Model of school competition

To fix ideas, I specify a simple model where schools attract students by making costly improvements to their quality. In this model, each school’s competitive incentives can be summarized by a conceptual measure — the change in enrollment in response to a change in its quality. My model is based on the model in Card et al. (2010), but I extend the model to better reflect the considerations for schools in my setting. First, I explicitly model the school’s objective as maximizing revenues from enrollments subject to the monetary costs of quality provision (as opposed to a general taste for market share subject to a non-monetary managerial effort cost). Second, I allow for potentially binding capacity constraints at the school level that dampen incentives for further improvement.

Assume that a student of type  $\theta_i \sim F(\theta)$  has utility for attending school  $j$  with

$$u_{ij} = V_j + \beta(\theta_i) M_j - d_{ij} + \epsilon_{ij}, \quad (3.1)$$

---

<sup>4</sup>The academic school year spans from the fall of a calendar year to the spring of the subsequent year. To avoid confusion, I refer to all academic years by their ending year (so I refer to SY2011-12 as 2012).

where  $V_j$  is a common school quality term,  $M_j$  are fixed school characteristics whose value to students may depend on student type  $\theta_i$ ,  $d_{ij}$  is the commuting cost for student  $i$  to school  $j$ , and  $\epsilon_{ij}$  is an idiosyncratic preference for school  $j$ . So, students trade off their taste for the school's observable and unobservable characteristics against the (distance) cost of attending. Since I focus on public schools within the same district, I assume that the primary cost for attendance is the commuting cost  $d_{ij}$ , which is normalized to 1 to serve as a numeraire.

Meanwhile, a school's objective is to maximize its revenues, so it solves

$$\begin{aligned} \max_{V_j} \quad & Rn_j(V_j, V_{-j}) - c(V_j)n_j(V_j, V_{-j}) \\ \text{s.t.} \quad & n_j(V_j, V_{-j}) \leq N_j \\ & Rn_j(V_j, V_{-j}) - c(V_j)n_j(V_j, V_{-j}) - C_j \geq 0, \end{aligned} \tag{3.2}$$

where  $R$  is the per-student revenue,  $n_j(V_j, V_{-j})$  is the total number of students who would enroll at school  $j$  given quality  $V_j$  at school  $j$  and quality  $V_{-j}$  at all other schools, and  $c(V_j)$  is a strictly increasing function that captures the marginal cost of providing quality  $V_j$  to each student.<sup>5</sup> Two constraints enter the school's optimization problem. The first is a capacity constraint, where the total enrollment of students at school  $j$  must be below the school's capacity  $N_j$ . The second constraint is a solvency constraint, which assumes that schools must collect enough revenue to meet their fixed costs ( $C_j$ ). By modeling the quality choice as  $V_j$ , the common component of utility that are equally valued by all students, I abstract from more complex considerations where schools may adjust their optimal quality choice based on the selection of students in their market (as in McMillan, 2004).

The equilibrium enrollments  $\{n_j\}$  are determined by student preferences, schools' quality choices, and the assignment mechanism. Under the assumptions laid out in the model, in any equilibrium where schools simultaneously choose their quality, a solvent school that is below the capacity constraint will satisfy the first order condition

$$[R - c(V_j)] \underbrace{\left( \frac{\partial n_j(V_j, V_{-j}) / \partial V_j}{n_j(V_j, V_{-j})} \right)}_{=\partial \log n_j(V_j, V_{-j}) / \partial V_j} = c'(V_j). \tag{3.3}$$

$[R - c(V_j)] > 0$  and is strictly decreasing for a solvent school, and  $c'(V_j)$  is strictly

---

<sup>5</sup>For simplicity, I've suppressed the dependence of  $n_j$  on fixed school characteristics and the distribution of student types.

increasing. It follows that schools whose relative enrollment is more responsive to quality increases ( $V_j$ ) will choose higher quality in equilibrium. In other words, holding all else equal, a school’s competitive pressure for enrollment is summarized by its semi-elasticity of enrollment with respect to quality ( $e_j \equiv \partial \log n_j(V_j, V_{-j}) / \partial V_j$ ), and schools with a higher  $e_j$  will choose a higher quality in equilibrium.

On the other hand, when capacity constraints bind, schools do not have enrollment-based incentives to further increase school quality. Instead, the optimal school quality is determined by

$$V_j^c = \min \{v_j : n_j(v_j, V_{-j}) \geq N_j\}. \quad (3.4)$$

This expression formalizes a common critique that the expected effects of competition may be low if school capacities are binding. Furthermore, the expression implies that capacity-constrained schools may even choose to *decrease* school quality in response to competition-enhancing policies if the increase in the semi-elasticity of enrollment is accompanied by an increase in the *level* of enrollment.

Although the setup is stylized, it is sufficiently flexible to also formalize an alternate theory of school quality provision. This “non-strategic” theory is that schools simply provide the highest level of quality that they are able to afford, so the realized school quality  $\tilde{V}_j$  satisfies

$$\left[ R - c(\tilde{V}_j) \right] n_j(\tilde{V}_j, V_{-j}) = C_j. \quad (3.5)$$

This model can capture a common concern with the introduction of school choice, which is that the introduction of competing schools may lower enrollment at incumbent schools. This lower enrollment in turn requires the school to devote a greater share of expenditures towards covering fixed costs, ultimately eroding the quality of instruction. This alternate model centers on the level of enrollment and generates no *a priori* relationship between school quality and the semi-elasticity of enrollment with respect to quality ( $e_j$ ).

### 3.4 Empirical model of the market

Given the richness of data available, it is straightforward to map the conceptual model from Section 3.3 to its empirical analogs. To ensure that my demand estimates are not confounded by subsequent changes that schools may have made following the centralized assignment mechanism, I focus on modeling the market from the first year of the match, which took place from Fall 2011 to Spring 2012. I focus on the market for entry grades at each school (0th grade for elementary schools, 6th grade for middle

schools, and 9th grade for high schools) since mechanism take-up is highest for those grades and the interpretation of the choice process is the most straightforward. The model has three components: student demand, school supply, and the assignment mechanism. I discuss the empirical implementation of each component in turn.

### 3.4.1 Demand estimation

The school district uses student-proposing deferred acceptance as its assignment mechanism. When rank-order lists are unrestricted in length, the algorithm is strategy-proof, so it is optimal for students to truthfully report their preferences over schools to the assignment mechanism (Roth, 1982). The students' reported rank-order lists can then be directly interpreted as ranked preference data, which are key for estimating flexible models of discrete choice that allow for rich heterogeneity in preferences (Berry et al., 2004). Although the district limits the number of schools that students can rank to five, Figure C-1 shows that most students rank fewer than five schools. I take truth-telling as a reasonable approximation of student behavior and treat students' submitted rank order lists as a partial preference ordering over the set of all schools. Since students are guaranteed an offer at their neighborhood school, I assume that all unranked schools are less preferred than the student's neighborhood school.

In principle, the student demand Equation 3.1 is non-parametrically identified from variation in commuting distances  $d_{ij}$  and student rank order lists (Agarwal and Somaini, 2019). In practice, the distribution of student preferences is a high dimensional object, and I face a tradeoff between analytical tractability and the flexibility to capture realistic substitution patterns. Based on the results from Pathak and Shi (2018), where a richly specified multinomial logit model performed as well, and occasionally better, than a more computationally intensive mixed multinomial logit model in predicting student choices following a policy change, I use a flexible multinomial logit model as my baseline approach for approximating student demand.

I model the utility of school  $j$  for student  $i$  as

$$u_{ij} = \beta(X_i) X_j + \gamma(X_i) M_{ij} + V_j + \varepsilon_{ij}, \quad (3.6)$$

where I allow for preferences for school-specific characteristics  $X_j$  and student-school characteristics  $M_{ij}$  to differ by student characteristics  $X_i$ . Specifically, I allow for the student's utility to depend on a school's average student achievement (differentially by a student's gifted status and income), the share of minority and ELL students (differentially by a student's race and ELL status, respectively), and the distance

between the student and the school (differentially by a student’s race and income). I also include dummies for whether a student has a sibling currently enrolled at the school, and whether the school is the student’s neighborhood school.  $V_j$  captures all of the components of the school’s utility that are common to all students. Finally,  $\varepsilon_{ij}$  is distributed independently and identically from a type I extreme-value distribution.

I estimate the empirical demand Equation 3.6 from students’ submitted rank order lists by exploded logit (Hausman and Ruud, 1987). I estimate the model separately for the elementary, middle, and high school entry grades, allowing the parameters to vary flexibly by grade. The model estimates are reported in Table 3.2. The parameters are generally precisely estimated and intuitive in sign and magnitude. Students are likely to rank schools that are their neighborhood schools and especially more likely to rank schools where they currently have a sibling, and they are less likely to rank schools that are far away. Hispanic and low income students find distance more costly, while black students find distance less costly. Furthermore, consistent with Goodreau et al. (2009); Idoux (2022), I find that black, Hispanic, and ELL students are relatively more likely to rank schools that have a higher share of similar students. Finally, I find that gifted students are relatively more likely to rank schools with high average achievement, whereas low income students are relatively less likely to do so. The estimates are also generally stable across specifications — when I vary the match-specific terms in the demand equation, the coefficients on the residential school, distance, and sibling controls remain quantitatively and qualitatively similar across specifications.

The primary concern with the multinomial logit model is that its independence of irrelevant alternatives (IIA) property would be a poor approximation of realistic substitution patterns. To assess whether the rich heterogeneity in Equation 3.6 alleviates the IIA concern, I directly assess the model’s ability to fit substitution patterns by comparing the empirical choice shares for students’ Top 2 choices to the predicted shares under Equation 3.6. Specifically, let

$$p_{jk} = P(u_{ij} > u_{ik} > u_{ix}) \quad \forall x \neq j, k$$

be the probability that a student ranks school  $j$  first and school  $k$  second. If the demand model is sufficiently flexible to capture students’ substitution patterns, then the predicted probability of ranking *both*  $j$  first and  $k$  second should closely match the realized rates. To avoid concerns with overfitting, I estimate demand parameters for this exercise on a random hold-out sample of students.<sup>6</sup> Figure 3-1 plots the correlation

---

<sup>6</sup>Given that the model is estimated from a subsample, it will also be subject to great estimation

between the empirical  $p_{jk}$  and the model's predictions. The R squared between actual choice shares and the model's predictions range between .756 for high schools and .819 for elementary schools. Furthermore, the slope between the two measures is close to 1. To assess the importance of microdata, I also replicate the exercise while removing all individual heterogeneity from the multinomial logit model. Consistent with the restrictiveness of the IIA property, the corresponding model estimates in Figure C-2 are generally unable to rationalize the most popular choice combinations observed in the data.

### 3.4.2 School supply and assignment mechanism

To fully characterize the model, I need to observe  $N_j$ , the capacity constraint for each school. However, capacities are difficult to observe, and may not be rigidly binding at some schools.<sup>7</sup> Instead, I calculate a lower bound for school capacities using the assignment data. If a school's offers were rationed in the assignment process, then that school's capacity is the total number of offers it gave. If seats were not rationed, I assume that the school's capacity is not binding at the margin. This assumption is sufficient for calculating the school's expected enrollment response for *small* changes in utility, which corresponds to our estimand of interest. But it is important to highlight that better data on capacities would be important for assessing any large counterfactual changes.

The assignment process that coordinates offers across schools is the standard student-proposing deferred acceptance algorithm. Each student is initially assigned priorities at each school based on the student's location and characteristics, as well as a random tiebreaking number (in the case two students have the same priority at a school). The algorithm proceeds in rounds. In each round,

1. All unassigned students apply to their most preferred option from the set of schools where they have not been rejected
2. Schools rank all current applications and tentatively accepted students. If the total number of students exceeds a school's capacity, the school rejects candidates with the lowest priorities and tentatively accepts the rest.

The algorithm ends when all students either are tentatively assigned to a school or exhausted their rank order lists. At that point, all tentative matches are then finalized, and any remaining students and seats remain unassigned.

---

error than the baseline demand estimates that are estimated on the entire sample of students.

<sup>7</sup>For example, neighborhood schools are required to serve all students in their area.



Since I observe all the inputs to the assignment process as well as the assignment algorithm, it is straightforward to replicate the match process under the submitted rank choices or any counterfactual rankings. In most years, I can nearly completely replicate the match based on the data that I observe. However, in the first year of the match, my data is less complete — I observe data from a combination of the main match and a supplementary round. I do not have the data necessary to separate the rankings and outcomes for the main round and the supplementary round. Nevertheless, by running the student-proposing deferred acceptance algorithm on the rank ordering and priorities that I do observe, I can still replicate realized offers for 90% of all students. As a result, I treat the imperfect data as a reasonable approximation of the true match process in 2012.

The assignment process is at the *program* level, and students may be matched to different programs within the school based on their eligibility, seat preferences, and program preferences. However, most of my data and outcomes are at the higher *school* level. As a result, I make an additional simplifying restriction where I abstract from the smaller programs within a school, and I treat the ranking and assignment process as occurring at the school level. Figure C-4 compares the share of student assignments that are replicated using school-level choices and priorities to the share of assignments that are replicated using program-level choices and priorities. The replication rate falls by less than 5% after aggregating the options to the school level.

### 3.4.3 Validation

I have made simplifying assumptions to implement each part of the market model. To assess whether the model can capture the core outcomes in the market, I test whether the model's predicted *offer* shares for each school match their empirical counterparts. This is more demanding than the individual tests that have been presented so far since it requires the *combination* of all three components of the market to be approximately correct. Moreover, the test implicitly places the most weight on fitting the features that are most relevant for outcomes (for example, predicting a student's top choice is particularly important if the student has a high priority, whereas predicting a student's later choices is particularly important if the student has a low priority). As a result, the test may be a particularly useful comprehensive assessment of the model's fit.

Figure 3-2 plots the empirical offer share for each school against the model's average predicted offer share after simulating the market 100 times. In each simulation, I draw idiosyncratic taste shocks for each student-school pair from the assumed type I

extreme value distribution. I then rank schools according to each student’s realized utilities and run the assignment algorithm while taking the student’s lottery number, school priorities, and school capacities as given. To ensure that my results are not driven by overfitting, I again estimate demand parameters from a randomly selected holdout sample, and I simulate the market using the remaining students.<sup>8</sup> I find that schools’ predicted offer shares are tightly correlated with schools’ actual offer shares, even for schools that are not at capacity, which is particularly reassuring that I have captured the key components of the assignment process. The notable exception is that elementary school students are less likely to be unassigned in my simulations than in reality, which may reflect some measurement error in the neighborhood priorities given to students. Even in the elementary school case, though, empirical offer shares closely match estimated offer shares, so the excess assigned students do not appear to benefit any school in particular.

### 3.5 Estimates of competitive pressure

The key estimand is each school’s semi-elasticity of enrollment with respect to its quality.<sup>9</sup> I estimate this semi-elasticity by using the empirical model from Section 3.4 to simulate enrollment changes when the school’s quality term increases.<sup>10</sup> Specifically, for each school  $j$ , I estimate  $e_j$  by increasing its quality measure  $V_j$  by a fixed increment  $\Delta$  while holding all other schools’ qualities fixed at their estimated levels. I then re-simulate the market’s offer shares using the same approach as in Section 3.4.3, where I draw students’ idiosyncratic taste shocks from the assumed distribution and run the match separately for students’ rank order lists under each simulation draw. It follows that the estimated semi-elasticity is then

$$\hat{e}_j = \frac{\hat{n}_j(\hat{V}_j + \Delta, \hat{V}_{-j}) - \hat{n}_j(\hat{V}_j, \hat{V}_{-j})}{\Delta}, \quad (3.7)$$

where  $\hat{n}_j(V_j, V_{-j})$  is school  $j$ ’s simulated offer share given school qualities  $(V_j, V_{-j})$ . The mechanism is discrete, so there is a bias-variance tradeoff between choosing a

---

<sup>8</sup>Removing students from the market simulation without adjusting school capacities would create excess capacity at otherwise capacity constrained schools. To address this issue, I reduce the capacities of each school by the number of realized offers given to students in the holdout sample. This ensures that the realized matches for the simulation sample are unaffected by the removal of other students as long as rank-order lists remain the same.

<sup>9</sup>Formally, this measure is defined in Equation 3.3 as  $e_j \equiv \partial \log n_j(V_j, V_{-j}) / \partial V_j$ .

<sup>10</sup>This approach is based on the approach from Bayer and McMillan (2005). I adapt their framework for simulating demand elasticities to the centralized assignment system in my setting.

large  $\Delta$  (where the local approximations for capacity become less valid) and choosing a small  $\Delta$  (where estimation and simulation error may be substantial relative to the actual change in offers shares). I set  $\Delta = 1$ , which can be interpreted as the preference cost of commuting an extra 1km to school for a white, non-FRL student.

Figure 3-3 plots the empirical CDF for the estimated school-level enrollment semi-elasticities. Competitive incentives for enrollment are generally low but can be substantial at some schools. I estimate that 32% of schools are already capacity constrained, and have no enrollment-based incentives to improve. These capacity constrained schools also constitute 32% of offers, so capacity constrained schools are not systematically larger or smaller than non-capacity constrained schools. On the other hand, more than half of all schools can expect to enroll at least 8.3% more students if their perceived quality increases by a 1km-equivalent utility increment, so enrollments are far from inelastic. There is a long tail of schools facing particularly high competitive pressure — the 90th percentile of schools can expect to enroll 19.3% more students for the same utility increase. Larger schools have generally less elastic demand, so the distribution of enrollment semi-elasticities are shifted to the left when schools are weighted by the number of offers, but otherwise, the patterns remain comparable.

Table 3.3 compares my estimates of competitive pressure to traditional concentration-based proxies. The results support the general intuition that competitive pressure is lower in areas with fewer competitors, but also highlight that concentration is not a sufficient statistic for competition. On average, schools with one more competitor within a 5km radius have 0.24pp higher enrollment semi-elasticities, and schools with a 10% higher market share among students living within a 5km radius of the school have 0.38pp lower enrollment semi-elasticities. However, while concentration is negatively correlated with competition, the two are distinct quantities. Even when I include both measures of concentration and allow the coefficients to vary arbitrarily by grade, the adjusted R-squared of the regression of  $e_j$  on concentration is less than .04, so more than 95% of the variation in competitive pressure exists conditional on concentration.

To further understand the variation in competitive pressure, I disaggregate the distribution of competitive pressure by school characteristics. Figure 3-4 plots the density of the competitive measure by the sector of the school. As expected, the traditional sector schools generally face the least competitive pressure, whereas alternate sectors like the charter and magnet schools face strong incentives to recruit additional students. Interestingly, although innovation schools were created in the district as alternatives to traditional schools, they do not necessarily face greater

competitive pressures than even traditional schools, and they are the most likely to be capacity-constrained. Similarly, Figure 3-5 plots the density of the competitive measure by the grade level of the school. Elementary schools tend to face more competitive pressure than high schools, whereas middle schools have the largest dispersion between schools with low competitive pressure and schools with high competitive pressure.

## 3.6 Relationship between competition and school outcomes

The demand model is agnostic about the content of the quality term  $V_j$ . As a result, the competitive pressure term  $e_j$  captures the school's competitive incentives to provide *any* characteristic that contributes to students' common utility  $V_j$ . To better understand the specific margins of school responses to competition, I use two complementary approaches to test for the effects of competition on school outcomes. The first approach makes cross sectional comparisons between schools, and the second approach makes within-school comparisons following a policy change. I find consistent conclusions across both approaches: competition alone does not lead to substantial increases in school effectiveness, but it does induce schools to devote discretionary funding towards instruction rather than administration.

### 3.6.1 Cross sectional variation

As a starting point, I consider the cross sectional difference in school inputs and outputs between observably similar schools that face different degrees of competitive pressure. To do so, I estimate the school level regression

$$y_j = \beta e_j + \theta X_j + \nu_j, \quad (3.8)$$

where  $e_j$  is the school's competitive pressure, and  $X_j$  is a set of controls for the size, average student composition, and sector of the school. To ensure that I am capturing the effects of competition, I focus on examining school outcomes after the introduction of centralized assignment in 2012. The key identifying assumption for the causal effect of competitive pressure on school outcomes is that

$$E[e_j \nu_j] = 0, \quad (3.9)$$

so a school’s competitive pressure is uncorrelated with other unobserved factors that also contribute to the outcome ( $\nu_j$ ). This is a strong assumption. For example, schools in competitive areas may also be more likely to enroll students whose parents closely monitor school principals and teachers. To assess the potential concerns with student sorting or other omitted variables, I use value-added as my measure of school effectiveness rather than test scores to control for student sorting. I also systematically probe the robustness of my results to including additional controls (in the spirit of Altonji et al., 2005; Oster, 2019).<sup>11</sup>

Table 3.4 reports estimates of  $\beta$  across a variety of specifications. I find that schools facing higher competition are not more effective at increasing test scores. Without controls, a 10 percentage point increase in competitive pressure  $e_j$  is correlated with a .015 (0.013 s.e.) standard deviation *decrease* in value-added.<sup>12</sup> This negative relationship becomes slightly smaller in magnitude after including controls for average student characteristics and the sector of the school but remains negative and statistically indistinguishable from 0.<sup>13</sup> If there is omitted variables bias in Equation 3.8 that is obscuring the effects of competitive pressure on school effectiveness, it would have to be mostly uncorrelated with the observable characteristics of the school and bias  $\beta$  downwards. In other words, schools facing *less* competition would need to have better unobserved determinants of school effectiveness.

On the other hand, I find that schools facing greater competitive pressure do spend more on instruction. Without controls, a 10 percentage point increase in competitive pressure  $e_j$  is correlated with a 2.1pp (0.9 s.e.) increase in the share of a school’s discretionary budget spent on instruction. This relationship also remains stable after including controls even though the adjusted  $R^2$  of the regression increases substantially. It is also worth highlighting that in the alternate model from Section 3.3, larger schools should face lower fixed costs and devote more money to instruction. Since larger schools also generally face less competitive pressure, this should bias the coefficient on

---

<sup>11</sup>Note, however, that the addition of controls in this setting is not innocuous. There is a separate concern that student composition or size reflect equilibrium outcomes, so controlling for them may also *introduce* bias (Angrist and Pischke, 2009).

<sup>12</sup>The results also remain similar when I include schools whose capacities bind, although the interpretation of the estimates becomes more complicated. In the model, the relationship between competitive pressure and school quality only holds for schools that are below capacity. Equation 3.4 highlights that capacity constrained schools may face strong incentives to raise quality to the point where their capacities bind, but no incentives to further increase quality.

<sup>13</sup>The district places some underperforming schools under “targeted interventions,” which provides additional support to try to turn around their performance. There is some uncertainty about whether these interventions are a confounder or an outcome of school competition. I drop all schools that receive interventions from the district at any point in the sample period in the last column, and my results remain similar.

competitive pressure ( $\beta$ ) downwards. Although the fact that the coefficient increased from .214 to .247 after controlling for school size provides some support for this explanation, the small magnitude of the change suggests that fixed costs are not the primary explanations for differences in schools' spending decisions.

### 3.6.2 Variation from centralized assignment

Although cross sectional comparisons are a useful baseline, it is ultimately difficult to reject the possibility that schools facing greater competitive pressure are different in unobservable ways. As a complementary exercise, I use the introduction of centralized assignment as exogenous variation in schools' competitive environments and consider the within-school change in inputs and outputs. My event-study estimation equation is

$$y_{jt} = \sum_t \beta_t (e_j \times I_t) + \alpha_j + \gamma_{gt} + \theta X_{jt} + \eta_{jt}, \quad (3.10)$$

where  $e_j$  is the school's estimated competition measure,  $I_t$  is an indicator for year  $t$ ,  $\alpha_j$  is a school fixed effect,  $\gamma_{gt}$  are time fixed effects (that vary by group  $g$ ), and  $X_{jt}$  are additional controls for the school's composition of students. In the baseline specification, I allow for differential trends by the sector of the school, but I do not control for the composition of students at the school (to allow for the possibility that student composition is an outcome of the policy). The key identifying assumption for  $\beta_t$  is that

$$E[e_j \eta_{jt}] = 0, \quad (3.11)$$

so schools facing greater competitive pressure are not experiencing any changes to their school outcomes that are not explained by the observable controls.<sup>14</sup>

I find results that are consistent with the cross sectional results. Figure 3-6 reports the event-study coefficients for the outcome of school value-added. High and low elasticity schools appear to be on similar trends before the introduction of the match, which bolsters the interpretation of  $\beta_t$  as the causal effect of increasing competitive pressure on schools. I do not find any evidence that schools facing more competition increased their effectiveness following the introduction of centralized assignment, and if anything, the point estimates are generally negative. A school with a 10p.p. higher

---

<sup>14</sup>Note that although the identification assumptions for this approach is weaker, interpreting  $\beta_t$  as the effects of competition also requires that centralized assignment particularly increased competitive pressure at high  $e_j$  schools. This additional assumption would be satisfied, for example, if opting out of neighborhood schools was uniformly difficult. On the other hand, this additional assumption may be violated if parents were able to freely choose competitive schools even before centralized assignment.

competitive pressure *lowered* its value-added by 0.012 (0.015 s.e.) standard deviations following the introduction of centralized assignment.

Table 3.6 summarizes the pooled differences in differences estimates across a variety of robustness checks that rule out other potential identification concerns. For example, the state exam that is the basis of my value-added estimates changed its format for high school students after 2014. Although I standardize test scores each year into standard deviations, there may be a concern that different schools are differentially affected by the new exam format. When I drop results after 2014, the estimate becomes positive, but remains small in magnitude and statistically indistinguishable from zero. The effects are also similar if I exclude targeted intervention schools, so my results are not explained by policies that the district may have introduced to turn around some underperforming schools.

On the other hand, I do find that the share of expenditures that are dedicated to instruction increases at schools facing high competitive pressure. Table 3.6 also reports the coefficients from estimating Equation 3.10 on the outcome of instructional expenditure shares at the school. I find that schools facing a 10p.p. higher competitive pressure increases their instructional spending shares by 1.8 percentage points (0.84 s.e.) following the introduction of centralized assignment. These effects also remain similar across specifications, so they do not appear to be driven by changes in student composition or subsequent policy changes.<sup>15</sup>

## 3.7 Conclusion

I use a novel method to directly measure schools' competitive incentives and draw two conclusions. First, the distribution of competitive pressure is dispersed within a public school district. A third of schools face no additional incentives to increase enrollment on the margin, but at least a quarter of schools, including those in the traditional sector, face substantial incentives to attract more students. Second, schools respond to competitive pressure by shifting expenditures from administration to instruction, but they do not appear to increase their school *effectiveness* on academic achievement.

The results in this paper focus on a single, large school district in the U.S., and may not be representative of effects in other institutional settings or contexts. However, it is worth highlighting that at least within public school districts in the U.S., the

---

<sup>15</sup>Given that the per-student funding for the school depends on student characteristics, the stability of estimates after I explicitly control for student composition is particularly reassuring that the results are mechanically determined by any enrollment changes.

district in question is unusually decentralized and flexible. The results are therefore likely to be close to the upper bound of plausible competitive effects within large U.S. districts where the possibility of expanding choice is especially relevant. Moreover, the approach in this paper can be implemented in other school districts with a centralized assignment mechanism, and it would be useful to assess the extent to which competitive pressures and their effects differ in other districts.


Finally, there are also several other fruitful avenues for future research. It would be useful to more comprehensively consider students' outside options, either to other school districts or private schools to assess the degree to which schools may also be responding to students' decisions to leave the district altogether. Another key direction is to better understand the school's supply function. Schools may respond through increasing instructional spending rather than school effectiveness because changing the latter is either difficult or simply unprofitable. Disentangling these two possible explanations would be crucial for the design of policy that can ensure that students ultimately benefit from competition.




# Tables and figures

Figure 3-1: Model fit: substitution patterns

(a) Elementary schools



(b) Middle schools



(c) High schools






Figure 3-2: Model fit: offer shares

(a) Elementary schools



(b) Middle schools



(c) High schools






Figure 3-3: Distribution of competitive pressure




Note: The figure plots the empirical CDF of the school-level estimates of Equation 3.7, the semi-elasticity of enrollment with respect to school quality. Intuitively, the measure captures the school's expected increase in enrollment after a utility increase equivalent to the costs of a 1km commute change. The blue series weights all schools equally, while the red series weights each school by the number of students it enrolls through the assignment process.

Figure 3-4: Competitive pressure by sector




Note: The figure plots the kernel density of the school-level estimates of Equation 3.7, the semi-elasticity of enrollment with respect to school quality, separately for each sector in the district.

Figure 3-5: Competitive pressure by grade



Note: The figure plots the kernel density of the school-level estimates of Equation 3.7, the semi-elasticity of enrollment with respect to school quality, separately by the school's entry grade. 9th grade is the entry grade for high schools. 6th grade is the entry grade for middle schools. 0th grade (i.e., kindergarten) is the entry grade to elementary schools.

Figure 3-6: Event study coefficients for school effectiveness



Note: The figure plots event study coefficients from estimating Equation 3.10 on the baseline sample of all schools with estimated competition semi-elasticities  $e_j$  and value-added estimates. The coefficients are differential changes in the school's annual value-added estimates by the school's competitive pressure  $e_j$ . The outcome is the school's value-added on state exams, estimated from a student-level value-added model. The scale of the value-added estimates is standardized to test score standard deviations. Standard errors are clustered at the school level, and bars indicate the 95% confidence interval.

Table 3.1: Summary statistics

	Sample period	
	2009-2017	2012
Number of students	181811	81870
Number of schools	256	186
Student shares:		
Black	0.143	0.145
Hispanic	0.567	0.580
ELL	0.234	0.225
Gifted	0.117	0.119
Free/Reduced Lunch	0.707	0.725
Female	0.491	0.493
School characteristics:		
Enrollment	440.4 [299.1]	440.2 [309.7]
Traditional	0.477	0.511
Innovation	0.151	0.135
Charter	0.192	0.163
Magnet	0.0793	0.0899

Table 3.2: Demand estimates

	Grade		
	0th	6th	9th
Neighborhood school	2.960*** (0.0354)	2.629*** (0.0410)	2.369*** (0.0391)
Sibling priority	5.617*** (0.116)	4.938*** (0.188)	4.524*** (0.178)
Distance (km)	-0.366*** (0.00925)	-0.222*** (0.00848)	-0.170*** (0.00870)
Black $\times$ Dist.	0.0696*** (0.0146)	0.0379*** (0.0135)	0.0921*** (0.00994)
Hispanic $\times$ Dist.	-0.0189 (0.0123)	-0.0535*** (0.0113)	-0.0240** (0.00957)
FRL $\times$ Dist.	-0.00422 (0.0112)	-0.0366*** (0.0108)	-0.0360*** (0.00812)
Black $\times$ Share Black	3.604*** (0.298)	0.626 (0.412)	1.468*** (0.389)
Hispanic $\times$ Share Hispanic	2.734*** (0.149)	2.340*** (0.159)	2.290*** (0.162)
ELL $\times$ Share ELL	1.575*** (0.277)	1.101*** (0.184)	1.858*** (0.288)
Gifted $\times$ Average achievement	0.920*** (0.210)	0.578*** (0.0838)	0.740*** (0.0994)
FRL $\times$ Average achievement	-1.297*** (0.0887)	-1.047*** (0.0918)	-0.864*** (0.0990)
Number of ranks	453635	135864	82777
Number of students	4985	3672	3599
Log-likelihood	-15921.3	-13576.4	-12069.8
Pseudo $R^2$	0.612	0.512	0.513

Note: The table reports parameter models from estimating Equation 3.1 using students' submitted ranked order lists. The model is an exploded logit model and the parameters are estimated by maximum likelihood. The first two rows are explanatory variables for whether the school is the student's assigned neighborhood school, and whether the student has a sibling attending the school. The final five rows are interactions between student characteristics and average school characteristics. The model is separately estimated by each entry grade. 9th grade is the entry grade for high schools. 6th grade is the entry grade for middle schools. 0th grade (i.e., kindergarten) is the entry grade to elementary schools. Standard errors are in parentheses. Stars indicate the level of significance: \* 10%, \*\* 5%, and \*\*\* 1%.



Table 3.3: Relationship between competitive pressure and concentration

	Outcome: competitive pressure ( $e_j$ )					
	(1)	(2)	(3)	(4)	(5)	(6)
Number of schools	0.00902*** (0.00310)	0.00238*** (0.000841)	0.000761** (0.000340)			
Enrollment share				-0.0383* (0.0220)	-0.185*** (0.0509)	-0.392*** (0.124)
Market radius	2km	5km	10km	2km	5km	10km
Number of obs.	151	151	151	151	151	151
Adjusted $R^2$	0.0476	0.0447	0.0262	0.00859	0.0347	0.0288

Note: The table reports bivariate regressions of competitive pressure on school-level measures of local concentration. The first row is the number of other schools within a given radius that are serving the same grades. The second row is the share of students living within a given radius and within the school's grade range that are currently attending the school. Stars indicate the level of significance: \* 10%, \*\* 5%, and \*\*\* 1%.

Table 3.4: Cross sectional effects of competitive pressure on school outcomes

(a) School effectiveness						
	Outcome: school value-added					
	(1)	(2)	(3)	(4)	(5)	(6)
Competitive pressure ( $e_j$ )	-0.152 (0.128)	-0.203 (0.130)	-0.159 (0.129)	-0.1000 (0.125)	-0.00977 (0.112)	-0.0291 (0.134)
Number of observations	579	579	579	579	277661	509
Number of schools	118	118	118	118	118	104
Adjusted $R^2$	0.00893	0.0302	0.0430	0.0826	0.118	0.0883
Controls:						
Size and grade		Y	Y	Y	Y	Y
Student composition			Y	Y	Y	Y
School sector				Y	Y	Y
Student-weighted					Y	
Dropping interventions						Y
(b) School spending						
	Outcome: share of expenditures on instruction					
	(1)	(2)	(3)	(4)	(5)	(6)
Competitive pressure ( $e_j$ )	0.214** (0.0904)	0.247*** (0.0845)	0.323*** (0.0622)	0.228*** (0.0660)	0.232*** (0.0642)	0.262*** (0.0755)
Number of observations	217	217	217	217	120974	191
Number of schools	112	112	112	112	112	99
Adjusted $R^2$	0.0339	0.165	0.339	0.428	0.490	0.420
Controls:						
Size and grade		Y	Y	Y	Y	Y
Student composition			Y	Y	Y	Y
School sector				Y	Y	Y
Student-weighted					Y	
Dropping interventions						Y

Note: The table reports coefficients from estimating Equation 3.8. The sample is all schools with both outcomes after 2012 and competitive pressure estimates. The school's competitive pressure is estimated from the empirical model of the market in 2012 using Equation 3.7. The school's value-added on state exams is estimated from a student-level value-added model and allowed to vary arbitrarily by year. Student composition controls are controls for the share of black, Hispanic, ELL, gifted, SPED, free/reduced lunch, and female students at the school. The last column drops schools that, at any point during the sample period, have received additional interventions from the district due to underperformance. Standard errors are in parentheses and clustered at the school-level. Stars indicate the level of significance: \* 10%, \*\* 5%, and \*\*\* 1%.

Table 3.6: Pooled differences in differences estimates

(a) School effectiveness						
	Outcome: school value-added					
	(1)	(2)	(3)	(4)	(5)	(6)
Competitive pressure ( $e_j$ ) $\times$ after match	-0.178 (0.149)	-0.118 (0.153)	-0.2184 (0.142)	0.0646 (0.142)	-0.113 (0.134)	-0.0310 (0.151)
Number of observations	1005	1005	906	661	437544	882
Number of schools	118	118	118	118	118	104
Adjusted $R^2$	0.243	0.253	0.277	0.343	0.397	0.278
School sector trends		Y	Y	Y	Y	Y
Student controls			Y			
Dropping post 2014				Y		
Student-weighted					Y	
Dropping interventions						Y

(b) School spending					
	Outcome: share of expenditures on instruction				
	(1)	(2)	(3)	(4)	(5)
Competitive pressure ( $e_j$ ) $\times$ after match	0.178* (0.0967)	0.181** (0.0842)	0.197** (0.0804)	0.204*** (0.0678)	0.158* (0.0856)
Number of observations	322	322	322	180941	283
Number of schools	109	109	109	113	96
Adjusted $R^2$	0.802	0.806	0.810	0.882	0.827
School sector trends		Y	Y	Y	Y
Student controls			Y		
Student-weighted				Y	
Dropping interventions					Y

Note: The table reports pooled coefficients from estimating Equation 3.10. The sample is all schools with competitive pressure estimates. The school's competitive pressure is estimated from the empirical model of the market in 2012 using Equation 3.7. The school's value-added on state exams is estimated from a student-level value-added model and allowed to vary arbitrarily by year. Student controls are controls for the log number of students at the school, as well as the share of black, Hispanic, ELL, gifted, SPED, free/reduced lunch, and female students. The last column drops schools that, at any point during the sample period, have received additional interventions from the district due to underperformance. Standard errors are in parentheses and clustered at the school-level. Stars indicate the level of significance: \* 10%, \*\* 5%, and \*\*\* 1%.

THIS PAGE INTENTIONALLY LEFT BLANK

# Appendix A

## Appendix for Chapter 1

### A.1 Additional Tables and Figures

Figure A-1: Prevalence of Occupations with Clear Promotion Tracks





Figure A-2: Distribution of firm promotion propensities



Note: The figure plots a kernel density estimate of the distribution of firm promotion propensities ( $\eta_j$ ) between 2004 and 2006. Each firm's promotion propensity is calculated as the residual firm promotion rates after controlling for differences in worker characteristics and occupational groups and averaged over the three years. Estimates in the figure have been winsorized at the 5th and 95th percentiles. For more details, see Section 1.4.1.

Figure A-3: Age profile of supervisors and promotions




Figure A-4: Visual intuition for IV first stage




Figure A-5: Impact of high promotion firms on a worker ever being promoted




Figure A-6: Impact of high promotion firms on worker promotions (by cohort)




Figure A-7: Impact of high promotion firms on formal labor market attachment (by cohort)

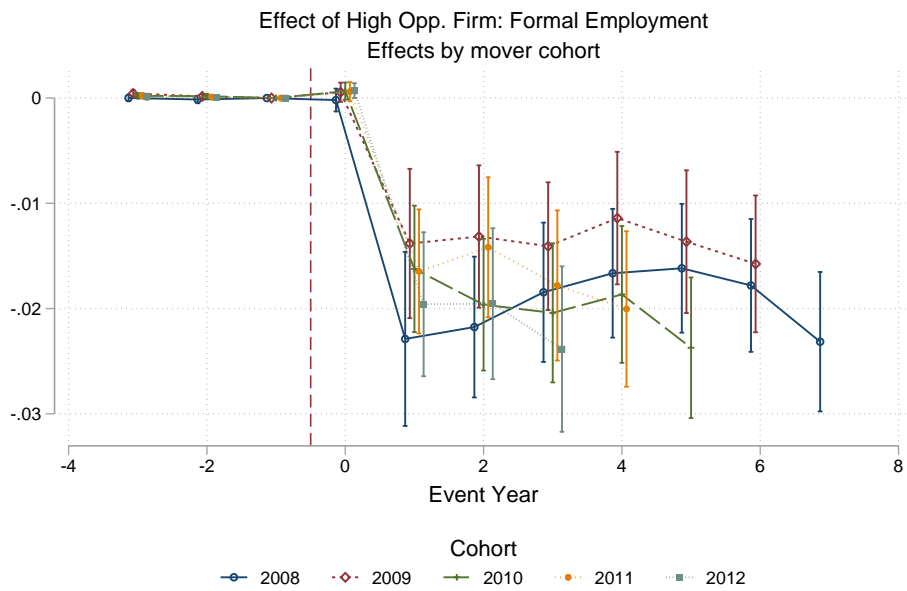




Figure A-8: Impact of high promotion firms on log earnings for employed workers (by cohort)




Figure A-9: Impact of high promotion firms on worker promotions (continuous measure)

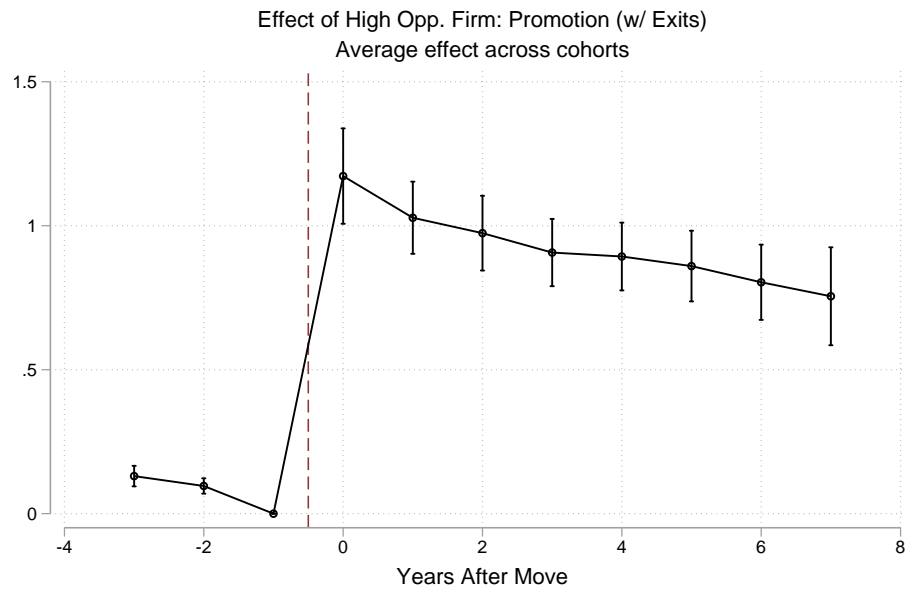


Figure A-10: Impact of high promotion firms on formal labor market attachment (continuous measure)




Figure A-11: Impact of high promotion firms on log earnings for employed workers (continuous measure)




Figure A-12: Impact of high promotion firms on worker promotions (mass layoffs)




Figure A-13: Impact of high promotion firms on formal labor market attachment (mass layoffs)

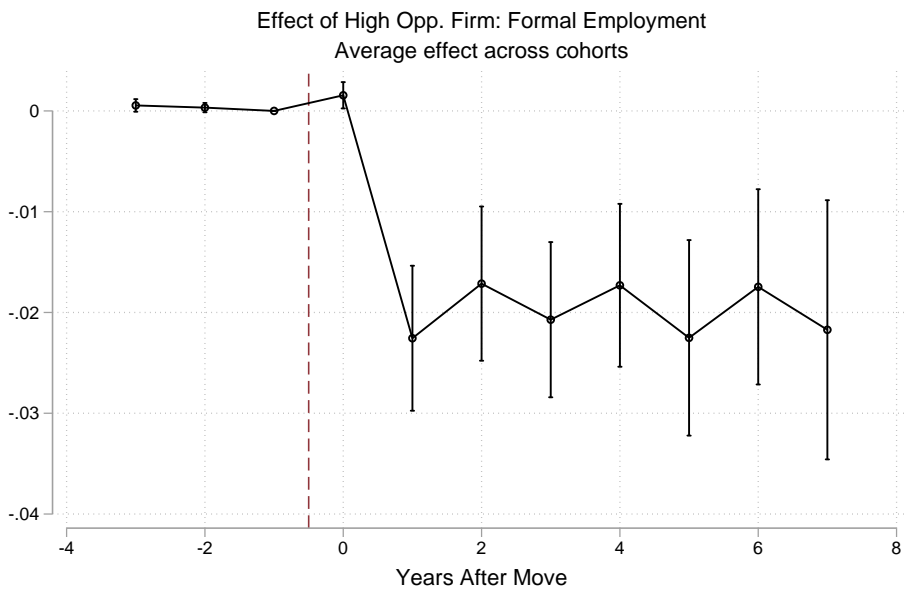


Figure A-14: Impact of high promotion firms on log earnings for employed workers (mass layoffs)



Figure A-15: Impact of high promotion firms on worker promotions (local hiring IV)




Figure A-16: Impact of high promotion firms on formal labor market attachment (local hiring IV)





Figure A-17: Impact of high promotion firms on log earnings for employed workers (local hiring IV)



Figure A-18: Comparison of main effects between alternative firm types

(a) Effects within 2 years of move



(b) Effects more than 2 years after move






Figure A-19: Robustness of model estimates to alternate assumptions



(a) Robustness to alternate kappa



(b) Robustness to alternate delta

Note: The figure shows the model estimates under alternative assumed values on  $\kappa$ , the degree of asymmetric information on the secondary market, and on  $\delta$ , the exogenous job separation rate. Adverse selection refers to the relative share of high ability workers in the pool of job-movers compared to the population share of high ability workers. Misallocation refers to the likelihood that a high ability worker is not working as a supervisor. The model can no longer exactly fit the sample means and treatment effect estimates at 50% of the baseline assumed value for  $\delta$ , so those estimates should be interpreted with some caution. The estimates under the baseline assumed value are indicated by dashed lines.

Table A.1: Correlation between measures of firm promotions

Measure	Correlation Coefficient			
Baseline measure	1			
Stayers only	0.925	1		
No controls	0.947	0.885	1	
All promotions	0.731	0.673	0.700	1

Note: This table reports the correlation matrix between the baseline measure of the firm’s promotion propensity  $\eta_j$  and alternate measures. All measures reflect the firm’s average  $\eta_{jt}$  over 2004-2006 and are winsorized at the 5th and 95th percentiles. The “stayers only” measure restricts the sample to workers who remained in the same firm, so the estimate reflects firms’ internal promotion rates. The “no controls” measure removes any controls and only considers the firms’ raw promotion rate. “All promotions” expands the baseline sample to also include workers who changed occupation groups.



Table A.2: Pooled estimates on the impact of high promotion firms on other worker outcomes

	(1)	(2)	(3)	(4)	(5)
Outcome:	Employed at destination firm	Promoted (cond. on formal emp.)	Earnings net of firm wage premia	Contractual salary	Earnings (adjusted for U.E.)
Effects within 2 years of moving	-0.0263*** (0.0041)	0.0136*** (0.00097)	-0.00625 (0.0036)	0.00964 (0.0071)	0.00799 (0.0055)
Effects more than 2 years after moving	-0.0447*** (0.0082)	0.0132*** (0.0010)	-0.00745 (0.0043)	0.0119* (0.0060)	-0.00665 (0.0058)
Number of observations	8279318	7378436	6640880	6390190	8279252
Number of dest. firms	45870	45870	45654	45496	45870
Adjusted $R^2$	0.617	0.304	0.795	0.776	0.775

Note: The table reports additional pooled event study coefficients from estimating Equation 1.6 on the baseline sample. The outcomes are whether the worker is employed at their original destination firm, whether the worker is promoted (conditional on the worker remaining in formal employment), the worker's log earnings after netting out the AKM firm effect, the worker's contractual salary, and a lower bound for worker earnings that include exiters (adjusting for the value of non-work time). Standard errors are clustered by the worker's destination firm. Stars indicate the level of significance: \* 5%, \*\* 1%, and \*\*\* 0.1%.

Table A.3: Worker potential subsample summary statistics

	Baseline	High Potential	Low Potential
Number of workers	1100590	198732	169589
Number of origin firms	162573	25758	26544
Number of destination firms	49785	26223	26347
Worker characteristics before move:			
Age	32.94 [6.761]	33.11 [6.613]	32.56 [6.774]
Female	0.262 [0.439]	0.284 [0.451]	0.234 [0.423]
Years of schooling	10.50 [2.941]	11.81 [2.716]	9.191 [2.966]
Monthly earnings (2010 Reals)	1242.9 [900.4]	1869.9 [1251.4]	874.5 [328.1]
Share to high promotion firm	0.417	0.435	0.406
Worker outcomes > 2 years following move:			
In formal employment	0.777	0.795	0.759
At destination firm	0.305	0.326	0.286
In supervisor occupation	0.0280	0.0476	0.0125

Note: The table reports summary statistics about the job-to-job movers in the baseline sample and the high and low potential worker subsamples. Pre-move worker characteristics refer to the snapshot of worker data from the year before the job-to-job transition ( $t = b - 1$ ). Outcomes more than two years following the move are averaged over all relevant years where the data are available. All statistics on characteristics and outcomes are in means, and standard deviations for continuous measures are in brackets.

Table A.4: World Management Survey questions and scoring criteria

Question	Score	Criteria
Developing Talent and Promoting High Performers	1	People are promoted primarily upon the basis of tenure
	5	We actively identify, develop and promote our top performers
Removing Poor Performers/Making Room for Talent	1	Poor performers are rarely removed from their positions
	5	We move poor performers out of the company or to less critical roles as soon as a weakness is identified

Note: The table reproduces the question and scoring rubric for the promotion and firing practice questions from the World Management Survey. All survey responses range from 1 to 5, with 5 indicating the most active firm practice and 1 indicating the least active. For more details, see Section 1.6.4.

Table A.5: Estimates of regional differences in labor market attachment

	Outcome: Leave formal labor market (next year)				
	(1)	(2)	(3)	(4)	(5)
High Firm Share ( $\beta_1$ )	0.0294*** (0.00315)	0.0154*** (0.00298)	0.0155*** (0.00298)	0.0213*** (0.00296)	0.00779* (0.00304)
High Firm Share $\times$ Super. ( $\beta_2$ )	-0.0385*** (0.00620)	-0.0382*** (0.00613)	-0.0285*** (0.00574)	-0.0176** (0.00567)	-0.0357*** (0.00602)
Number of observations	10489980	10489910	10489910	10489910	10127824
Number of municipalities	5489	5489	5489	5489	5487
Adjusted $R^2$	0.0233	0.0368	0.0371	0.0373	0.113
Controls:					
State-occupation-year FE	Y	Y	Y	Y	Y
Worker controls		Y	Y	Y	Y
Occupation-super.-year FEs			Y	Y	Y
Municipality controls				Y	
Firm FEs					Y

Note: The table reports estimates for municipal differences formal labor market attachment in the Brazilian formal employment sector (estimated using Equation 1.11). The sample is a 5% sample of all workers between the ages of 25 and 50 that work in occupational groups with a well-defined supervisor track in municipalities with at least 100 workers. The estimands of interest are the likelihood of leaving formal employment for all workers in municipalities with a high share of high promotion employers ( $\beta_1$ ) as well as the differential likelihood for promoted workers in municipalities with a high share of high promotion employers ( $\beta_2$ ). Worker controls are a quadratic in age interacted with gender along with indicators for the worker's education, race, and state. Municipality controls are controls for log employment and the overall hiring share of either high or low promotion firms that enter both linearly and as interactions with supervisor status. Standard errors are clustered by municipality. Stars indicate the level of significance: \* 5%, \*\* 1%, and \*\*\* 0.1%.

Table A.6: Sensitivity of model parameters to empirical moments ( $\Lambda$ )

Parameter	Sensitivity to $\eta$			
$\bar{Q}_H$	8.13	0.95	-19.32	-11.14
$\bar{Q}_L$	7.58	0.68	-19.02	-7.99
$\alpha$	-2.23	-1.08	24.59	11.96
$g$	7.31	3.55	-21.58	-10.49

Note: The ordering of the empirical moments are the average share of workers promoted, the average share of workers in formal employment, the effect of high promotion firms on worker promotions, and the effect of high promotion firms on formal employment. The interpretation of the sensitivity matrix is that for a local perturbation of the empirical moments that converges to  $\eta$ , the first-order asymptotic bias in the parameter estimates are  $\Lambda\eta$ . For more details, see Section A.3.3.

## A.2 Proofs of Propositions

### A.2.1 Partial Equilibrium Model

**Proposition.** *If information about job assignments on the secondary market is sufficiently weak (so  $\kappa$  is sufficiently small), then a unique Perfect Bayesian equilibrium exists. In this equilibrium, (i) job assignments for workers are efficient (given firms' information about workers) (ii) all turnover is involuntary (iii) wages are given by Equations 1.2 and 1.4.*

*Proof.* I characterize the equilibrium by solving the problem backward. I conjecture that the incumbent firms will promote workers that are revealed to be high ability, fire workers that are revealed to be low ability, and retain all workers whose abilities are still unobserved, and I show that this conjectured behavior is optimal given secondary market wages. I then prove uniqueness by showing that this conjectured behavior is the only one consistent with an equilibrium.

For separated workers who successfully convinced the secondary market that they were previously promoted, the secondary market will infer that they were high ability and offer the expected output

$$w_2^S = f_2(\theta_H).$$

Meanwhile, the likelihood that a low ability worker enters the unknown secondary market workers pool is

$$\delta + (1 - \delta) \underbrace{\left( \rho Q_H + (1 - \rho) Q_L \right)}_{\text{prob. fired}} \equiv \delta + (1 - \delta) \bar{Q},$$

whereas the likelihood that a high ability worker enters the unknown worker pool is  $\delta(1 - \bar{Q}\kappa)$ . By Bayes' rule,

$$\begin{aligned} \alpha' &\equiv Pr(\theta = \theta_H | \text{unknown quality}) \\ &= \frac{\alpha \delta (1 - \bar{Q}\kappa)}{\alpha \delta (1 - \bar{Q}\kappa) + (1 - \alpha) (\delta + (1 - \delta) \bar{Q})} \\ &< \alpha. \end{aligned}$$

Since low ability workers are both more likely to be fired and less likely to convince new employers that they are high ability, they are going to comprise a higher share of the secondary market workers of unknown quality than in the general population. So

the secondary market is adversely selected. The degree of adverse selection is higher when there is greater learning by initial employers ( $\bar{Q}$ ) or by the secondary market ( $\kappa$ ), and lower when there are more exogenous separations ( $\delta$ ). Assuming that it is still optimal for secondary market firms to assign workers of unknown quality to the low complexity occupation, offered wages for these workers will be their expected output, so

$$w_1^S = \alpha' f_1(\theta_H) + (1 - \alpha') f_1(\theta_L).$$

Next, I can then verify that the conjectured job assignment by incumbent firms is indeed optimal given the secondary market offers  $w_1^S, w_2^S$  and characterize the incumbent firms' wage offers.

The incumbent firms earn positive profits from promoted workers of known high ability regardless of  $g$  because the worker is not guaranteed to remain promoted in the secondary market. So, the incumbent firm needs to offer a promoted worker only their expected outside wage to retain them:

$$w_2^I = g [\kappa w_2^S + (1 - \kappa) w_1^S] < g f_2(\theta_H).$$

Similarly, the incumbent firm also benefits by retaining the workers whose ability it did *not* observe since the firm has private information that those workers are not adversely selected. For those workers, the probability that they are the high type is still  $\alpha$ , so the firm makes positive profits by retaining those unknown workers and offering them the expected outside wage for workers of unknown ability:

$$w_1^I = g [\alpha' f_1(\theta_H) + (1 - \alpha') f_1(\theta_L)] < g E[f_1(\theta) | \alpha].$$

Finally, for workers who are observed to be low ability,  $f_2(\theta_L) < f_1(\theta_L) < 0$ , so the firm is better off firing them rather than retaining them with any positive wage.

For this equilibrium to exist, two conditions need to hold. First, the expected productivity for unknown workers in the secondary market needs to be positive, so

$$\underbrace{\alpha' f_1(\theta_H) + (1 - \alpha') f_1(\theta_L)}_{=E[f_1(\theta)|\alpha']} \geq 0.$$

This condition also implies that the incumbent firm will also find it optimal to retain workers whose ability is unobserved, since  $\alpha > \alpha'$ . In addition, the wage gain upon promotion needs to be sufficiently low for the incumbent firm to still prefer to promote

workers of high ability, so

$$f_2(\theta_H) - w_2^I \geq f_1(\theta_H) - w_1^I.$$

Rearranging using the expression for wages,

$$\begin{aligned} f_2(\theta_H) - f_1(\theta_H) &\geq w_2^I - w_1^I \\ &= g[\kappa w_2^S + (1 - \kappa)w_1^S - w_1^S] \\ &= g\kappa(f_2(\theta_H) - E[f_1(\theta) | \alpha']). \end{aligned}$$

$\alpha'$  is decreasing in  $\kappa$  (since increasing the informativeness of job assignments reduces the likelihood that a high ability worker joins the unknown workers pool), so the right hand side of the inequality is also decreasing in  $\kappa$ . So, decreasing  $\kappa$  relaxes both inequality constraints, and both conditions clearly hold when  $\kappa \rightarrow 0$  as long as

$$E[f_1(\theta) | \alpha' = \alpha\delta / (\delta + (1 - \alpha)(1 - \delta)\bar{Q})] \geq 0,$$

so the equilibrium exists whenever  $\kappa$  is sufficiently small.

Finally, I show that if  $\kappa$  is sufficiently small such that the equilibrium exists, it is also the unique equilibrium. To do so, it suffices to show that for the incumbent firm, firing low ability workers, promoting high ability workers, and retaining unknown workers is the only strategy that is consistent with an equilibrium.

Clearly, there can be no equilibrium that exists where the incumbent firm retains all workers since the output of workers with  $\theta = \theta_L$  is strictly less than the outside option. The firm will also never retain low ability workers but fire unknown workers, since the market wages for those two types are identical and the expected productivity of the latter strictly dominates the former.

Similarly, there can be no equilibrium where the incumbent firm retains only high ability workers. Suppose the incumbent firms retain high ability workers and fire all other workers. This implies that expected outside wages exceed expected productivity for workers whose ability is unknown, which cannot be consistent with the competitive secondary market since the secondary market is adversely selected.

So, it suffices to consider whether an equilibrium exists where the incumbent firm treats high ability workers differently. There can be no equilibrium where the incumbent firm will fire high ability workers, since their expected outside wage will always be weakly less than  $f_2(\theta_H)$ . Similarly, if the incumbent firm keeps all high ability workers in the low complexity job, secondary market wages for workers of



unknown quality,  $w_1^S$ , rises, while secondary market wages for any promoted workers would remain the same. This strictly relaxes the incentive compatibility problem for promotions, so the incumbent firm would profitably deviate by promoting high ability workers instead.

It then follows that under the conditions where the conjectured equilibrium exists, the equilibrium where the incumbent firm promotes high ability workers, fires low ability workers, and retains workers of unknown ability in the low complexity job is also the unique equilibrium.  $\square$

**Proposition.** *In the equilibrium described in Proposition 1, workers initially employed at high learning firms are (i) more likely to be promoted and (ii) more likely to become unemployed than workers initially employed at low learning firms.*

*Proof.* This follows directly from the characterization of the equilibrium. The incumbent firm will promote all high ability workers and fire all low ability workers. Since the composition of workers and labor market parameters are the same between high and low learning firms, the results follow from differences in learning rates.

Specifically, the likelihood that a worker starting at firm  $f$  will be promoted by the end of the period is

$$\Pr(\text{Promoted}|Q_f) = \alpha Q_f (1 - \delta + \delta g \kappa),$$

and the likelihood that a worker will be be unemployed is

$$\Pr(\text{Unemployed}|Q_f) = (1 - g) [\delta + (1 - \alpha) (1 - \delta) Q_f].$$

Both likelihoods are increasing in  $Q_f$ .  $\square$

## A.2.2 General Equilibrium Model

**Proposition.** *Suppose (i) the marginal cost for vacancies  $c'(v)$  is log-concave and (ii) the elasticity of re-employment with respect to the number of high learning employer vacancies exceed the elasticity of expected secondary market worker output. Then, in more productive regions: (i) high learning employers post a greater share of vacancies, (ii) offered wages for incumbent workers are higher (iii) wage differentials between promoted and unpromoted workers are larger.*

*Proof.* Under the assumptions on labor market matching, the market level of high learning firm vacancies  $V_H$  and low learning firm vacancies  $V_L$  imply the labor market

parameters  $\rho$  and  $g$  as:

$$\rho = \frac{V_H}{V_H + V_L}$$

$$g = \frac{M(m^I, V_H + V_L - m^I)}{m^I},$$

where  $M(l, v)$  is the reduced form matching function that relates the level of job matches to the number of workers  $l$  and vacancies  $v$ , and  $m^I = (l, V_H + V_L)$  is the number of initially matched workers. Conditional on the number of vacancies (and their implied labor market parameters), the job assignment, offer, and secondary market decisions will follow the characterization in Proposition 1. So, the firm's expected profits from a vacancy filled with a new worker is

$$E[\pi_f | V_H, V_L] = (1 - \delta) \left\{ Q_f \alpha \left( \underbrace{f_2(\theta_H) - w_2^I}_{\text{profits from promoted}} \right) + (1 - Q_f) \left( \underbrace{E[f_1(\theta) | \alpha] - w_1^I}_{\text{profits from unpromoted}} \right) \right\},$$

where  $V_H, V_L$  are the total high and low learning firm vacancies in the market, respectively, and all other objects match their definition from Proposition 1. The firm profits only upon successfully retaining an initially assigned worker. Any possibility that a vacancy is filled in the secondary market is irrelevant since the secondary market is competitive.

In the vacancy creation problem, the firm solves

$$\max_v h(V_H, V_L) E[\pi_f | V_H, V_L] v - c(v),$$

where  $h(V_H, V_L) = \frac{M(l, V_H + V_L)}{V_H + V_L}$  is the probability that a vacancy will be initially matched with a worker. Note that the firm's choice of  $v$  does not affect  $h(V_H, V_L) E[\pi_f | V_H, V_L]$  under the assumption that each firm is atomistic. The first order condition to the vacancy problem equates the marginal cost of the vacancy to the expected equilibrium profits:

$$c'(v) = h E[\pi_f | V_H, V_L],$$

so the equilibrium share of vacancies from high learning firms is:

$$s = \frac{v(h E[\pi_H | V_H, V_L])}{v(h E[\pi_H | V_H, V_L]) + v(h E[\pi_L | V_H, V_L])},$$

where  $v(\pi)$  is the vacancy supply function given (expected) profits  $\pi$ .  $s$  is always

increasing in  $\psi$  if and only if

$$\frac{v'(\lambda\psi)\lambda}{v(\lambda\psi)} > \frac{v'(\psi)}{v(\psi)}$$

for all  $\lambda > 1$ . A sufficient (but not necessary) condition for this is if  $v$  is log-convex, which is satisfied if  $v = e^\pi$  (i.e.,  $c'$  is log), or if  $v(\pi) = \pi^{-a}$  for some  $a > 0$ .

Define firm's solution to the optimal vacancy problem as  $v_f^*(V_H, V_L)$ , where  $(V_H, V_L)$  are the total number of vacancies posted by high learning and low learning firms in the market. The equilibrium is the fixed point where  $v_f^*(v_H, v_L) = v_f \forall f \in \{H, L\}$ . To see that this fixed point exists under my assumptions, I can rewrite the maximization problem as

$$\max_v h(v_H + v_L) E[\pi(v_H, v_L)]v - c(v),$$

noting that the market level of vacancies imposes two negative externalities for the firm. First, excess vacancies increase the match rate on the secondary market, which drives up wage competition for incumbent workers. Second, excess vacancies lower the initial fill rate for new workers. There is also an offsetting force, where the share of firms that are high learning determine the degree of adverse selection on the secondary market. Expected profits are clearly decreasing in the total number of vacancies due to the first two mechanisms, so it suffices to ensure that the net effect of adding high learning employer vacancies does not induce unraveling (e.g., the adverse selection effect is not so strong as to push the slope of the best response curve for high learning vacancies above 1).

To be precise about this condition, notice that clearly  $h(v_H + v_L)$  is decreasing in the market  $v_H$ . So, a sufficient (but not necessary) condition for the best response function to be decreasing in  $v_H$  is if

$$\frac{\partial E[\pi(v_H, v_L)]}{\partial v_H} \leq 0.$$

Since the expected profits for each filled vacancy is the weighted average of the profits from a worker of known high ability and the profits from the average worker (and neither  $\alpha$  or the firm-level  $Q$  depends on the market characteristics), it suffices to consider whether

$$\frac{\partial \pi_1}{\partial v_H}, \frac{\partial \pi_2}{\partial v_H} \leq 0 \iff \frac{\partial w_1}{\partial v_H}, \frac{\partial w_2}{\partial v_H} \geq 0.$$

Secondary market profits are zero in any equilibrium, so the conditions that ensure firm profits are sufficiently well behaved are also exactly the conditions that ensure

offered wages are increasing. Now,

$$\begin{aligned}\frac{\partial w_1}{\partial v_H} &= \frac{\partial g}{\partial v_H} E_{\alpha'} [f_1(\theta)] + g \frac{\partial E [f_1(\theta) | \alpha']}{\partial v_H} \\ \frac{\partial w_2}{\partial v_H} &= \frac{\partial g}{\partial v_H} [\kappa f_2(\theta_H) + (1 - \kappa) E [f_1(\theta) | \alpha']] + g (1 - \kappa) \frac{\partial E [f_1(\theta) | \alpha']}{\partial v_H}.\end{aligned}$$

Observe that the wages of promoted workers will be more insulated from adverse selection in secondary market than the wages of unknown workers since their expected outside option includes the possibility of obtaining their true product, so a sufficient condition for both derivatives to be positive is for  $\frac{\partial w_1}{\partial v_H} \geq 0$ . A simple rearrangement of the derivative then yields the assumed condition in the proposition:

$$-\frac{\partial E [f_1(\theta) | \alpha']}{\partial v_H} / E [f_1(\theta) | \alpha'] \leq \frac{\partial g}{\partial v_H} / g.$$

Conditional on the equilibrium existing, it's straightforward to show that increasing the productivity term  $\psi$  increases the best response to any market-level vacancy, so the equilibrium number of vacancies will increase as well.

Similarly, the assumptions on secondary market matching are sufficient to ensure that the probability of receiving a secondary market offer,  $g$ , is increasing in the number of initial vacancies. So clearly, workers' likelihood of re-employment is higher in more productive regions. Meanwhile, incumbent wages increasing with high learning employer vacancies is exactly the sufficient condition that guarantees  $\partial E [\pi(v_H, v_L)] / \partial v_H \leq 0$ , so it follows that the same conditions that ensure the existence of the vacancy creation problem also ensures that occupational wages at the incumbent employers are weakly increasing as well.  $\square$

## A.3 Additional Empirical Details

### A.3.1 Data Construction

#### Annual worker panel

My primary data on worker earnings, job characteristics, and employer characteristics come from the universe of formal employment contracts in the RAIS data. Each observation in the raw data is a single employment contract within a state and year, so my first step is to construct an annual panel of workers' employment histories. Each observation in the annual panel is a worker's primary employment contract for

that year, and the resulting dataset serves as the basis for all subsamples and derived measures in my project.

To construct the unique worker-by-year panel of primary employment, I consider all employment contracts that covered at least six months over the year, entailed at least 20 contracted hours of work per week, and paid non-zero earnings. In the cases when there are multiple recorded employment contracts for the same worker and year that satisfy these selection criteria, I choose the employment contract covering the longest duration (in months), and I break any subsequent ties by selecting the contract with the highest average monthly earnings.

My preferred earnings measure is the average nominal monthly earnings over the employment contract. Where defined, this measure is highly correlated with the December monthly earnings measure that has been used in the literature (the correlation coefficient between these measures in logs is above .97), but average monthly earnings the additional advantage of being defined for partial employment spells in a year that ended before December.

I classify any years when a worker is not in the annual panel as years when the worker is out of the formal labor market. Correspondingly, I consider the worker to not be working as a supervisor in those years. I impute worker characteristics in years when the worker is not in the annual panel by using the last known observation (for gender, education, state, race, and birth year). Meanwhile, I impute a worker's counterfactual earnings by annually compounding the worker's last known earnings by the average wage growth for the worker's last known state-by-baseline-occupation-group. Finally, I assign a reason for separation to the out-of-formal-labor-market spell using the reason for separation field of the most recent employment contract. Employer-initiated separations are employer terminations with or without just cause (excluding contract expirations). Worker-initiated separations are voluntary worker separations. I combine separations for any other known reason (including contract expirations) into a single category.

### **Classifying promotions**

I use the CBO-02 occupation codes recorded in the RAIS for each employment contract from 2003 onwards to define promotions. The occupation codes follow a consistent hierarchical structure, so I can define supervisory jobs from the structure of the occupation codes themselves. As a check, I can also define supervisor jobs from the text of the job titles (by finding all job titles that contain the term *supervisores*). The results from these two classification methods are identical, which is reassuring about

the consistency of the occupation classification system.

The CBO-02 system classifies all occupations into a 6-digit occupational code. The first two digits of the occupation code indicate the main occupation group, which are generally broad classes of jobs like metalworkers, textile workers, or public services workers. Within production-level occupational groups, a third digit of “0” in the occupation code is reserved for the supervisors in the occupational group, whereas all other values refer to other sub-groups within the occupation that do not necessarily have a clear vertical interpretation relative to each other. Individual occupations are further differentiated by the three additional digits following these three base digits.

Within production-level occupation groups (CBO-02 codes starting with 41-99), all but four occupational groups contain supervisor occupations. On the other hand, none of the civil, managerial, professional, or technical-level occupational groups (CBO-02 codes starting with 01-39) contain supervisor occupations. So, occupational groups with observable lines of progression can be considered to be a proper (but nearly complete) subset of the production worker-level (*trabalhadores*) occupation groups.

### **Estimating promotion propensity**

As discussed in Section 1.4.1, I classify firms as high or low promotion firms based on their composition-adjusted promotion rates between 2004 and 2006. Specifically, for each of the three years  $t$ , I consider all workers from the annual panel between the ages of 25 and 50 who were in formal employment at a non-public sector firm in years  $t - 1$  and  $t$ . I further restrict the sample to all workers who remained in the same broad occupation group in both years, which bolsters the interpretation that these promotions reflect vertical job changes. I estimate  $\eta_{jt}$  on the sample using Equation 1.5 as the residual promotion rate for firm  $j$  at time  $t$  after adjusting for worker characteristics and differences in promotion rates in different occupation groups. Finally, to minimize measurement error or the contribution of year-specific shocks, I restrict the set of firms to those that had at least 10 workers in the estimation sample for each of the three years, and I define  $\eta_j = E[\eta_{jt}]$  as the firm’s average promotion residual over those three years.

It’s worth noting two additional details implicit in the baseline approach. I do not restrict the sample to workers who are at the same firm in both years, so some promotions in the data are external promotions where a worker was working as a line worker in a firm in year  $t - 1$  and as a supervisor in a different firm in year  $t$ . Furthermore, in the case of these external promotions, I attribute the promotion to the firm where the worker received the promotion (i.e., the firm in year  $t$ ), which

is consistent with my interpretation and model. However, although the inclusion of external promotions slightly increases power, these two details are inconsequential for my results. Table A.1 compares the correlation in firms’ promotion rates across a variety of alternate classification methods. The measures are highly correlated with each other. As an additional check, Figure A-18 compares the treatment effect estimates when I classify firms based on their internal promotion rates – the effects on promotions and turnover are slightly attenuated but otherwise similar.

### Measuring mass layoffs

I follow the literature on using linked employer-employee data to identify mass-layoff events. I first compile a firm-level panel of employment counts by aggregating the worker-level annual panel to the firm-year level, and I identify large employment drops or firm closures using the criteria from Schmieder et al. (2020). Mass-layoff events are when a firm with at least 50 employees experience at least a 30% drop in employment or disappear from the data altogether in the following year. Firm identifiers are not always longitudinally consistent, so reorganizations or spinoffs may be mistakenly classified as layoff events. I follow the literature to exclude these alternate scenarios by dropping any layoff event where at least 20% of displaced workers go to the same firm in the following year.

### Constructing the local hiring IV

I construct a panel of total employment and new hiring at the firm-by-municipality level by aggregating the worker-level annual panel. Since employment contracts specify both the firm and location of the employment establishment, the mapping from workers to the firm-municipality is clear. New hires are defined as the total number of workers who are working in the firm-municipality and were working at a different firm in the year prior. So, this measure excludes within firm transfers across municipalities, as well as any brief employment spells that would not be classified as the worker’s primary employment for the year. To ensure that the hiring shares are informative, I restrict my attention to municipalities that have at least 1000 workers and at least 200 new hires each year.

For a worker  $i$  in mover cohort  $c$  and municipality  $m$ , I calculate their jack-knife local hiring share instrument as

$$z_{mc} = \frac{G_{mc}^H - \sum_{J(i')=J(i)} H_{i'}}{G_{mc}^H + G_{mc}^L - \sum_{J(i')=J(i)} 1},$$

where  $G_{mc}^H, G_{mc}^L$  are the total new hires by high and low promotion firms, respectively, and  $i'$  are other movers in the same cohort. Since workers in my analysis sample are included in the new hire totals, I avoid the reflection problem that would arise from this functional dependence by excluding *all* new hires from worker  $i$ 's destination firm from the numerator and denominator. As a result, the interpretation of  $z_{mc}$  is the local hiring share by high promotion firms *excluding* the worker's destination firm. Technically,  $z_{mc}$  varies by destination firm due to the jack-knife procedure. However, this variation is minor, so I slightly abuse notation to focus on the main source of variation.

### A.3.2 Calculating Average Treatment Effects

I combine workers from multiple cohorts to increase the precision of my estimates and to ensure that I am capturing an average treatment effect that is representative across cohorts. However, researchers have cautioned that pooling treatment effects in designs with staggered treatment timing may yield unintuitive and potentially negative weighting of the underlying treatment effects.

To address these concerns and make the relevant comparisons clear, I allow all coefficients in the estimation equation to vary arbitrarily with the worker's cohort. This clearly emphasizes that all identification of treatment effects over time come solely from comparisons between workers who are in the same cohort (e.g., I compare workers who moved to a high promotion firm in 2009 to workers who moved to a low promotion firm in 2009). Furthermore, I can combine estimated treatment effects across cohorts using explicitly specified weights to estimate an average treatment effect across all mover cohorts.

I define my estimand of interest as the average effect across all cohorts. Correspondingly, the combined estimate of the cohort specific treatment effect  $\beta_{c\tau}$  at event time  $\tau$  is

$$\beta_{\tau} = \left[ \sum_{c=2008}^{2012} \beta_{c\tau} \right] / 5.$$

In the data, the number of workers in each cohort varies slightly, so an alternative is to weigh each cohort's treatment effect by the number of workers in the cohort. However, using uniform weighting across cohort years is more straightforward, and ensures that any differences in estimates across subgroups are not driven by any differences in the composition of workers across cohorts.

In practice, the difference between all of the possible approaches is small. Estimating



Equation 1.6 by pooling the treatment effect coefficients yield similar estimates. This is due to two reasons. First, cohort-specific treatment effects are already reasonably similar. In addition, the share of workers moving to high versus low promotion firms each year is also stable, so the OLS weights are roughly comparable to the uniform weights.

### A.3.3 Structural Quantification

#### Estimation method

The model yields nonlinear expressions for the means and treatment effects of high learning firms on promotions and turnover. I match these expressions to their empirical analogs using classical minimum distance, which requires numerically optimizing the nonlinear objective function in Equation 1.12. The objective function is straightforward to compute given that the model expressions are in closed form, yet numerical optimization can run the risk of hitting local rather than the global minima of the objective function.

I use the following algorithm to calculate my model parameters. First, I draw a starting guess for the parameters  $\theta = (\bar{Q}_H, \bar{Q}_L, \alpha, g)$  randomly from a uniform distribution. I then numerically minimize the objective function using the Nelder Mead algorithm from the R package `nloptr` with a stopping criterion for the relative change of  $10^{-12}$  and a constraint that each parameter lies on the interior between 0 and 1. I repeat the process 100 times, drawing a new random starting value each time, and I select the solution with the smallest objective across all the starting value draws.

The nonlinear model performs reasonably well in my setting. Across the 100 different starting values, the numerical algorithm reaches an objective below  $10^{-5}$  in 68 cases. The maximum standard deviation for any parameter estimate across these 68 cases is approximately  $10^{-13}$  and the mean objective is approximately  $10^{-22}$ .

#### Sensitivity of parameter estimates to moments ( $\Lambda$ )

Although the estimation of the parameters requires optimizing a nonlinear function, the choice of the minimum distance estimator ensures that all identification for the model parameters ultimately comes from the four empirical moments (conditional on the calibrated parameters). To help increase the transparency of the model estimates, I report the *sensitivity* of the model parameters to the matched moments, as defined by Andrews et al. (2017), in Table A.6.

Formally, the sensitivity measure in the classical minimum distance estimator is

$$\Lambda = (G'WG)^{-1}G'W,$$

where  $W$  is the chosen weighting matrix and  $G$  is the Jacobian of the model equations  $h(\theta)$  from Equation 1.12. The interpretation is that for a local perturbation of the empirical moments that converges to  $\eta$ , the first-order asymptotic bias in the estimated parameters is

$$E[\tilde{\theta}] = \Lambda\eta.$$


For more details, see Proposition 2 of Andrews et al. (2017). As expected, the estimated parameters are most sensitive to the treatment effect estimates of high promotion firms, particularly their effects on promotions.

# Appendix B

## Appendix for Chapter 2

### B.1 Appendix Figures


Figure B-1: Number of Unique Case Numbers Across Datasets versus NLRB Annual Reports



*Note:* This figure plots the total number of unique NLRB election cases each year in our data set and in the annual NLRB reports. These include all case types (e.g., ‘RC’ cases and non-RC cases) Our data set is from combining union election datasets from Henry Farber, J.P. Ferguson, and Thomas Holmes and publicly available data from the NLRB and picking one observation for each NLRB case number. See Appendix B.3 for details on our data construction process.

Figure B-2: Election Vote-Share Histogram, 50 + Vote Elections


### Density of Union Certification Elections




*Note:* This plots the vote-share histogram of elections with more than 50 total voters. Given the discreteness of the running variable and the fact that our sample includes elections with a small number of votes, it is difficult to detect manipulation from the vote-share density figure for the entire sample so we restrict the sample to elections with at least 50 votes. The figure was constructed using external union election data (e.g., not our final sample matched to the Census) but the sample was constructed to mirror the overall sample construction (see Appendix B.3 for details). Note, there may still be a small bias from the “integer problem” described in DiNardo and Lee (2004) that could lead to an excess mass of elections right below 50 % but simulations suggest that it is quite small with at least 50 votes.

Figure B-3: Log Employment and Payroll Estimates, 20-80 % Vote-Share Elections

Panel A. All Union Elections




Panel B. Manufacturing Union Elections



Note: This figure plots estimates from the *Flexible Controls* specification presented in Figure 2-3 Panel B. and Figure 2-4 Panel B. The log employment estimates are identical to the estimates in Figures 2-3 and 2-4 but the log payroll estimates are not otherwise reported.


Figure B-4: DHS Employment Estimates, 20-80 % Vote-Share Elections, 10 Yr Pre- and Post-Periods




*Note:* This figure plots the same DHS employment growth rate estimates as in Figure 2-3, Panel C and Figure 2-4, Panel C but includes the -10 to -5 pre-period estimates and the 6 to 10-year post-period estimates. Note, the panel is balanced from -5 years pre-election to 10 years post-election but not from -10 to -5 years pre-election. Consequently, each of the -5 to -10 point estimates average over slightly different cohorts.

Figure B-5: Election Win Rates and Challenged Vote Rates by Delay Time

Panel A. All Elections




Panel B. Manufacturing



*Note:* This figure plots the relationship between pre-election delay times, election win rates, and challenged votes in elections. Pre-election delay times are defined as the number of days between the election petition being filed and the election date. We then take the within-year percentiles of the election delay distribution and plot this on the x-axis. The share of elections with a challenged vote is defined as an indicator for any vote in the election being challenged. The sample of elections includes all elections in our “external elections dataset” described in Appendix B.3. The conditional regression coefficients are from regressing the election win indicator (or challenged vote indicator) on deciles of the number of eligible voters in the election, four-digit NAICS industry fixed effects, and election state fixed effects.

Figure B-6: Establishment-Level Total Factor Productivity Heterogeneity, Multi-Establishment Firms



*Note:* This figure plots the same estimates as in Figure 2-11 except restricting the TFP comparison to only be between multi-establishment firms in different TFP terciles. As with the other heterogeneity tests, the sample includes all manufacturing elections and pools the controls across the entire sample.



## B.2 Appendix Tables

Table B.1: Union Election Matched Sample Construction

	All Elections		Winning Elections	
	Elections	Eligible Voters	Elections	Eligible Voters
<b>Panel A: NLRB Election Sample</b>				
All Election, 1981-2005	94,824	5,991,865	44,492	2,461,138
Representation Elections (RC)	77,349	5,111,675	39,397	2,071,859
> 5 Eligible Voters	69,789	5,084,061	34,247	2,053,210
Non-Contested Elections	66,353	4,590,121	31,378	1,668,877
<b>Panel B: Final NLRB Sample Industry Shares</b>				
Manufacturing	0.307	0.408	0.253	0.307
Other	0.266	0.186	0.263	0.177
Services	0.426	0.405	0.484	0.515
<b>Panel C: Matched Census Sample</b>				
Elections Matched to Census Establishments	46,000			
Final Establishment-Level Outcome Sample	27,000			
20-80 % Election Sample	19,000			

*Note:* This table illustrates how our specific sample restrictions change the number of elections and eligible voters we have in our sample. Panel A plots the total number of elections and eligible voters for all elections and specifically for winning elections. The first row in Panel A. includes all unique NLRB cases with filing dates between 1981-2005 (the main years in our sample). The second row only includes representation (RC) elections. The third row drops elections without more than five eligible voters. The fourth row only includes non-contested elections (e.g., elections with one union on the ballot). Panel B. presents the industry composition of the remaining elections from the fourth row of Panel A. Note we use the NLRB election industry codes here rather than the LBD industry codes but the overall industry shares are reassuringly similar to the industry shares in Table 2.4. The three columns represent the total shares of elections and eligible voters for all elections and winning elections. Panel C shows our final sample sizes from the matched Census data. The sample restrictions between "Elections Matched to Census Establishments" and "Final Establishment-Level Outcome Sample" include keeping (1) the first election at each establishment, (2) at least three years of pre-election survival, (3) non-missing employment, payroll, and other controls at event time  $t = -1$ .

Table B.2: Post-Election Outcome Trends by Vote Share, 20-80 % Vote-Share Elections, Employment and Industry Ctrls.

Industry Group: Outcome:	All Industries		Manufacturing	
	DHS Emp	Survival	DHS Emp	Survival
<i>3-Year Post Election</i>				
Event-Time $\times$ 0-50 % Vote Share	-0.134 (0.100)	0.021 (0.037)	-0.181 (0.159)	0.042 (0.059)
Event-Time $\times$ 50-100 % Vote Share	-0.361*** (0.126)	-0.052 (0.051)	-0.543** (0.250)	-0.035 (0.099)
<i>5-Year Post Election</i>				
Event-Time $\times$ 0-50 % Vote Share	-0.119 (0.116)	-0.009 (0.047)	-0.150 (0.187)	0.023 (0.076)
Event-Time $\times$ 50-100 % Vote Share	-0.450*** (0.141)	-0.085 (0.060)	-0.537* (0.275)	-0.033 (0.116)
<i>10-Year Post Election</i>				
Event-Time $\times$ 0-50 % Vote Share	-0.218 (0.133)	-0.052 (0.057)	-0.186 (0.224)	-0.020 (0.097)
Event-Time $\times$ 50-100 % Vote Share	-0.354** (0.157)	-0.107 (0.070)	-0.676** (0.309)	-0.209 (0.140)
Exclude 50 % Elections	X	X	X	X
Industry + Employment Ctrls. Flexible Ctrls.	X	X	X	X
Number of Elections	19,000	19,000	6,000	6,000

*Note:* This table presents the same estimates as in Tables 2.3 but only includes the baseline industry and employment controls. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table B.3: Employment and Survival Bargaining Unit Share Interaction, 20-80 %  
Vote-Share Elections

Outcome:	DHS Employment	Survival
3-Year Post Election $\times$ Bargaining Unit Share	-0.109** (0.044)	-0.046*** (0.017)
5-Year Post Election $\times$ Bargaining Unit Share	-0.132*** (0.051)	-0.041* (0.021)
10-Year Post Election $\times$ Bargaining Unit Share	-0.057 (0.057)	-0.015 (0.025)
Industry + Employment Ctrl.	X	X
Flexible Ctrl.	X	X

*Note:* This table presents estimates from the same specification as Figure 2-3 for DHS employment growth rates except that we add (1) an interaction between the event-time  $\times$  win indicators with the share of the establishment's employment covered by the bargaining unit and (2) an interaction just between event-time indicators and the bargaining unit share. We report the interactions in (1) for three, five, and ten years post-election. Consequently, this specification estimates how treatment effects increase with the bargaining unit share, accounting for overall post-election trends across all elections by bargaining unit share. A survival estimate of -0.05 means that increasing the share of the establishment covered by the bargaining unit by 10 % leads to an additional 0.5 pct. pct. increase in establishment exit.

Table B.4: Manufacturing versus Services Employment and Survival Estimates, Robustness Checks

Specification:	Baseline		Pooled Controls		Good Matches		> 25 % Barg Unit Share		30-70 %	
	DHS Emp	Survival	DHS Emp	Survival	DHS Emp	Survival	DHS Emp	Survival	DHS Emp	Survival
5-Year Difference	-0.117*** (0.037)	-0.021 (0.016)	-0.118*** (0.035)	-0.022 (0.015)	-0.144*** (0.044)	-0.024 (0.018)	-0.121*** (0.041)	-0.028* (0.017)	-0.132*** (0.043)	-0.026 (0.018)
10-Year Difference	-0.172*** (0.043)	-0.058*** (0.019)	-0.159*** (0.04)	-0.05*** (0.018)	-0.196*** (0.05)	-0.061*** (0.022)	-0.193*** (0.047)	-0.06*** (0.021)	-0.171*** (0.049)	-0.054** (0.022)
Industry + Employment Ctrls.	X	X	X	X	X	X	X	X	X	X
Pooled Ctrls.			X	X						
Flexible Ctrls.	X	X			X	X	X	X	X	X

*Note:* This table presents robustness results for the differences between the service-sector and manufacturing results in Table 2.4. Specifically, it presents the differences between the five- and ten-year DHS employment growth rate and survival estimates for various alternative specifications. The first two columns present the differences for the estimates presented in Table 2.4. The "Pooled Controls" columns pool the controls across all cohorts as described in Section 2.4. The "Good Matches" columns restrict to election matches which we give a 95 % rating (see Appendix B.3 for details). The "Barg Unit Share" columns restrict to elections where the bargaining unit is at least 25 % of the total establishment employment. The 30-70 % columns restrict to elections with 30-70 % of the vote share. For all specifications with restrictions, we still use the entire sample for controls but restrict the treated variables to be estimated from the restricted sample. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table B.5: Single- Versus Multi-Establishment Firm Heterogeneity, Robustness Checks

Specification:	Baseline		Pooled Controls		30-70 %	
	DHS Emp	Survival	DHS Emp	Survival	DHS Emp	Survival
Outcome:						
5-Year Difference	-0.068 (0.058)	-0.061** (0.024)	-0.063 (0.054)	-0.053** (0.023)	-0.034 (0.065)	-0.057** (0.027)
10-Year Difference	-0.149** (0.066)	-0.093*** (0.03)	-0.13** (0.062)	-0.085*** (0.028)	-0.067 (0.075)	-0.06* (0.034)
Industry + Employment Ctrls.	X	X	X	X	X	X
Pooled Ctrls.			X	X		
Flexible Ctrls.	X	X			X	X

*Note:* This table presents robustness results for the differences between single- and multi-establishment firms presented in Figure 2-7. Specifically, it presents the differences between the five- and ten-year DHS employment growth rate and survival estimates for various alternative specifications. The first two columns present the differences for the estimates presented in Figure 2-7. The "Pooled Controls" columns pool the controls across all cohorts as described in Section 2.4. The 30-70 % columns restrict to elections with 30-70 % of the vote share. For all specifications with restrictions, we still use the entire sample to estimate controls but restrict the treated variables to be estimated from the restricted sample. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table B.6: Unionized versus Non-Unionized Firm Heterogeneity, Robustness Checks

Specification:	Baseline		Pooled Controls		Contracts since 1990		30-70 % Elections	
	DHS Emp	Survival	DHS Emp	Survival	DHS Emp	Survival	DHS Emp	Survival
Outcome:								
5-Year Difference	-0.187** (0.095)	-0.03 (0.041)	-0.139 (0.089)	-0.018 (0.039)	-0.179 (0.112)	-0.057 (0.049)	-0.197* (0.108)	-0.029 (0.046)
10-Year Difference	-0.336*** (0.104)	-0.108** (0.048)	-0.287*** (0.097)	-0.09** (0.045)	-0.412*** (0.121)	-0.149** (0.058)	-0.305** (0.119)	-0.098* (0.055)
Industry + Employment Ctrls.	X	X	X	X	X	X	X	X
Pooled Ctrls.			X	X				
Flexible Ctrls.	X	X			X	X	X	X

*Note:* This table presents robustness results for the differences between multi-establishment firms with and without any unionized establishments presented in Figure 2-9. Specifically, it presents the differences between the five- and ten-year DHS employment growth rate and survival estimates for various alternative specifications. The first two columns present the differences for the estimates presented in Figure 2-9. The "Pooled Controls" columns pool the controls across all cohorts as described in Section 2.4. The "Contracts since 1990" column only classifies firms as unionized versus non-unionized starting in 1990. This gives all firms at least five years of pre-election FMCS contract data that we can use to define the firms' unionization status. The 30-70 % columns restrict to elections with 30-70 % of the vote share. For all specifications with restrictions, we still use the entire sample to estimate controls but restrict the treated variables to be estimated from the restricted sample. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## B.3 Data and Matching Details Appendix

### NLRB Union Election Data

**Union Election Data Sources** We combine datasets on NLRB elections from Henry Farber, J.P. Ferguson, and Thomas Holmes and publicly available data from the NLRB to give us a near-complete set of union elections from 1961-2019. Internet links for the Ferguson, Holmes, and NLRB are available. For more details about the sources of these data see JP Ferguson’s website [here](#).

**NLRB Election Case Numbers** The ID variable in the election data is an NLRB **Case ID Number**. This case number is assigned after an election petition is first filed. A single case number, however, could include multiple different vote counts. For example, there might be (1) multiple different tallies of the same election or (2) multiple elections for the same case number.<sup>1</sup> Additionally, there might be separate elections for multiple different bargaining units filed under the same case number (e.g., if a union initially filed for a petition for one bargaining unit but the NLRB then split bargaining unit). Consequently, it is important to pick the vote count that actually corresponds to the outcome of the certification election. Finally, since the different data sources cover overlapping time periods, we have multiple observations of the same case number in different datasets.

We deal with multiple observations per case number within datasets somewhat differently for the different data sources. For the public NLRB data (the “Public Data”) there is information indicating why there are multiple observations for a single case number. Consequently, for a given bargaining unit, we pick the final tally of the last election for each case number. This ensures that we take the vote tally that determines the unions’ certification for cases where there are multiple counts of the same election or multiple ordered elections for the same bargaining unit. Within each case number, we then take the results from the election at the largest bargaining unit in cases where there are distinct bargaining units for a single case. For the other datasets, there is somewhat less clarity about why there are duplicate observations within the same case number. For these datasets, we first pick the observation with the last election date and then the observation with the largest bargaining unit size.

---

<sup>1</sup>There could be multiple tallies for the same election due to challenged votes (e.g., the first tally would not include challenged votes while the final tally would include challenged votes that were determined to be valid). There could be multiple elections for the same case number if an NLRB director orders a second election due to objections to the first election

This leaves us one observation per case number within each data set but duplicates across datasets. We take one observation per case number across datasets. For picking a single case number per dataset, we deprioritize observations in the Farber data given data irregularities in those data. Additionally, we prioritize the public data because we have more confidence that we are picking the correct observation across duplicates within the same case number.

**Variables in the Election Dataset** We define the following from the union election data that we use for our analysis and for our matching algorithm

- **Election City, State, and Address:** The data contain the city and state of the election that we use to match each election to an establishment in the LBD. For many observations, we also observe a street address that we also use for the matching.

For the “public data”, we observe an address for the employer and for the election site. There are two conceptual reasons why these addresses might be different. First, the election might not be held at the employers’ location.<sup>2</sup> This suggests that the employers’ address is better for name and address matching to Census establishments. Second, the listed address for the employer might be a corporate headquarters rather than the establishment where the bargaining unit works. This suggests that the election address is better for name and address matching. Since it is not conceptually clear which address to use, we check which address is more likely to match the text in the bargaining unit description (e.g., “all warehousemen at its Louisville, KY facility”). We find that the election site address is more likely to match address information in the bargaining unit description and consequently use the election site addresses when they disagree.

- **Election Vote Shares:** We define election vote shares as the number of votes for the union divided by the total number of votes in the election. This differs from the adjusted vote shares constructed in DiNardo and Lee (2004) and Frandsen (2021) to address the “integer problem” with constructing vote shares<sup>3</sup>. We do not apply this adjustment for two reasons. First, the integer problem is

---

<sup>2</sup>For example, when strikes, pickets, or lockouts are in progress, the election may be held at a neutral location (NLRB, 2020). As another example, when the employers’ location is different than the employees’ worksite (e.g., security guards), the election might be held at the work site

<sup>3</sup>The integer problem refers to the fact that since vote shares are based on a discrete number of votes, there will be a mechanical discontinuity in the number of elections with exactly 50 % vote shares



especially problematic for regression discontinuity designs but less of an issue with our difference-in-differences design. Second, since we don't impose any restrictions on the number of votes cast in the election, the adjustment proposed in DiNardo and Lee (2004) would lead to larger changes in our vote shares (e.g., a six-person election would be adjusted from 50 % to 41.7 %).

- **Contested Elections:** We define contested elections as elections with multiple unions on the ballot. We drop these elections for two reasons. First, these elections are often “union raids” where one union already represents a specific bargaining unit and another union challenges that union for representation (Sandver and Ready, 1998). Consequently, a winning election, in this case, would not lead to a switch from the establishment being non-unionized to unionized but instead just a switch in which union represents the bargaining unit. Second, the reported vote totals for multi-union elections may not actually represent the workers' support for the union. In particular, for multi-union elections, if none of the options (e.g., “union 1”, “union 2”, or “no union”) receive the majority of the votes, a runoff election is held between the highest two options (Fraundorf, 1990). Consequently, the unions' true support (the union vote share from the first election) may be different than the unions' support in the observed runoff election results.
- **Election Industry:** The election data contain industry codes indicating the industry of the election analysis. For our main analysis, we use the Census industry codes for the establishments we match each election to. For some of our analysis of the unmatched NLRB data (e.g., Figures 2-2 and B-5 and Tables B.1), we use the election industry codes to split up manufacturing and non-manufacturing elections. Since the industry codes in the election data come from different vintages (e.g., SIC versus NAICS industry codes), we use the modal employment-weighted industry crosswalks from Eckert et al. (2020) to crosswalk the industry codes to consistent NAICS 2012 industry codes.
- **Bargaining Unit Size and Share of Total Employment:** We define the `bargaining unit size` as the number of eligible voters from the NLRB election data. We define the `bargaining unit share of total employment` as the bargaining unit size divided by the establishment-level employment one year for the union election. Since we do not impose that the bargaining unit is smaller than the establishment, we cap the share at one.

- **Election Filing Date:** We define treatment timing based on the date that the election was filed. To maximize the number of observations that we observe election filing dates for, we pull the dates across case numbers when some observations are missing from one dataset (e.g., if the filing date is only available for a case in the Ferguson data but not the Farber data, we pull date from the Ferguson to Farber data). For five % of elections, we do not observe the filing date and instead use the election or case closing date.
- **Election Delay Time:** We define delay time as the number of days between the date the election petition was filed to the NLRB and the date the election was held. The availability of exact dates for these two concepts varies somewhat across time and datasets. Both dates are missing from the Farber data which is one reason why we prioritize the other datasets when duplicates across case numbers are available. However, as described above, we pull both dates across datasets when they are missing for some observations. For the Ferguson and Holmes data, the delay time is missing for cases that closed in 1982 and we only have a monthly measure for 1981 and part of 1983. These differences over time motivate our checks that the heterogeneity by delay time holds using both variation within-years (e.g., the within-year tercile measures) and across years (e.g., the continuous log specification). Additionally, there may have been some institutional changes over time that we do not want to include (e.g., the “Quickie Election Rule” decreased delay times but is not in our sample of elections).

## FMCS Contract Data

We combine contract data from Thomas Holmes for 1984-2003 and from the FMCS for 1997-2019. The Homes data are available here and the FMCS data are partially available here and the rest were obtained via a FOIA request. They include both notices of *initial contracts* (i.e., first-contract negotiation after an election) and *contract renegotiation or reopening* for existing contracts. There are two reasons that these contract notices likely underrepresent the universe of unionized establishments in the U.S. First, these “notices of bargaining” are provided to the FMCS so it can be ready to provide mediation. Although filing is legally incentivized, underreporting is possible. For example, an employer changing the terms of employment or a union striking without first filing a notice could be violating labor law. Second, some contract notices may represent a contract covering multiple establishments be we always only match

each contract to one establishment.<sup>4</sup>

There are duplicate observations both across the Holmes versus FMCS datasets and within each dataset.<sup>5</sup> However, unlike the NLRB election data, we have no IDs to restrict the dataset to unique observations. Consequently, to deal with duplicates, we match all contract observations to the Census establishments in the LBD and drop duplicates when multiple contract observations match to the same Census establishment.

We use the contract data to define

- **Previous contract at an establishment:** for each election establishment, we define an indicator for whether the establishment has a previous FMCS contract ever matched to the same establishment (e.g., indicating that another bargaining unit was already unionized at this establishment). To avoid contract matches related to the union election, we only include matched contracts starting one year before the election.
- **Unionized versus Non-Unionized Firms:** we define a firm as being (partially) unionized if at time  $t$  any of the establishments in the same FIRMID had an FMCS contract match in the current or previous five years. For the unionized versus non-unionized firm heterogeneity check, we also include elections at establishments with a previous contract (defined above) as unionized firms.

## FMCS Works' Stoppage Data

For Figure 2-2 Panel D., we use works' stoppage data from the FMCS from 1984-2005. The data are available [here](#). They include both strikes and employer-initiated lockouts. We match the works stoppages to the election data based on exact company names and cities rather than the *Soft TF-IDF* algorithm we use for the main analysis. Prior to matching, we use the same cleaning algorithms described below to clean the employer and city names in the FMCS works' stoppage data.

---

<sup>4</sup>Sometimes the FMCS contract notices explicitly mention that they apply to multiple locations (e.g., the address indicating various locations). In these cases, we will still only match the contract notice to one establishment if there is alternative location data available.

<sup>5</sup>The across-dataset duplicates come from the fact that the datasets overlap. The within-dataset duplicates could come from an employer and union submitting an FMCS notice for the same contract.

## Longitudinal Business Database

In Section 2.3, we mention potential concerns with how the LBD allocates employment across establishments at multi-establishment firms that could bias our results. To be more precise about the issue, while the LBD is an establishment-level dataset, some of the employment and payroll input data are received at higher levels of aggregation (e.g., at the EIN level). For example, one source used to construct the LBD is *IRS form 941s* that provide annual employment and payroll at the EIN-level which can cover multiple establishments. The Census uses an imputation model to allocate these EIN-level measures across establishments. This model primarily imputes employment changes across establishments based on their past employment. Consequently, employment changes at an establishment part of a multi-establishment firm might initially be allocated across all establishments. Thus, the LBD would initially underestimate the establishment-level decrease in employment. To correct some of these mistakes, the Census receives establishment-level information from the Company Organization Survey (COS), Economic Censuses, and Annual Survey of Manufacturers (ASM) that provide more accurate measures of establishment-level employment and survival. These alternative surveys are not, however, conducted for all establishments annually (e.g., the Economic Census is only conducted every five years). So there might be a few years lag before the LBD reports the correct establishment employment and exit. This lag mirrors the spike in establishment births and deaths every five years during the economic census years when the Census has establishment-level data for each establishment (Jarmin and Miranda, 2002). See Chow et al. (2021) for details about these issues with the LBD construction.

We use the LBD to define the following establishment-level variables

- **Employment:** total number of employees who received wages or other compensation during the pay period that included March 12th.
- **Payroll:** total “wages, tips, and other compensation” for employees over the entire year.
- **Establishment Survival:** indicator for whether the establishment has positive employment for at least one year in the future and in the past. Consequently, an establishment that has 50 employees one year, 0 employees the next, and 50 employees the following year would be defined as a “survivor” in the intermittent year. Since the LBD only measures March 12 employment, these establishments could be true survivors (e.g., seasonal businesses).

- **Establishment-Level NAICS Codes:** We classify each establishment into a 2012 NAICS industry using the Fort and Klimek (2016) NAICS codes.

## Plant-Level TFP from the Annual Survey of Manufacturers

We define plant-level productivity using inputs and outputs from the Annual Survey of Manufacturers (ASM) and TFP measures calculated by Foster et al. (2016). To classify each election into different terciles of the plant productivity distribution, we first take all ASM observations, with and without union elections, with non-missing TFP and calculate year by NAICS 6 industry TFP percentiles. For each of our manufacturing union elections, we then assign the election the plant's most recent TFP percentile in the previous five years (e.g., if the establishment was sampled by the ASM in year  $E_i - 2$  but not  $E_i - 1$ , we assign the establishment its  $E_i - 2$  productivity rank). Based on the election observations with defined TFP, we then classify the elections into within-year terciles based on these rankings.

## Matching Elections, Contracts, and LBD Establishments

Our data on union elections and contract notices contain information on the name and location of the employer, but no unique identifiers (like EIN) that could use to directly link the establishments to administrative Census data firms. We instead use a fuzzy-matching algorithm to link each election or contract to its corresponding Census record from the Standard Statistical Establishment List/Business Register. The algorithm is based on the name and geographic similarity of establishments. Our algorithm is based upon the Soft TF-IDF approach used by Kline et al. (2019), but extends their approach to incorporate the additional address data.

**Name and Address String Cleaning:** We start by standardizing and cleaning the name and address strings. Our cleaning procedure builds on the `stnd_compname` and `stnd_address` Stata name standardization programs (Wasi and Flaaen, 2015). We clean addresses as follows:

1. Remove most symbols, non-numeric or letter characters, and non-standard ASCII characters.
2. Removed PO boxes, building/suite/room numbers, and company names at the start of addresses (e.g., **GENERAL SUPPLY COMPANY** 2651 1ST STREET.)

3. Standardize common address and city name strings (e.g., `ST`  $\Rightarrow$  `STREET`, `TWENTY FIRST`  $\Rightarrow$  `21ST`, and `LIC`  $\Rightarrow$  `LONG ISLAND CITY`) and correct common address and city misspellings.

We clean the employer names as follows

1. Remove most symbols, non-numeric or letter characters, and non-standard ASCII characters.
2. Remove the portion of company names in parentheses. The union election data often contain supplemental information in the parentheses portion of the name (e.g., `(wage employees only)`).
3. Remove the portion of company names following `DOING BUSINESS AS (DBA)` or `A DIVISION OF`
4. Combine consecutive singleton letters and symbols separated by spaces (e.g., `A T & T`  $\Rightarrow$  `AT&T` and `D R HORTON`  $\Rightarrow$  `DR HORTON`).
5. Remove company entity types (e.g., `CORP`, `INC`, etc.), articles, and standard common company names (e.g., `MANUFACTURERS`  $\Rightarrow$  `MANUFACTURING`).

**Election, Contract, and Census Address Geocodes:** We geocode all addresses. This allows us to construct measures of address similarity based on the geographic distance between two addresses. We use geographic distance rather than string distance to measure address similarity because there may be addresses with very similar strings that are very different addresses (e.g., `100 Main St.` may be very far away from `10 Main St.`).

For the election and contract data, we first try to geocode all addresses with the Census Bureau's Geocoding API because these geocodes are the most likely to match the Census's internal geocodes. For the observations where the Census's geocoder cannot find a geocode, we try the geocodio geocoder. When an observations' street address is missing or we cannot geocode it, we take the city/state geocode or the zip-code geocode.

For the Census data, we use the geocodes in the SSEL/Business Register (DeSalvo et al., 2016). These geocodes, however, are only available since 2002 (Akee et al., 2017). For observations where we do not have a geocode we first try to match it to a geocoded address. If the same address was not geocoded from 2002-2016, we instead take the average geocode of all addresses we see in 2002-2016 in the same city/state or zip code.

**Matching Algorithm** We implement a matching algorithm based on the string similarity of the cleaned employer names and the geographic distance between geocoded addresses. The standard Soft TF-IDF algorithm computes a match score between two firm names that is increasing in their string similarity. The algorithm is particularly suitable for our application since it overweights similarities in uncommon words between the two names and discounts similarities in common words. Although it's possible to match the unionization records to the Census data based on employer name similarity alone, the procedure is likely to generate false establishment matches (especially given that establishments at multiunit firms may all share the same name, like "CVS" or "Starbucks"). Consequently, we instead also incorporate the geography information to distinguish between these potential matches.

We implement our matching algorithm as follows

1. For each election, we take all Census establishments in the same state that share at least one common word.<sup>6</sup>
2. For each election-establishment pair, we calculate the Soft TF-IDF similarity measure between the employer name strings. Specifically, let  $A_j$  be the set of all words in the election name string and  $B_k$  be the set of all words in the establishment name string. The total number of election names is  $J$  and the total number of Census names is  $K$ . The Soft TF-IDF distance is defined as

$$s_{jk} = \text{Soft TF-IDF}(A_j, B_k) = \sum_{w \in A_j} \text{weight}(w, A_j) \times \text{m-score}(w, B_k) \quad (\text{B.1})$$

where  $\text{weight}(w, A_j)$  is defined as

$$\text{weight}(w, A_j) = \frac{\text{TF}(w, A_j) \times \text{IDF}(w, A, B)}{\left[ \sum_{w' \in A_j} (\text{TF}(w', A_j) \times \text{IDF}(w', A, B))^2 \right]^{1/2}}, \text{ where} \quad (\text{B.2})$$

$$\text{TF}(w, A_j) = \frac{\text{freq}(w, A_j)}{\sum_{w' \in A_j} \text{freq}(w', A_j)}, \text{ and} \quad (\text{B.3})$$

$$\text{IDF}(w, A, B) = -1 \times \log \left( \frac{\sum_{j'} \mathbb{1}[w \in A_{j'}] + \sum_{k'} \mathbb{1}[w \in B_{k'}]}{J + K} \right). \quad (\text{B.4})$$

Intuitively, the TF portion of the weight gives higher weights to words part of

---

<sup>6</sup>We require that the establishments share at least one common word because this vastly reduces the number of string and distance calculations we need to make. For single-word companies, we only require that the potential matches share the same first letter. This allows us to match single-word establishments even with misspellings

shorter names. The IDF portion of the weight gives higher weights to less common words relative to all words included in any election or Census establishment name. We give higher weights to less common words because two names sharing a common word (e.g., **manufacturing**) is less likely to indicate a correct match than two words sharing a less common word (e.g., **wanaque**).

The  $\text{m-score}(w, B_k)$  is defined as follows

$$\text{m-score}(w, B_k) = \bar{m}(w, B_k) \times \text{weight}(\bar{w}, B_k) \times \mathbb{1}[\bar{m}(w, B_k) > \theta] \quad (\text{B.5})$$

where  $\bar{m}(w, B_k)$  is the highest Jaro-Winkler distance between the word  $w$  and any word in the name  $B_k$

$$\bar{m}(w, B_k) = \max_{w' \in B_k} \text{Jaro-Winkler}(w, w') \quad (\text{B.6})$$

and  $\bar{w}$  is the word in  $B_k$  that maximizes the Jaro-Winkler string distance.  $\theta$  is a threshold below which the m-score is defined as zero. The Jaro-Winkler string distance is a measure of how similar two strings are. It considers the number of matching characters in the strings and the number of transpositions necessary to get the strings to match (e.g., **Boston** and **Bostno** require one transposition). Finally, it also places a higher weight on matching characters at the beginning of strings. See Kline et al. (2019) for details.

3. We calculate the Haversine distance between the election and Census establishment geocoordinates as follows

$$d_{j,k} = \min(\text{Haversine Distance}(\text{geo\_coord}_j, \text{geo\_coord}_k), \bar{d}). \quad (\text{B.7})$$

where  $\bar{d}$  is our distance top code (e.g., distances above a certain threshold are unlikely to be informative).

4. We combine the string similarity measure and the distance measure for each pair of elections and establishments as follows

$$\text{match score}_{jk} = (1 - \beta) \cdot s_{jk} + \beta [1 - (d_{jk}/\bar{d})^\gamma] \quad (\text{B.8})$$

where  $\beta$  is the relative weight placed on distance versus string name similarity.  $\gamma$  is the relative weight placed on very close versus farther away matches (e.g., a very concave  $\gamma$  places much more weight on exact geographic matches than



matches that are even slightly farther away).

5. For each election, we pick the Census establishment with the highest match score $_{jk}$ . This yields a potential match for each election but these matches may be very low quality or incorrect.
6. We only keep matches where match score $_{jk}$  is above a minimum threshold  $p$ .

The matching algorithm has several tuning parameters that determine the relative weights placed on each component of the final match score. For the parameters used to calculate the Soft TF-IDF score and the final match score (e.g.,  $\theta$ , the  $p$  parameter in the JW string distance, and  $\gamma$ ), we use details about our institutional setting to optimize these parameters in a principled manner. We first optimize the Soft TF-IDF parameters by matching each election record to at most one contract record. We then choose the parameters that maximize the discontinuity in the likelihood that an election record has a matching contract record across the 50% vote-share threshold.

To pick the minimum match score  $p$ , we exploit the fact that the size of the election bargaining unit in the election data and the number of employees at the Census establishment give us information about whether or not the match is correct. In particular, having a larger bargaining unit than the number of workers at the establishment indicates an incorrect match.<sup>7</sup> We first directly calculate the probability that an election record was matched correctly to a Census record (as a function of the records' match score) by comparing the bargaining unit size to the number of workers at the Census establishment. For a matched set of records with match score  $s$ , we define the average likelihood that the matched Census employment is at least as high as the number of recorded votes  $m(s)$ . On the other hand, the likelihood that the employment at random Census establishment is at least as high as the number of recorded votes is  $\underline{m}$ . We assume records where the name and geographic location match exactly are "true" matches, which correspondingly allows us to estimate that a pair of records with a match score of  $s$  is matched correctly with probability:

$$p(s) = \frac{m(s) - \underline{m}}{m(1) - \underline{m}}. \tag{B.9}$$

We include all record matches where the correct match probability  $p(s)$  is at least 75%, and we select the geography weight that maximizes the number of elections that

---

<sup>7</sup>There may be cases of larger bargaining unit sizes than establishment employment that actually are correct matches. For example, there may be data mistakes in the bargaining unit size or the measures may cover different time periods.

are matched in this process. We then use the same parameters to also match contract notices to the Census records.

# Appendix C

## Appendix for Chapter 3

### C.1 Additional Tables and Figures

Figure C-1: Number of schools ranked in application

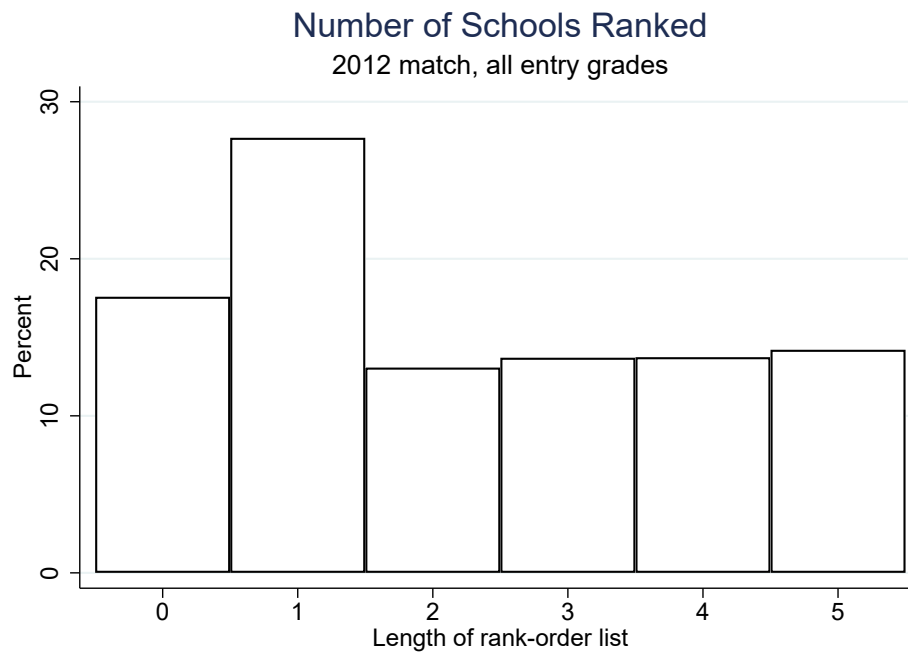





Figure C-2: Substitution patterns (no heterogeneity)



(a) Elementary schools



(b) Middle schools



(c) High schools

Figure C-4: Match replication rate






Figure C-5: Event study coefficients (other specifications)

(a) Basic controls



(b) Additional controls

# Bibliography

- Abadie, Alberto, Susan Athey, Guido Imbens, and Jeffrey Wooldridge (2017) “When Should You Adjust Standard Errors for Clustering?” *arXiv:1710.02926 [econ, math, stat]*.
- Abadie, Alberto, Matthew M. Chingos, and Martin R. West (2018) “Endogenous Stratification in Randomized Experiments,” *The Review of Economics and Statistics*, 100 (4), 567–580.
- Abaluck, Jason, Mauricio Caceres Bravo, Peter Hull, and Amanda Starc (2021) “Mortality Effects and Choice Across Private Health Insurance Plans,” *The Quarterly Journal of Economics*, 136 (3), 1557–1610.
- Abdulkadiroglu, Atila, Nikhil Agarwal, and Parag A. Pathak (2017a) “The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match,” *American Economic Review*, 107 (12), 3635–3689.
- Abdulkadiroglu, Atila, Parag A. Pathak, Jonathan Schellenberg, and Christopher R. Walters (2017b) “Do Parents Value School Effectiveness?” Working Paper 23912, National Bureau of Economic Research.
- Abowd, John M., Francis Kramarz, and David N. Margolis (1999) “High Wage Workers and High Wage Firms,” *Econometrica*, 67 (2), 251–333.
- Acemoglu, Daron and Jorn-Steffen Pischke (1998) “Why Do Firms Train? Theory and Evidence,” *The Quarterly Journal of Economics*, 113 (1), 79–119.
- Addario, Sabrina L. Di, Patrick M. Kline, Raffaele Saggio, and Mikkel Sølvsten (2021) “It Ain’t Where You’re From, It’s Where You’re At: Hiring Origins, Firm Heterogeneity, and Wages,” Working Paper 28917, National Bureau of Economic Research.
- Agarwal, Nikhil and Paulo Somaini (2019) “Revealed Preference Analysis of School Choice Models,” Technical Report w26568, National Bureau of Economic Research, Cambridge, MA.
- Akee, Randall, Elton Mykerezzi, and Richard M Todd (2017) “Reservation Employer Establishments: Data from the U.S. Census Longitudinal Business Database,” 35.

- Allende, Claudia (2019) “Competition Under Social Interactions and the Design of Education Policies,” 104.
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber (2005) “Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools,” *Journal of Political Economy*, 113 (1), 151–184.
- Altonji, Joseph G. and Charles R. Pierret (2001) “Employer Learning and Statistical Discrimination,” *The Quarterly Journal of Economics*, 116 (1), 313–350.
- Alvarez, Jorge, Felipe Benguria, Niklas Engbom, and Christian Moser (2018) “Firms and the Decline in Earnings Inequality in Brazil,” *American Economic Journal: Macroeconomics*, 10 (1), 149–189.
- Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro (2017) “Measuring the sensitivity of parameter estimates to estimation moments,” *The Quarterly Journal of Economics*, 132 (4), 1553–1592.
- (2020) “Transparency in Structural Research,” *Journal of Business & Economic Statistics*, 38 (4), 711–722.
- Angrist, Joshua D. and Jörn-Steffen Pischke (2009) *Mostly Harmless Econometrics: An Empiricist’s Companion*, Princeton: Princeton University Press, 1st edition.
- Angrist, Joshua D. and Miikka Rokkanen (2015) “Wanna Get Away? Regression Discontinuity Estimation of Exam School Effects Away From the Cutoff,” *Journal of the American Statistical Association*, 110 (512), 1331–1344.
- Arellano-Bover, Jaime (2020) “Career Consequences of Firm Heterogeneity for Young Workers: First Job and Firm Size,” 90.
- Arellano-Bover, Jaime and Fernando Saltiel (2020) “Differences in On-the-Job Learning Across Firms,” 65.
- Arnold, David (2019) “Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes,” *Working Paper*.
- Autor, David H. (2001) “Why Do Temporary Help Firms Provide Free General Skills Training?” *The Quarterly Journal of Economics*, 116 (4), 1409–1448.
- Bagger, Jesper, François Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin (2014) “Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics,” *American Economic Review*, 104 (6), 1551–1596.
- Baker, George, Michael Gibbs, and Bengt Holmstrom (1994) “The Internal Economics of the Firm: Evidence from Personnel Data,” *The Quarterly Journal of Economics*, 109 (4), 881–919.
- Baron, James N. and William T. Bielby (1986) “The Proliferation of Job Titles in Organizations,” *Administrative Science Quarterly*, 31 (4), 561–586.



- Bayer, Patrick and Robert McMillan (2005) “Choice and Competition in Local Education Markets,” Technical Report w11802, National Bureau of Economic Research, Cambridge, MA.
- Becker, Gary S. (1964) *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education, First Edition*: NBER.
- Bennett, Magdalena (2020) “How Far is Too Far? Estimation of an Interval for Generalization of a Regression Discontinuity Design Away from the Cutoff,” 50.
- Benson, Alan, Danielle Li, and Kelly Shue (2019) “Promotions and the Peter Principle,” *The Quarterly Journal of Economics*, 134 (4), 2085–2134.
- Berger, David, Kyle Herkenhoff, and Simon Mongey (2021) “Minimum Wages and Welfare,” *Working Papers*, 49.
- Berry, Steven, James Levinsohn, and Ariel Pakes (2004) “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market,” *Journal of Political Economy*, 112 (1), 68–105.
- Blau, Francine D. and Jed Devaro (2007) “New Evidence on Gender Differences in Promotion Rates: An Empirical Analysis of a Sample of New Hires,” *Industrial Relations: A Journal of Economy and Society*, 46 (3), 511–550.
- Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen (2019) “What Drives Differences in Management Practices?” *American Economic Review*, 109 (5), 1648–1683.
- Bloom, Nicholas, Renata Lemos, Raffaella Sadun, Daniela Scur, and John Van Reenen (2014) “The New Empirical Economics of Management,” Technical Report w20102, National Bureau of Economic Research, Cambridge, MA.
- Bluestone, Barry and Bennett Harrison (1982) *The Deindustrialization of America: Plant Closings, Community Abandonment, and the Dismantling of Basic Industry*, New York: Basic Books, Inc., Publishers.
- Boehm, Christoph E., Aaron B. Flaaen, and Nitya Pandalai-Nayar (2019) “Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tohoku Earthquake,” *Review of Economics and Statistics*, 101 (1).
- Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa (2019) “A Distributional Framework for Matched Employer Employee Data,” *Econometrica*, 87 (3), 699–739.
- Bradley, Daniel, Incheol Kim, and Xuan Tian (2017) “Do Unions Affect Innovation?” *Management Science*, 63 (7), 2251–2271.
- Bresnahan, Timothy F. (1989) “Chapter 17 Empirical studies of industries with market power,” in *Handbook of Industrial Organization*, 2, 1011–1057: Elsevier.

- Bronfenbrenner, Kate (1996) “Final Report: the Effects of Plant Closing or Threat of Plant Closing on the Right of Workers to Organize,” *Report Submitted to the The Labor Secretariat of North American Commission for Labor Cooperation*.
- (1997) “The Role of Union Strategies in NLRB Certification Elections,” *ILR Review*, 50 (2), 195–212.
- (2000) “Uneasy Terrain: The Impact Of Capital Mobility On Workers, Wages, And Union Organizing,” *Report Submitted to the U.S. Trade Deficit Review Commission*.
- (2001) “Uneasy Terrain: The Impact of Capital Mobility on Workers, Wages, and Union Organizing. Part II: First Contract Supplement.”
- (2002) “Overcoming the Challenges to Organizing in Manufacturing,” *Report Submitted to the AFL-CIO*.
- Bronfenbrenner, Kate and Tom Juravich (1998) “It Takes More Than House Calls: Organizing to Win with a Comprehensive Union-Building Strategy,” in *Organizing to Win: New Research on Union Strategies*, Kate Bronfenbrenner and Sheldon Friedman, et al. (eds.), 19–36, Ithaca, NY: ILR Press.
- Brown, Jenny (2011) “Store Workers Say: Whose Walmart? OUR Walmart!,” *Labor Notes*.
- Burdett, Ken and Melvyn Coles (2003) “Equilibrium Wage-Tenure Contracts,” *Econometrica*, 71 (5), 1377–1404.
- Butani, Shail, George Werking, and Vinod Kapani (2005) “Employment dynamics of individual companies versus multicorporations,” *Monthly Labor Review*, 13.
- Caldwell, Sydnee and Nikolaj Harmon (2019) “Evidence from Coworker Networks,” 105.
- Cameron, Andrew Tangel and Doug (2020) “WSJ News Exclusive \textbar Boeing to Move All 787 Dreamliner Production to South Carolina,” *Wall Street Journal*.
- Campos, Christopher and Caitlin Kearns (2022) “The Impact of Neighborhood School Choice: Evidence From Los Angeles’ Zones of Choice,” 96.
- Card, David (1996) “The Effect of Unions on the Structure of Wages: A Longitudinal Analysis,” *Econometrica*, 64 (4), 957–979.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline (2018) “Firms and Labor Market Inequality: Evidence and Some Theory,” *Journal of Labor Economics*, 36 (S1), S13–S70.
- Card, David, Ana Rute Cardoso, and Patrick Kline (2016) “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women \*,” *The Quarterly Journal of Economics*, 131 (2), 633–686.

- Card, David, Martin D. Dooley, and A. Abigail Payne (2010) “School Competition and Efficiency with Publicly Funded Catholic Schools,” *American Economic Journal: Applied Economics*, 2 (4), 150–176.
- Card, David, Jörg Heining, and Patrick Kline (2013) “Workplace Heterogeneity and the Rise of West German Wage Inequality,” *The Quarterly Journal of Economics*, 128 (3), 967–1015.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019) “The Effect of Minimum Wages on Low-Wage Jobs,” *The Quarterly Journal of Economics*, 134 (3), 1405–1454.
- de Chaisemartin, Clément and Xavier D’Haultfoeuille (2020) “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 110 (9), 2964–2996.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff (2014) “Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates,” *American Economic Review*, 104 (9), 2593–2632.
- Chodorow-Reich, Gabriel (2014) “The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008-9 Financial Crisis,” *The Quarterly Journal of Economics*, 129 (1), 1–59.
- Chow, Melissa, Teresa Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White (2021) “Redesigning the Longitudinal Business Database,” *NBER Working Paper* (w28839), w28839.
- Clark, Kim B (1980) “The Impact of Unionization on Productivity: A Case Study,” *ILR Review*, 33 (4), 20.
- Cocci, Matthew D and Mikkel Plagborg-Møller (2019) “Standard Errors for Calibrated Parameters,” 26.
- Cornwell, Christopher, Ian M. Schmutte, and Daniela Scur (2021) “Building a Productive Workforce: The Role of Structured Management Practices,” *Management Science*.
- Crane, Leland D. and Ryan A. Decker (2019) “Business Dynamics in the National Establishment Time Series (NETS),” *Finance and Economics Discussion Series Divisions of Research & Statistics and Monetary Affairs Federal Reserve Board*, 2019 (034).
- CRS (2013) “The National Labor Relations Act (NLRA): Union Representation Procedures and Dispute Resolution,” *Congressional Research Service Report*, 38.
- Cullen, Zoë and Ricardo Perez-Truglia (2021) “The Old Boys’ Club: Schmoozing and the Gender Gap,” 50.

- Dafny, Leemore S. (2005) “How Do Hospitals Respond to Price Changes?” *American Economic Review*, 95 (5), 1525–1547.
- Davis, Steven J, John Haltiwanger, Kyle Handley, Josh Lerner, and Javier Miranda (2014) “Private Equity, Jobs, and Productivity,” *American Economic Review*, 104, 3956–3990.
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh (1996) *Job Creation and Destruction*, Cambridge, MA, USA: MIT Press.
- DeSalvo, Bethany, Frank Limehouse, and Shawn D. Klimek (2016) “Documenting the Business Register and Related Economic Business Data,” *CES Working Papers* 16-17.
- DiNardo, John and David S. Lee (2004) “Economic Impacts of New Unionization on Private Sector Employers: 1984-2001,” *Quarterly Journal of Economics*, 119 (4), 1383–1441.
- Dinlersoz, Emin, Jeremy Greenwood, and Henry Hyatt (2017) “What Businesses Attract Unions? Unionization over the Life Cycle of U.S. Establishments,” *ILR Review*, 70 (3), 733–766.
- Dix-Carneiro, Rafael, Pinelopi Koujianou Goldberg, Costas Meghir, and Gabriel Ulyssea (2021) “Trade and Informality in the Presence of Labor Market Frictions and Regulations,” 71.
- Dix-Carneiro, Rafael and Brian K. Kovak (2017) “Trade Liberalization and Regional Dynamics,” *American Economic Review*, 107 (10), 2908–2946.
- Doeringer, Peter B. and Michael J. Piore (1971) *Internal labor markets and manpower analysis*: Lexington, Mass., Heath [1971].
- Dube, Arindrajit, Ethan Kaplan, and Owen Thompson (2016) “Nurse Unions and Patient Outcomes,” *ILR Review*, 69 (4), 803–833.
- Duggan, Mark G. (2000) “Hospital Ownership and Public Medical Spending,” *The Quarterly Journal of Economics*, 115 (4), 1343–1373.
- Dunlop, John (1994) “Dunlop Commission on the Future of Worker-Management Relations: Final Report.”
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson (1989) “The Growth and Failure of U. S. Manufacturing Plants,” *The Quarterly Journal of Economics*, 104 (4), 671.
- Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge (2020) “Reallocation Effects of the Minimum Wage,” *Working Paper*, 60.

- Eckert, Fabian, Teresa Fort, Peter Schott, and Natalie Yang (2020) “Imputing Missing Values in the US Census Bureau’s County Business Patterns,” *NBER Working Paper* (w26632), w26632.
- Eeckhout, Jan and Philipp Kircher (2011) “Identifying Sorting-In Theory,” *The Review of Economic Studies*, 78 (3), 872–906.
- Epple, Dennis, Richard E. Romano, and Miguel Urquiola (2017) “School Vouchers: A Survey of the Economics Literature,” *Journal of Economic Literature*, 55 (2), 441–492.
- Estlund, Cynthia (1993) “Economic Rationality and Union Avoidance: Misunderstanding the National Labor Relations Act,” *Texas Law Review*, 71 (5).
- Evans, Stephen and Roy Lewis (1989) “Union Organisation, Collective Bargaining And The Law: An Anglo- American Comparison Of The Construction Industry,” *Comparative Labor Law Journal*, 33.
- Fan, Ying (2013) “Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market,” *American Economic Review*, 103 (5), 1598–1628.
- Farber, Henry S (2005) “Nonunion Wage Rates and the Threat of Unionization,” *Industrial and Labor Relations Review*, 58 (3), 19.
- Farber, Henry S. and Robert Gibbons (1996) “Learning and Wage Dynamics,” *The Quarterly Journal of Economics*, 111 (4), 1007–1047.
- Farber, Henrys (2001) “Union Success in Representation Elections: Why Does Unit Size Matter?” *ILR Review*, 54 (2), 329–348.
- Ferguson, John-Paul (2008) “The Eyes Of The Needles: A Sequential Model Of Union Organizing Drives, 1999–2004,” *Industrial And Labor Relations Review*, 62 (1), 19.
- Figlio, David and Cassandra M. D. Hart (2014) “Competitive Effects of Means-Tested School Vouchers,” *American Economic Journal: Applied Economics*, 6 (1), 133–156.
- Flanagan, Robert (2007) “Has Management Strangled U.S. Unions?” in *What Do Unions Do? A Twenty Year Perspective Edited by James Bennett and Bruce Kaufman*, New York: Routledge.
- Fort, Teresa C and Shawn Klimek (2016) “The Effects Of Industry Classification Changes On Us Employment Composition,” *CES Working Paper 18-28*, 24.
- Foster, Lucia, Cheryl Grim, and John Haltiwanger (2016) “Reallocation in the Great Recession: Cleansing or Not?” *Journal of Labor Economics*, 34 (S1), S293–S331.
- Foster, Lucia, John Haltiwanger, and Chad Syverson (2008) “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?” *American Economic Review*, 98 (1), 394–425.

- Foulkes, Fred (1980) *Personnel Policies in Large Nonunion Companies*, Englewood Cliffs, N.J.: Prentice-Hall.
- Frandsen, Brigham (2021) “The Surprising Impacts of Unionization: Evidence from Matched Employer-Employee Data,” *Journal of Labor Economics*, 39 (4), 56.
- Frandsen, Brigham R. (2017) “Party Bias in Union Representation Elections: Testing for Manipulation in the Regression Discontinuity Design when the Running Variable is Discrete,” in Cattaneo, Matias D. and Juan Carlos Escanciano eds. *Advances in Econometrics*, 38, 281–315: Emerald Publishing Limited.
- Fraudorf, Martha Norby (1990) “The effect of voting procedures on union representation elections: The multi-union case,” *Journal of Labor Research*, 11 (3), 323–335.
- Freedman, Audrey (1979) “Managing Labor Relations,” *Conference Board Report*.
- Freeman, Richard B and Morris M Kleiner (1990a) “Employer Behavior in the Face of Union Organizing Drives,” *ILRR*, 43 (4), 16.
- Freeman, Richard B. and Morris M. Kleiner (1990b) “The Impact of New Unionization on Wages and Working Conditions,” *Journal of Labor Economics*, 8 (1, Part 2), S8–S25.
- Freeman, Richard and Morris Kleiner (1999) “Do Unions Make Enterprises Insolvent?” *ILR Review*, 52 (4).
- Freeman, Richard and James Medoff (1984) *What Do Unions Do?*, N.Y.: Basic Books.
- Friedman, Milton (1951) “Some Comments on the Significance of Labor Unions for Economic Policy,” in *The Impact of the Union*, 204–259, New York: Harcourt Brace.
- (1962) *Capitalism and Freedom*: University of Chicago Press.
- Friedrich, Benjamin (2020) “Internal Labor Markets and the Competition for Managerial Talent,” 84.
- Gale, D. and L. S. Shapley (1962) “College Admissions and the Stability of Marriage,” *The American Mathematical Monthly*, 69 (1), 9–15.
- Ganong, Peter and Pascal Noel (2020) “Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession,” *American Economic Review*, 110 (10), 3100–3138.
- General Accounting Office, United States (2002) “Collective Bargaining Rights: Information on the Number of Workers with and without Bargaining Rights,” *U.S. General Accounting Office Report GAO-02-835*.

- Gerard, François, Lorenzo Lagos, Edson Severnini, and David Card (2018) “Assortative Matching or Exclusionary Hiring? The Impact of Firm Policies on Racial Wage Differences in Brazil,” 94.
- Gibbons, Robert and Lawrence Katz (1992) “Does Unmeasured Ability Explain Inter-Industry Wage Differentials,” *The Review of Economic Studies*, 59 (3), 515.
- Gibbons, Robert and Michael Waldman (1999) “A Theory of Wage and Promotion Dynamics inside Firms,” *The Quarterly Journal of Economics*, 114 (4), 1321–1358.
- Gilraine, Michael, Uros Petronijevic, and John D. Singleton (2021) “Horizontal Differentiation and the Policy Effect of Charter Schools,” *American Economic Journal: Economic Policy*, 13 (3), 239–276.
- Giroud, Xavier and Holger M. Mueller (2017) “Firm Leverage, Consumer Demand, and Employment Losses during the Great Recession,” *The Quarterly Journal of Economics*, 132 (1), 271–316.
- Glaeser, Edward L. and Andrei Shleifer (2001) “Not-for-profit entrepreneurs,” *Journal of Public Economics*, 81 (1), 99–115.
- Golan, Limor (2005) “Counteroffers and Efficiency in Labor Markets with Asymmetric Information,” *Journal of Labor Economics*, 23 (2), 373–393.
- Goncalves, Felipe (2021) “Do Police Unions Increase Misconduct?,” 76.
- Gonzaga, Gustavo, William F. Maloney, and Alejandra Mizala (2003) “Labor Turnover and Labor Legislation in Brazil,” *Economía*, 4 (1), 165–222.
- Goodman-Bacon, Andrew (2021) “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 225 (2), 254–277.
- Goodreau, Steven M., James A. Kitts, and Martina Morris (2009) “Birds of a feather, or friend of a friend? using exponential random graph models to investigate adolescent social networks\*,” *Demography*, 46 (1), 103–125.
- Greenhouse, Steven (2011) “Labor Board Tells Boeing New Factory Breaks Law,” *The New York Times*.
- Greenwald, Bruce C. (1986) “Adverse Selection in the Labour Market,” *The Review of Economic Studies*, 53 (3), 325–347.
- Gregory, Victoria (2019) “Firms as Learning Environments: Implications for Earnings Dynamics and Job Search,” 57.
- Guajardo, Jose A., Morris A. Cohen, and Serguei Netessine (2016) “Service Competition and Product Quality in the U.S. Automobile Industry,” *Management Science*, 62 (7), 1860–1877.

- Haanwinckel, Daniel and Rodrigo R Soares (2020) “Workforce Composition, Productivity, and Labor Regulations in a Compensating Differentials Theory of Informality,” 93.
- Haltiwanger, John, Ron S. Jarmin, and Javier Miranda (2013) “Who Creates Jobs? Small versus Large versus Young,” *Review of Economics and Statistics*, 95 (2), 347–361.
- Harasztosi, Peter and Attila Lindner (2018) “Who Pays for the Minimum Wage?” *Working Paper*, 110.
- Harju, Jarkko, Simon Jäger, and Benjamin Schoefer (2021) “Voice at Work,” *Working Paper*, 62.
- Hastings, Justine S., Thomas J. Kane, and Douglas O. Staiger (2008) “Heterogeneous Preferences and the Efficacy of Public School Choice.”
- Hatton, Erin (2014) “Temporary Weapons: Employers’ Use of Temps against Organized Labor,” *ILR Review*, 67 (1), 86–110.
- Hausman, Jerry A. and Paul A. Ruud (1987) “Specifying and testing econometric models for rank-ordered data,” *Journal of Econometrics*, 34 (1), 83–104.
- Heneman, Herbert G and Marcus H Sandver (1983) “Predicting The Outcome Of Union Certification Elections: A Review Of The Literature,” *Industrial And Labor Relations Review*, 36 (4), 24.
- Herkenhoff, Kyle, Jeremy Lise, Guido Menzio, and Gordon M. Phillips (2018) “Production and Learning in Teams,” Working Paper 25179, National Bureau of Economic Research.
- Hoxby, Caroline M. (2000) “Does Competition among Public Schools Benefit Students and Taxpayers?” *The American Economic Review*, 90 (5), 1209–1238.
- Hoxby, Caroline Minter ed. (2003) *The economics of school choice*, A National Bureau of Economic Research conference report, Chicago: University of Chicago Press.
- Hsieh, Chang-Tai and Miguel Urquiola (2006) “The effects of generalized school choice on achievement and stratification: Evidence from Chile’s voucher program,” *Journal of Public Economics*, 90 (8), 1477–1503.
- Ichniowski, Casey, Kathryn Shaw, and Giovanna Prennushi (1997) “The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines,” *American Economic Review*, 87 (3), 291–313.
- Idoux, Clemence (2022) “Integrating New York City Schools: The Role of Admission Criteria and Family Preferences,” 100.



- Ito, Koichiro and James M. Sallee (2018) “The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel Economy Standards,” *The Review of Economics and Statistics*, 100 (2), 319–336.
- Jarmin, Ron S. and Javier Miranda (2002) “The Longitudinal Business Database,” *SSRN Electronic Journal*.
- Jarosch, Gregor (2015) “Searching for Job Security and the Consequences of Job Loss,” 56.
- Jarosch, Gregor, Ezra Oberfield, and Esteban Rossi-Hansberg (2019) “Learning from Coworkers,” Technical Report w25418, National Bureau of Economic Research, Cambridge, MA.
- Jovanovic, Boyan (1979) “Job Matching and the Theory of Turnover,” *Journal of Political Economy*, 87 (5), 972–990.
- Jovanovic, Boyan and Yaw Nyarko (1997) “Stepping-stone mobility,” *Carnegie-Rochester Conference Series on Public Policy*, 46, 289–325.
- Kahn, L. B. and F. Lange (2014) “Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures,” *The Review of Economic Studies*, 81 (4), 1575–1613.
- Kahn-Lang, Ariella and Kevin Lang (2020) “The Promise and Pitfalls of Differences-in-Differences: Reflections on 16 and Pregnant and Other Applications,” *Journal of Business & Economic Statistics*, 38 (3), 613–620.
- Kleiner, Morris M. (2001) “Intensity of management resistance: understanding the decline of unionization in the private sector,” *Journal of Labor Research*, 22 (3), 519–540.
- Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar (2019) “Who Profits from Patents? Rent-Sharing at Innovative Firms,” *The Quarterly Journal of Economics*, 134 (3), 1343–1404.
- Knepper, Matthew (2020) “From the Fringe to the Fore: Labor Unions and Employee Compensation,” *The Review of Economics and Statistics*, 102 (1), 98–112.
- Kochan, Thomas A, Katz Harry, and McKersie Robert (1986a) *The Transformation of American Industrial Relations*, New York: Basic Books.
- Kochan, Thomas A, Robert McKersie, and John Chalykoff (1986b) “The Effects of Corporate Strategy and Workplace Innovations on Union Representation,” *ILR Review*, 16.
- Kornai, János, Eric Maskin, and Gérard Roland (2003) “Understanding the Soft Budget Constraint,” *Journal of Economic Literature*, 41 (4), 1095–1136.

- Lagos, Lorenzo (2019) “Labor Market Institutions and the Composition of Firm Compensation: Evidence from Brazilian Collective Bargaining.”
- LaLonde, Robert, Gerard Marschke, and Kenneth Troske (1996) “Using Longitudinal Data on Establishments to Analyze the Effects of Union Organizing Campaigns in the United States,” *Annales d’Économie et de Statistique* (41/42), 155.
- Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler (2019) “Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market,” *SSRN Electronic Journal*.
- Lange, Fabian (2007) “The Speed of Employer Learning,” *Journal of Labor Economics*, 25 (1), 1–35.
- Lazear, Edward P. (1979) “Why Is There Mandatory Retirement?” *Journal of Political Economy*, 87 (6), 1261–1284.
- (2009) “Firm-Specific Human Capital: A Skill-Weights Approach,” *Journal of Political Economy*, 117 (5), 914–940.
- Lee, David S. and Alexandre Mas (2012) “Long-Run Impacts of Unions on Firms: New Evidence from Financial Markets, 1961–1999,” *The Quarterly Journal of Economics*, 127 (1), 333–378.
- Lee, David S., Justin McCrary, Marcelo J. Moreira, and Jack R. Porter (2021) “Valid t-ratio Inference for IV,” Working Paper 29124, National Bureau of Economic Research.
- Levitt, Martin and Terry Conrow (1993) *Confessions of a Union Buster*, New York: Crown Publishers.
- Lise, Jeremy and Jean-Marc Robin (2017) “The Macrodynamics of Sorting between Workers and Firms,” *The American Economic Review*, 107 (4), 1104–1135.
- Logan, John (2002) “Consultants, lawyers, and the ‘union free’ movement in the USA since the 1970s,” *Industrial Relations Journal*, 33 (3), 197–214.
- Luca, Dara Lee and Michael Luca (2019) “Survival Of The Fittest: The Impact Of The Minimum Wage On Firm Exit,” *NBER Working Papers*.
- Maestas, Nicole, Kathleen J. Mullen, David Powell, Till von Wachter, and Jeffrey B. Wenger (2018) “The Value of Working Conditions in the United States and Implications for the Structure of Wages,” Working Paper 25204, National Bureau of Economic Research.
- Manski, Charles F. (1993) “Identification of Endogenous Social Effects: The Reflection Problem,” *The Review of Economic Studies*, 60 (3), 531–542.

- Mas, Alexandre and Amanda Pallais (2019) “Labor Supply and the Value of Non-Work Time: Experimental Estimates from the Field,” *American Economic Review: Insights*, 1 (1), 111–126.
- McCue, Kristin (1996) “Promotions and Wage Growth,” *Journal of Labor Economics*, 14 (2), 175–209.
- McMillan, Robert (2004) “Competition, incentives, and public school productivity,” *Journal of Public Economics*, 88 (9), 1871–1892.
- Menezes-Filho, Naércio Aquino, Marc-Andreas Muendler, and Garey Ramey (2008) “The Structure of Worker Compensation in Brazil, with a Comparison to France and the United States,” *Review of Economics and Statistics*, 90 (2), 324–346.
- Milgrom, Paul and John Roberts (1992) *Economics, Organization and Management*, Englewood Cliffs, N.J: Pearson, 1st edition.
- Milne, Cynthia N (1985) “Double-Breasted Operations: A Management Perspective,” *The Labor Lawyer*, 1 (4), 12.
- Mincer, Jacob A. (1974) *Schooling, Experience, and Earnings*: NBER.
- Moscarini, Giuseppe and Fabien Postel-Vinay (2018) “The Cyclical Job Ladder,” *Annual Review of Economics*, 10 (1), 165–188.
- Munger, Peter, Stephen X Mungertt, and Thomas J Mungerttt (1988) “Plant Closures And Relocations Under The National Labor Relations,” *Georgia State University Law Review*, 5 (1), 41.
- Muralidharan, Karthik and Venkatesh Sundararaman (2015) “The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India \*,” *The Quarterly Journal of Economics*, 130 (3), 1011–1066.
- Murphy, Kevin J. (1985) “Corporate performance and managerial remuneration: An empirical analysis,” *Journal of Accounting and Economics*, 7 (1), 11–42.
- Neilson, Christopher A (2021) “Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students.”
- Nelson, Douglas (1996) “The Political Economy of U.S. Automobile Protection,” in *The Political Economy of American Trade Policy*, 133–196: University of Chicago Press.
- Neumark, David and Michael L. Wachter (1995) “Union Effects on Nonunion Wages: Evidence from Panel Data on Industries and Cities,” *Industrial and Labor Relations Review*, 49 (1), 20–38.
- Newey, Whitney K. and Daniel McFadden (1994) “Chapter 36 Large sample estimation and hypothesis testing,” in *Handbook of Econometrics*, 4, 2111–2245: Elsevier.

- NLRB (2020) “National Labor Relations Board: Casehandling Manual: Part Two Representation Proceedings.”
- OECD (2019) “Higher education in Brazil,” in *Rethinking Quality Assurance for Higher Education in Brazil*, 65–90: OECD.
- OECD-IDB (2014) “Detailed description of Employment Protection Legislation in LAC countries.”
- O’Flaherty, Brendan and Aloysius Siow (1995) “Up-or-Out Rules in the Market for Lawyers,” *Journal of Labor Economics*, 13 (4), 709–735.
- Oreopoulos, Philip, Till von Wachter, and Andrew Heisz (2012) “The Short- and Long-Term Career Effects of Graduating in a Recession,” *American Economic Journal: Applied Economics*, 4 (1), 1–29.
- Oster, Emily (2019) “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business & Economic Statistics*, 37 (2), 187–204.
- Pastorino, Elena (2019) “Careers in Firms: the Role of Learning and Human Capital.”
- Pathak, Parag A. and Peng Shi (2018) “How Well Do Structural Demand Models Work? Counterfactual Predictions in School Choice,” SSRN Scholarly Paper ID 3356092, Social Science Research Network, Rochester, NY.
- Perry, Guillermo E., Omar Arias, Pablo Fajnzylber, William F. Maloney, Andrew Mason, and Jaime Saavedra-Chanduvi (2007) *Informality: Exit and Exclusion*: The World Bank.
- Postel-Vinay, Fabien and Jean-Marc Robin (2002) “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 70 (6), 2295–2350.
- Prendergast, Canice (1993) “The Role of Promotion in Inducing Specific Human Capital Acquisition,” *The Quarterly Journal of Economics*, 108 (2), 523–534.
- Prince, Jeffrey T. and Daniel H. Simon (2017) “The Impact of Mergers on Quality Provision: Evidence from the Airline Industry,” *The Journal of Industrial Economics*, 65 (2), 336–362.
- Ridley, Matthew and Camille Terrier (2018) “Fiscal and Education Spillovers from Charter School Expansion,” Working Paper 25070, National Bureau of Economic Research.
- Roomkin, Myron and N Block (1981) “Case Processing Time And The Outcome Of Representation Elections: Some Empirical Evidence,” *University Of Illinois Law Review* (4), 25.
- Roth, Alvin E. (1982) “The Economics of Matching: Stability and Incentives,” *Mathematics of Operations Research*, 7 (4), 617–628.

- Roth, Jonathan and Pedro H. C. Sant’Anna (2021) “When Is Parallel Trends Sensitive to Functional Form?” *arXiv:2010.04814 [econ, stat]*.
- Rothstein, Jesse (2007) “Does Competition Among Public Schools Benefit Students and Taxpayers? Comment,” *American Economic Review*, 97 (5), 2026–2037.
- Sandver, Marcus Hart and Kathryn J. Ready (1998) “Trends in and determinants of outcomes in multi-union certification elections,” *Journal of Labor Research*, 19 (1), 165–172.
- Schmieder, Johannes F., Till von Wachter, and Joerg Heining (2020) “The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany.”
- Schmitt, John and Ben Zipperer (2009) “Dropping the Ax: Illegal Firings During Union Election Campaigns, 1951-2007,” *CEPR Working Papers*, 22.
- Selten, Reinhard (1978) “The chain store paradox,” *Theory and Decision*, 9, 127–159.
- Slichter, Sumner, James Healey, and Robert Livernash (1960) *The Impact of Collective Bargaining on Management*, Washington DC: Brookings Institution.
- Sojourner, Aaron J., Brigham R. Frandsen, Robert J. Town, David C. Grabowski, and Min M. Chen (2015) “Impacts of Unionization on Quality and Productivity: Regression Discontinuity Evidence from Nursing Homes,” *ILR Review*, 68 (4), 771–806.
- Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter (2019) “Firming Up Inequality,” *The Quarterly Journal of Economics*, 134 (1), 1–50.
- Sorkin, Isaac (2018) “Ranking Firms Using Revealed Preference,” *The Quarterly Journal of Economics*.
- Sun, Liyang and Sarah Abraham (2020) “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 53.
- (2021) “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 225 (2), 175–199.
- Sweeting, Andrew (2013) “Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry,” *Econometrica*, 81 (5), 1763–1803.
- Taber, Christopher and Rune Vejlin (2020) “Estimation of a Roy/Search/Compensating Differential Model of the Labor Market,” *Econometrica*, 88 (3), 1031–1069.
- Taschereau-Dumouchel, Mathieu (2020) “The Union Threat,” *The Review of Economic Studies*, 87 (6), 2859–2892.

- Topel, Robert H. and Michael P. Ward (1992) "Job Mobility and the Careers of Young Men," *The Quarterly Journal of Economics*, 107 (2), 439–479.
- Traxler, Franz (1994) "Collective Bargaining: Levels and Coverage," in *OECD Employment Outlook*: OECD.
- Urquiola, M. (2016) "Competition Among Schools," in *Handbook of the Economics of Education*, 5, 209–237: Elsevier.
- Verma, Anil (1985) "Relative Flow of Capital to Union and Nonunion Plants Within a Firm," *Industrial Relations*, 24 (3), 395–405.
- Waldman, Michael (1984) "Job Assignments, Signalling, and Efficiency," *The RAND Journal of Economics*, 15 (2), 255–267.
- Walters, Christopher R. (2018) "The Demand for Effective Charter Schools," *Journal of Political Economy*, 126 (6), 2179–2223.
- Wamsley, Laurel (2017) "Billionaire Owner Shuts Down DNAinfo, Gothamist Sites A Week After Workers Unionize," *NPR*.
- Wasi, Nada and Aaron Flaaen (2015) "Record Linkage Using Stata: Preprocessing, Linking, and Reviewing Utilities," *The Stata Journal: Promoting communications on statistics and Stata*, 15 (3), 672–697.
- Weiler, Paul (1983) "Promises to Keep: Securing Workers' Rights to Self-Organization under the NLRA," *Harvard Law Review*, 96 (8), 1769.
- Whitehurst, Grover J. "Russ" (2017) "The 2016 Education Choice and Competition Index," March.
- Zimmerman, Anne (2000) "Pro-Union Butchers at Wal-Mart Win a Battle, but Lose the War," *Wall Street Journal*.