

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Essays on Local Labor Markets

Permalink

<https://escholarship.org/uc/item/68q9011f>

ISBN

9798189275906

Author

Jin, Richard

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Essays on Local Labor Markets

by

Richard Jin

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Card, Chair

Professor Enrico Moretti

Professor Jesse Rothstein

Spring 2026

Essays on Local Labor Markets

Copyright 2026
by
Richard Jin

Abstract

Essays on Local Labor Markets

by

Richard Jin

Doctor of Philosophy in Economics

University of California, Berkeley

Professor David Card, Chair

This dissertation studies how institutions of higher education and local labor markets impact the labor market outcomes of college-educated workers in the United States. Chapter 1 focuses on the interaction between colleges and local labor markets, and in particular how they shape geographic mobility. Chapter 2 highlights how local labor market conditions impact a wide range of individual labor market outcomes. Chapter 3 analyzes the effects of changes at the school level on various outcomes, including geographic mobility.

In Chapter 1, which I co-author with Heidi Artigue, we quantify the impact of alumni networks on the geographic mobility of job seekers for nearly 1,400 US colleges and universities. We use detailed employment and education information on LinkedIn users to isolate college-educated workers who faced an exogenous job separation in a mass layoff or firm closure. Using a nested logit model of location choice, we compare the migration decisions of job seekers who were displaced in the same city and who attended different but similar and geographically proximate universities. We find that a 1% increase in the number of co-alumni in the city of displacement increases a job seeker's odds of staying there by 0.4%. Conditional on moving, a 1% increase in a potential destination's number of co-alumni increases the odds of choosing that city over another by 0.9%. Co-alumni may both impact job search and provide local amenities. Using data on the presence or absence of co-alumni at new jobs, we conclude that the job search channel is particularly important. Co-alumni from the same or neighboring graduating class have much larger impacts on location choice, indicating true network effects rather than idiosyncratic matches between alumni of certain colleges and jobs in certain cities. We also find strong impacts of having more local co-alumni who work in the same industry.

In Chapter 2, I study the labor market impacts of a recession on high-skill workers by quantifying the medium and long-run effects of the 2001 Dot-Com bubble burst on employment and geographic mobility. I find that college-educated workers in the US who were exposed to larger local unemployment shocks were significantly less likely to be employed in the

years following the recession, with eventual recovery in the long run. Workers who became re-employed ended up at smaller firms, but did not systematically switch into lower-paying industries. Impacts on mobility were particularly large. Individuals in more negatively affected local labor markets were not only more likely to move, but also relocated to lower-paying and cheaper places. Effects on employment and migration outcomes were primarily concentrated among workers in the tech sector.

Chapter 3 studies the impact of widespread expansions in Computer Science and Engineering majors across US colleges and universities during the 21st century on individual labor market outcomes. Using rich data on individual educational and employment histories, I compare the outcomes of students who entered their undergraduate institution as a rapid expansion in the production of Computer Science or Engineering degrees took place to cohorts at the same institution who had completed major declarations just prior to the expansion. I find that relative to pre-expansion counterparts, students in first post-expansion cohorts were less likely to obtain an advanced degree and more likely to be employed in higher-paying, high-technology industries and smaller firms after graduation. I also find strong short-run effects on geographic mobility—individuals in post-expansion cohorts were more likely to stay in the same location as their undergraduate institution immediately following graduation. However, five years after graduation, this reduced out-mobility largely disappears, suggesting that the rampant increase in the production of local STEM graduates did not directly lead to a larger local college-educated workforce beyond the short run.

Altogether, I show that both the locations and educational institutions of college graduates meaningfully shape their post-graduation economic outcomes. These impacts are far-reaching, spanning a broad range of locations and schools and encompassing millions of individuals.

Contents

Contents	i
List of Figures	iii
List of Tables	v
1 College Alumni Networks and Mobility Across Local Labor Markets	1
1.1 Introduction	1
1.2 Data	7
1.3 Theoretical Framework and Empirical Strategy	12
1.4 Identification	18
1.5 Results	23
1.6 The Role of Peers and Same-Industry Networks	32
1.7 Conclusion	36
1.8 Figures and Tables	38
2 The Dynamics of High-Skill Worker Reallocation: Evidence from the Dot-Com Bubble	56
2.1 Introduction	56
2.2 Data	58
2.3 Empirical Strategy	60
2.4 Results	62
2.5 Conclusion	66
2.6 Figures and Tables	68
3 The Labor Market Impacts of 21st Century STEM Expansions in US Higher Education	81
3.1 Introduction	81
3.2 Data	83
3.3 Empirical Strategy	85
3.4 Results	87
3.5 Conclusion	89

3.6	Figures and Tables	91
Bibliography		97
A	Additional Material for College Alumni Networks and Mobility Across Local Labor Markets	104
A.1	Appendix Figures and Tables	104
A.2	Data Appendix	123
A.3	Model Derivations	125
A.4	Additional Notes for Empirical Exercises	127
B	Additional Material for The Dynamics of High-Skill Worker Reallocation: Evidence from the Dot-Com Bubble	131
B.1	Appendix Figures	131
C	Additional Material for The Labor Market Impacts of 21st Century STEM Expansions in US Higher Education	135
C.1	Appendix Tables	136

List of Figures

1.1	College Alumni Network Size Across Local Labor Markets, by School	38
1.2	Differences in Alumni Network Size Across Local Labor Markets, by School Pair	39
1.3	Probability of Relocation, Large vs. Small Local Alumni Network	40
1.4	Relationship Between Probability of Staying in City of Displacement and log Alumni Network Size	41
1.5	Concentration of and Re-Employment at Local In-Network Firms	42
1.6	Effect of Local Alumni Network Size on Mobility, by Age	43
1.7	Effect of Same-Cohort and Non-Cohort Co-Alumni Network Size on Out-mobility and Directed Migration	44
1.8	Effect of Same-Industry Co-Alumni Network Size on Out-mobility	45
1.9	Relationship Between Changes in Local Alumni Network Size and Changes in Citywide Earnings Premia	46
2.1	Local Dot-Com Recession Shocks, by Commuting Zone	68
2.2	Impact of Dot-Com Local Shocks on Employment	69
2.3	Impact of Dot-Com Local Shocks on Size of Future Employers	70
2.4	Impact of Dot-Com Local Shocks on Change in Industry Earnings Premium Relative to Initial Industry	71
2.5	Impact of Dot-Com Local Shocks on Geographic Mobility	72
3.1	Impact of Partial and Full Exposure to CS and Engineering Expansions on Completion of CS and Engineering Degrees	91
3.2	Impact of Partial and Full Exposure to CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation	92
3.3	Impact of Partial and Full Exposure to CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation	93
A.1	Pooled 2010-2019 Industry Shares for Prime-Age College-Educated Workers, Revelio vs. ACS	104
A.2	Average Yearly Degrees Awarded from 2010-2019, Revelio vs. IPEDS	105
A.3	A College-Educated Job Seeker on LinkedIn Can Access Many Pieces of Information about a Local Alumni Network	106
A.4	Job Seeker's Location Choice Decision in Work History Timeline	107

A.5	Job Seeker's Location Decision Tree	108
A.6	Predictive Accuracy of School Strata in Explaining School Major Share	109
B.1	Impact of Dot-Com Local Shocks on Employment, Longer Time Horizon	131
B.2	Impact of Dot-Com Local Shocks on Employment, Within-School Variation	132
B.3	Impact of Dot-Com Local Shocks on Employment, No Pre-Recession Adjustment	133
B.4	Impact of Dot-Com Local Shocks on Geographic Mobility, Within-School Variation	134

List of Tables

1.1	Summary Statistics for Job Seekers at Time of Displacement	47
1.2	Effect of Alumni Network Size in City of Displacement on Probability of Staying	48
1.3	Effect of Alumni Network Size in City of Displacement on Probability of Staying — Logit and Saturated Linear Probability Models	49
1.4	Effect of Alumni Network Size in City of Displacement on Probability of Staying — Control Function Approach with Instrumental Variable	50
1.5	Effect of Alumni Network Size in City of Displacement on Probability of Staying and Re-Employment with or without Co-Alumni	51
1.6	Effect of Alumni Network Size in Destination City on Directed Migration	52
1.7	Effect of Alumni Network Size in Destination City on Directed Migration — Control Function Approach with Instrumental Variable	53
1.8	Effect of Destination Alumni Network Size on Directed Migration for Job Seekers Who Move and are Re-employed with Versus without Co-alumni at New Firm	54
1.9	Effect of Same-Industry Alumni Network Size in Destination City on Directed Migration	55
2.1	Worker Summary Statistics	73
2.2	Impact of Dot-Com Local Shocks on Change in Commuting Zone Earnings Pre- mium Relative to Initial Location	74
2.3	Impact of Dot-Com Local Shocks on Change in Commuting Zone Log Median Rent Relative to Initial Location	75
2.4	Impact of Dot-Com Local Shocks on 2006 Employment Relative to Pre-2000 Average for Tech and Non-Tech Workers	76
2.5	Impact of Dot-Com Local Shocks on Change in Size of Future Employers for Tech and Non-Tech Workers	77
2.6	Impact of Dot-Com Local Shocks on Change in Industry Earnings Premium Rel- ative to Initial Industry for Tech and Non-Tech Workers	78
2.7	Impact of Dot-Com Local Shocks on Moving out of Initial Location for Tech and Non-Tech Workers	79
2.8	Impact of Dot-Com Local Shocks on Change in Commuting Zone Earnings Pre- mium Relative to Initial Location for Tech and Non-Tech Workers	80
3.1	Impact of CS and Engineering Expansions on Major Completions	94

3.2	Impact of CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation	95
3.3	Impact of CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation	96
A.1	Variation in School Major Share Explained by School Strata	110
A.2	Effect of Alumni Network Size in Destination City on Location Choice from Multinomial Logit Framework	111
A.3	Effect of Destination-Origin Difference in Alumni Network Size on Directed Migration	112
A.4	Effect of Alumni Network Size (of Bachelor's Institution) in City of Displacement on Probability of Staying	113
A.5	Effect of Alumni Network Size (of Bachelor's Institution) in Destination City on Directed Migration	113
A.6	Effect of Alumni Network Size in City of Displacement on Probability of Staying for Workers Displaced Outside of Home City	114
A.7	Effect of Destination Alumni Network Size on Directed Migration for Workers Displaced Outside of Home City	114
A.8	Effect of Alumni Network Size in Destination City on Directed Migration for Workers who are Displaced Outside of Home City and Do Not Return to Home City	115
A.9	Estimated Impact of Alumni Network Size in City of Displacement on Probability of Staying is Robust to Including Location-Specific Demographic Heterogeneity	116
A.10	Inclusion of Location-Specific Demographic Heterogeneity Does Not Impact Estimated Effect of Alumni Network Size in Destination City on Directed Migration	117
A.11	Estimated Impact of Alumni Network Size in City of Displacement on Probability of Staying is Robust to Exploiting Narrower Variation Within Firm-Strata	118
A.12	Inclusion of Firm-Specific Heterogeneity Does Not Impact Estimated Effect of Alumni Network Size in Destination City on Directed Migration	119
A.13	Controlling for Major Does Not Affect Estimated Effect of Alumni Network Size in City of Displacement on Probability of Staying	120
A.14	Inclusion of Major-Specific Heterogeneity Does Not Impact Estimated Effect of Alumni Network Size in Destination City on Directed Migration	121
A.15	Effect of Alumni Network Size (2010-2019 Yearly Average) in City of Displacement on Probability of Staying	122
A.16	Effect of Alumni Network Size (2010-2019 Yearly Average) in Destination City on Directed Migration	122
C.1	Top 50 Computer Science and Engineering Expansions	136
C.2	Impact of CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation, No Great Recession Cohorts	139

C.3	Impact of CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation, No Great Recession Cohorts	139
C.4	Impact of Partial and Full Exposure to CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation, No Great Recession Cohorts . . .	140
C.5	Impact of Partial and Full Exposure to CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation, No Great Recession Cohorts . .	140
C.6	Impact of Largest CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation	141
C.7	Impact of Largest CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation	141
C.8	Impact of Partial and Full Exposure to Largest CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation	142
C.9	Impact of Partial and Full Exposure to Largest CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation	142

Acknowledgments

I have been so fortunate to have had many wonderful teachers and mentors at Peddie, Williams, Wharton and Berkeley. Without their dedication and belief in me, this dissertation would not have been possible.

I would like to especially thank my committee members. To Enrico, thank you for always providing valuable feedback on my work and for your support. It has been such a privilege to work with one of my research idols from before the Ph.D. even started. Jesse, thank you for being a fantastic teacher in the Labor field course, for the opportunity to work on a variety of really cool projects as a GSR, and for your valuable feedback on my work and support. And to Dave, thank you for everything— from GSR-ing after my first year, to all the helpful talks about research, to chats about life in general. Your mentorship and support have meant more to me than you'll ever know, and I am eternally grateful.

Chapter 1 is co-authored with Heidi Artigue, who has been a wonderful collaborator. I also thank the team of brilliant research assistants that helped with this project— Emily Fan, Ari Pereira Gomes, Stephen Liu and Rocky Zhao.

I would like to thank my friends. You are all individuals with such immense talents and of exceptional character. Included among this group, are my office mates in Evans 666. It has been such a joy and honor to progress through this journey alongside you.

Finally, above all else, thank you to my parents. You have been my strongest support system during my hardest moments, and my biggest fans during my wins. I am and will always be in awe of you both. You have accomplished so much in your lives and overcome so many obstacles in the process— hardships that I will never have to endure because your capacity to give love and nourish our family is limitless. Out of the many blessings I've had in this life, being your son is the greatest of them all. I'm excited to join you both in the Ph.D. club, and can finally invoke one of my favorite movie quotes, from *Step Brothers*: "This is a house of learned doctors!"

Chapter 1

College Alumni Networks and Mobility Across Local Labor Markets

1.1 Introduction

An individual's choice of where to live and work affects every aspect of their life. Places vary in both their economic opportunities and the local amenities that they provide to workers, and the co-location of workers gives rise to agglomeration economies that benefit both workers and the firms that employ them (Glaeser and Gottlieb 2009). The location choices of college-educated workers are particularly policy relevant, given evidence of the externalities from a better-educated local workforce (Moretti 2004a; Shapiro 2006) and the rising concentration of these workers in relatively few cities, both in the US (Moretti 2012; Diamond and Gaubert 2022) and other high-income countries (Iammarino, Rodriguez-Pose, and Storper 2019). These economic consequences are magnified by the fact that college graduates are highly mobile across local labor markets (Molloy, Smith, and Wozniak 2011; Jia et al. 2023).

College alumni networks are an obvious mechanism that help *push* and *pull* the graduates from a particular college to a particular set of cities. Nearly all college graduates have expansive networks of co-alumni in many local labor markets, with some cities having tens of thousands of alumni from a single school. Such networks arguably play an important role in job search. Indeed, the online platform LinkedIn shows job seekers the employers of not only direct contacts but also other alumni from their *alma maters* — presumably because personal contacts can aid in the job matching process. Alumni networks can also help employers learn about the graduates from particular schools, improving the precision of signals in the matching process (Dustmann et al. 2016). Although there is a lot of work on how interpersonal networks affect job search,¹ and job search is by far the most

¹Existing work on network effects looks at neighbors (Bayer, Ross, and Topa 2008; Hellerstein, Kutzbach, and Neumark 2019), family members (Kramarz and Skans 2014), and classmates (Kramarz and Thesmar 2013; Eliason et al. 2023; Ost, Pan, and Webber 2024; Zhu 2025).

commonly-stated reason for moving (Jia et al. 2023), there is surprisingly little evidence on how interpersonal networks of any kind, and alumni networks in particular, directly impact location choice.

In addition to providing information that can be helpful in the job search and matching process, the presence of local co-alumni can directly impact the value of a particular city if graduates value social interactions with their college friends and the friends of their friends. Many alumni networks also have local branches that organize social events, reinforcing the strength of alumni connections and promoting local ties.

This paper provides the first causal evidence on how local college alumni networks affect mobility and hiring for workers across the US. Using LinkedIn data, we identify nearly 20,000 college graduates who are observed on a job in the period 2010-2018 and subsequently are displaced from a mass layoff or firm closure. We follow these workers in the year after displacement and observe whether they remain in the same original metro area (their “origin”) or move to a new metro area (a new “destination”). We model the decision of both *whether* to move and *where* to move using a nested logit structure in which a displaced job seeker first chooses to stay in the origin or leave, and then conditional on leaving, chooses among all possible new destinations. Accordingly, we estimate both the impact of origin alumni network size on the probability of remaining in the origin after a displacement event, and the impact of destination alumni network size on the choice over alternative cities. In our framework, local alumni networks increase the number of job leads for local opportunities and also provide potential local amenities. We highlight the importance of the job search channel by distinguishing between the probability of re-employment at local firms that employ other alumni from the same college, and at firms in destination cities with and without coworkers from the same college.

Our approach overcomes several substantial challenges to understanding how college networks impact the labor market. First, we observe a panel of workers with detailed information on both their *alma maters* and their locations and the identities of their employers over time. In the US, administrative data typically has either individual education histories or worker locations obtained from employment histories, but not both.² As a result, data availability has hindered research on the impacts of networks on geographic mobility.

Second, even with specialized data, the researcher must contend with several threats to identification of the effects of local alumni network size. Most prominently, worker location choice is endogenous. Workers who voluntarily leave their firms to change jobs and/or locations are not randomly selected. We isolate a sample of workers who were exogenously separated from their jobs and consequently had to make a new decision about whether to remain in their city of displacement or move to a new location.

Next, one must contend with the endogeneity of college alumni networks. It is natural to expect that graduates from the same institution could have correlated preferences for location characteristics. For example, NYU graduates may tend to favor large cities, or

²Some countries outside of the US have administrative datasets that can link education, employment and location histories, but have insufficiently many local labor markets to study geographic mobility.

UCSD graduates coastal cities. If so, we would expect to see workers move to places with a large alumni network, even in the absence of any treatment effect. We address concerns about correlated preferences among alumni by leveraging LinkedIn’s broad coverage of college-educated US workers to facilitate narrow worker comparisons. To identify and estimate the effect of alumni network size in the origin on the probability of finding re-employment in the origin, we compare job seekers who were displaced in the same city and went to observably similar schools in the same area, but who have differently-sized alumni networks in their city of work at the time of displacement. For example, we compare two workers who were both laid off due to mass layoffs in Denver, CO, one of whom attended Williams College and the other Amherst College, both highly selective small liberal arts colleges located in rural Massachusetts that also have similar distributions of majors. Our identifying assumption is that any unobserved location preferences are absorbed by a combination of fixed effects for the city of displacement, the location of the *alma mater*, and the type of *alma mater* as defined by size, selectivity and other institutional characteristics.

Another source of endogeneity in alumni networks is the possibility of idiosyncratic match effects between the alumni of certain colleges and the jobs in a certain city. We address this concern in several ways. First, we further refine our comparisons between displaced workers in the same location who attended different but similar and geographically adjacent schools. Specifically, we additionally compare such workers who were also displaced in the same industry or even the same firm. Next, we show that estimated effects are largest for networks of alumni from the same or neighboring cohorts — groups of students with whom a job seeker is most likely to have interacted. We also directly absorb match effects across all city-school pairs in extended analyses of the impact of more granular co-alumni networks on location choice. Finally, we implement a control function approach using an instrumental variable for a school’s local alumni network size based on a combination of the historical share of in-state students at the college and the distance from the college to a given city.

Our first key finding is that an increase in the size of a worker’s local college alumni network decreases their likelihood of relocating to a different local labor market one year after displacement. Specifically, increasing the network size in the city of displacement by 1% (the average number of co-alumni in the city of displacement is approximately 17,000) increases a worker’s odds of staying by 0.39%. This result is driven primarily by increases in the probability of remaining in the city of displacement and finding a new job at a firm that employs co-alumni: increasing the local college alumni network size by 1% increases a displaced worker’s odds of staying in the same local labor market and finding re-employment at a firm with at least one co-alumni by an even larger 0.55%. The latter effect is surprising given that for the average job seeker in our sample, only 1 in 15 firms in the displacement city have a co-alum. The effects are also larger for older job seekers, who are generally less likely to move than younger job seekers.

We then investigate whether, conditional on moving to a new destination, an increase in a potential destination’s alumni network size increases the job seeker’s odds of choosing that destination as their place of work and residence. We derive and estimate a specification similar to a “gravity equation” that relates the number of job seekers from a school who are

displaced in a particular origin and choose a specific destination to the number of co-alumni in each destination, as well as other determinants of job seekers' utility in each destination. Our identifying assumption is that, outside of the city of displacement, correlated preferences for distant locations do not differ among the alumni of a sufficiently narrowly defined set of geographically adjacent schools. For instance, we assume that Williams and Amherst alumni who are displaced in Denver do not meaningfully differ in their latent preferences for living and working in Dallas, but allow arbitrary differences in preferences between Williams alumni and UT Austin alumni.

We test for relevant threats to identification to strengthen the credibility of our results. Our results are robust to incorporating location-varying heterogeneity in demographic characteristics in our model, which accounts for correlated preferences in a richer way. We rule out labor demand factors common to job seekers of the same major, industry, or firm, as well as locational preferences for and return migration to home locations. We estimate that, conditional on moving, a 1% increase in the number of co-alumni working in an alternative local labor market increases the odds of a displaced job seeker also choosing to work and live there, relative to any other place, by 0.91%. Thus, an increase in a destination's alumni network size substantially improves a job seeker's likelihood of locating there.

We also show that destination network size has a larger impact on the joint flow of job seekers to the destination and to firms with other co-alumni, than on the joint flow of job seekers to the destination and to firms without other co-alumni. This finding is consistent with local alumni networks impacting geographic mobility in large part through their assistance in job finding. We also show that conditional on leaving the city of displacement, destination choices among younger job seekers are far more responsive than those of their older counterparts to differences in the alumni network sizes in alternative cities. Whereas older job seekers may have stronger networks across different cities that originated outside of college, younger job seekers are less likely to have built these additional networks and more likely to rely on college alumni networks.

To lend further credibility to our estimates, we construct instrumental variables for origin and destination alumni network size based on the interaction between a school's historical share of in-state students (calculated in 1972) and the inverse distance between the school and different cities. This approach isolates variation in network size that comes from historically driven differences in the "home bias" of graduates from a given school, combined with relative distance to different cities. We use these instrumental variables in a control function approach within our nested logit model to re-estimate the impacts of local alumni network size on both out-mobility and directed migration. The resulting estimates of the effect of network size are very similar to our baseline estimates, and show significant effects of alumni network size on the decision to stay in or leave the city of displacement, and on the choice of alternative destinations for those who move.

Finally, we examine additional mechanisms by which local alumni networks affect location choice. We find that same- and neighboring-cohort alumni networks, which are more likely to consist of direct peers, influence location choice more strongly than the vastly larger networks of co-alumni from distant cohorts with whom a job seeker did not overlap during school. We

also show that industry-specific human capital can affect job seekers' geographic mobility. If workers have accrued significant experience in a particular industry, then switching industries can be costly (Parent 2000). Indeed, we find that displaced job seekers are more likely to locate in cities with larger networks of co-alumni who work in their industry of displacement, consistent with Moretti and Yi (2025).

In these latter analyses, we construct extremely granular alumni networks, which allows us to absorb additional threats to identification that vary across city-school pairs in our baseline analyses. Altogether, we show that local college alumni networks, which are spread across local labor markets and are salient social networks for tens of millions of college-educated workers, have a significant impact on geographic mobility.

Our work contributes to a broad literature on interpersonal networks and labor market outcomes, which has focused on hiring as the primary outcome of interest. Previous studies have found positive effects on hiring from residential networks (Bayer, Ross, and Topa 2008; Hellerstein, Kutzbach, and Neumark 2019), family networks (Kramarz and Skans 2014), and educational networks (Kramarz and Thesmar 2013; Zimmerman 2019; Eliason et al. 2023; Ost, Pan, and Webber 2024; Zhu 2025). In these studies, networks facilitate hiring to a specific firm via direct referrals from known individuals.

We contribute to this literature by providing the first evidence on how interpersonal networks *jointly* affect both geographic mobility and job search. Recent work has demonstrated the impact of smaller personal networks on mobility, including from family members (Huttunen, Møen, and Salvanes 2018), cell phone contacts (Büchel et al. 2020; Blumenstock, Chi, and Tan 2025), birthplaces (Stuart and Taylor 2021; Buggle et al. 2023), Facebook friends (Koenen and Johnston 2024), and randomly-assigned World War II Navy shipmates (Green 2024). We link this body of work to the literature on interpersonal networks and job search. Our setting is advantageous since we have rich individual-level data on employment histories and various labor market outcomes. This information allows us to directly conceptualize geographical areas as not just places but rather as local labor markets in which large social networks actively help job seekers find re-employment. We further emphasize this role by focusing exclusively on exogenously separated job seekers who are forced to both search for a new job and potentially relocate.

Our second key contribution is to extend the scope of relevant networks far beyond individual-level connections to the entire universe of alumni. Widely used job search platforms such as LinkedIn provide information and suggestions to job seekers of firms with co-alumni who are not direct contacts. Estimated treatment effects from our networks definition also pick up signaling benefits that one's *alma mater* provides to firms. If a city has a large network of alumni from a particular school spread across many firms, and those firms in turn learn that the school's alumni are a good match, then the firms may be more likely to hire additional alumni. This channel may be particularly relevant for younger college graduates without additional networks from previous employers (Cingano and Rosolia 2012; Hensvik and Skans 2016). In this regard, we additionally contribute to a more recent literature on the impact of online networks on labor market outcomes, e.g. individual Facebook friends (Gee, Jones, and Burke 2017; Koenen and Johnston 2024). We study even larger

online networks and provide the first evidence on how an individual’s set of all co-alumni (observed online) from their *alma mater* can shape mobility.

Our focus on college alumni networks, most directly related to the aforementioned studies on educational networks’ labor market impacts, also moves beyond the existing literature’s firm-specific network measures, e.g. the number or share of classmates at a firm. In our setting, if job search manifests across an entire local labor market, then our treatment mechanism has a scope beyond a single firm. Our data and approach are uniquely suited to capture this phenomenon, with detailed information on educational histories and locations of over 50 million workers per year on average. We compute alumni network sizes for nearly every 4-year US college/university in each US local labor market, solidifying our study’s external validity. By contrast, previous studies of firm-level educational networks are difficult to scale up to the local labor market level since they use data from settings covering a narrow set of schools and geographies (Eliason et al. 2023; Ost, Pan, and Webber 2024; Zhu 2025).

Altogether, our results suggest a multifaceted value of local college alumni networks for job seekers. First, the network of alumni from a displaced college-educated worker’s *alma mater* in the local labor market of displacement provides insurance against displacement’s negative labor market impacts by facilitating re-employment in the same local labor market. Co-alumni networks can thus shield workers from having to relocate across metros, which can entail high costs of moving (Kennan and Walker 2011). At the same time, job seekers are also drawn to large co-alumni networks in distant cities, which themselves provide consumption amenities but also facilitate re-employment in a new place.

College alumni networks’ impacts on geographic mobility have direct implications for job seekers’ economic well-being. Local labor markets have earnings premia that are not solely attributable to worker sorting (Card, Rothstein, and Yi 2025), so that location choice directly affects earnings. Moreover, in the presence of human capital spillovers, an increased flow of college-educated workers to any particular city raises the wages of both college-educated and non-college-educated workers in the city (Moretti 2004b, 2004a). Additionally, an increased number of college graduates working in a city’s tradeable sector creates “local multipliers” (Moretti 2010) by boosting employment in the non-tradeable sector.

The remainder of this paper is organized as follows. In Section 1.2, we describe our various data sources and present descriptive facts about the geographic variation in network sizes and characteristics of our sample of job seekers. In Section 1.3, we present a model of location choice, which we use to estimate treatment effects on our two margins of mobility. We address potential identification concerns in Section 1.4. Section 1.5 presents our main results. We provide additional evidence in Section 1.6, where we consider more granular definitions of local co-alumni networks, based on cohort and industry, and evaluate their effects on mobility. Section 1.7 concludes.

1.2 Data

We combine data on individual job histories, self-reported educational histories, and monthly flows into and out of firms. Our data covers workers in the US from 2010-2019, and is provided by the company Revelio Labs.³ Revelio Labs creates a standardized database of employment records by collecting information from online public profiles, e.g. LinkedIn, resumes, and job postings. We use these data files to create a sample of exogenously displaced job seekers for whom we can observe educational histories and both pre- and post-displacement job locations. We now discuss relevant aspects of the data in more detail.

Individual Job and Educational Histories

Our empirical exercise requires observing an individual’s employing firm(s), location(s) of employment, and timing of spells of employment and job search. Thus, our primary data file consists of 226 million job spells for 86 million individuals working in the US between 2010 and 2019. Each spell is linked to a worker, so that we can track workers over time. For each spell, we observe the position start date and end date (if the spell has been completed and its end recorded). For each position, we also observe occupation title, a measure of occupational seniority, predicted salary, and the state and metropolitan area of employment.⁴ For 80 percent of the job spells in our sample, we also know the associated company and industry defined by 2-digit NAICS code. We restrict our data to before the COVID-19 pandemic to minimize the role of remote workers in our sample, and treat the metro area of employment as the metro area of residence for all workers. We leverage the observed locations of each position and the panel structure of the data to track worker mobility across local labor markets in the US. There are 102 metropolitan areas delineated in the data, as well as 50 non-metropolitan areas, one for each state. We define each metropolitan area and state-specific non-metropolitan area as a distinct local labor market, a term we use interchangeably with city, for a total of 152 local labor markets. We define a “firm” as a company and local labor market pair.⁵

In order to measure a college-educated worker’s exposure to the local college alumni network of their *alma mater(s)*, we require data on workers’ individual educational histories. We obtain individual self-reported educational histories from public LinkedIn profiles and include education spells up to 2019. We capture not only younger LinkedIn users who obtained Bachelor’s or advanced degrees during the latter portion of our 2010-2019 analysis time frame, but also older users who graduated prior to 2010 and remain in the labor force between 2010 and 2019.⁶ When reported, the data also contains information on the degree

³Data access via Wharton Research Data Services (WRDS): <https://wrds-www.wharton.upenn.edu/>

⁴Actual salaries are not observed for most of the job spells in the data. Rather, Revelio Labs generates predicted salaries based on a model that takes into account job title, firm, location and other characteristics. State and metro area are recorded for 96 percent of positions during the period we study.

⁵We cannot distinguish different establishments of the same firm in the same local labor market. However, this is not a first-order concern — we only study migration patterns across metro areas.

⁶We only keep prime-age workers age 22-55. Thus the earliest graduation dates are in the 1980s.

obtained for a given education spell (71% coverage), e.g. Bachelor's, MBA, etc., and the field of degree (60% coverage).⁷

Validation of Revelio Labs Data

We validate the representativeness of our data for college-educated job seekers. First, we show that the nationwide industry mix for college-educated workers in the Revelio Labs data is quite similar to that observed in the American Community Survey (ACS), and remarkably so in some cities. In Figure A.1, we plot the pooled 2010-2019 2-digit NAICS shares, among prime-age college-educated workers for both the Revelio Labs data and ACS, for the entire US and Washington DC. We find a strong, positive relationship between Revelio and ACS nationwide industry shares, reassuring us of the representativeness of our data.

We also validate that the self-reported educational histories from LinkedIn that we observe accurately capture the true number of alumni from these schools. The National Center for Education Statistics' Integrated Postsecondary Education Data System (IPEDS) records the number of degrees awarded in each year for each 4-year US college/university. In Figure A.2, we plot the average yearly number of degrees from 2010-2019 observed for each 4-year US college/university in the Revelio Labs data against the same number reported in IPEDS. Reassuringly, there is a strong, positive relationship between the average number of degrees awarded in the 2010s in our data and the IPEDS data, with a slope of approximately two-thirds. Accordingly, our data is well-suited to measure the sizes of local college alumni networks and investigate their impact on geographic mobility.

Calculating Local College Alumni Network Size

We combine the data on self-reported educational histories and individual job histories to measure local college alumni network size for each local labor market. For each city and year, we first identify the set of all workers who held an active position in that city and year. We then match each worker to the set of college(s) and/or universit(ies) from which they obtained any degree of a Bachelor's level or higher. For a job seeker i , we define local alumni network size based on the *alma mater* for their highest degree, denoted as s . We construct two primary network size measures for i : the log number of workers in a city, either in the year of displacement (1) or averaged across the years of displacement among job seekers from s displaced in the same city (2), who have any degree from s .⁸ We use the alumni network for a worker's highest degree in order to focus on the effect of more-recent alumni networks, and to avoid contaminating the treatment effect by the choice of additional education for

⁷For users who report spending four or more years at a Bachelor's degree-granting educational institution, but who are missing information on precise degree obtained, we impute Bachelor's degree attainment to cover well above 71% of degrees.

⁸We do not calculate city-school-major network sizes given incomplete coverage of majors.

workers who return to education after some time in the labor force. As such, we also drop job spells that occurred prior to completion of a worker’s highest degree.⁹

The main limitation we face in constructing these measures of local alumni network size is that we observe only those alumni who use LinkedIn and who have updated profile information on education history and current location. Thus, we cannot observe the in-person alumni networking effects or those manifesting through other platforms. Nevertheless, LinkedIn is the largest job-related social media platform in the US, and heavily utilized by college-educated job seekers, our population of interest.¹⁰ Although we do not observe individual-level connections, because LinkedIn highlights for users the presence of co-alumni who are *not* individual connections, our measure of network size is closely related to what a platform user would actually observe. As an example, Figure A.3 displays a present-day snapshot of detailed information that a college-educated job seeker might see about their local alumni network when looking for jobs in a city, including the number of co-alumni at those jobs. Furthermore, for a given job seeker, these platforms may suggest new connections from the individual’s college alumni network spread across myriad local labor markets.

Mass Layoffs and Monthly Flows Data

We identify workers who have been exogenously separated from their jobs, either through a mass layoff or firm closure, between 2010 and 2018. We supplement our worker panel data with Revelio Labs’ Workforce Dynamics data, from which we observe monthly inflows and outflows of workers for each firm-by-job-type in the database, and define a set of conditions based on drops in employment at the firm level that are consistent with a mass layoff or firm closure.¹¹ Because we observe the start and end dates of job spells as well as the associated firm, we characterize exogenously separated workers as those whose end dates coincide with these aforementioned declines in employment at the same firm. Our definitions of mass layoff and firm closure events closely follow those of Schmieder, Von Wachter, and Heining (2023). More details can be found in section A.2 of our Data Appendix.

Sample of Job Seekers

Our sample of job seekers consists solely of prime-age workers, i.e. age 22-55, with at least a Bachelor’s degree from a 4-year US college/university and were displaced in a mass layoff or firm closure event between 2010 and 2018. We keep the job spell of displacement as well as subsequent active job spells between 2010 and 2019. In some analyses, we also

⁹However, for completeness, we also calculate network sizes based only on Bachelor’s degrees, and find that our quantitative results are robust to this definition.

¹⁰In our data, over 80 percent of workers for whom we observe both job histories and educational degree attainment report earning a Bachelor’s degree or higher, much higher than the national average. Although our data is therefore not representative of the US workforce as a whole, we are interested exclusively in the location choice of college graduates.

¹¹We define a job type as a job category and job seniority level. In the Workforce Dynamics files, there are seven aggregated job categories, as well as seven distinct seniority levels.

incorporate individual-level heterogeneity from Revelio Labs’ data on demographic characteristics imputed from LinkedIn user names. We briefly describe this additional data in Data Appendix A.2. Our final sample consists of 17,680 displaced workers whom we observe both at displacement and post-displacement and for whom we can trace educational histories.

Descriptive Facts about Networks and Job Seekers

Here, we characterize both the college alumni networks that we study and our sample of displaced college-educated job seekers. We begin by documenting that the alumni networks of US colleges and universities are dispersed widely across local labor markets. In Figure 1.1, we show heat maps of alumni network sizes in each local labor market for three schools of different types: a) a highly selective small liberal arts college, Williams College; b) a highly selective mid-sized private university, Northwestern University; and c) a large public university, California State University (CSU) Fullerton.¹² Despite wide differences in their fundamental characteristics, all three schools share the common feature that their alumni networks are spread across local labor markets in the US. While the largest local co-alumni networks are present in the local labor markets closest to the schools themselves (e.g. New York City and Boston for Williams, Chicago for Northwestern, and Los Angeles for CSU Fullerton), these schools’ alumni networks expand well beyond their home regions.

One may presume that the distributions observed in Figure 1.1 are simply characteristics of different school types. If this were the case, then we would expect to see similar concentrations of alumni in each local labor market between two highly similar and geographically proximate schools. Yet we find that pairs of observably comparable schools may have substantially different local alumni network sizes. In Figure 1.2, we show the log differences in alumni network size between each school from Figure 1.1 and an observably similar, geographically proximate school. Panel (a) compares Williams and Amherst College. Both schools are top-ranked liberal arts colleges nationwide, of the same size, and are located only 50 miles apart. But across the US, some cities have more alumni from one school versus the other. Cleveland has 1.4 times as many alumni from Williams as from Amherst, but neighboring Pittsburgh has more Amherst alumni. In panels (b) and (c), we find similar differences for our other examples, comparing Northwestern University to the University of Chicago and comparing CSU Fullerton to CSU Long Beach. To identify the effects of network size on geographic mobility, we use these relative differences in local network sizes between different schools of the same type and same location. In Section 1.4, we rule out endogenous sources of variation in network size and describe possible sources of exogenous variation.

We show descriptively that there is a strong correlation between the network size in a job seeker’s city of displacement and their post-displacement relocation probability. We start with a simple characterization of cities with large and small networks by defining, for each school, cities in the population-weighted top and bottom terciles of local alumni network size.

¹²We take the yearly average (2010-2019) log number of alumni, across all degree types, working in the local labor market. We exclude non-metropolitan areas here, but include them in our main analyses.

For a given school, the top (bottom) tercile consists of the cities with the largest (smallest) number of alumni from that school who account for the first (bottom) third of the school's nationwide alumni network. For example, New York City is in the top tercile of local alumni network size for Williams College, whereas cities with fewer alumni from Williams such as Atlanta and Phoenix are in the bottom tercile.¹³

We first compare the characteristics of job seekers in large and small local co-alumni networks, defined as top and bottom terciles of within-school local alumni network size, for our sample of workers displaced in mass layoffs or firm closures. We report summary statistics in Table 1.1. Our sample of college-educated job seekers is majority male majority White, and have, on average, over a decade of labor market experience. They are displaced in local labor markets with large co-alumni networks — often exceeding thousands of co-alumni. Reassuringly we see that the demographic composition of workers in large and small markets is highly similar, and even the predicted salary at displacement is similar across both groups at slightly above \$80k/year. Thus, at least on observables, our sample of job seekers appears well-balanced across different network size groups.

Next, we investigate whether job seekers who are displaced in local labor markets that have larger co-alumni networks are more likely to remain in the same local labor market. In Figure 1.3, we plot the cumulative share of workers who have ever relocated by each quarter after displacement, separately for workers in the top and bottom terciles of local co-alumni network size for their respective *alma maters*. For completeness, we also show that the share of workers who, in quarters prior to displacement, are in a different metro than the metro of displacement is near-zero for both groups. In the post-period, however, a sizable gap in the probability of relocation emerges — within 5 quarters post-displacement, the cumulative share of workers in the bottom tercile who have relocated (40%) is twice as high as that of workers in the top tercile (20%). We then consider the full distribution of alumni network size across schools and cities. In Figure 1.4, we find a strong, positive relationship between the log number of alumni in the city from which a job seeker is displaced, discretized into many bins, and the average probability for each bin of staying in the same city 1 year post-displacement.

The significantly lower post-displacement relocation probabilities of job seekers displaced in cities with large versus small networks does not necessarily point to a causal relationship between network size and geographic mobility. Larger (smaller) cities such as New York City (Milwaukee) are mechanically more (less) likely to be classified as having large (small) alumni networks for all schools. If workers are generally less likely to leave big cities, then we would misattribute the gap in relocation probability to alumni network size. Another challenge in identifying the causal effect of alumni network size on geographic mobility is that workers from different schools have latent location preferences that impact their geographic

¹³In our tercile characterization, large cities such as New York City are systematically more likely to show up in the top tercile. An alternative characterization is to define, within the same city, top and bottom terciles of network sizes based on schools. This alternative, by contrast, would systematically place small schools in the bottom tercile and large schools in the top tercile. Regardless of the tercile definition, our negative correlation between network size and relocation probability holds.

mobility. Yet another factor influencing geographic mobility is the possibility that firms in specific cities have strong idiosyncratic match effects with alumni from specific schools.

1.3 Theoretical Framework and Empirical Strategy

Model of Worker Location Choice

We present a static model of location choice in which a displaced job seeker chooses a city in which to live and work in two steps, following a nested logit structure. In any city, the size of a worker’s co-alumni network affects their utility from living there. This effect comes through both job-related and non-job-related channels. First, through influence on job search, a larger network increases a job seeker’s utility from work. Second, workers assign social or amenity value to having more co-alumni nearby. Our model predicts that an increase in a city’s alumni network size increases the probability that a worker chooses to live there. To identify network size’s impact on choice probabilities, we assume equivalent structures for match effects and correlated locational preferences among alumni of sufficiently similar and geographically proximate schools.

We consider the location choice of a recently displaced job seeker as an adult, taking their choice of school as given. Among J local labor markets, workers choose one location in which to both live and work.¹⁴ A random utility-maximizing job seeker i who graduated from school s and is displaced from their previous job in origin city o first chooses whether to stay in o or leave. If the job seeker chooses to leave, they then choose a city d among all other destinations in the set $J \setminus \{o\}$. We illustrate this location decision in Figure A.5.

In any city m , the job seeker gains utility from work y_{ms} , e.g. earnings, from local amenities B_{ms} , and from idiosyncratic taste shocks ϵ_{ioms} .¹⁵ We assume that ϵ_{ioms} follows a generalized extreme value (GEV) distribution (McFadden 1978), which yields a nested logit framework in which the two nests are the choice to stay in o , *Stay*, and the choice over all alternatives $d \neq o$ conditional on leaving, *Leave*. Workers incur disutility from cost of living r_m , and incur moving costs C_{om} . We assume that moving costs depend on individual-specific factors, X'_i , and factors that vary across origins and destinations, $\tilde{C}_{om}$, where $\tilde{C}_{om} \neq 0$ if and only if $m \neq o$. We allow X'_i to enter negatively in C_{oo} with parameter η — for instance, older individuals are more likely to be married and may be less likely to move, i.e. lower “cost” of staying in the origin. For all other cities $m \neq o$, X'_i enters positively in C_{om} with parameter κ .¹⁶ The indirect utility in both the origin o and an alternative destination $d \in J \setminus \{o\}$ are

¹⁴To illustrate this, in Figure A.4 we present a timeline of a job seeker’s labor market history and highlight, in red, the point in time at which they are displaced and forced to re-optimize their job and location choice.

¹⁵Work utility y_{ms} is a reduced-form way to capture job search and matching technology, including the benefit of shorter unemployment duration, increased match probability, or better expected match quality.

¹⁶For simplicity, we assume these factors do not vary with the actual city of origin, but we can extend our framework to allow for such heterogeneity. We show that our results are robust to allowing characteristics to vary across all locations in Section 1.5. In all other destinations, X'_i does not impact destination choice probabilities because they are not destination-specific.

sums of the observable components, denoted as V_{ioos} and V_{iods} , and the taste shocks:

$$U_{ioos} = \underbrace{y_{os} + B_{os} - r_o + X'_i \eta}_{\equiv V_{ioos}} + \epsilon_{ioos} \quad (1.1)$$

$$U_{iods} = \underbrace{y_{ds} + B_{ds} - r_d - X'_i \kappa - \tilde{C}_{od}}_{\equiv V_{iods}} + \epsilon_{iods} \quad (1.2)$$

We allow the size of the worker's local co-alumni network in any city m to directly influence the work utility in m . After a job seeker is displaced and the previous worker-job match is destroyed, they search for a new job and seek to maximize their expected work utility from some new job. For a job seeker from school s , we allow the work utility for some job h in city m , y_{msh} , to systematically vary across cities and schools, captured by μ_{ms} as well as idiosyncratically, with an idiosyncratic match component ν_h that is i.i.d. with CDF $F(\nu_h)$:

$$y_{msh} = \mu_{ms} + \nu_h \quad (1.3)$$

We assume that in each city, a job seeker has a finite number of job draws proportional to the number of co-alumni from school s in that city, σ_{ms} , by a factor of $a < 1$. This assumption does not require that all draws come from jobs where co-alumni are present, and instead simply reflects the idea that job seekers are more likely to sample more jobs in a city with many co-alumni.¹⁷ We assume that $F(\nu_h)$ is Type-1 Extreme Value (McFadden 1978), and allow for different scale parameters in the origin, ω^{Stay} , and alternative destinations, ω^{Leave} . Following Kuhn and Shen (2013), if we define y_{ms} , the job seeker's work utility in city m , as the expected value of the job with highest work utility in a sample of $a\sigma_{ms}$ draws, then :

$$y_{ms} = \begin{cases} \underbrace{\mu_{os} + \omega^{Stay}(\gamma + \log(a))}_{\equiv A_{os}} + \omega^{Stay} \log(\sigma_{os}), & \text{if } m = o \\ \underbrace{\mu_{ds} + \omega^{Leave}(\gamma + \log(a))}_{\equiv A_{ds}} + \omega^{Leave} \log(\sigma_{ds}), & \text{if } m = d \text{ where } d \in J \setminus \{o\} \end{cases} \quad (1.4)$$

where γ is Euler's constant, and the expected value of ν_h is $\omega^{Stay}\gamma$ and $\omega^{Leave}\gamma$ in the origin and alternative destinations, respectively, yielding a scaled city-school work utility A_{ms} for all m . We now have the network returns $\omega^{Stay} \log(\sigma_{os})$ and $\omega^{Leave} \log(\sigma_{ds})$, which nest many channels through which the local alumni network increases work utility in a location. First, if firms have experience with previous alumni from school s , then the network provides a signal to employers about the job seeker's productivity and facilitates better matches between job seekers in those networks and their new firms (Dustmann et al. 2016). Additionally, employees may pass along information about jobs in their respective local areas to other network members, potentially shifting the distribution of wage offers that a job seeker can

¹⁷Our setup implies that in cities with no co-alumni, the job seeker has zero job draws. This rarely binds in practice, as the schools in our sample have positive alumni counts in every or almost every city.

receive (Calvo-Armengól and Jackson 2004; Beaman 2012). Yet another possibility is that having more employed members from the alumni network decreases the duration of job search, or that they directly influence hiring other co-alumni. We allow the aforementioned as well as other network-related mechanisms to affect y_{ms} , and measure their composite effects. To separately identify $\omega^{Stay} \log(\sigma_{os})$ from A_{os} (and $\omega^{Leave} \log(\sigma_{ds})$ from A_{ds} for $d \neq o$), we require additional structure on A_{ms} . We test several possibilities in Section 1.3, and discuss the validity and implications of our assumptions in more detail in Section 1.4.

As is the case for work utility, we also assume that for any city m , the alumni network size in m enters into amenities in a consistent manner. We decompose the city amenity into a common component experienced across schools and a school-specific component:

$$B_{ms} = B_m^{City} + \tilde{B}_{ms} \quad (1.5)$$

The city amenity B_m^{City} captures a broad set of local non-wage amenities that all college graduates value equally, as in Diamond (2016). We further decompose the school-specific component $\tilde{B}_{ms}$ into two channels, only one of which the alumni network directly enters:

$$\tilde{B}_{ms} = g(\sigma_{ms}) + Z_{ms} \quad (1.6)$$

The alumni network affects local amenities via $g(\sigma_{ms})$, whereby workers value having a larger network for non job-related reasons, e.g. they like to socialize with friends and acquaintances from college. We assume that $g(\sigma_{ms})$ is log-linear with parameters ρ^{Stay} and ρ^{Leave} for each nest. That is, $g(\sigma_{os}) = \rho^{Stay} \log(\sigma_{os})$ and $g(\sigma_{ds}) = \rho^{Leave} \log(\sigma_{ds})$ if $d \neq o$. The second component of school-specific local amenities, Z_{ms} , captures correlated preferences for a location among a school's alumni. For example, Boston may be more latently attractive to workers from some college in New England than those from a school in Minneapolis. It is not possible to exactly identify Z_{ms} without further structure. We make the assumption that for sufficiently similar schools in a set K , defined as a *strata*, $Z_{ms} = Z_{mK(s)}$ for all $s \in K$. We discuss this assumption and its implications in more detail in Section 1.4.

Given our nested logit structure, we can write both the probability that worker i , displaced in city o , chooses to stay in o (denote $Pr(stay_{ios})$) and the probability, conditional on leaving o , of choosing to work in some alternative city d (denote $Pr(h_{iods}|Leave)$). Starting with $Pr(stay_{ios})$:

$$Pr(Stay_{ios}) = \frac{\exp(V_{ioos})}{\exp(V_{ioos}) + \underbrace{\exp(\lambda \log(\sum_{m=1}^{J \setminus \{o\}} \exp(\frac{V_{iom s}}{\lambda})))}_{\equiv IV_{os}}} \quad (1.7a)$$

$$= \frac{\exp \left[\underbrace{(\omega^{Stay} + \rho^{Stay})}_{\equiv \beta^{Stay}} \log(\sigma_{os}) + (A_{os} + Z_{os}) + (B_o^{City} - r_o) + X'_i \eta \right]}{\exp \left[\underbrace{(\omega^{Stay} + \rho^{Stay})}_{\equiv \beta^{Stay}} \log(\sigma_{os}) + (A_{os} + Z_{os}) + (B_o^{City} - r_o) + X'_i \eta \right] + \exp(\lambda IV_{os})} \quad (1.7b)$$

Here, $\beta^{Stay} \equiv \omega^{Stay} + \rho^{Stay}$ captures the total impact of origin network size on the probability of staying in the origin. The parameter λ , which comes from our assumption of GEV-distributed idiosyncratic taste shocks, is the measure of independence in unobserved utility across alternative destinations. We also define an essential object in our nested logit structure: IV_{os} , the inclusive value of the *Leave* nest.¹⁸ The inclusive value multiplied by the GEV parameter, λIV_{os} , captures the expected utility across alternative destinations in the *Leave* nest. For $Pr(h_{iods}|Leave)$, we have:

$$Pr(h_{iods}|Leave) = \frac{\exp(\frac{V_{iods}}{\lambda})}{\sum_{m=1}^{J \setminus \{o\}} \exp(\frac{V_{ioms}}{\lambda})} \quad (1.8a)$$

$$= \frac{\exp \left[\frac{1}{\lambda} \left(\underbrace{(\omega^{Leave} + \rho^{Leave})}_{\equiv \beta^{Leave}} \log(\sigma_{ds}) + (A_{ds} + Z_{ds}) + (B_d^{City} - r_d) - \tilde{C}_{od} \right) \right]}{\sum_{m=1}^{J \setminus \{o\}} \exp \left[\frac{1}{\lambda} \left(\underbrace{(\omega^{Leave} + \rho^{Leave})}_{\equiv \beta^{Leave}} \log(\sigma_{ms}) + (A_{ms} + Z_{ms}) + (B_m^{City} - r_m) - \tilde{C}_{om} \right) \right]} \quad (1.8b)$$

where $\frac{\beta^{Leave}}{\lambda} \equiv \frac{1}{\lambda}(\omega^{Leave} + \rho^{Leave})$ captures destination network size's impact on location choice conditional on moving. For simplicity, we denote $\tilde{\beta}^{Leave} \equiv \frac{\beta^{Leave}}{\lambda}$ as our parameter of interest.¹⁹ Under this framework, we study two related outcomes: 1) the job seeker's decision of whether to remain in their city of displacement (out-mobility), and 2) conditional on deciding to leave, the job seeker's location choice post-displacement (directed mobility).

We show in Appendix A.3 that if $\beta^{Stay} > 0$, then Equation 1.7b predicts that an increase in the number of co-alumni in the city of displacement increases the probability of staying. Likewise, we show in Appendix A.3 that if $\tilde{\beta}^{Leave} > 0$, Equation 1.8b predicts that an increase in the number of co-alumni in an alternative destination increases the probability of choosing that destination, conditional on leaving.²⁰ We allow for this impact to come both from channels related to job search, i.e. ω^{Stay} and ω^{Leave} , or from networks' value as a local consumption amenity, i.e. ρ^{Stay} and ρ^{Leave} . Although we do not aim to quantify the relative magnitudes of these two broad channels in influencing mobility and location choice, we conduct empirical tests that provide suggestive evidence on which channel is stronger.

If $\omega^{Stay} > 0$ and $\omega^{Leave} > 0$, then we should additionally expect that increasing local co-alumni network size increases the probability of not only choosing to live and work in that city but also working at local firms with other co-alumni. Intuitively, the many job search-related channels that are nested in the two parameters capture mechanisms by which

¹⁸This term varies at the origin-school level since the observable component of indirect utility in any alternative destination does not have destination-varying individual-specific components

¹⁹We are able to identify both $\tilde{\beta}^{Leave}$ and β^{Leave} , but focus on the former since it aligns more directly with our causal effect of interest. The latter, meanwhile, captures the impact on unconditional, rather than conditional, destination choice probability.

²⁰Theoretically, the same condition should hold for β^{Leave} , since theoretically λ is between 0 and 1. Our estimated values of λ indeed fall in this range.

employed members of the local co-alumni network increase the probability that job seekers from their network are re-employed at their *incumbent firms*. For instance, co-alumni provide direct referrals and are more likely to transmit information about openings at their incumbent firms rather than other firms. A positive value of ρ^{Stay} and ρ^{Leave} , by contrast, does not deliver the same prediction, since a job seeker's amenity valuation of the local alumni network is independent of whether they work at the same firm as a co-alumni.²¹

Impact of Local Alumni Networks on Out-mobility

We first test the hypothesis that college-educated workers who experience exogenous job separation in a mass layoff or firm closure are more likely to remain in a local labor market with a higher concentration of alumni from their *alma mater*. From Equation 1.7b, we can estimate β^{Stay} in two steps. First, we obtain an estimate of the inclusive value of leaving, IV_{os} , which we denote as $\widehat{IV}_{os}$, from estimating the parameters in the conditional destination probability expression in Equation 1.8b. We discuss our estimation of those parameters, including $\widehat{\beta}^{Leave}$, in Section 1.3. In the second step, we include $\widehat{IV}_{os}$ on the right-hand side and estimate β^{Stay} via a binary logit model immediately derived from Equation 1.7b.

To identify β^{Stay} , we must absorb the impact of origin-school work utility A_{os} and origin-specific correlated preferences among a school's alumni Z_{os} . To do so, we invoke our identifying assumption that correlated preferences $Z_{os} = Z_{oK(s)}$ for alumni of all schools s belonging to the same narrowly defined set K . As a baseline, we make a parallel assumption that $A_{os} = A_{oK(s)}$, which rules out systematic differences in origin work utility for alumni of schools within the same strata. Accordingly, we estimate the following binary logit:

$$\log\left(\frac{Pr(stay_i)}{1 - Pr(stay_i)}\right) = \alpha + \beta^{Stay} \log(\sigma_{ose_i}) + X'_i \eta + \delta_{e_i} + \phi_{oK(s)} - \lambda \widehat{IV}_{os} + \epsilon_{ios} \quad (1.9)$$

The logit dependent variable is an indicator for job seeker i staying in the city of displacement o 1 year post-displacement. Here, e_i refers to the year of displacement. Accordingly, $\log(\sigma_{ose_i})$ is the co-alumni network size in o in the year of displacement. We also include year fixed effects δ_{e_i} , and other individual characteristics X'_i . Our origin $\times$ strata fixed effects $\phi_{oK(s)}$ absorb the remaining components of observable indirect utility in the origin — correlated preferences $Z_{oK(s)}$, fixed work utility $A_{oK(s)}$, amenities B_o^{City} , and cost of living r_o .

To show that our results are robust to different relaxations of our baseline identifying assumptions, we additionally estimate linear probability models (LPM) as reasonable approximations of Equation 1.9 that can handle more saturated fixed effects. One relaxation to our assumption on A_{os} , the city-school component of work utility, is to include school fixed effects that absorb the possibility of schools having different relative value-added in earnings (Chetty et al. 2020; Mountjoy and Hickman 2021). Another relaxation to assuming within-strata equivalence of city-school components of work utility and correlated locational

²¹Job seekers may be re-employed with other co-alumni by random chance, but we show that the rate of re-employment with co-alumni far exceeds the rate from random chance.

preferences is to only require this equivalence for workers in the same industry, major or firm. We explore all of these possibilities.

Impact of Local Alumni Networks on Directed Mobility

Next, we aim to estimate $\tilde{\beta}^{Leave}$, which captures the impact of network size in an alternative destination on a job seeker's probability of choosing that destination, conditional on moving.²² Using the choice probability expression in Equation 1.8b, it is technically possible, using absorbing fixed effects, to estimate $\tilde{\beta}^{Leave}$ via conditional logit. However, with over 100 possible destination cities and an exhaustive set of fixed effects, it is not feasible in our setting. Instead, we recover consistent estimates of $\tilde{\beta}^{Leave}$ using the result in Guimaraes, Figueirdo, and Woodward (2003) that a conditional logit framework in which the indirect utility components vary at most at the group $\times$ alternative level, where a group refers to a set of individuals, can be reformulated as a Poisson-structured gravity equation with collapsed group $\times$ alternative level data and group fixed effects.²³ In our setting, we define groups as origin-school pairs, and we estimate the following econometric specification via Poisson pseudo-maximum likelihood (PPML) for *movers only*:

$$P_{ods} = \exp(\tilde{\beta}^{Leave} \log(\sigma_{ds}) + \phi_{od} + \phi_{os} + \phi_{dK(s)}) \varepsilon_{ods} \quad (1.10)$$

where P_{ods} is the number of job-seeking alumni from school s who are displaced in a mass layoff or firm closure in city o and move to city d 1-year post-displacement. Within origin-school groups, we calculate co-alumni network size $\log(\sigma_{ds})$ by averaging across years of displacement. We include origin metro $\times$ school fixed effects as our group fixed effects, which absorbs the denominator of our choice probability expression in Equation 1.8b. We also include a set of origin metro $\times$ destination metro fixed effects ϕ_{od} , which absorb citywide amenities B_d^{City} , cost of living r_d , and moving costs $\tilde{C}_{od}$. Invoking our identifying assumption that correlated locational preferences are the same across alumni of schools in the same strata K , i.e. $Z_{dK(s)} = Z_{ds}$, we also include destination metro $\times$ school strata fixed effects $\phi_{dK(s)}$. These fixed effects also allow for several different assumptions about the structure of the city-school component of work utility A_{ds} . If A_{ds} primarily reflects matches between firms in d with job seekers from the same narrow class of schools, which our strata capture, then $\phi_{dK(s)}$ largely absorbs A_{ds} . We also directly absorb any location-invariant city-specific and city-strata factors. And if A_{ds} also includes any location-invariant, school-specific determinants, e.g. school value-added, these are already factored out in Equation 1.8b. We provide an extended discussion of A_{ds} and identification in Section 1.4.

²²We can also obtain the impact on the unconditional probability, β^{Leave} itself, by first estimating $\tilde{\beta}^{Leave}$, and then using our estimated coefficient on λ in Equation 1.9 to scale the estimate.

²³Conceptually, our Poisson structure is additionally advantageous given the well-documented issues with estimating gravity equations via OLS in the presence of heteroskedasticity and zero flows (Silva and Tenreyro 2006), e.g. we observe very few moves from Seattle to Nebraska.

1.4 Identification

We aim to estimate $\beta^{Stay} \equiv \omega^{Stay} + \rho^{Stay}$, the effect of (log) alumni network size in the city of displacement on the odds of staying, as well as $\tilde{\beta}^{Leave} \equiv \frac{1}{\lambda}(\omega^{Leave} + \rho^{Leave})$, the effect of (log) alumni network size in an alternative destination on the odds of choosing that destination, conditional on moving. These total impacts combine ω^{Stay} and ω^{Leave} , the composite effect of log alumni network size on work utility, with ρ^{Stay} and ρ^{Leave} , which governs the alumni network's impact on a city's amenity value. Identification of β^{Stay} and $\tilde{\beta}^{Leave}$ requires overcoming two threats: endogeneity of worker location choice and endogeneity of local college alumni network size. In this section, we discuss how our empirical strategy overcomes these concerns.

Endogeneity of Worker Location Choice

Our theoretical framework targets the direct influence of local alumni networks on the location decisions of job seekers who were forced into actively making them. The inclusion of all job seekers would otherwise pose a selection threat to our causal channel, as workers who voluntarily change jobs and/or locations are not randomly selected. Voluntary job switchers may have unobserved characteristics that directly influence their mobility across local labor markets. This selection is problematic if these unobserved characteristics are both systematically correlated with local network size and unrelated to the network's direct influence on job search and local amenities.²⁴ Furthermore, if voluntary job switchers move to another city because they have already secured another job offer in that city prior to separation, then we have a potential reverse causality concern whereby mobility is predetermined.

Our exclusive focus on workers who are displaced in mass layoffs or firm closures overcomes these concerns. This sample restriction follows other studies on the impact of networks (Hellerstein, Kutzbach, and Neumark 2019; Ost, Pan, and Webber 2024) or local labor market size (Moretti and Yi 2025) on labor market outcomes. When comparing the pre-displacement raw relocation probabilities of job seekers in local labor markets with the smallest vs. largest alumni networks for their schools, as in Figure 1.3, we find that these probabilities are zero for both groups. This finding of no pre-displacement mobility suggests that our sample of job seekers did in fact experience displacement as a sudden event that forced a re-optimization of location choice. Thus, we find little evidence that workers select into large or small networks based on their inherent propensity to move.

Endogeneity of Local College Alumni Network Size

Local college alumni networks vary in size in nonrandom ways. For instance, they are larger in large metros and in metros near the college's campus. To identify β^{Stay} and $\tilde{\beta}^{Leave}$,

²⁴For example, voluntary job switchers may have higher unobserved ability and in turn be more inherently mobile as they preemptively seek out new job opportunities in locations with co-alumni who are more willing to hire them.

we must isolate the component of variation that is uncorrelated with a job seeker’s preferences and mobility. We classify the endogenous determinants of location choice into labor supply factors driven by correlated preferences among alumni and labor demand factors driven by productivity matches with local firms.

On the labor supply side, we must account for unobserved location preferences. At a high level, some cities may be more attractive to college-educated job seekers in general (high B_m^{City} and/or high city-specific component of A_{ms}). Graduates from some colleges may also generally be more productive (high school-specific component of A_{ms}). Additionally, migration costs may systematically vary in ways that make some cities more or less difficult to depart from or migrate to (variation in $\tilde{C}_{od}$). Our rich fixed effects account for all these possibilities: in models that estimate β^{Stay} we net out the impact of city-invariant school-specific work utility determinants and include city fixed effects,²⁵ and our directed mobility analysis includes fixed effects for origin-school pairs and origin-destination pairs. These origin-destination pair fixed effects account not only for flexible migration costs, but also allow the relative attractiveness of a destination to vary by origin.²⁶ For example, if residents laid off in New York City have a high average presence for large metros, we account for their increased likelihood to relocate to a city like Chicago relative to a resident laid off in Little Rock. And to whatever extent these broad location preferences vary by individual characteristics, we show that our results are robust to controlling for such preferences for each potential destination within the same gender, age bracket, and race.

The next labor supply concern is the unobserved preference structure for cities that varies by school (variation in Z_{ms}). If a school’s alumni may have correlated preferences for locations (Manski 1993), such preferences would be positively correlated with both local network size and geographic mobility. For example, if University of California San Diego (UCSD) graduates have latent preferences for coastal living, then we would expect both a large UCSD alumni network in Los Angeles and a UCSD job seeker to be relatively likely to choose Los Angeles, even in the absence of a network treatment effect. To address this concern, we group schools into *strata*, made up of observationally very similar schools in the same location. We assume that for a school s in a strata K , $Z_{ms} = Z_{mK(s)}$. Of course, this assumption is reasonable only if the strata are quite narrowly defined. We use coarsened exact matching (CEM) to match schools by geographic proximity to each other, size, admissions rate, 75th percentile composite SAT score, and share of enrolled freshman who are in-state. We provide additional details in section A.2 of our Data Appendix.

Our sample of job seekers comes from 1,399 schools in 402 strata — an average of only 3.5 schools per strata.²⁷ Since schools in our strata are matched both on location and

²⁵Our saturated specifications include fixed effects for origin cities interacted with fixed effects for the school strata, resulting in a flexible structure that also accounts for arbitrary productivity matches between cities and strata, which we discuss further below.

²⁶If any of these factors are time-varying, then our fixed effects would not fully account for them. This threat is primarily relevant for a longer time horizon, which we minimize by restricting our time frame of analysis to only a decade that both follows the Great Recession and predates the COVID-19 Pandemic.

²⁷We define 431 strata for 1,840 4-year colleges and universities, for an average size of 4.3 schools per

key institutional dimensions, we maximize the plausibility of our identifying assumption. For example, Williams College and Amherst College are in the same strata, along with Middlebury College, Bowdoin College, and Wesleyan University. Thus, our identification assumption would require that a Williams alum and an Amherst alum have the same average preference for living in a city such as Dallas (in practice, we interact our strata with other fixed effects to relax this assumption in multiple ways). Other examples of strata include large public universities in Arizona (University of Arizona and Arizona State University), and California State University schools in Northern California (Cal State Chico, Cal State Fresno, and Cal State Sacramento). Across strata, we allow for arbitrary variation in preferences. For example, even for two schools as similar as UCSD and San Diego State, we make no assumption on the preferences of one school’s alumni compared to the other, since one school is slightly more selective and thus in a different strata.

We use these strata to account for preference correlation in several ways. In our out-mobility analysis, we estimate models with fixed effects for origin city-industry-strata triples, so that our identification assumption need hold only for workers in the same strata, city, and industry of employment. In our directed mobility analysis, we include fixed effects for destination-strata pairs, as well as fixed effects for origin-*school* pairs, so that our identification assumption need only hold for distant metros. We thus allow for different preferences, even within strata, for the metro of residence at displacement, with which the worker is presumably relatively familiar.

Our identifying assumption is not very restrictive, as it still allows for a wide range of correlated preferences to exist without threatening identification. For instance, consider an alternative assumption that the alumni of any two schools that are located in the same geographical area, e.g. New York City, have the same latent locational preferences. We estimate models allowing for the possibility that alumni from different schools in New York City who all lost their jobs in the same city and industry may still have systematically different preferences for New York City, or elsewhere. This possibility is only ruled out if the two schools are highly similar across many dimensions.

We fortify our identifying assumption’s plausibility by matching schools on the share of enrolled freshmen who are in-state. Geographic proximity to home is the top predictor of college choice (Avery, Howell, and Page 2014), and standard location choice models often include proximity to home as a source of locational preference, e.g. Diamond (2016). Furthermore, nearly half of graduates from US 4-year colleges and universities reside in the nearest metro to their *alma mater* (Conzelmann et al. 2025). Within a small set of observably similar schools in the same location whose student bodies feature the same degree of geographic locality, this preference is unlikely to differ across the schools’ respective alumni. To rule out any remaining confounding home bias, we estimate specifications on a restricted sample of job seekers who are displaced outside of the metro of their *alma mater* (“home metro”).²⁸ We additionally exclude workers who return to their home metro after being

strata, but only 1,399 schools are represented in our final estimation sample.

²⁸We have information on high school attendance for a very small fraction of job seekers. Thus we cannot

displaced elsewhere.

Next, we consider potential confounding determinants of mobility on the labor demand side that vary across both cities and schools. There are two endogeneity concerns that we address. First, we have assumed that the component of work utility that is common to some city-school pair $m \times s$ can be decomposed into a city-strata component and separate city and school components, with remaining city-school match effects coming through the network size σ_{ms} . Our concern is with a potential lingering non-zero idiosyncratic match effect between city m firms and school s alumni that increases work utility y_{ms} . A positive match would both increase a graduate's probability of living in m and be positively correlated with σ_{ms} . Because job search is typically very local in nature (Manning and Petrongolo 2017), this threat is most acute in the city of displacement.²⁹ Thus, in our analysis of conditional destination choice, whose estimates we use to also analyze the stay vs. leave margin, we include origin $\times$ school fixed effects to directly account for the possibility of local city-school productivity matches.

We further assuage concerns about city-school match effects by showing that our results are robust to leveraging within-firm (and strata) variation in network size. We effectively compare the location decisions of job seekers who are alums of schools in the same strata and were displaced at the *same firm*. Additionally, if firms imperfectly observe their productivity matches with schools, and instead learn about match effects through the co-alumni they employ, then these matches would constitute a treatment mechanism rather than a threat to identification, paralleling Montgomery (1991) and Dustmann et al. (2016). Finally, any residual match effects are less likely to impact our sample of job seekers with, on average, over a decade of labor market experience.³⁰

The second, related concern for labor demand is that if firms in a given city and/or industry systematically prefer to hire specific majors, then differences in schools' major composition, even within a set of highly comparable schools, could lead to overstating β^{Stay} and β^{Leave} . This is only a minor concern if major's productivity signal sharply attenuates with labor market experience, as in the literature on employer learning from workers' schooling. Nevertheless, to the extent that differences in school major composition manifest as differences in industrial composition of workers, we estimate models that absorb any unobserved factors varying across both cities and industries. To provide even more reassurance, we show that school strata alone explain at least 40% of the variation in school major shares for each of the six most common majors in our data.³¹ And as a final check, we estimate models that

reliably infer home location based on pre-college education.

²⁹On LinkedIn, job search is online and more widespread, and thus networks have a broader geographic scope. However, over 70% of displaced workers in our sample still remain in their city of displacement.

³⁰Such match effects are likeliest to manifest upon graduation (such as via on-campus recruiting), as the productivity signal that schooling sends to employers quickly diminishes as workers gain experience (Altonji and Pierret 2001; Lange 2007). The comparison of schools within the same area is also significant. Although a high proportion of large firms recruit new graduates on-campus (Forsythe and Weinstein 2025), recruiting remains local in nature (Weinstein 2022a). Our school strata consist of schools whose graduates draw attention from very similar sets of recruiting firms.

³¹In Table A.1, for each major recorded in the Revelio Labs data, we report the R-squared from a

directly control for factors varying across destination-major pairs and major-strata pairs.

While we cannot rule out all unobserved determinants of mobility that are correlated with local network size, we also reinforce our findings against all factors varying at the city-school pair level in Section 1.6. Here, we recalculate alumni networks within each industry of employment, and estimate models with city $\times$ school fixed effects that directly absorb the labor supply and labor demand threats discussed above.

Finally, we briefly discuss possible sources of exogenous variation in local alumni network sizes. Because we do not observe these networks forming in the years after a school opens, we are not able to empirically differentiate between possible origins, so our commentary is relatively speculative in nature. One possibility is that many geographically proximate schools that are highly similar today have substantially differed on key characteristics in the past, potentially resulting in student bodies with different post-graduation migration patterns. On one hand, such schools may have admitted very different students in the past — on various dimensions such as race, gender, and pre-college geography. The schools’ own characteristics could have also differed, including major offerings, enrollment, or perceived quality.³² Another possibility is idiosyncratic timing across schools in firm recruiting. Because we analyze exogenously displaced job seekers in the 2010s as opposed to fresh college graduates from earlier time periods, their location decisions are not likely to be directly impacted by the aforementioned factors other than through these factors’ influence on network size.

Motivated by this intuition, we construct two instruments for local alumni network size, one each for the origin and a potential destination based on historical variation in schools’ shares of in-state students. Concretely, we interact a school’s in-state enrollment share in 1972 with the inverse distances between the school and the origin or a potential destination. Our instruments are $W_{os} \equiv \frac{Share_{s,1972}}{1+dist_{os}}$ and $W_{ds} \equiv \frac{Share_{s,1972}}{1+dist_{ds}}$, where $dist_{os}$ and $dist_{ds}$ are the great circle distances (in miles) between the school and the origin and destination, respectively. We employ a control function approach using these instruments, following

regression of the major’s share at each school, from IPEDS in 2010, on school strata fixed effects. For each of the six most common majors, at least 40% of the variation in school major share is explained by school strata alone. For six of the eleven majors, the explained variation is at least 48%. We corroborate school strata’s explanatory power by plotting a histogram of standardized residuals for each major in Figure A.6. For all majors, the standardized residuals are overwhelmingly concentrated at zero, with little spread in the distribution. Thus the average major share within strata predicts the school’s major share with high accuracy for all majors.

³²Take Pomona College and Claremont McKenna College (CMC), for example. Both schools have been neck-and-neck in US News rankings during the 21st century, but in the late 1980s, which predates the graduation year of nearly all job seekers in our sample, Pomona’s ranking (~ 5 th) was far ahead of CMC’s (~ 21 st) (Source:<https://andyreiter.com/datasets/>). If Pomona’s higher ranking drew students with higher unobserved skill who then sorted to higher-wage cities, e.g. NYC (Combes, Duranton, and Gobillon 2008; Card, Rothstein, and Yi 2025), then we might expect more Pomona alums than CMC alums from older cohorts, but a smaller gap for new graduates from present-day cohorts. Indeed, Pomona has 1.5 times as many alumni as CMC in NYC, yet we see a remarkable similarity between the two schools in the number of new graduates who locate in NYC — between 2016-2019, the average number of CMC alumni from the 2016-2019 graduating cohorts in New York City was 169, compared to 168 for Pomona.

Terza, Basu, and Rathouz (2008), to isolate arguably exogenous sources of variation in network size and re-estimate β^{Stay} and $\tilde{\beta}^{Leave}$. We provide more details in Appendix A.4.

Our instruments are strongly positively correlated with local alumni network sizes. Intuitively, if historical alumni of a school with a high in-state enrollment share in 1972 were more likely to choose nearby cities (high inverse distance), then we would expect larger alumni networks in nearby cities in the present day. We indeed find that schools have large alumni networks in nearby cities, and our instruments isolate the variation in present-day network sizes that comes from these historical differences. Exclusion requires that differences in the 1972 in-state enrollment shares between two neighboring schools that are highly similar today, including having similar present-day in-state enrollment shares, are not systematically correlated with unobserved determinants of present-day job seekers' mobility. Because our sample of job seekers all graduated well after 1972, we minimize the plausibility that contemporaneous unobserved factors affect the location decisions of modern job seekers.

As an example of our targeted variation in network size, consider Boston University and Northeastern University, two schools in the same strata with a similar in-state share in 2010. Both schools are large private research universities in Boston with near-identical rankings and similar cross-yields.³³ But one school historically had a very local student body and the other did not. In 2010, only 27% of Northeastern University freshmen were from Massachusetts, compared to a similar 20% for Boston University. Yet in 1972, 81% of Northeastern University students were from Massachusetts, compared to only 48% for Boston University. In the present day, we see that compared to Boston University, Northeastern University has larger alumni networks in both Boston and throughout New England. While other factors, many of which we have discussed in this section, contribute to these differences in network sizes, we control for those factors and solely leverage variation that comes from these differences in home backgrounds among alumni who predate our sample of job seekers.

1.5 Results

Impact of Local Alumni Network Size on Out-Mobility

Table 1.2 presents our baseline second-step estimates of β^{Stay} in Equation 1.9. In all of our logits, we include the log co-alumni network size in the city of displacement, the inclusive value of leaving, individual controls, school controls, and fixed effects for the year of displacement.³⁴ To account for sampling variability in estimating the inclusive value of the *Leave* nest in our first step, we report bootstrapped standard errors based on resampling clusters at the city $\times$ school level. In column 1, we add origin fixed effects, and then in column 2, fixed effects for the origin interacted with the metro of a job seeker's *alma mater*. Our

³³Source: <https://www.parchment.com/c/college/tools/college-cross-admit-comparison.php?compare=Eastern+University&with=Boston+University>

³⁴School controls are absorbed by any strata fixed effects, since within strata, schools are matched on observables.

preferred specification in column 3 includes origin $\times$ school strata fixed effects. This within-strata comparison of co-alumni network size is consistent with our identifying assumption that correlated locational preferences and city-school work utility matches within a school's nationwide alumni network, the primary confounders to identifying the effect of local alumni network size, are the same across alumni networks of sufficiently similar and geographically proximate schools in the same strata.

Across all specifications, we report both our estimates of β^{Stay} , which has a direct elasticity of odds interpretation, and the associated average marginal effects.³⁵ We find that increasing a job seeker's alumni network size in the city of displacement by 1% yields a 0.4% - 0.7% increase in the odds of staying, including a 0.4% increase in our preferred specification in column 3 (p-value < 0.01). Based on the associated average marginal effect, the magnitude of this estimate is highly similar to that of the estimated impact of local city $\times$ industry size on the relocation of displaced workers in Moretti and Yi (2025). Our estimates of λ , captured by the coefficient on the inclusive value of leaving the city of displacement, range from -0.5 to -0.37, indicating a reasonable amount of correlation between alternative destinations within the leave nest.

We also estimate more saturated models via linear approximations of our logit in Equation 1.9.³⁶ These saturated models allow for more flexible assumptions about the structure of correlated locational preferences Z_{os} and city-school components A_{os} of work utility in our model. In Table 1.3, in column 1 we again report our estimates from our preferred logit model, and show in column 2 that a linear probability model with the same variables yields a near-identical estimated average marginal effect. With reassurance that the linear probability model is a sensible approximation, we show that if we add school fixed effects in column 3, which absorbs any potential location-invariant school-specific components of work utility, our estimated average marginal effects remain similar in magnitude and statistically significant at the 1% level. In column 4, we instead include fully interacted origin $\times$ strata $\times$ industry of displacement fixed effects. Thus our assumption that city-school correlated preferences and work utility are captured by city-strata factors only needs to hold for workers displaced in the *same industry*. We effectively compare the out-mobility of a Williams alum who is laid off in the Finance industry in Denver to an Amherst alum who is also laid off in the Finance industry in Denver. We estimate a similar average marginal effect of 0.05 — doubling the origin network size yields a 3.5 percentage point increase in the probability of staying (p-value < 0.01), which is a 4.5% increase given the $\sim 77\%$ baseline probability of staying.

Our finding of a strong, positive effect of origin alumni network size on the probability of staying holds when using a control function approach with an instrumental variable that isolates arguably exogenous variation in origin network size. We re-introduce the instru-

³⁵The interpretation of β^{Stay} is that a 1% increase in the origin alumni network size yields a $\beta^{Stay}\%$ increase in the odds of staying. This comes from a log-log interpretation for our logit, which relates log odds to log network size.

³⁶These more saturated models do not converge when estimating logits, given the high dimensionality of fixed effects.

mental variable from Section 1.4, $W_{os} \equiv \frac{Share_{s,1972}}{1+dist_{os}}$, where $Share_{s,1972}$ is the school's in-state enrollment share in 1972 and $dist_{os}$ is the great circle distance (in miles) between the origin and school. We provide more details about our estimation procedure in Appendix A.4. The F-statistic from the first stage is 349.3, signifying the instrument's relevance. We report our two-stage residual inclusion (2SRI) estimate of β^{Stay} in column 1 of Table 1.4, alongside our baseline estimate in column 2. The 2SRI estimate of 0.478 is close to our baseline estimate of 0.389, and is statistically significant at the 5% level.³⁷ Our estimate of λ , the dissimilarity parameter and coefficient on the inclusive value of the *Leave* nest, is -0.327 (p-value < 0.05), which is also similar to the baseline estimate of -0.367.

Thus we show that increasing the network size in a job seeker's city of displacement meaningfully increases the probability of staying in that metro. This effect combines the job search channel ω^{Stay} and amenity channel ρ^{Stay} . Descriptively, only a tiny fraction of firms in a given city have co-alumni coworkers for the average job seeker, so in the absence of a treatment effect driven by ω^{Stay} , by random chance we would expect very few job seekers to become employed at "in-network" firms. In Panel A of Figure 1.5, we plot a histogram of the share of firms in a job seeker's metro 1 year post-displacement across all displaced job seekers in our sample. For the vast majority of job seekers, close to 0% of local firms are in-network firms. Even for very large schools, only 10% of local firms are in-network. Nevertheless, more than 60% of workers are re-employed in-network, as shown in Panel B.

This evidence is suggestive that the job search channel is dominant in the treatment effect β^{Stay} . Formally, we re-estimate Equation 1.9 on two new outcomes. First, we define $1(Stay \& alum)_i$ as an indicator equal to 1 if worker i stays in metro o and is re-employed at a firm with a co-alum from school s 1-year post-displacement. Next, we define $1(Stay \& no alum)_i$ as an indicator equal to 1 if worker i stays in metro o and is re-employed 1-year post-displacement at a firm without any co-alumni from school s . We separately re-estimate Equation 1.9 as a logit with these two new dependent variables and denote the estimated coefficients on $\log(\sigma_{ose_i})$ as $\beta^{Stay,alum}$ and $\beta^{Stay,noalum}$, respectively.³⁸ A finding of $\beta^{Stay,alum} > \beta^{Stay,noalum}$ would strongly suggest that job search channels are dominant in the treatment effect.

The results are presented in Table 1.5. We re-report our baseline estimate of β^{Stay} , 0.389, in column 1. Relative to this baseline, we find that increasing the alumni network size in

³⁷From bootstrapped standard errors based on 400 reps and re-sampling clusters at the origin $\times$ school level.

³⁸To maintain algebraic consistency with the log-odds equation for $Pr(stay_i)$ that is our econometric specification in Equation 1.9, we would need an additional predictor for the log conditional probability of being re-hired at a firm with (without) a co-alumni conditional on staying in the metro of displacement. However, accurately measuring this log conditional probability would require additional modeling of the separate firm and job-level hiring outcome that is beyond the scope of this paper. As a reasonable approximation, we assume that this log conditional probability is the sum of (1) the local log co-alumni network size in the city of displacement, and (2) some function that is identical for workers in the same city o who attended geographically proximate and highly similar schools in the same strata $K(s)$. The effect of (1) is directly captured by $\beta^{Stay,alum}$ or $\beta^{Stay,noalum}$, and the effect of (2) is absorbed by city $\times$ school strata fixed effects $\phi_{oK(s)}$.

the city of displacement has an even more pronounced effect on the joint probability of staying and finding re-employment at a firm with another co-alumni (estimated elasticity of jointly staying and finding in-network re-employment, $\beta^{Stay,alum}$, of 0.55, p-value < 0.01). By contrast, in response to increasing the local alumni network size, job seekers are actually less likely (estimated $\beta^{Stay,noalum}$ of -0.34, p-value < 0.01) to stay and be re-employed without a co-alumni at their new firm. In parallel, we see that the negative impact of a higher inclusive value of the *Leave* nest on a job seeker's probability of staying comes largely from a decreased likelihood of finding new employment at local in-network firms. Meanwhile, a higher inclusive value has no discernible impact on the likelihood of staying in the city of displacement and finding employment outside of the alumni network. This result comes from the offset of two opposite-sign effects: 1) a smaller inclusive value pushes job seekers to local in-network firms and away from out-of-network firms, and 2) a larger inclusive value pushes job seekers away to other cities. Altogether, our results suggest that the local co-alumni network's positive impact on a job seeker's propensity to stay in the city of displacement comes largely from the network's value in directing job seekers to local in-network firms. This finding is perhaps unsurprising given descriptive evidence that job-related factors are the primary drivers of geographic mobility in the US (Jia et al. 2023), yet novel given a lack of focus on job search in previous studies on social networks' impact on geographic mobility.

We also test whether co-alumni network size in the origin differentially impacts out-mobility for job seekers across age groups. In addition to local networks of co-alumni, older job seekers with more labor market experience have larger networks of former colleagues who can provide both information about job opportunities (Cingano and Rosolia 2012) and signals of these job seekers' productivity to their incumbent employers (Hensvik and Skans 2016). By contrast, younger job seekers with less labor market experience may more heavily rely on their college alumni networks in the job search process. Furthermore, younger job seekers may be more likely to value local co-alumni networks as a local amenity. Friends and acquaintances can greatly impact location choice.³⁹ For younger job seekers, it is reasonable to expect that a higher share of their immediate social networks come from school.

Thus we might expect that a larger network in the origin would make a younger job seeker more likely to stay than an older one. Yet the baseline inter-metro mobility rate among younger individuals is far greater, indicating a higher propensity to leave (Jia et al. 2023). We repeat our empirical exercises from Table 1.5, which estimate the effect of co-alumni network size in the city of displacement on a job seeker's probability of staying as well as jointly staying and finding re-employment at a firm with or without a co-alumni, separately for older (35 and older) and younger workers (40 and under). We plot our estimates in Panel I of Figure 1.6.⁴⁰

³⁹For instance, Koenen and Johnston (2024) document that on average, 78% of a Facebook user's friends live within 100 miles of their home CZ. Büchel et al. (2020) find that individuals strongly prefer neighborhoods where more of their phone contacts reside.

⁴⁰We use slightly overlapping samples, i.e. workers between the ages of 35 and 40 are part of both groups, in order to preserve more of our sample for reliability. Conceptually, it is also unclear whether workers with more than 10 but fewer than 20 years of experience should be classified as old or young.

We find that the impact of co-alumni network size on a job seeker’s probability of staying in the same local labor market is greater for older workers (estimated β^{Stay} of 0.621, p-value < 0.01) than for younger workers (estimated β^{Stay} of 0.338, p-value < 0.01). We also find larger estimated impacts on jointly staying in the same metro and finding re-employment at an in-network firm for older workers (estimated $\beta^{Stay,alum}$ of 0.653, p-value < 0.01) than for younger workers (estimated $\beta^{Stay,alum}$ of 0.551, p-value < 0.01), albeit with a smaller gap. Therefore, while we may have expected younger workers, who are more likely to leverage college alumni networks, to also be more likely to stay in the origin in response to an increase in network size, we cannot conclude this to be the case. Instead, we largely capture the generally higher likelihood of older workers to not move.

Our estimated treatment effect of origin network size on the probability of staying is robust to a variety of additional tests. First, while we define alumni networks based on the worker’s highest degree obtained, it is possible that workers with post-Bachelor’s education are nevertheless influenced by their Bachelor’s alumni network. We re-estimate Equation 1.9 using only Bachelor’s-level alumni networks and present results in Table A.4, and find that effects are slightly smaller but similar in magnitude to our baseline specification.

Additionally, one might worry that job seekers are especially attached to the metro where they went to school in a way that is not driven by the alumni network. Because the alumni network is almost always largest in this “home metro”, such an attachment could bias our estimate. Therefore, in Table A.6, we exclude all workers who are displaced in the same metro in which they attended school, estimating only on workers displaced in other metros.⁴¹ Our estimate of β^{Stay} is 0.354, which is very similar to our baseline estimate of 0.389. Thus our results are not sensitive to the widespread inclusion of job seekers who remain in their home metro and this source of potential home bias.

Beyond home bias, we show that our results are robust to controlling for location preferences that vary across individual demographic characteristics. We define 12 demographic groups based on gender (male or female), age group (22-30, 31-39, 40+), and race (white or non-white), and include destination $\times$ demographic group fixed effects in our estimation of the *Leave* nest’s inclusive value and add origin $\times$ demographic group fixed effects to our baseline logit in Equation 1.9. We even estimate models that additionally interact these fixed effects with school strata. Thus we rule out threats as narrow as workers of the same gender, age and race as well as from the same origin or strata having common location preferences. Our estimates of β^{Stay} , reported in Table A.9, are larger than baseline but remain statistically significant at the 1% level.

We also alleviate concerns about idiosyncratic matches between firms in specific cities and schools. Instead of comparing job seekers who are displaced in the same industry and location and who went to schools in the same strata, we estimate a model in which we compare the out-mobility of job seekers who are alumni of schools in the same strata and were displaced from the *same firm*. Our estimates in Table A.11 are slightly larger in magnitude than our baseline estimates, and remain statistically significant at the 1% level.

⁴¹We also estimate the inclusive value in a first step on this same sample.

We then address concerns about differences in the composition of majors across schools, even within the same strata. Concretely, instead of isolating variation in network size within the origin $\times$ industry of displacement (and additionally school strata), we isolate variation in network size within a narrow origin-major-strata cell. Thus even if similar schools differ in their production of certain majors, we only compare graduates displaced in the same city who studied the same major. We report estimates in Table A.13 that are very similar to our baseline results, with even larger estimated elasticities and average marginal effects.

Finally, we consider an alternative measure of alumni network size that does not depend on a job seeker's year of displacement. We repeat our analysis after redefining the alumni network to its average value across the 2010-2019 time period. Table A.15 shows that there is virtually no change to our estimates.

Impact of Local Alumni Network Size on Directed Mobility

We turn to quantifying the impact of destination alumni network size on directed migration to that destination. Table 1.6 presents results from estimating variations of our gravity equation in Equation 1.10 via PPML. Since this gravity equation yields equivalent estimates of $\tilde{\beta}^{Leave}$ as implied by the multinomial logit framework of the *Leave* nest (Guimaraes, Figueirdo, and Woodward 2003), and the alumni network size is expressed in logs, we interpret our estimates of $\tilde{\beta}^{Leave}$ as the elasticity of odds of moving to a destination city d , relative to another city, with respect to the co-alumni network size in d . In all specifications, we cluster standard errors three-ways at the origin $\times$ destination, origin $\times$ school, and destination $\times$ school levels, paralleling Moretti and Wilson (2017). We also include origin $\times$ school fixed effects in all specifications, as consistent with Guimaraes, Figueirdo, and Woodward (2003). Our simplest specification in column 1 only includes origin $\times$ school fixed effects and destination fixed effects, and recovers large and statistically significant (at the 1% level) estimated elasticities of the odds of moving to destination city d with respect to the number of co-alumni in d . Our preferred specification in column 5 includes both destination $\times$ strata fixed effects and origin $\times$ destination fixed effects, and exactly corresponds to our econometric model in Equation 1.10. The former absorbs correlated locational preferences and distant city-specific work utility matches among alumni of highly similar neighboring schools, and the latter absorbs heterogeneous moving costs across origin-destination pairs. Our preferred estimate of 0.907 indicates that a 1% increase in the number of co-alumni in a destination d yields a 0.91% increase (p-value < 0.01) in the odds, conditional on moving, of moving to d over another city — a fairly substantial impact.

We can also re-frame our elasticity of odds interpretation that comes directly from our estimated $\tilde{\beta}^{Leave}$ into a probability-based interpretation. Consider two cities d_1 and d_2 , where conditional on moving, a job seeker's probability of choosing d_1 is 5% whereas their probability of choosing d_2 is 50%. Our estimated $\tilde{\beta}^{Leave}$ of 0.91 implies that doubling the network size in d_1 raises the probability of choosing to live and work there, conditional on moving, from 5% to 9%.⁴² For d_2 , the city that already has a very high baseline conditional

⁴²Calculation is as follows. Doubling the network size in d_1 increases the log odds of choosing d_1 by

choice probability, doubling its network size raises the same probability from 50% all the way to 65%, a substantial increase.

We obtain similar estimates of $\tilde{\beta}^{Leave}$ using a control function approach that isolates arguably exogenous variation in destination network size using the instrumental variable introduced in Section 1.4, $W_{ds} \equiv \frac{Share_{s,1972}}{1+dist_{ds}}$. Here, $Share_{s,1972}$ is the school's in-state enrollment share in 1972 and $dist_{ds}$ is the great circle distance (in miles) between the destination and school. We provide more details about our estimation procedure in Appendix A.4. We establish the instrument's relevance by noting, in column 1 of Table 1.7, a large F-statistic of 1,117.5 in our first stage.⁴³ We report our two-stage residual inclusion (2SRI) estimate of $\tilde{\beta}^{Leave}$ in column 1 of Table 1.7. The estimate of 0.999 is very close to our baseline estimate of 0.907, and is statistically significant at the 1% level.⁴⁴

Next, we evaluate the relative contributions of the job search channel and the amenity channel. That is, we evaluate the relative importance of ω^{Leave} and ρ^{Leave} to $\tilde{\beta}^{Leave}$. Concretely, we define P_{ods}^{alum} as the number of displaced job-seeking alumni from school s who move from city o to d and are re-employed 1-year post-displacement at firms with at least one other co-alumni. Likewise, we define P_{ods}^{noalum} as the number of o to d movers from school s who are re-employed at firms without any co-alumni. We then separately re-estimate Equation 1.10 with P_{ods}^{alum} and P_{ods}^{noalum} as our dependent variables, with $\tilde{\beta}^{Leave,alum}$ and $\tilde{\beta}^{Leave,noalum}$, the respective coefficients on the log network size in d , as the corresponding parameters of interest. While our estimate of $\tilde{\beta}^{Leave}$ informs us about the total effect of an increase in the size of the local alumni network in d on the likelihood of moving to d , our estimate of $\tilde{\beta}^{Leave,alum}$ provides indirect evidence on the strength of the network's job search-related channels in affecting this location choice. Likewise, a large, positive estimate of $\tilde{\beta}^{Leave,noalum}$ would suggest that job seekers may still value local alumni networks for non job-related reasons. We re-emphasize that even for large schools with the largest networks across cities, only a small fraction of firms in those cities have co-alumni from those schools (10% average share in city of displacement for job seekers who went to schools with more than 20,000 students). Thus we do not expect that $\tilde{\beta}^{Leave,alum}$ would exceed $\tilde{\beta}^{Leave,noalum}$ for purely mechanical reasons.

Table 1.8 presents the decomposition exercise, where we re-estimate our most preferred specification of Equation 1.10 with the two alternate outcomes P_{ods}^{alum} and P_{ods}^{noalum} in columns 2 and 3, respectively. Similar to our out-mobility results in Section 1.5, in cities with an increase in local alumni network size, we find an increased coincidence of directed migration

0.91 log(2). This change in log odds implies a multiplicative change in odds of $\exp(0.91 \log(2))$. The original odds of choosing d_1 are $\frac{0.05}{1-0.05} \approx 0.053$, meaning the new odds are $0.053 \exp(0.91 \log(2))$. Denoting the new conditional choice probability of d_1 as p_{d_1} , and using the odds expression where new odds are equal to $\frac{p_{d_1}}{1-p_{d_1}}$, we obtain that $p_{d_1} \approx 0.09$. We can do the same calculation for d_2 .

⁴³Several factors contribute to this large F-statistic. First, there are 399,131 observations in the first stage regression, at the origin $\times$ school $\times$ destination level. Additionally, college graduates tend to locate close to school, so it is unsurprising to see a strong correlation between inverse distance and network size.

⁴⁴From bootstrapped standard errors based on 400 reps and re-sampling clusters at the origin $\times$ school level.

and in-network re-employment that is greater than the increase in mobility alone. Compared to our baseline estimate of $\hat{\beta}^{Leave}\%$ of 0.91%, which we re-report in column 1 of Table 1.8, in column 2 we estimate an even larger effect of the destination log alumni network size on the relative odds of jointly moving to the destination and finding re-employment at a firm with another co-alumni (1.12%, p-value < 0.01). By contrast, we find a much smaller effect on the relative likelihood of moving to the destination and finding re-employment at a firm without any co-alumni present (0.56%, p-value < 0.01).

We also find that, in contrast to our out-mobility analysis, these large effects of destination network size on the location choice of movers are even more pronounced for younger job seekers. In Panel II of Figure 1.6, we re-estimate our gravity equation in Equation 1.10 separately for older (≥ 35) and younger (≤ 40) workers. Our dependent variable is now the number of o to d movers from school s within each age group. Conditional on moving, we find that younger workers are more likely to choose a destination city that experiences a 1% increase in its co-alumni network size, with a large estimated elasticity of odds (1.1%, p-value < 0.01), than older workers (estimated elasticity of 0.64%, p-value < 0.01). The lower propensity to stay in the origin post-displacement and the very large mobility response of younger job seekers to increases in destinations' alumni network sizes suggests that younger job seekers are quite willing to move to where their networks are.

Altogether, despite our focus on more broadly defined networks that primarily include individuals with whom job seekers in our sample have had no prior interaction, we find strong impacts of alumni network size on both out-mobility and directed migration. These findings strengthen a burgeoning literature on social networks and migration that has focused on more personalized networks (Büchel et al. 2020; Green 2024; Koenen and Johnston 2024).⁴⁵ In our unique setting in which the social networks we study are more directly related to job search, we additionally find even greater impacts of social networks on joint mobility and in-network re-employment. In light of job search's outsized importance in determining inter-metro migration, this finding suggests that college alumni networks play a significant role in influencing the geographic mobility of college-educated workers in the US.

We note, however, that other factors may still have a greater impact on geographic mobility. For instance, our estimated elasticities of the odds of migration with respect to the co-alumni network size in a destination are smaller than estimated elasticities of migration with respect to net-of-tax rate (Kleven et al. 2014) or tax rate differentials (Moretti and Wilson 2017).⁴⁶ This result is perhaps unsurprising, since earnings are a key component of indirect utility in many models of location choice, including ours to the extent that work utility nests earnings, and changes in tax rates directly and saliently change earnings. By comparison, changes in alumni network size impact earnings only by indirect means, through

⁴⁵As a point of comparison, Koenen and Johnston (2024) estimate that one additional Facebook friend in a commuting zone increases the probability of locating there by 0.33 percentage points. In our setting, one additional co-alumni does not have the same impact on mobility, but any change in alumni network size beyond a small percentage increase would.

⁴⁶Our elasticities with respect to local alumni network size are between a third to half as large as those in Kleven et al. (2014).

various job search-related mechanisms.

Our directed mobility results are robust to the same tests as for the out-mobility results detailed above, as well as additional tests. First, we consider a non-nested multinomial logit structure in which job seekers choose among all possible locations, including the origin, non-sequentially. We provide more details in Appendix A.4. In Table A.2, we re-estimate Equation 1.10 but for both movers and stayers, and obtain an estimated elasticity of location choice with respect to network size that is slightly larger than our estimate of β^{Stay} but smaller than our estimate of $\tilde{\beta}^{Leave}$. This is sensible given that the multinomial logit structure delivers a singular β that governs both the out-mobility and directed migration margins.

From this framework, we also derive and estimate an alternative gravity equation based on Moretti and Wilson (2017). This specification, Equation A.11 in Appendix A.4, is a log-odds gravity equation estimated via OLS. This approach unfortunately drops many observations from our sample due to zero within-school flows between many origins and destinations that are more effectively handled via PPML estimation. But since identification relies on destination-origin differences in alumni network size, we can directly include destination $\times$ school fixed effects to absorb any remaining concerns about within-school correlated preferences and city-school determinants of work utility in non-origin cities. We present OLS estimates in Table A.3. Reassuringly, even with destination $\times$ school fixed effects, our estimates of the elasticity (odds) of choosing a destination over staying in the origin, with respect to a change in the destination-origin alumni network size difference, are between our estimates of β^{Stay} and $\tilde{\beta}^{Leave}$.

Next, we re-estimate Equation 1.10 using Bachelor’s-degree-only alumni networks and present results in Table A.5. We find slightly smaller estimated effects to baseline, but still similar in magnitude, providing reassurance that our results are not driven by workers with post-Bachelor’s education.

Next, we address the concern that workers may be especially attached to the metro they attended school in. As in the robustness checks in Section 1.5, we repeat our analysis using only workers who experience job loss outside their school’s metro and present results in Table A.7. We find similar estimated elasticities as in our baseline specification. One might additionally worry that our results are driven by “return migration”, or workers moving back to the metro they went to school in, where they tend to have large alumni networks. Only 6% of job seekers who are displaced outside their “home metro” actually make this choice. Nevertheless, in Table A.8 we exclude such individuals. Unsurprisingly, the resulting treatment effect is slightly smaller, but remains both large and statistically significant.

We address the possibility of location preferences varying across demographic groups by estimating expanded versions of our gravity equation that collapses counts of movers to the origin-destination-school-demographic group level (in comparison to origin-destination-school in our baseline analysis). We define demographic groups in the same manner as our parallel exercise for out-mobility in Section 1.5. Table A.10 reports estimates that are very similar to our baseline estimates.

We also estimate an expanded version of our gravity equation that includes firm-specific heterogeneity. We relate counts of movers collapsed to the origin-destination-school-firm

level to destination alumni network size, and directly control for factors that vary across school-firm pairs and destination-firm pairs. Table A.12 reports our estimated elasticities from these models, which are very similar to our baseline estimates.

Likewise, we estimate an expanded version of our gravity equation that includes major-specific heterogeneity. We relate counts of movers collapsed to the origin-destination-school-major level to destination alumni network size, and directly control for factors that vary across school-major pairs and destination-major pairs. Our estimates, as reported in Table A.14, are very similar to our baseline estimated elasticities.

Finally, we consider an alternative measure of destination alumni network size that does not depend on job seekers' years of displacement. We calculate network size based on the yearly average across 2010-2019, rather than the average across years of displacement within an origin-school pair. Table A.16 shows virtually no effect on our estimated effects.

1.6 The Role of Peers and Same-Industry Networks

In this section, we analyze two mechanisms by which local co-alumni networks can influence geographic mobility. First, we consider the role of strong, weak, and potential weak ties. Specifically, we estimate the impact of both same-cohort and non-cohort co-alumni network size in a potential destination on location choice. Finally, we analyze to what extent a city's concentration of co-alumni in the same industry from which a job seeker was displaced affects that job seeker's geographic mobility.

Same-Cohort Co-Alumni Networks and the Role of Direct Peers

Our theoretical framework and primary analyses define a local co-alumni network as the entire set of alumni from a particular school in a given city. This broad definition nests strong ties — a job seeker's close personal network from school; weak ties — loose connections with other co-alumni that the job seeker knows personally; and potential ties — co-alumni that may show up on LinkedIn but are outside of the job seeker's personal network (Granovetter 1973). A job seeker's co-alumni network in a given city largely consists of potential ties from cohorts with whom they had no interaction during school. However, this does not necessarily mean that these potential connections from a job seeker's *alma mater* exert a stronger influence on geographic mobility than direct peers — strong, weak, or potential ties from the same or adjacent cohorts in the local co-alumni network. If the local co-alumni network's value in job search primarily comes from direct referrals and the transmission of job information from personal LinkedIn connections, or if job seekers strongly prefer to reside near friends and acquaintances from school, then the network's overall impact on geographic mobility instead stems from the job seeker's smaller network of peers from the same or adjacent cohorts.

To test this possibility, we quantify the impact of both same-cohort and non-cohort local co-alumni network size on the location choice of job seekers. We define the same-cohort

alumni network size as the log number of co-alumni from same and neighboring cohorts of the same highest degree.⁴⁷ Consequently, the non-cohort alumni network size is the log number of co-alumni who either did not complete the same degree as the job seeker or do not belong to the same or a neighboring cohort.

To test the relative strength of the two networks, we decompose the impact of total network size into separate effects of same-cohort and non-cohort alumni network size, in both the city of origin and alternative destinations, on a job seeker's indirect utility. Thus we estimate four parameters of interest: $\beta^{Stay,c}$ and $\beta^{Stay,-c}$ rather than a singular β^{Stay} ; as well as $\tilde{\beta}^{Leave,c}$ and $\tilde{\beta}^{Leave,-c}$ rather than a singular $\tilde{\beta}^{Leave}$. Here, c refers to same-cohort and $-c$ refers to non-cohort. We estimate $\beta^{Stay,c}$ and $\beta^{Stay,-c}$ using the appropriately modified form of our binary logit in Equation 1.9:

$$\begin{aligned} \log\left(\frac{Pr(stay_i)}{1 - Pr(stay_i)}\right) = & \alpha + \beta^{Stay,c} \log(\sigma_{ose_i,c}) + \beta^{Stay,-c} \log(\sigma_{ose_i,-c}) \\ & + X'_i \eta + \delta_{e_i} + \phi_{oK(s)} - \lambda \widehat{IV}_{osc} + \epsilon_{iosc} \end{aligned} \quad (1.11)$$

Here, $\log(\sigma_{ose_i,c})$ and $\log(\sigma_{ose_i,-c})$ are same-cohort and non-cohort network sizes, respectively. The inclusive value $\widehat{IV}_{osc}$ is estimated from a modified version of Equation 1.10, using origin $\times$ school $\times$ cohort fixed effects in lieu of origin $\times$ school fixed effects and using both same-cohort and non-cohort log network size in lieu of total network size in a destination d . This same gravity equation yields consistent estimates of $\tilde{\beta}^{Leave,c}$ and $\tilde{\beta}^{Leave,-c}$, as follows:

$$P_{odsc} = \exp(\tilde{\beta}^{Leave,c} \log(\sigma_{ds,c}) + \tilde{\beta}^{Leave,-c} \log(\sigma_{ds,-c}) + \phi_{osc} + \phi_{od} + \phi_{dK(s)}) \varepsilon_{odsc} \quad (1.12)$$

Here, our dependent variable is the number of job seekers from school s and graduating in cohort c , defined as the graduation year and degree of highest completion, e.g. Bachelor's, MBA, etc., who are displaced in city o and move to city d one year post-displacement. Our primary explanatory variables are $\log(\sigma_{ds,c})$ and $\log(\sigma_{ds,-c})$, the log number of same-cohort and non-cohort alumni, respectively, in d .

We display our estimates of $\beta^{Stay,c}$ and $\beta^{Stay,-c}$ in Panel I of Figure 1.7, and our estimates of $\tilde{\beta}^{Leave,c}$ and $\tilde{\beta}^{Leave,-c}$ in Panel II. We find that for both the stay vs. leave decision and the choice over all alternative destinations conditional on moving, the estimated impact of same-cohort network size ($\beta^{Stay,c}$ and $\tilde{\beta}^{Leave,c}$) on our two mobility margins is greater than that of non-cohort network size. In the case of directed migration, our estimate of $\tilde{\beta}^{Leave,c}$ (0.69, p-value < 0.01) is over twice as large as our estimate of $\tilde{\beta}^{Leave,-c}$ (0.27, p-value < 0.01).

Thus our results suggest particularly strong roles for same-cohort alumni in impacting the geographic mobility of college-educated job seekers. This impact is notable given that same-cohort networks are much smaller (non-cohort networks in the city of displacement

⁴⁷For Bachelor's degrees, the set of neighboring cohorts for some graduating year t is $[t - 3, t + 3]$. For Master's degrees and MBA degrees, the set of neighboring cohorts is $[t - 1, t + 1]$. For doctorate degrees, the set of neighboring cohorts is $[t - 4, t + 4]$.

are over 6 times as large, on average).⁴⁸ Our estimates with respect to both same-cohort and non-cohort network size imply that one additional co-alumni with whom a job seeker overlapped during school exerts the same impact on mobility as several additional co-alumni from non-adjacent cohorts.

Our significantly higher estimated impacts of same-cohort alumni network size also suggest that our estimated impacts of local alumni network size on both out-mobility and directed migration capture a true network effect. Idiosyncratic match effects between alumni from a specific school and firms in a specific city are not tied to direct or potential connections between network members. Consequently, such effects are common to different cohorts from the same school. By contrast, co-alumni who graduated around the same time as a job seeker are the most likely network members to interact with the job seeker during the job search process. Our finding of even higher effects from these members speaks to a phenomenon that extends beyond match effects that are far more cohort-invariant.

Co-Alumni Networks and Industry-Specific Human Capital

In this section, we consider a more granular definition of a local co-alumni network, specifically the number of co-alumni in a local labor market who work in the same industry from which a job seeker is displaced.⁴⁹ We ask whether a higher concentration of co-alumni who work in a job seeker’s city and industry of displacement increases the probability that the job seeker stays in that city. We acknowledge that workers are mobile across industries, but our paper’s focus is on geographic mobility specifically, and we do not model a job seeker’s joint location and industry choice.

We motivate our focus on same-industry co-alumni networks by first noting that our sample of job seekers has, on average, more than a decade of labor market experience and thus consists of many individuals who have accrued significant industry-specific human capital (Neal 1995; Parent 2000). For such job seekers, displacement can be very costly if they cannot find new employment in the same industry. To capture the relationship between industry-specific human capital and earnings, we now assume that the number of job draws a job seeker can have in a particular city is proportional to the number of local same-industry alumni rather than all alumni. Thus work utility y_{ds} is log-linear in a city’s number of same-industry alumni. We discuss this modification in greater detail in Appendix A.4.

We continue to assume that the amenity value of having more co-alumni in a particular city is constant across a school’s alumni and not tied to a job seeker’s industry of displacement. To the extent that job seekers differentially value same-industry versus different-industry co-alumni for non-job reasons, our estimates of local industry co-alumni network size on mobility could partially capture this mechanism in addition to job search channels.

We begin by testing whether an increase in the number of same-industry co-alumni in a job seeker’s city of displacement o increases their likelihood of remaining in the same city.

⁴⁸The mean same-cohort network size across job seekers’ cities of displacement is about 2,400 co-alumni, whereas the mean non-cohort network size is over 14,000 co-alumni.

⁴⁹As in previous sections, our industry definition is a 2-digit NAICS code.

We estimate a modified version of Equation 1.9:

$$\log\left(\frac{Pr(stay_i)}{1 - Pr(stay_i)}\right) = \alpha + \beta^{Stay} \log(\sigma_{osje_i}) + X_i' \eta + \delta_{e_i} + \phi_{os} - \lambda \widehat{IV}_{osj} + \epsilon_{iosj} \quad (1.13)$$

Our primary explanatory variable is now $\log(\sigma_{osje_i})$, the log number of alumni in metro o and industry j in the year of displacement who went to worker i 's *alma mater* s . We can now identify β^{Stay} , without imposing any identifying assumptions about city-school work utility A_{ms} or correlated preferences Z_{ms} for any city m , by including destination $\times$ school fixed effects in the estimation of the inclusive value $\widehat{IV}_{osj}$ and origin $\times$ school fixed effects ϕ_{os} .⁵⁰ Crucially, by including ϕ_{os} (and destination $\times$ school fixed effects in the estimation of the inclusive value), we fully absorb lingering threats to identification at the city-school level. We also absorb the impact of total network size in the city of origin.

Our second-step estimate of β^{Stay} is 0.456 (p-value < 0.01). In Figure 1.8, we show that this estimate is larger than our baseline estimate of the impact of total log alumni network size on out-mobility from Equation 1.9. We note that the two estimates are not perfectly comparable since our estimated effect of same-industry origin network size comes from a model that already absorbs the impact of total network size. Nevertheless, the large relative magnitude of our estimated coefficient on same-industry network size suggests that the presence of co-alumni who work in the industry from which a job seeker was most recently separated and had accrued meaningful industry-specific human capital influences out-mobility in a way that extends beyond the mere presence of all co-alumni.

Next, we show that increasing the number of same-industry co-alumni in a potential destination city increases the odds that a job seeker chooses to live and work in that city over another, conditional on moving. We incorporate even more industry-specific heterogeneity into our model by allowing for additional city $\times$ industry and school $\times$ industry work utility components.⁵¹ This inclusion additionally controls for factors such as UC Berkeley graduates potentially having strong productivity matches with firms in the Tech sector, or workers in the New York City Finance sector being more productive. We describe these additional modifications in Appendix A.4, and accordingly estimate the following extension of our primary gravity equation used for our directed mobility analysis:

$$P_{odsj} = \exp(\widetilde{\beta}^{Leave} \log(\sigma_{dsj}) + \phi_{osj} + \phi_{od} + \phi_{ds} + \phi_{dj}) \varepsilon_{odsj} \quad (1.14)$$

We estimate Equation 1.14 on data collapsed at the origin $\times$ destination $\times$ school $\times$ industry level, where P_{odsj} is the number of job seekers who went to school s and were displaced in industry j who move from o to d . We include origin $\times$ school $\times$ industry fixed effects

⁵⁰We estimate $\widehat{IV}_{osj}$ using a modified version of Equation 1.10 at the origin $\times$ school $\times$ industry (of displacement) $\times$ destination level. We replace origin $\times$ school fixed effects with origin $\times$ school $\times$ industry fixed effects, and include same-industry network size in destination d instead of total network size. We replace destination $\times$ strata fixed effects with destination $\times$ school fixed effects, and estimate parameters via PPML.

⁵¹Recall in our baseline model, work utility is defined at the city-school level as y_{ds}

ϕ_{osj} to absorb school $\times$ industry work utility factors and the combined attractiveness of all alternative destinations (see denominator in Equation 1.8b). We also include origin $\times$ destination fixed effects ϕ_{od} to absorb moving costs. Our inclusion of destination $\times$ school fixed effects ϕ_{ds} absorbs citywide co-alumni network size, correlated preferences and city-school work utility factors without requiring additional identifying assumptions. Finally, we include ϕ_{dj} , which are destination $\times$ industry fixed effects that control for location $\times$ industry-specific factors that impact work utility in each location.

We present our estimates in Table 1.9. In our most preferred specification in Column 3 that coincides exactly with Equation 1.14, our estimated elasticity of the odds of choosing a destination city, relative to another, with respect to the same-industry co-alumni network size in that destination, is 0.79% (p-value < 0.01). This estimate is similar to our estimated elasticity of the odds of choosing a destination city with respect to the total alumni network size, irrespective of industry, in that city. Since our inclusion of destination $\times$ school fixed effects already controls for the impact of total network size on location choice, our large estimate of the effect of local same-industry co-alumni network size on destination choice suggests that, as was the case for the stay vs. leave margin, same-industry co-alumni networks play a significant role in influencing the directed mobility of job seekers.

1.7 Conclusion

We show that local college alumni networks have a significant impact on the geographic mobility of exogenously displaced college-educated job seekers. Job seekers are both more likely to stay in the same local labor market post-displacement if they are locally surrounded by a higher concentration of co-alumni, and more likely to move to other local labor markets that have more co-alumni working there. The overall effect of local college alumni network size on mobility comes from both a network's value in job search and its direct value as a consumption amenity, with our evidence suggesting that the job search channel plays a strong role. College alumni networks in alternative destinations have a particularly strong impact on the directed migration of younger job seekers who are less likely to have built other social networks post-graduation. Finally, within a local labor market's stock of all graduates from a particular institution, those who were in a job seeker's same or neighboring cohorts as well as those employed in the same industry from which a job seeker was displaced have a particularly strong impact on their location choice.

Our study advances understanding of internal migration in the US, and potentially other large developed countries, by emphasizing a key determinant of the inter-metro mobility of the most mobile subpopulation — college-educated workers. If internal migration improves individual economic outcomes, then college alumni networks provide an indirect pathway for college-educated job seekers to potentially increase earnings and find better job matches, both locally and far away. In Figure 1.9 panel A, we show that for job seekers who move, the change in their alumni network size (new destination minus origin) is positively correlated

with the change in citywide place effects on earnings from origin to destination.⁵² In panel B of the same figure, when we regress this change in earnings premia (either an indicator for positive change or the actual magnitude) on the change in network size, conditioning on origin $\times$ strata fixed effects, we find a strong, positive relationship between moving to a metro with more co-alumni and moving to a metro with larger place effects. While this relationship is largely descriptive in nature, it suggests that networks drive not only mobility in itself but also mobility to areas of higher economic opportunity. Quantifying these gains with more precise earnings data would be insightful in future research.


We also underscore potential micro-foundations of agglomeration economies and economic spillovers in local labor markets. A strong presence of alumni from a particular school in a given city attracts more job seekers from the same school, and those job seekers overwhelmingly end up at the same firms where co-alumni already work. To the extent that this coincidence of connected college-educated workers at the same local firms exerts a meaningful influence on knowledge spillovers, labor market pooling and matching externalities, college alumni networks directly contribute to the existence of agglomeration economies. Furthermore, in the presence of human capital externalities and local multipliers that predict wage increases for workers without a college degree in response to an increase in a city's stock of college-educated workers, these networks can indirectly improve the economic outcomes of non college-educated workers.

These economic implications highlight many important considerations for policymakers. Local governments often incentivize firms to relocate nearby. This relocation could simply reallocate college-educated workers who were already in the local area. Yet if these firms have a large presence of co-alumni from any particular school(s), then network-induced directed mobility could drive a possibly larger influx of college-educated workers from those schools and the ensuing economic impacts of this migration. Other local policies such as Idaho's House Bill 718 might provide subsidies for students conditional on those students remaining in-state for some time post-graduation (Idaho 66th Legislature 2022), which could encourage other co-alumni to also return to the local area in the future. From a broader welfare perspective, however, tradeoffs between aggregate economic efficiency and spatial inequality must be taken into consideration. While a local government may prioritize policies that capitalize on alumni networks' influence on geographic mobility to encourage more in-migration of college-educated workers, this inflow may come at the expense of a more productive local labor market that in turn loses these workers. Evaluating these tradeoffs is of great policy importance, and a promising line of future research.


⁵²We use estimates of commuting zone earnings premia from Card, Rothstein, and Yi (2025), and map commuting zones to our metro areas using county-level population weights.

1.8 Figures and Tables


Figure 1.1: College Alumni Network Size Across Local Labor Markets, by School



(a) Williams College




(b) Northwestern University




(c) CSU Fullerton

Notes: The heat maps show the 2010-2019 yearly average log number of alumni (across all degree types) from the schools in panels a-c that are working in each metro area. Counts of alumni are obtained from self-reported educational histories on LinkedIn profiles observed in the Revelio Labs data. Our three panels correspond to different types of schools: a) a small, highly selective liberal arts college; b) a mid-sized, highly selective private university; and c) a large, non-highly selective public university.


Figure 1.2: Differences in Alumni Network Size Across Local Labor Markets, by School Pair



(a) Amherst College (red) and Williams College (blue)




(b) University of Chicago (red) and Northwestern University (blue)



(c) CSU Long Beach (red) and CSU Fullerton (blue)


Notes: The heat maps show the difference in 2010-2019 yearly average log number of alumni (across all degree types) that are working in each metro area between pairs of highly similar, geographically proximate schools in panels a-c. Counts of alumni are obtained from self-reported educational histories on LinkedIn profiles observed in the Revelio Labs data. Our three panels correspond to pairs of 3 different types of schools: a) highly selective small liberal arts colleges in Massachusetts; b) elite mid-sized universities in the Chicago metro area; and c) large universities in the Los Angeles metro area that belong to the California State University System.

Figure 1.3: Probability of Relocation, Large vs. Small Local Alumni Network



Notes: This plot compares the relocation probabilities of workers displaced in cities with large vs. small alumni networks. We plot the share of workers in our sample who have ever worked in a different metro from their metro in the quarter of displacement $\tau = 0$, for each quarter τ after displacement. For quarters prior to displacement, we plot the share of workers who are in a different metro from their metro at $\tau = 0$. We split the sample by top and bottom population-weighted terciles of local college alumni network size. In other words, for a given school, we define the top tercile as the cities with the largest number of alumni from that school who account for the first third of the entire alumni population. Similarly, the bottom tercile is the set of cities with the smallest number of alumni from the school who account for the final third of the school's alumni population.


Figure 1.4: Relationship Between Probability of Staying in City of Displacement and log Alumni Network Size



Notes: This plot shows the relationship between alumni network size in the city of displacement and the probability of staying post-displacement. We partition the log yearly average (2010-2019) number of alumni in each job seeker's city of displacement into bins. For each bin, we plot the average (across job seekers) probability of staying in the same city 1 year post-displacement.

Figure 1.5: Concentration of and Re-Employment at Local In-Network Firms

Panel A: Share of Local Firms with Co-Alumni Across Job Seekers




Panel B: Share of Workers Re-Employed at In-Network Firms

	Share of workers w/ co-alum at firm	Share of firms in metro w/ co-alum (mean)	Share of firms (10+ employees) in metro w/ co-alum (mean)
All	0.66	0.07	0.22
By experience			
0 to 5-yr post-grad	0.69	0.07	0.23
6 to 10-yr post-grad	0.68	0.07	0.22
11 to 20-yr post-grad	0.66	0.07	0.22
20-yr+ post-grad	0.61	0.06	0.20
By school type			
Public school grad	0.70	0.09	0.26
Private school grad	0.59	0.03	0.14
By school size			
Less than 5,000 students	0.60	0.03	0.11
5,000-20,000 students	0.66	0.05	0.17
20,000+ students	0.73	0.10	0.30


Notes: In Panel A, we plot, across all displaced job seekers in our sample, a histogram of the share of firms in their respective cities 1 year post-displacement that have at least one co-alumni (in-network firm). In Panel B, we report the mean of this share in the middle column. In the right-most column, we report the mean share of in-network firms among local firms with 10+ employees. In the left-most column, we report the share of workers who are re-employed at an in-network firm 1 year post-displacement. We report these statistics for various subpopulations of our sample. In the top row, we include all displaced job seekers who are re-employed 1 year post-displacement. The next block splits the sample into groups of workers who have different amounts of labor market experience. In the third block, we split the sample into workers who matriculated from public vs. private schools. In the final block, we split job seekers by the size of their *alumni* networks.

Figure 1.6: Effect of Local Alumni Network Size on Mobility, by Age

I) Effect of Origin Alumni Network Size on Probability of Staying




II) Effect of Destination Alumni Network Size on Directed Migration




Notes: We plot estimates of the impact of network size in the city of displacement (origin)/destination on out-mobility/directed migration, separately by age group. Panel I estimates Equation 1.9 separately for workers 40 years old and younger (a) and 35 and over (b). Baseline estimates for all workers are in Table 1.5. We estimate logits with three dependent variables: indicators for whether a displaced worker stays in the origin 1 year post-displacement (1(Stay)); jointly stays and has a co-alum at their new firm (1(Stay+co-alum)); and jointly stays but without any co-alums at their new firm (1(Stay+no co-alum)). All models include the inclusive value of leaving (estimated from Equation 1.10); individual controls: gender, race, age and age-squared; year of displacement fixed effects; and origin $\times$ school strata fixed effects, with strata defined in Data Appendix A.2. Standard errors are bootstrapped, with 400 reps resampling clusters at the origin $\times$ school level. Panel II plots estimates from Equation 1.10 via Poisson pseudo-maximum likelihood (PPML), separately by age group. Dependent variable is the number of displaced alumni from school s in city o in each age group who move to city d 1 year post-displacement. We include origin $\times$ destination fixed effects, origin $\times$ school fixed effects, and destination $\times$ school strata fixed effects. Standard errors are clustered three-ways by origin $\times$ destination, origin $\times$ school, and destination $\times$ school. 95% confidence intervals are drawn for each estimate.

Figure 1.7: Effect of Same-Cohort and Non-Cohort Co-Alumni Network Size on Out-mobility and Directed Migration



Notes: We plot estimates of the impact of same-cohort and non-cohort network size on out-mobility and directed migration. Panel I presents estimates from Equation 1.11 for out-mobility — a logit with the outcome as an indicator for whether a displaced worker stays in the same metro (origin) 1 year post-displacement. We plot the estimated coefficients on the log number of co-alumni in the origin who are and are not in the job seeker’s graduating or neighboring cohort, respectively. We define cohort as the graduation year $\times$ degree of highest completion. We include the inclusive value of leaving; individual controls: gender, race, age and age-squared; year of displacement fixed effects; and origin $\times$ school strata fixed effects. School strata are defined as in Data Appendix A.2. Standard errors are bootstrapped, with 400 reps resampling clusters at the origin $\times$ school $\times$ cohort level. Panel II plots estimates of same (and neighboring)-cohort network size and non-cohort network size’s impacts on destination choice, conditional on moving, from Equation 1.12 estimated via PPML. Dependent variable is the number of displaced alumni from school s and cohort c in city o who move to city d 1 year post-displacement. We include origin $\times$ destination fixed effects, origin $\times$ school $\times$ cohort fixed effects, and destination $\times$ school strata fixed effects. Standard errors are clustered three-ways by origin $\times$ destination, origin $\times$ school, and destination $\times$ school. 95% confidence intervals are drawn for each estimate.

Figure 1.8: Effect of Same-Industry Co-Alumni Network Size on Out-mobility



Notes: We plot the coefficient on same-industry log alumni network size in our out-mobility logit from Equation 1.13 (left), alongside the coefficient on total log alumni network size in our separate baseline out-mobility logit from Equation 1.9 (right). Outcome is an indicator for whether a displaced worker stays in the same metro 1 year post-displacement. Our estimate of the same-industry network size effect comes from a model that already includes origin $\times$ school fixed effects, and whose estimated inclusive value comes from a gravity equation, Equation 1.10, that directly includes destination $\times$ school fixed effects and same-industry network size in the destination. The origin $\times$ school fixed effect thus absorbs the impact of total network size in the origin on out-mobility. Also included are the same additional controls and fixed effects in column 3 of Table 1.2. Standard errors are bootstrapped, with 400 reps resampling clusters at the city $\times$ school $\times$ industry level (city $\times$ school level for total network size analysis). On the right side, we plot our baseline estimate of origin total network size's impact on the probability of staying, from column 3 of Table 1.2. See table notes for Table 1.2 for more details. We examine the impact of same-industry network size on directed migration in Table 1.9.

Figure 1.9: Relationship Between Changes in Local Alumni Network Size and Changes in Citywide Earnings Premia

Panel A: Raw Correlation Across Movers



Panel B: Correlation Within Origin and Strata

	Outcome	
	$1(\Delta \text{ Place Effect} > 0)$	$\Delta \text{ Place Effect}$
	(1)	(2)
1-Year Change in log Number of Co-Alumni	0.016** (0.007)	0.010*** (0.002)
Observations	2,067	2,067
Origin x school strata FE	X	X

Notes: We show descriptively that job seekers, conditional on moving, relocate to cities with both higher network sizes and higher earnings premia. In Panel A, we partition the difference in log alumni network size between a job seeker's city of displacement (origin) and their city 1 year post-displacement (destination) into bins, for movers only. For each bin, we plot the average destination-origin difference in citywide earnings premia from Card, Rothstein, and Yi (2025). The earnings premia calculated in Card, Rothstein, and Yi (2025) are at the commuting zone level, which we map to our metro areas based on county population weights. In Panel B, we regress, for movers only, an indicator for a positive destination-origin difference in earnings premia (column 1) and the actual difference (column 2) on the difference in log alumni network size between destination and origin. We include origin $\times$ school strata fixed effects. School strata are defined as in Data Appendix A.2. Standard errors are clustered at the origin $\times$ school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table 1.1: Summary Statistics for Job Seekers at Time of Displacement

	<u>All</u>		<u>Top Tercile</u>		<u>Bottom Tercile</u>	
	mean	sd	mean	sd	mean	sd
Number of Co-Alumni in Metro	17371.47	27393.98	33874.02	33268.70	3848.16	12580.25
Age	35.56	9.25	34.64	8.84	36.76	9.68
Female	0.35	0.48	0.37	0.48	0.34	0.48
White	0.81	0.39	0.80	0.40	0.81	0.39
Asian	0.07	0.25	0.07	0.26	0.06	0.24
Black	0.07	0.25	0.07	0.25	0.07	0.25
Hispanic	0.05	0.22	0.06	0.23	0.06	0.23
Other	0.00	0.06	0.00	0.07	0.00	0.05
Predicted Salary of Laid-Off Position	83585.23	41288.72	82982.27	41776.13	80908.28	38427.85
Observations	20459		8807		6904	

Notes: We present summary statistics for our sample of displaced job seekers at the time of displacement. We split the sample by top and bottom population-weighted terciles of local college alumni network size, within school and across metros. In other words, for a given school, we define the top tercile as the cities with the largest number of alumni from that school who account for the first third of the entire alumni population. Similarly, the bottom tercile is the set of cities with the smallest number of alumni from the school who account for the final third of the school's alumni population. Age is imputed based on year of entry to the worker's Bachelor's degree program. Female and ethnicity is predicted from a user's name on their LinkedIn profile. Predicted salary comes from a model, developed by Revelio Labs, taking firm, location, occupation, etc. as inputs.

Table 1.2: Effect of Alumni Network Size in City of Displacement on Probability of Staying

	Outcome: 1(Stay metro 1-year post-displacement)		
	(1)	(2)	(3)
log Number of Co-Alumni	0.325*** (0.012)	0.426*** (0.061)	0.389*** (0.072)
Inclusive Value	-0.487*** (0.045)	-0.386*** (0.089)	-0.367*** (0.110)
Average Marginal Effect of log Number of Co-Alumni	0.054*** (0.002)	0.067*** (0.009)	0.060*** (0.011)
Observations	17,435	12,943	10,743
Individual controls	X	X	X
Institutional controls	X	X	
Layoff year FE	X	X	X
Origin FE	X		
Origin x school metro FE		X	
Origin x school strata FE			X

Notes: Estimates from second-step logit in Equation 1.9. Dependent variable is an indicator for whether a displaced worker stays in the same metro 1 year post-displacement. Estimation proceeds in two steps. In a first step, we estimate, via Poisson pseudo-maximum likelihood (PPML), the parameters in our choice probability model (Equation 1.8b) using our gravity equation for destination choice (conditional on moving) in Equation 1.10. We use these estimates to construct the inclusive value. In our second step, we estimate logits at the worker level. All models include individual controls: gender, race, age and age-squared. Columns 1 and 2 include institutional controls, namely school size. We include fixed effects for the calendar year of displacement in all columns. In column 1 we include origin fixed effects, which are replaced in column 2 by origin $\times$ school metro fixed effects, and finally replaced by origin $\times$ school strata fixed effects in column 3. School strata are defined as in Data Appendix A.2. Standard errors are bootstrapped, based on 400 reps and cluster resampling at the city $\times$ school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table 1.3: Effect of Alumni Network Size in City of Displacement on Probability of Staying
— Logit and Saturated Linear Probability Models

	Outcome: 1(Stay metro 1-year post-displacement)			
	Logit	LPM		
	(1)	(2)	(3)	(4)
log Number of Co-Alumni	0.389*** (0.072)	0.061*** (0.012)	0.070*** (0.023)	0.050*** (0.014)
Inclusive Value	-0.367*** (0.110)	-0.060*** (0.018)	-0.084 (0.053)	-0.042** (0.020)
Average Marginal Effect of log Number of Co-Alumni	0.060*** (0.011)	0.061*** (0.012)	0.070*** (0.023)	0.050*** (0.014)
Observations	10,743	11,135	11,012	9,221
Individual controls	X	X	X	X
Layoff year FE	X	X	X	X
School FE			X	
Origin x school strata FE	X	X	X	
Origin x school strata x industry FE				X

Notes: Estimates from second-step logit and linear probability model (LPM) approximations in Equation 1.9. Dependent variable is an indicator for whether a displaced worker stays in the same metro (origin) 1 year post-displacement. Estimation proceeds in two steps. First, we estimate, via Poisson pseudo-maximum likelihood (PPML), the parameters in our choice probability model (Equation 1.8b) using our gravity equation for destination choice (conditional on moving) in Equation 1.10. We use these estimates to construct the inclusive value of leaving. In our second step, we estimate worker-level logits and LPMs. All models include individual controls: gender, race, age and age-squared; and year of displacement fixed effects. Column 1 reports our baseline logit estimate, with column 2 as the LPM equivalent. School strata are defined as in Data Appendix A.2. In column 3, we add school fixed effects to the LPM in column 2. In column 4, we replace origin $\times$ strata fixed effects and school fixed effects with fully interacted origin $\times$ strata $\times$ industry of displacement fixed effects. Standard errors are bootstrapped, with 400 reps and cluster resampling at the city $\times$ school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table 1.4: Effect of Alumni Network Size in City of Displacement on Probability of Staying
— Control Function Approach with Instrumental Variable

	Outcome: 1(Stay metro 1-year post-displacement)	
	CF+IV	Baseline
	(1)	(2)
log Number of Co-Alumni	0.478** (0.206)	0.389*** (0.072)
Inclusive Value	-0.327** (0.138)	-0.367*** (0.110)
Observations	10,068	10,743
First Stage	1.422*** (0.076)	
F-statistic	349.3	
Individual controls	X	X
Layoff year FE	X	X
Origin x school strata FE	X	X

Notes: Column 1 reports our two-stage residual inclusion (2SRI) estimate of β^{Stay} , the effect of alumni network size in the city of displacement (origin) on the probability of staying, using a control function approach with an instrumental variable, alongside our baseline estimate of β^{Stay} from Table 1.2 column 3. Our instrument is $W_{os} = \frac{Share_{s,1972}}{1+dist_{os}}$, where $Share_{s,1972}$ is a school's share of in-state students in 1972 and $dist_{os}$ is the great circle distance between origin and school in miles. In a first stage, origin log network size in a given year is regressed on W_{os} , $dist_{os}$, the estimated inclusive value of leaving, origin $\times$ school means of individual controls (gender, race, age and age-squared), year of displacement fixed effects, and origin $\times$ school strata fixed effects, where strata are defined as in Data Appendix A.2. The inclusive value comes from 2SRI estimates, via Poisson pseudo-maximum likelihood (PPML), of parameters in our choice probability model plus a first stage residual (Equation A.7). Our first stage is estimated at the origin $\times$ school $\times$ year of displacement level, weighted by number of individuals in each cell, with 10,424 observations. First stage standard errors are clustered at the origin $\times$ school level. We then estimate a binary logit of the probability of staying in the origin on origin log network size, the first stage residual, individual-level controls in lieu of averages from the first stage, and keep all other predictors the same. We report bootstrapped standard errors with 400 reps and cluster resampling at the origin $\times$ school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table 1.5: Effect of Alumni Network Size in City of Displacement on Probability of Staying and Re-Employment with or without Co-Alumni

	Outcome		
	1(Stay)	1(Stay + co-alum at firm 1-yr post-displacement)	1(Stay + no co-alum at firm 1-yr post-displacement)
	(1)	(2)	(3)
log Number of Co-Alumni	0.389*** (0.072)	0.545*** (0.062)	-0.343*** (0.075)
Inclusive Value	-0.367*** (0.110)	-0.229** (0.095)	-0.081 (0.118)
Average Marginal Effect of log Number of Co-Alumni	0.060*** (0.011)	0.111*** (0.012)	-0.048*** (0.010)
Observations	10,743	10,245	9,633
Individual controls	X	X	X
Layoff year FE	X	X	X
Origin x school strata FE	X	X	X

Notes: Estimates of the impact of log alumni network size in the city of displacement (origin) on the joint probability of staying and finding re-employment with or without a co-alumni. We estimate the second-step logit in Equation 1.9. Column 1 reports our baseline estimate of β^{Stay} . Dependent variables in columns 2 and 3 are indicators for jointly staying in the same metro and having or not having, respectively, a co-alum at their new firm. Estimation proceeds in two steps. First, we estimate, via Poisson pseudo-maximum likelihood (PPML), the parameters in our gravity equation for destination choice (conditional on moving) in Equation 1.10. We use these estimates to construct the inclusive value. In our second step, we estimate logits at the worker level. All models include individual controls: gender, race, age and age-squared; fixed effects for the calendar year of displacement; and origin $\times$ school strata fixed effects, where origin refers to the city in which the worker was displaced. School strata are defined as in Data Appendix A.2. Standard errors are bootstrapped, based on 400 reps and cluster resampling at the city $\times$ school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table 1.6: Effect of Alumni Network Size in Destination City on Directed Migration

	Outcome: Number of o to d movers				
	(1)	(2)	(3)	(4)	(5)
Log Alumni Network Size	0.819*** (0.010)	0.837*** (0.032)	0.612*** (0.015)	0.790*** (0.045)	0.907*** (0.060)
Observations	459,361	51,360	87,142	31,422	17,431
Origin x School FE	X	X	X	X	X
Destination FE	X				
Destination x School Metro FE				X	
Destination x Strata FE		X			X
Origin x Destination FE			X	X	X

Notes: Estimates of $\tilde{\beta}^{Leave}$, the effect of destination alumni network size on destination choice conditional on moving, from gravity equation (Equation 1.10) via Poisson pseudo-maximum likelihood (PPML). Dependent variable is the number of displaced alumni from school s in city o who move to a different city d 1 year post-displacement. Primary independent variable is the log alumni network size (for school s) in d . Column 1 includes origin $\times$ school fixed effects and destination fixed effects. Column 2 adds destination $\times$ school strata fixed effects, where strata are defined as in Data Appendix A.2. Column 3 replaces destination fixed effects and destination $\times$ strata fixed effects with origin $\times$ destination fixed effects. Column 4 adds destination $\times$ school metro fixed effects, where school metro is the metro of a job seeker's *alma mater*. Column 5 replaces destination $\times$ school metro fixed effects with destination $\times$ school strata fixed effects. Standard errors are clustered three-ways at the origin $\times$ destination, origin $\times$ school, and destination $\times$ school levels. Significance levels: * 10%, ** 5%, and *** 1%.

Table 1.7: Effect of Alumni Network Size in Destination City on Directed Migration — Control Function Approach with Instrumental Variable

	Outcome: Number of o to d movers	
	CF+IV	Baseline
	(1)	(2)
Log Alumni Network Size	0.999*** (0.164)	0.907*** (0.060)
Observations	15,649	17,431
First Stage	2.594*** (0.078)	
F-statistic	1,117.5	
Origin x School FE	X	X
Destination x Strata FE	X	X
Origin x Destination FE	X	X

Notes: In column 1, we report our two-stage residual inclusion (2SRI) estimate of $\tilde{\beta}^{Leave}$, the effect of destination network size on destination choice conditional on moving, using a control function approach with an instrumental variable. We report this estimate alongside our baseline estimate of $\tilde{\beta}^{Leave}$ Table 1.6 column 5. Our instrument is $W_{ds} = \frac{Share_{s,1972}}{1+dist_{ds}}$, where $Share_{s,1972}$ is a school's share of in-state students in 1972 and $dist_{ds}$ is the great circle destination-school (excluding origin) distance in miles. In a first stage, destination log network size, averaged across job seekers' years of displacement within origin $\times$ school, is regressed on W_{ds} , $dist_{ds}$, origin $\times$ school fixed effects, origin $\times$ destination fixed effects and destination $\times$ strata fixed effects, where strata are defined as in Data Appendix A.2. There are 399,131 observations at the origin $\times$ destination $\times$ school level in this first stage. First stage standard errors are clustered three ways by origin $\times$ school, origin $\times$ destination, and destination $\times$ school. We then estimate a gravity equation via Poisson pseudo-maximum likelihood (PPML) relating the number of displaced alumni from school s in city o who move to city d 1 year post-displacement to destination log network size, the first stage residual, and the same fixed effects and $dist_{ds}$ from the first stage. PPML drops many separated observations from the first stage. We report bootstrapped standard errors based on 400 reps and cluster resampling at the origin $\times$ school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table 1.8: Effect of Destination Alumni Network Size on Directed Migration for Job Seekers Who Move and are Re-employed with Versus without Co-alumni at New Firm

	Outcome		
	Number of o to d movers	Number of o to d movers w/ co-alum	Number of o to d movers w/o co-alum
	(1)	(2)	(3)
Log Alumni Network Size	0.907*** (0.060)	1.168*** (0.123)	0.564*** (0.153)
Observations	17,431	6,100	2,034
Origin x Destination FE	X	X	X
Origin x School FE	X	X	X
Destination x Strata FE	X	X	X

Notes: Estimates of $\tilde{\beta}^{Leave}$, the effect of destination alumni network size on destination choice conditional on moving, from gravity equation (Equation 1.10) via Poisson pseudo-maximum likelihood (PPML), separately for workers re-employed with or without a co-alumni at their new firm. Dependent variable in column 1 is the number of displaced alumni from school s in city o who move to a different city d 1 year post-displacement. Dependent variables in columns 2 and 3 are the number of alumni from school s who move from o to d and are also re-employed with or without, respectively, a co-alumni at their firm 1 year post-displacement. Primary independent variable is the log alumni network size (for school s) in city d . All columns include origin $\times$ destination pair fixed effects, origin $\times$ school fixed effects, and destination $\times$ school strata fixed effects, where strata are defined as in Data Appendix A.2. Standard errors are clustered three-ways at the origin $\times$ destination, origin $\times$ school, and destination $\times$ school levels. Significance levels: * 10%, ** 5%, and *** 1%.

Table 1.9: Effect of Same-Industry Alumni Network Size in Destination City on Directed Migration

	Outcome: Number of o to d movers		
	(1)	(2)	(3)
Log Alumni Network Size	0.636*** (0.017)	0.873*** (0.088)	0.787*** (0.267)
Observations	70,458	10,417	7,386
Origin x School x Industry FE	X	X	X
Origin x Destination FE	X	X	X
Destination x Industry FE	X		X
Destination x School FE		X	X

Notes: Estimates of $\tilde{\beta}^{Leave}$, the effect of same-industry alumni network size in a destination on the destination choice conditional on moving. We estimate our gravity equation (Equation 1.14) via PPML, with data collapsed to the origin $\times$ destination $\times$ school $\times$ industry level. Dependent variable is the number of alumni from school s who are displaced in city o and industry j and move to a different city d 1 year post-displacement. Primary independent variable is the log number of alumni (for school s) in d for a given industry j . All models include origin $\times$ school $\times$ industry fixed effects and origin $\times$ destination pair fixed effects. Column 1 also adds destination $\times$ industry fixed effects. Column 2 replaces destination $\times$ industry fixed effects with destination $\times$ school fixed effects. Column 3 includes both destination $\times$ industry fixed effects and destination $\times$ school fixed effects. Standard errors are clustered three-ways at the origin $\times$ destination, origin $\times$ school, and destination $\times$ school levels. Significance levels: * 10%, ** 5%, and *** 1%.

Chapter 2

The Dynamics of High-Skill Worker Reallocation: Evidence from the Dot-Com Bubble

2.1 Introduction

In recent times, unemployment rates for college graduates have increased, with rates for young college graduates in particular rising to levels that historically have only been observed during recessions (Thompson 2025). It has been well documented that recessions adversely impact the immediate and long-run employment outcomes of new college graduates.¹ But there is much less evidence on how recessions that disproportionately impact high-skill labor affect the future labor market outcomes of college graduates who are already in the labor force. Accounting for these additional effects meaningfully informs economic and policy discussions, as there are nearly 70 million prime-age college-educated workers in the US today (Federal Reserve Bank of St. Louis 2025). Thus any widespread unemployment among the college-educated labor force can meaningfully reshape the broader labor market.

In this paper, I analyze the medium-run and long-run labor market impacts of a recession on college-educated workers, focusing on the 2001 Dot-Com Bubble. The recession that emerged from the collapse of the Dot-Com Bubble lasted from March to September of 2001, and disproportionately affected high-skill workers in the US. This latter feature, paired with the fact that the recession occurred during the present century yet sufficiently long ago, makes the Dot-Com recession a uniquely suitable historical setting to study how economic downturns affect college-educated workers in the long run. Leveraging both variation across local labor markets in the severity of unemployment shocks induced by the recession and data that traces college graduates' post-recession employment histories over a long period of time, I quantify the recession's impacts on several key outcomes, including subsequent employment and mobility across industries and locations.

¹See Oreopoulos, Wachter, and Heisz (2012); Rothstein (2023).

My approach exploits the fact that the Dot-Com crash induced a national economic downturn (National Bureau of Economic Research 2001), yet local labor markets differed in the severity of these downturns. Given the national scope of the recession, rather than reflecting differences in the timing of shocks, differences in local severity of unemployment shocks more plausibly reflected pre-existing differences in potential exposure, e.g. local industry mixes vary widely across space. Identifying the labor market impacts of the Dot-Com recession also requires overcoming concerns about worker selection. To minimize the plausibility that differences in local Dot-Com shocks reflected underlying differences between workers in potential outcomes, I make narrow comparisons between workers within observably similar, narrowly defined groups. These groups isolate workers from the same state and initial industry who are also of the same sex, level of education, and graduating cohort.

I first document that college-educated workers who were initially exposed to larger local unemployment shocks from the Dot-Com recession faced immediate negative effects on subsequent employment up to 6 years post-recession. Concretely, a 1 percentage point larger increase in the local unemployment rate during the recession decreased a worker's probability of employment, relative to their pre-recession baseline probability, in each of the next few years by almost 0.3 percentage points. For individuals who were subsequently re-employed, those who were initially located in more negatively affected areas were more likely to end up at smaller companies following the recession. However, these workers did not end up in systematically lower-paying industries, with little evidence of industry-switching both immediately after the Dot-Com crash and over a decade afterwards.

Next, I examine the recession's impact on geographic mobility. While mobility responses were weak in the first few years after the recession, workers who were initially in local labor markets with larger Dot-Com shocks were significantly more likely to have moved elsewhere by 2005 and beyond. This mobility response is consistent with the mechanism documented in Blanchard and Katz (1992) in which individuals move out of regions that face adverse employment shocks. Surprisingly, however, rather than move to places with larger place-specific pay premia, workers moved to regions that, on average, had lower-paying firms. These places also had lower median rents. Given the negative impacts of the recession on employment, one possible explanation is that immediate earnings losses from the loss of employment drove negatively impacted individuals to lower cost of living areas.

Finally, I show that the Dot-Com recession's impacts on the labor market outcomes of college-educated workers were primarily concentrated in the technology sector. Estimated effects on employment were nearly five times as large in magnitude for tech and tech-adjacent workers as for non-tech workers. Tech workers in more adversely affected regions were also more likely to move after the recession, and more likely to choose destinations with lower average pay premia. Given the sheer size of the US college-educated workforce and the large scale of recent layoffs to high-skill workers in technology-adjacent industries (West 2025), these findings can provide insights on the potential outcomes of workers facing difficult economic conditions in the present and near future.

This paper contributes to a broad literature on the negative labor market impacts of recessions, and their effects on high-skill workers in particular. Recent work has documented

negative effects of entering the labor market during the Great Recession on earnings in the long run (Oreopoulos, Wachter, and Heisz 2012; Rothstein 2023). These effects also generalize to other recessions more broadly (Kahn 2010; Schwandt and Wachter 2019). I complement these studies by focusing on the medium and long-run outcomes of college graduates who were already in the workforce, as opposed to new labor market entrants. While individuals who were already in the labor force do not suffer the immediate negative shock to initial job quality induced by these recessionary environments, a systemic economic downturn that disproportionately affects high-skill workers can still have lasting effects on their career trajectories.

This focus on the labor market outcomes of individuals who were already in the labor force adds to a large literature on the impacts of negative employment shocks. Much of the existing literature focuses on the effects of job displacement on future earnings (Jacobson, LaLonde, and Sullivan 1993; Wachter, Manchester, and Song 2009; Schmieder, Von Wachter, and Heining 2023). I complement these studies by analyzing a broader set of labor market outcomes and focusing on the outcomes of college-educated workers in particular. The Dot-Com bubble burst is a uniquely suitable period to study because unlike traditional mass layoff events that more heavily affect less-educated workers, this recession had an immediate impact on high-skill labor.

Surprisingly, there is little evidence on how the Dot-Com recession affected the medium and long-run labor market outcomes of high-skill workers, despite a high potential relevance today. Some work has documented positive impacts on earnings of entering the tech sector during the Dot-Com bubble itself (Hombert and Matray 2019), but does not isolate post-bubble labor market impacts of the recession. Other work examines the impacts of the Dot-Com crash on human capital investments, and specifically the completion of degrees fields tied to sectors that were most affected by the bubble burst (Weinstein 2022b), but not on the outcomes of individuals who were already in the labor force. I complement these studies by analyzing the effects of the bubble burst on a comprehensive set of individual outcomes encompassing both employment and geographic mobility.

The rest of the paper proceeds as follows. Section 2.2 describes the data sources I use to measure local Dot-Com recession shocks and trace individuals' post-recession employment. Section 2.3 describes the empirical strategy I use to identify the recession's labor market impacts. I present results in Section 2.4. Section 2.5 concludes.

2.2 Data

I combine measures of local unemployment rates at the time of the Dot-Com recession with individual-level data on the educational histories and employment histories (both pre- and post-recession). The individual-level data comes from the company People Data Labs, which collects and harmonizes information from various private and public sources such as LinkedIn profiles, and covers job spells from prior to the 21st century through the present day. I now discuss the data sources in more detail.

Measuring Local Unemployment Shocks

The 2001 recession induced by the burst of the Dot-Com Bubble was a nationwide economic downturn whose severity varied widely across local areas. To measure these associated local downturns, I quantify local unemployment shocks at the commuting zone (CZ) level, a geographic unit that reasonably captures a local labor market.² A key advantage of using CZs as my local labor market definition is that they span the entire US, thus capturing employment in both urban and rural areas. I start with historical data on monthly unemployment rates at the county level provided by the Bureau of Labor Statistics (BLS) Local Area Unemployment Statistics (LAUS) series.

Because counties are embedded within CZs, I aggregate these unemployment rates to calculate CZ-level unemployment rates as the average unemployment rate of the counties within the same CZ, weighting each county by the size of its labor force. To smooth over any idiosyncratic monthly fluctuations, I compute average quarterly unemployment rates for each CZ. I then define a CZ's Dot-Com shock as the change, in percentage points, in its quarterly unemployment rate between 2000 Q4 and 2001 Q3. Q4 of 2000 immediately predates the beginning of the recession, whereas the end of 2001 Q3 coincides with the recession's end.

Figure 2.1 displays a heat map of Dot-Com shocks across CZs. CZs varied widely in their local exposure to the Dot-Com recession. Certain areas in the Southeastern US saw the local unemployment rate increase by more than 4 percentage points in the span of less than a year, whereas other areas, including many CZs in the Great Plains, saw very little change in their unemployment rates. Even within states, CZs were differentially exposed to the Dot-Com recession. San Jose's unemployment rate increased by more than 2 percentage points from 2000 Q4 to 2001 Q3, but in nearby Sacramento, the unemployment rate only increased by 0.5 percentage points during the same time period. I leverage variation in these local shocks to identify the recession's impacts on individuals' post-recession employment outcomes.

Individual Employment Histories

I link my measure of local Dot-Com unemployment shocks to data on the employment histories of college-educated workers to identify these workers' exposure to the shocks and their post-recession labor market outcomes. First, I use data on self-reported educational histories to identify individuals who have completed at least a Bachelor's degree by the end of 2000. For each education spell, I observe the start and end dates, as well as the degree type and institution.

I combine this information with data recording individuals' job spells to isolate college graduates who were employed at the time of the bubble burst. I have detailed information on the location of each job spell, including the city name and state, which I map to the relevant county and, in turn, CZ. Thus the local Dot-Com shock to which an individual is

²I use 1990 CZ definitions.

exposed is based on their job location between 2000 Q4 and 2001 Q3. For workers observed in multiple locations during this time frame, I take the location in which the individual spent the longest amount of time.

To trace their subsequent employment over time, I observe characteristics of their jobs and firms before and at the time of the bubble burst, in addition to these same characteristics for years afterwards. I observe characteristics of both job spells and unemployment spells, such as duration; job characteristics such as occupation title and a measure of occupational seniority; and firm characteristics such as number of employees and industry. Given the rich data on job and firm characteristics, I am able to study how the recession impacted workers' labor market outcomes over a long time horizon.

The main limitation of this data source is that individuals who do not have an updated online presence, e.g. LinkedIn, are unlikely to appear in the data. If individuals who experienced particularly negative effects of the Dot-Com recession on their future labor market outcomes were more likely to drop out of the labor force or fail to update their employment histories online, then any estimated impacts of the recession may be understated. To account for the pattern that LinkedIn usage tends to decrease with age,³ I only compare outcomes for workers with the same highest degree level and graduation year. Nevertheless, my results should be interpreted cautiously as capturing the recession's impacts on workers who remained in the labor force long enough to overlap with and utilize online platforms that could trace their employment histories.

Table 2.1 present means and standard deviations for two primary worker outcomes in 2006 (five years post-recession) and worker characteristics in 2000. The vast majority of workers observed in 2000 were still employed for at least three months in 2006. Among workers who were still employed, 14% had moved to a different CZ by 2006. In 2000, just before the recession, individuals were observed at fairly large firms (on average, nearly 2,200 employees at their company).⁴ About a third of the sample is female. Most of the sample has no more than a Bachelors degree, but over a quarter have a highest degree that is a Masters and 12% have a Doctorate. At the time of the recession, workers on average had over a decade of labor market experience.

2.3 Empirical Strategy

To identify the effects of the Dot-Com Bubble on the medium- and long-run labor market outcomes of high-skill workers, I treat the bubble burst as an exogenous shock to high-skill employment and isolate ensuing variation in local unemployment shocks across CZs. This approach parallels Yagan (2019), which leverages variation in Great Recession severity across CZs in the US, and is similar in spirit to Autor, Dorn, and Hanson (2013), which likewise leverages variation across CZs in import exposure.

³Source: <https://www.linkedhelper.com/blog/linkedin-demographics/>

⁴I use company, firm, and employer interchangeably. Here, the term aggregates all establishments. For example, Microsoft would be a company, including all of its offices across different cities.

I quantify the recession’s labor market impacts by estimating both static and dynamic econometric models. The static model relates an individual’s labor market outcome in a single post-recession year to their local Dot-Com unemployment shock in 2001, as follows:

$$y_{i,post} = \beta \Delta U_{c(i),2001} + \theta_{g(i),2000} + \varepsilon_i \quad (2.1)$$

where $y_{i,post}$ captures a post-recession labor market outcome for individual i , where *post* refers to 5 or 15 years post-recession, e.g. 2006 and 2016. I study two broad classes of outcomes, employment and geographic mobility, to comprehensively quantify the recession’s labor market impacts. Employment outcomes of interest include employment status (relative to pre-recession average from 1997-2000), the change in firm size between post-recession and pre-recession employer, and the change in earnings premia between a worker’s post-recession and initial industry. I complement these outcomes with mobility outcomes—whether the individual moved to a different CZ, and both the change in earnings premia and difference in median rent between their post-recession and initial CZ. These mobility outcomes reflect a key margin of potential adjustment after economic downturns, as highlighted in Blanchard and Katz (1992).

The local Dot-Com shock in i ’s CZ in 2001 is denoted by $\Delta U_{c(i),2001}$. Thus β captures the medium or long-run effects of the local Dot-Com shock, as well as underlying economic forces related to the recession. To isolate comparisons of outcomes for similar workers, I include a set of group fixed effects $\theta_{g(i),2000}$. Here, $g(i)$ denotes a group of individuals who, in 2000, work in the same industry and share the same demographic and educational characteristics as individual i , including graduation year, gender, and highest degree completed.⁵ In my preferred econometric specifications, I further restrict these groups to only include workers residing in the same state in 2000. Another refinement I use is to restrict these groups to only include workers who matriculated from the same college or university.

I complement this econometric specification with a dynamic model that more flexibly traces the Dot-Com recession’s effects on outcomes over many consecutive years, rather than at only two points in time. Concretely, I estimate:

$$y_{it} = \sum_{\substack{k=-4 \\ k \neq 0}}^M \beta_k [\Delta U_{c(i),2001} \times 1(k = t - 2001)] + \theta_{g(i),2000} + \delta_t + \varepsilon_{it} \quad (2.2)$$

Now, y_{it} captures an individual’s labor market outcome in year t . The coefficients β_k are the recession’s effects on outcomes in each post-recession year, and in several pre-period years. I define M as the maximum number of post-recession years, and consider both $M = 6$ (up to 2007) and $M = 15$ (up to 2016). I additionally include year fixed effects δ_t , which absorb any time-varying unobserved determinants of labor market outcomes.

The assumption required to identify β and β_k is that within the same narrowly defined group of individuals who are of the same education level and cohort, and gender and who

⁵Highest degree completed and graduation year effectively serve as a proxy for age.

are both working in the same industry and residing in the same state in 2000, individuals are as good as randomly assigned across CZs. My narrow group definition overcomes various concerns about selection into differentially shocked local labor markets. Workers of the same education level and cohort entered the labor force at the same time and faced a similar macroeconomic labor market. By isolating individuals in the same state in 2000, I account for the possibility that people in the same geographic region may have different location preferences that could impact key mobility outcomes. And although variation across CZs in their local Dot-Com unemployment shocks may reflect differential pre-recession exposure given unique local industry composition, I rule out industry-specific determinants of post-recession outcomes by isolating workers in the same initial industry.⁶ Empirically, I assuage concerns about selection by showing that in the dynamic econometric model from Equation 2.2, estimated coefficients on the local Dot-Com shocks in the pre-recession period are statistically indistinguishable from zero. This suggests that more negatively affected CZs were not already trending downward in various economic outcomes.

2.4 Results

Impact of Local Dot-Com Shocks on Employment

The Dot-Com recession had immediate negative impacts on employment. Figure 2.2 plots estimated effects on employment from Equation 2.2. Here, I define my outcome as “relative employment”, defined as an indicator for at least three months of employment in a given year minus the yearly average value of this indicator for the four pre-recession years from 1997 to 2000. Reassuringly, the estimated effects for all years in the pre-recession period are statistically indistinguishable from zero.⁷ However, in 2002, the year immediately following the recession, I estimate that a 1 percentage point increase in a CZ’s local unemployment shock decreased a worker’s probability of employment, relative to their pre-recession baseline, by more than 0.2 percentage points. In the next two years, the estimated effects correspond to a nearly 0.3 percentage point decrease in relative employment. Although employment exhibited a gradual recovery around 2005, there was still no full recovery by 2007. Even six years post-recession, I estimate that a 1 percentage point increase in a CZ’s local unemployment shock still made a worker over 0.1 percentage points less likely to be employed. This finding suggests that the Dot-Com recession had lasting effects on employment in the medium-run.

I extend the time horizon and plot estimated effects for up to 15 years post-recession in Figure B.1. In the long run, workers in initially more adversely affected CZs are no less likely to be employed, suggesting long-term recovery. Nevertheless, I interpret the long-run

⁶One possible concern is that the earnings premia associated with specific industries vary widely across cities. Card, Rothstein, and Yi (2025) find that location-specific earnings premia and industry earnings premia are largely additively separable, suggesting that any interactive effects between location and industry are unlikely to pose a selection threat.

⁷I cluster standard errors at the level of CZ in 2000 for all models.

impacts on employment with caution. The Great Recession began six years after the Dot-Com recession, and it is possible that CZs with larger Dot-Com shocks subsequently had larger local unemployment shocks from the Great Recession, which had even larger impacts on employment (Yagan 2019). Indeed, my estimated effect of a larger local Dot-Com shock on relative employment crawls back to near-zero by 2008, but drops somewhat sharply in 2009, which coincided with the end of the Great Recession.

I additionally show estimates from a dynamic model that leverages variation in local Dot-Com shocks within groups of individuals not just from the same initial industry, sex, highest degree level and graduation year but also from the same alma mater. This specification addresses potential selection concerns by controlling for differences in ability or specialized human capital across schools that could otherwise drive differences in employment outcomes. I plot estimates in Figure B.2 that are similar in magnitude to my baseline estimates that impose a same-state but not same-school restriction within groups.

I also consider a simpler definition of employment that does not adjust for pre-recession employment status. I re-estimate Equation 2.2 after defining a new employment outcome as an indicator for at least three months of employment in a given year. I plot estimates in Figure B.3 and find the same pattern of immediate negative impacts of local Dot-Com shocks on employment, with gradual recovery over the next few years.

Next, I examine the recession's impact on the characteristics of post-recession employment. I first study the impact on changes in company size of subsequent employers. Even if workers' employment eventually re-stabilized after the Dot-Com crash, it is plausible that they ended up at lower-paying firms. Given the canonical model of Burdett and Mortensen (1998) in which higher-paying firms are larger, I use a size-based measure to proxy for earnings. Concretely, I measure the change in the log number of employees at a worker's company in a given post-recession year, relative to the average log number of employees at their employers in the four years prior to the recession. I re-estimate Equation 2.2 to analyze dynamic effects on this company size measure. I again empirically validate the identifying assumption with estimated effects for the pre-recession period that are statistically indistinguishable from zero. For the post-recession period, I show in Figure 2.3 that workers who were exposed to larger local Dot-Com shocks systematically end up at smaller firms, with all estimated effects statistically significant at the 5% level. Thus workers in more adversely affected areas were not only less likely to be employed for several years afterwards, but were also more likely to end up at smaller firms, which tend to be lower-paying, conditional on employment.

Another possible margin of change in post-recession employment is whether workers who were exposed to larger unemployment shocks ended up in lower-paying industries. Using weighted averages of estimated establishment-level pay premia, as in Abowd, Kramarz, and Margolis (1999) and Card, Heining, and Kline (2013), within industries, Card, Rothstein, and Yi (2024) compute measures of earnings premia across 4-digit NAICS industries.⁸ I use

⁸These measures better capture industry quality than earnings alone, as they net out the contribution of sorting of workers of higher quality into higher-paying industries.

these premia, denoted as “industry effects”, to quantify the impact of Dot-Com shocks on the change in the associated industry effects between a worker’s post-recession industry and their pre-recession industry in 2000.⁹ I plot estimates from Equation 2.2 in Figure 2.4, and find no evidence that workers who were exposed to larger local unemployment shocks ended up in industries with smaller industry effects after the recession. One potential explanation is that industry switching is quite costly, and workers may be unwilling to leave industries in which they have accrued meaningful industry-specific human capital (Neal 1995; Parent 2000), even if it means going to smaller, potentially lower-paying firms.

Impact of Local Dot-Com Shocks on Geographic Mobility

Having shown strong negative impacts of the Dot-Com recession on employment and post-recession employer size as a proxy for earnings, I now turn to analyzing the recession’s impacts on geographic mobility. A core prediction of Blanchard and Katz (1992) is that regions with adverse unemployment shocks should see out-migration of workers. To test this prediction, I estimate the dynamic econometric specification in Equation 2.2 where I define the outcome as an indicator for whether a worker is observed in a different CZ from their CZ in 2000.¹⁰ In Figure 2.5, I plot estimated effects, along with 95% confidence intervals, for up to 15 years after the recession. While mobility is slow to adjust, I find that by 2005, a clear pattern emerges whereby workers in more adversely shocked CZs are more likely to move elsewhere. Specifically, a 1 percentage point increase in the local Dot-Com shock makes a worker 1 percentage point more likely to move elsewhere five years later.

As in my analysis of impacts on employment, to assuage additional concerns about selection, I estimate a dynamic model that leverages variation in local Dot-Com shocks within groups of individuals not just from the same initial industry, sex, highest degree level and graduation year but also from the same alma mater. Estimated effects on mobility in Figure B.4 are similar in magnitude to my baseline estimates that impose a same-state but not same-school restriction within groups.

A natural follow-up question is not just whether the recession impacted whether workers moved, but also *where* they moved. The directed migration of college-educated workers is highly responsive to local labor demand (Wozniak 2010). Thus one might expect that after the recession, workers in more adversely affected CZs would migrate to other CZs with greater economic opportunity. To capture economic opportunity in a CZ, I use estimated CZ-level earnings premia (denoted as “CZ effects”) from Card, Rothstein, and Yi (2025),

⁹One caveat is that the estimated industry effects come from data that spans 2010-2018, which occurs after the time period that I study. To the extent that inter-industry earnings differentials changed significantly between 2001 and 2010-2018, these estimated industry effects may not be an accurate measure of industry quality. Nevertheless, early evidence from Krueger and Summers (1988) from the late 20th century suggests these differentials are quite stable over time.

¹⁰I define the 2000 CZ based on where people work, and not where they reside. Data on location of residence is not available for that time frame.

which are constructed in a parallel way to the industry effects used previously.¹¹ I estimate the impact of the Dot-Com recession on moving to a higher-opportunity location using the static econometric specification in Equation 2.1, where the outcome is the change in CZ effects between a worker's CZ in 2006 or 2016 and their pre-recession CZ in 2000.

I present estimates in Table 2.2, with effects for 2006 in columns 1-3 and effects for 2016 in columns 4-6. All columns include group fixed effects, where groups isolate individuals with the same characteristics in 2000. My preferred estimates in columns 2 and 5 define a group as a set of workers in the same 4-digit NAICS industry, same state, and of the same sex, degree level and graduation year in 2000. I find that a 1 percentage point increase in the local Dot-Com shock led workers to locate in a CZ in 2006 with a 1.1% lower earnings, net of worker sorting, relative to their 2000 CZ.¹² This effect doubles when examining the change between 2000 and 2016. Together, these results suggest that workers in more adversely affected local labor markets moved to fundamentally lower-earnings CZs after the recession, and did not necessarily move to economic opportunity.

One possible explanation for this phenomenon is that since high-wage cities also tend to have higher cost of living (Moretti 2013), workers instead moved to cheaper cities after the recession. In Table 2.3, I show this is indeed the case. A 1 percentage point increase in the local Dot-Com unemployment shock led workers to, on average, live in a CZ in 2006 whose median rent was 3% cheaper than that of their CZ in 2000, and a CZ in 2016 whose median rent was 5% cheaper.¹³ Given the strong negative impacts on employment for several years following the recession, and the observation that workers moved to lower-earnings CZs, it is plausible that any ensuing earnings losses forced displaced workers to move out of higher-earnings but expensive CZs.

Impact of Local Dot-Com Shocks on Tech Workers

Given the Dot-Com bubble's origins in the tech sector, it is reasonable to speculate that the strong impacts of the recession on both employment and geographic mobility would be concentrated among workers in the tech sector. Table 2.4 presents estimated effects on individual employment relative to average employment status from 1997-2000, from Equation 2.1, separately for workers in the tech sector versus those in non-tech sectors in 2000.¹⁴ I

¹¹The CZ effects capture the average pay premium in log quarterly earnings (for 2010-2018, in dollars) at establishments in that CZ relative to pay at all nationwide establishments in the restaurant sector.

¹²The estimated coefficient is -1.1, and since the CZ effect is expressed in log quarterly earnings, the dependent variable is a change in log earnings, where earnings are net of worker sorting and capture the pure CZ earnings premium. Thus the percentage point change in the local Dot-Com shock corresponds to a percentage change in earnings.

¹³Data on county-level rents come from the 2000 decennial Census. I take a CZ-level weighted average, using the number of housing units in each county as weights, to construct a CZ-level measure of median rent in 2000 dollars.

¹⁴I define the tech sector as all industries within the 2-digit NAICS code 51 (information), all industries within the 3-digit NAICS code 334 (computer and electronics production), and NAICS code 5415 (computer systems design).

find that the strong overall negative effects of the recession on employment are concentrated among workers in the tech sector, as the estimated effect on relative employment in 2006 is over four times as large for tech workers as for non-tech workers, and statistically significant at the 5% level. Conditional on employment in 2006, tech workers in CZs with larger Dot-Com shocks were far likelier than their non-tech counterparts to end up in firms that were smaller than their initial pre-recession firms, as shown in Table 2.5. This discrepancy grew even larger 15 years post-recession in 2016. However, I do not find any evidence that workers in either the tech sector or non-tech sectors switched to industries with lower earnings premia, as shown in Table 2.6.

Next, I examine the impacts of local Dot-Com shocks on geographic mobility, separately for tech and non-tech workers. In Table 2.1, I find that in both 2006 and 2016, tech workers in CZs with larger local unemployment shocks in 2001 were slightly more likely than their non-tech counterparts to move out of their initial CZ in 2000. However, the estimated effects are not statistically significant at the 10% level and thus should be interpreted with caution. When looking at directed migration, Table 2.8 shows that estimated effects of the local Dot-Com shocks on the change in a worker's CZ earnings premia are larger for tech workers than for non-tech workers. A 1 percentage point larger local Dot-Com shock led tech workers to locate in CZs with, on average, 1.4% lower earnings in 2006 and 2.9% lower earnings in 2016, with both estimates statistically significant at the 1% level. By contrast, the same increase in the local Dot-Com shock led non-tech workers to locate in CZs with, on average, 1% lower earnings in 2006 and 2.0% lower earnings in 2016.

These findings of larger effects of the recession on tech workers are perhaps unsurprising given that the Dot-Com bubble originated in the tech sector. Interestingly, despite little evidence of strongly negative impacts on employment for non-tech workers, the Dot-Com recession still induced strong migration responses from non-tech workers in adversely affected areas. This finding underscores the importance of studying mobility responses to recessions, even if the underlying sources of the recession are more concentrated.

2.5 Conclusion

The Dot-Com recession of 2001 had negative impacts on the employment and geographic mobility of college-educated workers in both the medium run and long run. Workers in hard-hit areas were immediately less likely to be employed, with recovery delayed until over half a decade later. Conditional on re-employment, larger local unemployment shocks also induced individuals to end up at smaller firms, which tend to be lower-paying. Individuals did not, however, systematically switch into lower-paying industries. These impacts on employment were accompanied by strong mobility responses. In the long run, workers who were initially in local labor markets that were more adversely affected were more likely to move elsewhere. Rather than move to high-wage cities following these negative shocks, such workers were more likely to end up in cities with lower pay premia and lower costs of living. Consistent with the Dot-Com bubble's origins in the tech sector, the recession had the largest impacts,


both for employment and mobility, on workers in the tech sector. While the recession did not have quite the same impact on employment for non-tech workers, it still induced their migration towards lower-earnings local labor markets.

In a post-COVID environment with a difficult labor market for young college graduates and with layoffs in the tech sector, the Dot-Com recession provides an important historical precedent for how poor labor market conditions or a recession that particularly hits high-skill labor impacts outcomes in both the medium run and long run. While one might expect that college graduates with high levels of human capital may be insured against negative unemployment shocks, the Dot-Com recession highlights this may not necessarily be the case. Recovery in employment took years, and when additionally considering mobility, effects were more persistent for many years afterwards.

As the career trajectories of today's young college-educated workers unfold, there will be great interest in understanding how the present day labor market conditions impact them going forward. The rise of artificial intelligence (AI) calls into question whether structural changes in the economy will unfold in a manner that does not parallel the time period that followed the Dot-Com crash over two decades earlier. On one hand, automation has long been a concern well before the existence of AI, yet had thus far not seemed to dramatically reduce the availability of high-skill jobs (Autor 2015). On the other hand, there remain questions about what kinds of tasks AI will replace, with diverging implications for aggregate employment (Autor and Thompson 2025). Thus, while the Dot-Com crash provides many valuable insights into how recessions that are especially relevant for high-skill workers impact their labor market outcomes, there is scope for future research to deliver additional insights for a changing landscape.


2.6 Figures and Tables

Figure 2.1: Local Dot-Com Recession Shocks, by Commuting Zone




Notes: This map plots the Dot-Com recession shock for each commuting zone (CZ), calculated as the change (in percentage points) in the CZ's unemployment rate between 2000 Q4 and 2001 Q3. CZ-level unemployment rates are computed by aggregating monthly county-level counts of unemployed individuals and the size of the labor force as documented in the BLS LAUS series. For visual convenience, Alaska and Hawaii are omitted.

Figure 2.2: Impact of Dot-Com Local Shocks on Employment




Notes: This plot presents estimated effects of the Dot-Com Local Unemployment Shocks on relative employment, following Equation 2.2. The outcome is the difference between an individual's yearly employment status (indicator equal to 1 if the individual is employed for at least 3 months in a given year) and their average employment status from 1997-2000. The Dot-Com local unemployment shock is defined as the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's commuting zone (CZ) of residence in 2000, which I interact with dummy variables for each year (relative to 2001) equal to 1 if the present year is equivalent to the relative year. Years relative to 2001 (relative year 0) are marked on the horizontal axis. I plot 95% confidence intervals around each estimated effect, clustering standard errors at the level of CZ in 2000. Here, I restrict the sample of workers to include only those who are observed as employed in 1997.

Figure 2.3: Impact of Dot-Com Local Shocks on Size of Future Employers




Notes: This plot presents estimated effects of the Dot-Com Local Unemployment Shocks on switching to larger companies, following Equation 2.2. The outcome is the difference between an individual's current company's log number of employees and the yearly average log number of employees at their employers from 1997-2000. The Dot-Com local unemployment shock is defined as the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's commuting zone (CZ) of residence in 2000, which I interact with dummy variables for each year (relative to 2001) equal to 1 if the present year is equivalent to the relative year. Years relative to 2001 (relative year 0) are marked on the horizontal axis. I plot 95% confidence intervals around each estimated effect, clustering standard errors at the level of CZ in 2000.

Figure 2.4: Impact of Dot-Com Local Shocks on Change in Industry Earnings Premium Relative to Initial Industry



Notes: This plot presents estimated effects of the Dot-Com Local Unemployment Shocks on switching to higher or lower paying industries, following Equation 2.2. The outcome is the difference between an individual's current industry's earnings premium and the earnings premium at their initial industry in 2000. The industry earnings premia are from Card, Rothstein, and Yi (2024), and capture the premium in log quarterly earnings (for 2010-2018, in dollars) in a given industry (a 4-digit NAICS code) relative to the quarterly earnings in the restaurant sector. The Dot-Com local unemployment shock is defined as the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's commuting zone (CZ) of residence in 2000, which I interact with dummy variables for each year (relative to 2001) equal to 1 if the present year is equivalent to the relative year. Years relative to 2001 (relative year 0) are marked on the horizontal axis. I plot 95% confidence intervals around each estimated effect, clustering standard errors at the level of CZ in 2000.

Figure 2.5: Impact of Dot-Com Local Shocks on Geographic Mobility



Notes: This plot presents estimated effects of the Dot-Com Local Unemployment Shocks on geographic mobility, following Equation 2.2. The outcome is an indicator equal to 1 if, in a given year, an individual is observed in a different commuting zone (CZ) from their primary CZ in 2000. The Dot-Com local unemployment shock is defined as the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's CZ of residence in 2000, which I interact with dummy variables for each year (relative to 2001) equal to 1 if the present year is equivalent to the relative year. Years relative to 2001 (relative year 0) are marked on the horizontal axis. I plot 95% confidence intervals around each estimated effect, clustering standard errors at the level of CZ in 2000.

Table 2.1: Worker Summary Statistics

	mean	sd
<u>2006 Outcomes</u>		
Employed	0.97	0.16
Changed CZ	0.14	0.35
<u>2000 Characteristics</u>		
log Employer Size	7.69	2.57
Female	0.32	0.47
Bachelors	0.62	0.48
Masters	0.26	0.44
Doctorate	0.12	0.32
Years Since Graduation	12.12	8.17
Observations	1,120,582	

Notes: Table of worker summary statistics, with means and standard deviations. The top panel displays two primary outcomes in 2006, an indicator for whether a worker is employed in 2006 and an indicator for whether they are observed in a different commuting zone (CZ) from their primary CZ in 2000, just prior to the recession. In the bottom panel, I show summary statistics for individual characteristics in 2000. These include the log number of employees at their primary employer in 2000, female, three levels of highest degree (Bachelors, Masters, Doctorate), and years since graduation from their institution of highest degree.

Table 2.2: Impact of Dot-Com Local Shocks on Change in Commuting Zone Earnings Premium Relative to Initial Location

	Outcome: Difference Between 2006 or 2016 CZ Earnings Premium and 2000 CZ Earnings Premium					
	2006			2016		
	(1)	(2)	(3)	(4)	(5)	(6)
Dot-Com Shock	-0.552* (0.300)	-1.073*** (0.221)	-0.753* (0.387)	-1.073* (0.582)	-2.124*** (0.402)	-1.462** (0.744)
Observations	1,110,523	957,362	386,000	1,030,695	884,700	354,621
Industry x Sex x Degree x Cohort FE	X			X		
Industry x Sex x Degree x Cohort x State FE	X			X		
Industry x Sex x Degree x Cohort x School FE				X		

Notes: Estimates from Equation 2.1. Dependent variable is the difference between an individual's current commuting zone's (CZ) earnings premium and the earnings premium in their initial CZ in 2000. The CZ earnings premia are from Card, Rothstein, and Yi (2025), and capture the average pay premium in log quarterly earnings (for 2010-2018, in dollars) at establishments in that CZ relative to pay at all nationwide establishments in the restaurant sector. Current CZ is defined as the primary CZ of work in 2006 (columns 1-3) or 2016 (columns 4-6). The primary explanatory variable is the Dot-Com local shock, which is the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's CZ of residence in 2000. All models include group fixed effects. Groups in columns 1 and 4, denoted as "base groups", are defined as industry $\times$ sex $\times$ degree $\times$ cohort cells in 2000, where industry is the worker's 4-digit NAICS in 2000, degree is whether the highest degree attained is Bachelors, Masters or Doctorate, and cohort is the graduation year. Relative to base groups, groups in columns 2 and 5 further restrict to cells of individuals in the same state in 2000, while groups in columns 3 and 6 instead isolate individuals who went to the same college or university. Standard errors are clustered at the level of CZ in 2000. Significance levels: * 10%, ** 5%, and *** 1%.

Table 2.3: Impact of Dot-Com Local Shocks on Change in Commuting Zone Log Median Rent Relative to Initial Location

	Outcome: Difference in log(Median Rent) Between 2006 or 2016 CZ and Initial CZ in 2000					
	2006			2016		
	(1)	(2)	(3)	(4)	(5)	(6)
Dot-Com Shock	-1.557*** (0.580)	-2.703*** (0.458)	-2.229*** (0.794)	-3.146*** (1.125)	-5.418*** (0.812)	-4.373*** (1.504)
Observations	1,112,973	958,416	387,008	1,033,169	885,820	355,641
Industry x Sex x Degree x Cohort FE	X			X		
Industry x Sex x Degree x Cohort x State FE	X			X		
Industry x Sex x Degree x Cohort x School FE				X		

Notes: Estimates from Equation 2.1. Dependent variable is the difference between an individual's current commuting zone's (CZ) log median rent (measured in 2000) and the log median rent in their initial CZ in 2000. Rents are calculated as CZ-level averages of county median rents (weighted by number of housing units in each county) from the 2000 decennial Census. The primary explanatory variable is the Dot-Com local shock, which is the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's CZ of residence in 2000. All models include group fixed effects. Groups in columns 1 and 4, denoted as "base groups", are defined as industry $\times$ sex $\times$ degree $\times$ cohort cells in 2000, where industry is the worker's 4-digit NAICS in 2000, degree is whether the highest degree attained is Bachelors, Masters or Doctorate, and cohort is the graduation year. Relative to base groups, groups in columns 2 and 5 further restrict to cells of individuals in the same state in 2000, while groups in columns 3 and 6 instead isolate individuals who went to the same college or university. Standard errors are clustered at the level of CZ in 2000. Significance levels: * 10%, ** 5%, and *** 1%.

Table 2.4: Impact of Dot-Com Local Shocks on 2006 Employment Relative to Pre-2000 Average for Tech and Non-Tech Workers

	Outcome: 2006 Employment Relative to 1997-2000 Average	
	Tech	Non-Tech
	(1)	(2)
Dot-Com Shock	-0.614** (0.258)	-0.135 (0.167)
Observations	198,377	1,016,233
Industry x Sex x Degree x Cohort x State FE	X	X

Notes: Estimates from Equation 2.1, separately for Tech workers (NAICS codes 51XX, 334X and 5415) and Non-Tech workers. Dependent variable is the difference between an individual's employment status in 2006 (indicator equal to 1 if the individual is employed for at least 3 months) and their average employment status from 1997-2000. The primary explanatory variable is the Dot-Com local shock, which is the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's commuting zone (CZ) of residence in 2000. All models include group fixed effects, where groups are defined as industry $\times$ sex $\times$ degree $\times$ cohort $\times$ state cells in 2000. Here, industry is the worker's 4-digit NAICS in 2000, degree is whether the highest degree attained is Bachelors, Masters or Doctorate, and cohort is the graduation year. Standard errors are clustered at the level of CZ in 2000. Significance levels: * 10%, ** 5%, and *** 1%.

Table 2.5: Impact of Dot-Com Local Shocks on Change in Size of Future Employers for Tech and Non-Tech Workers

	Outcome: Change in log(Company Size) Between 2006 and 1997-2000 Average			
	Tech	Non-Tech	Tech	Non-Tech
	(1)	(2)	(3)	(4)
Dot-Com Shock	-5.694* (3.016)	-2.878*** (0.706)	-16.124*** (5.281)	-6.117*** (1.308)
Observations	174,308	893,114	162,762	813,464
Industry x Sex x Degree x Cohort x State FE	X	X	X	X

Notes: Estimates from Equation 2.1, separately for Tech workers (NAICS codes 51XX, 334X and 5415) and Non-Tech workers. Dependent variable is the difference between an individual's current company's log number of employees and the yearly average log number of employees at their employers from 1997-2000. The primary explanatory variable is the Dot-Com local shock, which is the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's commuting zone (CZ) of residence in 2000. All models include group fixed effects, where groups are defined as industry $\times$ sex $\times$ degree $\times$ cohort $\times$ state cells in 2000. Here, industry is the worker's 4-digit NAICS in 2000, degree is whether the highest degree attained is Bachelors, Masters or Doctorate, and cohort is the graduation year. Standard errors are clustered at the level of CZ in 2000. Significance levels: * 10%, ** 5%, and *** 1%.

Table 2.6: Impact of Dot-Com Local Shocks on Change in Industry Earnings Premium Relative to Initial Industry for Tech and Non-Tech Workers

	Outcome: Difference Between 2006 Industry Earnings Premium and 2000 Industry Earnings Premium	
	Tech	Non-Tech
	(1)	(2)
Dot-Com Shock	0.151 (0.226)	0.076 (0.046)
Observations	160,325	835,577
Industry x Sex x Degree x Cohort x State FE	X	X

Notes: Estimates from Equation 2.1, separately for Tech workers (NAICS codes 51XX, 334X and 5415) and Non-Tech workers. Dependent variable is the difference between an individual's 2006 industry's earnings premium and the earnings premium at their initial industry in 2000. The industry earnings premia are from Card, Rothstein, and Yi (2024), and capture the premium in log quarterly earnings (for 2010-2018, in dollars) in a given industry (a 4-digit NAICS code) relative to the quarterly earnings in the restaurant sector. The primary explanatory variable is the Dot-Com local shock, which is the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's commuting zone (CZ) of residence in 2000. All models include group fixed effects, where groups are defined as industry $\times$ sex $\times$ degree $\times$ cohort $\times$ state cells in 2000. Here, industry is the worker's 4-digit NAICS in 2000, degree is whether the highest degree attained is Bachelors, Masters or Doctorate, and cohort is the graduation year. Standard errors are clustered at the level of CZ in 2000. Significance levels: * 10%, ** 5%, and *** 1%.

Table 2.7: Impact of Dot-Com Local Shocks on Moving out of Initial Location for Tech and Non-Tech Workers

	Outcome: 1(Changed CZ)			
	2006		2016	
	Tech	Non-Tech	Tech	Non-Tech
	(1)	(2)	(3)	(4)
Dot-Com Shock	1.183 (0.960)	0.862 (0.583)	1.275 (1.915)	0.899 (0.973)
Observations	151,872	806,469	143,404	742,294
Industry x Sex x Degree x Cohort x State FE	X	X	X	X

Notes: Estimates from Equation 2.1, separately for Tech workers (NAICS codes 51XX, 334X and 5415) and Non-Tech workers. Dependent variable is an indicator equal to 1 if, in 2006 (columns 1-2) or 2016 (columns 3-4), an individual is observed in a different commuting zone (CZ) from their primary CZ in 2000. The primary explanatory variable is the Dot-Com local shock, which is the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's CZ of residence in 2000. All models include group fixed effects, where groups are defined as industry $\times$ sex $\times$ degree $\times$ cohort $\times$ state cells in 2000. Here, industry is the worker's 4-digit NAICS in 2000, degree is whether the highest degree attained is Bachelors, Masters or Doctorate, and cohort is the graduation year. Standard errors are clustered at the level of CZ in 2000. Significance levels: * 10%, ** 5%, and *** 1%.

Table 2.8: Impact of Dot-Com Local Shocks on Change in Commuting Zone Earnings Premium Relative to Initial Location for Tech and Non-Tech Workers

	Outcome: Difference Between 2006 or 2016 CZ Earnings Premium and 2000 CZ Earnings Premium			
	2006		2016	
	Tech	Non-Tech	Tech	Non-Tech
	(1)	(2)	(3)	(4)
Dot-Com Shock	-1.437*** (0.495)	-1.002*** (0.186)	-2.915*** (0.928)	-1.967*** (0.335)
Observations	151,771	805,484	143,294	741,242
Industry x Sex x Degree x Cohort x State FE	X	X	X	X

Notes: Estimates from Equation 2.1, separately for Tech workers (NAICS codes 51XX, 334X and 5415) and Non-Tech workers. Dependent variable is the difference between an individual's current commuting zone's (CZ) earnings premium and the earnings premium in their initial CZ in 2000. The CZ earnings premia are from Card, Rothstein, and Yi (2025), and capture the average pay premium in log quarterly earnings (for 2010-2018, in dollars) at establishments in that CZ relative to pay at all nationwide establishments in the restaurant sector. Current CZ is defined as the primary CZ of work in 2006 (columns 1-2) or 2016 (columns 3-4). The primary explanatory variable is the Dot-Com local shock, which is the difference in unemployment rate between Q4 2000 and Q3 2001 in the individual's CZ of residence in 2000. All models include group fixed effects, where groups are defined as industry $\times$ sex $\times$ degree $\times$ cohort $\times$ state cells in 2000. Here, industry is the worker's 4-digit NAICS in 2000, degree is whether the highest degree attained is Bachelors, Masters or Doctorate, and cohort is the graduation year. Standard errors are clustered at the level of CZ in 2000. Significance levels: * 10%, ** 5%, and *** 1%.

Chapter 3

The Labor Market Impacts of 21st Century STEM Expansions in US Higher Education

3.1 Introduction

The dramatic shift towards STEM (science, technology, engineering and mathematics) fields in higher education during the 21st century has significantly changed the composition of the global high-skill workforce. In the US in particular, the number of degrees in STEM fields awarded at postsecondary institutions has more than doubled since 2000.¹ Policy-driven expansions in the production of STEM degrees at these institutions have spurred much of this sharp increase. Yet there is limited evidence on how these systematic nationwide expansions in STEM degree production impacted both students' post-graduation labor market outcomes and the ensuing implications for retention of STEM graduates across local labor markets.

While there exists substantial evidence on the positive impacts of STEM major choice on post-graduation outcomes such as earnings (Altonji, Blom, and Meghir 2012; Kirkeboen, Leuven, and Mogstad 2016; Deming and Noray 2020), this paper quantifies the broad labor market impacts of widespread changes in STEM major completions, specifically in Computer Science (CS) and Engineering, across institutions. Beyond these expansions' effects on individual post-graduation outcomes, their effects on local labor market outcomes are of significant economic and policy relevance. Given positive externalities from a large local college-educated workforce (Moretti 2004a; Shapiro 2006), and from STEM graduates in particular (Winters 2014), as well as longstanding policy concerns about brain drain from local regions (Gottlieb 2004), I investigate impacts on geographic mobility to probe the extent to which schools' expansions in CS and Engineering degree production could have increased the concentration of STEM graduates in the surrounding area.

¹Source: https://nces.ed.gov/programs/digest/d23/tables/dt23_322.10.asp

I combine information on completed degrees in CS and Engineering throughout the 21st century with detailed data on education and employment histories for millions of college-educated workers in the US. I identify an expansion event as a sharp single-year increase in degree production in these fields at a single institution, and quantify the widespread effects of these expansions on a wide range of individual outcomes, including advanced degree attainment, future employment in high-technology industries, future employer size, and geographic mobility. To identify these effects, I compare the post-graduation outcomes of individuals at the same school and in neighboring cohorts, where the older cohort had already declared their majors by the time the expansion event began but the younger cohort entered school right as the expansion began. This approach minimizes possible confounding impacts of differences across schools on future outcomes that could otherwise be related to the decision to expand CS and Engineering degrees, as well as possible selection of different types of students into these schools. By comparing neighboring cohorts, I also mitigate the potential impact of graduating into different macroeconomic job search environments on future outcomes.

I first document that CS and Engineering degree expansions positively increased the individual probability of obtaining a CS or Engineering undergraduate degree. Next, I find that the expansions had significant impacts on various post-graduation outcomes. Students in cohorts that were exposed to an expansion event were less likely to obtain an advanced degree than those in slightly older cohorts at the same school who were not exposed. They were also more likely, both immediately after graduation and five years out, to work in high-technology industries that are immediately adjacent to CS and Engineering and to work for smaller companies. Finally, exposed cohorts were less likely to leave the local labor market in which they went to school in the short run, but were no less likely to leave after five years.

Together, these results suggest that widespread expansions in these STEM fields yielded potential benefits for both students and, at least in the short run, the local areas surrounding the institutions where expansions took place. Relative to advanced degrees in law or medicine, advanced degrees in CS and Engineering offer a lower return in future earnings (Altonji, Humphries, and Zhong 2023). Thus the lower rates of advanced degree attainment among exposed cohorts does not necessarily point to a deleterious effect, especially since students instead shifted to employment in higher-paying, high-technology industries upon graduation. My finding that individuals in cohorts exposed to expansions were more likely to stay in the local area upon graduation also suggests that these institution-level policies could successfully align, within a short time horizon, with the interests of local policymakers. Despite potential concerns that the production of more STEM graduates would not yield any meaningful change in local retention of college graduates if students could simply move after graduation, in this context, higher production directly led to higher immediate retention. Nevertheless, these effects appear to be limited to a narrow window of time.

This paper contributes to a large literature on the impacts of STEM major choice on post-graduation labor market outcomes. Studies consistently find that students who choose STEM majors earn more, on average, than their non-STEM counterparts (Altonji, Blom, and Meghir 2012; Deming and Noray 2020), and that this gap does not merely reflect the selection of individuals with higher earnings ability into these majors (Kirkeboen, Leuven,

and Mogstad 2016). While this evidence highlights that choosing a STEM field can significantly affect a college student’s post-graduation career, there is much less evidence on how the access to such majors at the institutional level impacts the same outcomes. In one study, Bleemer and Mehta (2022) analyze the impact on earnings of access to the Economics major at a single institution. I expand this scope by analyzing how the expansion of STEM access across a wide range of institutions affects students’ future employment outcomes.

This paper also relates to a literature on the labor market impacts of expansions in post-secondary education. Most studies have focused on expansions via new college openings, rather than expansions in fields within existing schools. Studies on mass openings of new schools in France during the 1990s find that the expansions had strong positive impacts on individuals’ employment and wages (Ferhat 2022; Nimier-David 2023). Parallel evidence for the US is less prevalent given the relative lack of new college openings over the last few decades. I complement these studies by analyzing the individual-level labor market impacts of field-specific expansions across schools in the US. Closely related work in Switzerland finds that the establishment of new STEM-specific schools increased job opportunities for high-skill workers (Schultheiss, Backes-Gellner, and Pfister 2023). I widen this scope to include expansions in multiple STEM fields across a wider range of schools in a much larger setting.

Finally, I contribute to a policy-relevant literature that analyzes the relationship between a local area’s production of new college graduates and its overall concentration of college-educated workers. Local policymakers often express concerns about brain drain (Gottlieb 2004), prompting questions about whether colleges and universities can directly increase a local area’s stock of college graduates simply by producing more of them each year. Early evidence from Bound et al. (2004) suggests only a modest positive relationship between a local area’s production of new college graduates and its stock of all college graduates, with later evidence from Abel and Deitz (2012) corroborating this finding. I complement this analysis by studying the impact of STEM expansions on geographic mobility. Some descriptive evidence from England suggests that excessive STEM specialization harms universities’ abilities to retain graduates locally (Kitagawa et al. 2022). I focus on the US and on the causal impact of university-level policies.

The rest of the paper proceeds as follows. Section 3.2 describes the data sources I use to identify expansions in CS and Engineering and trace individuals’ educational histories and post-graduation employment. Section 3.3 describes the empirical strategy I use to identify the labor market impacts of these expansions. I present results in Section 3.4. Section 3.5 concludes.

3.2 Data

I use data on observed changes in major-specific completions at various schools from 2000-2019 to identify schoolwide expansions in CS and Engineering degree production. To identify college graduates who were potentially exposed to these expansions and their ensuing labor market outcomes, I use rich individual-level data on educational histories and post-

graduation employment histories. The data come from the company People Data Labs, which collects and harmonizes information from various private and public sources such as LinkedIn profiles. I now discuss the data sources in more detail.

STEM Expansions

To identify potential CS and Engineering expansions, I use data from the National Center for Education Statistics' Integrated Postsecondary Education Data System (IPEDS), which reports the yearly number of completed degrees for each major at nearly all 4-year colleges and universities in the US. For each field and institution, I identify an expansion event as a sharp single-year increase in the number of completed degrees. In order to classify an event, I require that both the number of degrees in the relevant field increased by 20 and that the percentage increase exceeds 30%. I allow the same institution to have multiple events, but exclude any events that are less than 10 years apart, in order to only keep cohorts who were exposed to either a single event or none at all.

In total, I document 597 expansions across CS and Engineering that fit such criteria. I display the largest events in Table C.1. Because I observe completed degrees and not enrollments, for any event I count the first treated cohort as the cohort entering four years prior to the year in which the spike in completed degrees is observed. For instance, if I observe a 50% jump and increase of 30 degrees in CS at a particular institution from 2016-2017, then the first treated cohort is the one entering the school in 2013. The intuition is that on average, the students from this cohort graduate in 2017, in line with the observed spike in completed degrees. In my empirical strategy, I compare the outcomes of these students to those who have just completed major declarations, typically by the end of sophomore year, when the first treated cohort enters.² I discuss this approach in more detail in Section 3.3.

Educational Histories and Post-Graduation Outcomes

To identify workers who attended schools that implemented CS and Engineering expansions during the relevant time frame, I use data on self-reported educational histories linked to job spells for hundreds of millions of individuals. From educational histories, for a given individual I know their institution, degree level and field of study. I also know the start and end dates of each education spell. Thus based on the timing of these spells and the timing of expansions that are unique to institutions and fields, I identify individuals within the same institution who were and were not affected by such expansions.

To study these individuals' post-graduation outcomes, I use data on their post-graduation job spells. For each spell, I observe the position start date and end date, using the start date to isolate spells that began after college graduation. I also observe job characteristics such as firm characteristics, including the number of employees and industry; and granular worker locations that I map to the commuting zone (CZ) level. I identify 2.4 million workers who

²While major declaration timelines are not universal, declaration by the end of sophomore year is common across schools.

attended a school with an expansion event and whose entry cohort was within two years on either side of the event. For instance, if a documented spike in CS degrees occurred in 2017, I keep the first treated cohort entering in 2013, as well as cohorts entering in 2011, 2012, and 2014.

One limitation of this data source is that individuals who do not have an updated online presence, e.g. LinkedIn, are unlikely to appear in the data. Given that LinkedIn usage is higher among younger college graduates in later cohorts, this limitation in itself is unlikely to yield a massive loss in coverage. Another concern is that the data source is over-represented in more technical fields, i.e. STEM fields, and other fields that require advanced degrees such as law or medicine. By comparison, workers who majored in the humanities or some social sciences are under-represented. To the extent that individuals are more likely to substitute between STEM and STEM-adjacent fields instead of between STEM and humanities fields, my sample may consist of more compliers who could be induced into majoring in CS or Engineering after an expansion, as opposed to students who would not choose a CS or Engineering major regardless of whether or not an expansion occurred. In this respect, my estimates could be closer to upper bounds.

3.3 Empirical Strategy

I quantify the impact of CS and Engineering expansions on individual labor market outcomes by comparing the outcomes of students at the same school who are in neighboring cohorts, where one cohort is immediately affected by an expansion event whereas the other declares majors just before the expansion takes place. For my running example of the school with a documented spike in CS degrees in 2017, the first treated cohort would enter in 2013, with the majority graduating in 2017. Meanwhile, the cohort entering in 2011 would have largely declared their majors before the CS expansion took place.

I implement this design by estimating the following econometric model, which I define as the “baseline” specification:

$$y_{i,post} = \beta D_{sc(i)} + X_i' \gamma + \theta_s + \varepsilon_i \quad (3.1)$$

where $y_{i,post}$ denote post-graduation labor market outcomes for individual i , with $post$ either capturing the first job post-graduation or five years post-graduation. The primary coefficient of interest is β , associated with the explanatory variable $D_{sc(i)}$ — an indicator equal to 1 if i belongs to the first cohort c that experienced a CS or Engineering expansion event at their school, where a school-expansion event pair is denoted as s . I include individual controls X_i' , which in practice is a single dummy variable for whether the individual is female. To identify β , I include fixed effects for each school-expansion event pair, θ_s , and estimate Equation 3.1 only on individuals belonging to either the first treated cohort, i.e. completion event year minus four, or the last non-treated cohort, i.e. completion event year minus six for their school’s expansion event. This restriction ensures that I compare the outcomes of students at the same school who are only two years apart, where the first treated cohort sees a spike in

CS or Engineering completions upon their graduation while the last non-treated cohort are instead juniors who have largely declared their majors just prior to the first treated cohort entering school.

I complement this specification with a model that allows for both partial and full exposure to a CS or Engineering expansion event at a given school. I denote this as the “exposure” specification. I expand the sample to include not only the first treated cohort and the last treated cohort, but also the second treated cohort (entering one year after the first) and the partially treated cohort that immediately precedes the first treated cohort. The partially treated cohort enters school five years before the documented spike in CS or Engineering completions, and thus are sophomores who are still able to declare their major when the school’s expansionary policies begin. Concretely, I estimate:

$$y_{i,post} = \beta_{partial} \mathbf{1}\{c(i) = y(s) - 1\} + \beta_{full} \mathbf{1}\{c(i) \geq y(s)\} + X_i' \gamma + \theta_s + \varepsilon_i \quad (3.2)$$

Here, $\mathbf{1}\{c(i) = y(s) - 1\}$ is an indicator for whether individual i ’s cohort year $c(i)$ is in the cohort that precedes the first year $y(s)$ of the CS or Engineering expansion at school-event s , with the sharp increase in completed degrees observed four years after $y(s)$. This cohort is the partially treated cohort, for which my estimate of $\beta_{partial}$ captures the impact of the expansion for cohorts who were first exposed as sophomores. By comparison, β_{full} captures the impact on the outcomes of fully treated cohorts, i.e. $c(i) \geq y(s)$, who entered school in or the year after the expansion.

Identifying β , $\beta_{partial}$ and β_{full} requires that within the same school and time period, students who are exposed to CS and Engineering expansions do not systematically differ in unobserved ways from those who are not. My empirical approach overcomes two key identification concerns. First, schools differ systematically on a variety of observed and unobservable characteristics, such as average student ability, program quality, and proclivity to select into specific majors. Exploiting variation in expansions across schools can overstate their impacts if schools that expanded their CS or Engineering programs attract students who were more likely to choose those majors to begin with or to have different underlying propensities to choose high-tech industries or employment at smaller firms or another labor market outcome of interest. By restricting any comparison of outcomes to students within the same school and from adjacent cohorts, I isolate students who are on average highly similar on observed dimensions, including standardized test scores, family background, etc., and plausibly similar on unobserved dimensions.³

Second, even within the same school, graduates from different cohorts vary in the macroeconomic labor market conditions that they face upon graduation. These conditions can directly influence major choice (Blom, Cadena, and Keys 2021), such that pre-expansion cohorts could have been systematically more or less likely than post-expansion cohorts to have chosen CS or Engineering majors even in the absence of any expansion. Differences

³While it is theoretically possible for two adjacent cohorts at the same school to be different on a wider range of characteristics, this is quite rare in practice. Schools do not typically experience sudden changes in quality, resources or admissions criteria that are sharp enough to affect the entire student body.

in these conditions can also affect different cohorts' post-graduation labor market outcomes themselves. By isolating neighboring cohorts that are only two years apart, I avoid comparing students who may have faced different macroeconomic environments upon graduation or during school. Since all expansion events that I study occurred before 2019, I eliminate a sudden change in the broader labor market such as COVID that otherwise caused even adjacent cohorts to graduate into very different job search environments. To overcome concerns about any rapid changes in macroeconomic conditions from the Great Recession, I repeat my analysis on a subsample that excludes Great Recession cohorts from 2007-2009.

3.4 Results

The widespread CS and Engineering expansions across US colleges and universities in the 21st century had immediate impacts on the individual likelihood of completing a CS or Engineering degree. In Table 3.1, I separately estimate the impact of exposure to a CS expansion on CS degree attainment, Engineering expansion on Engineering degree attainment, and a concurrent expansion in both fields on obtaining a degree in either field. By including fixed effects for a school-event pair, I compare individuals in first treated cohorts to individuals at the same school who entered two years prior. I show that exposure to a CS expansion increased the probability of completing a CS degree by 0.7 percentage points, exposure to an Engineering expansion increased the probability of an Engineering degree by 0.18 percentage points, and exposure to a joint expansion increased the probability of a CS or Engineering degree by 0.22 percentage points. I complement this analysis by estimating Equation 3.2 and quantifying the impacts of partial exposure, i.e. sophomore at the time of expansion versus full exposure, i.e. freshman during or after expansion. In Figure 3.1, I show that for each of CS, Engineering, and joint expansions, the probability of completing a degree in that field for a partially exposed cohort increased relative to an untreated cohort. But for a fully treated cohort, the estimated increase in probability is significantly larger. Altogether, I confirm that the CS and Engineering expansions indeed increased individual probabilities of completing degrees in those fields.

Next, I quantify the impacts of the expansions on post-graduation labor market outcomes. I present estimated effects from Equation 3.1 on various outcomes associated with individuals' first post-graduation job in Table 3.2. First, I show that relative to cohorts within the same school who narrowly missed the expansion, individuals in the first cohort under a CS or Engineering expansion are on average, 1.4 percentage points less likely to obtain a Masters or Doctorate degree (p-value < 0.01). This effect does not necessarily indicate a reduction in earnings. In fact, such individuals are more likely to have their first post-graduation job in a high-technology industry.⁴ These are industries that are not only adjacent to CS and Engineering but also have higher relative earnings premia, i.e. average pay premium of firms in the industry independent of worker ability (Card, Rothstein, and Yi 2024). I also find

⁴The high-tech industries are NAICS 51XX– Information; NAICS 334X– Computer and Electronic Manufacturing; NAICS 5413– Engineering; and NAICS 5415– Computer Systems Design

that cohorts exposed to CS or Engineering expansions at their alma maters tend to end up in smaller firms, on average, which is broadly consistent with the prevalence of start ups in tech relative to other industries. This finding of employment in high-technology industries and smaller firms holds when looking at individuals' jobs five years post-graduation (Table 3.3). In fact, individuals in cohorts exposed to a CS or Engineering expansion are even more likely to be in a high-technology industry five years post-graduation than immediately upon graduation.

Across these outcomes, I then additionally compare individuals who were exposed to CS or Engineering expansions as sophomores, representing partially treated cohorts, to untreated cohorts the year prior. I present estimated effects on first-job outcomes for partial treatment and full treatment from Equation 3.2 in Figure 3.2. Unsurprisingly, the magnitudes of the estimated effects on fully treated cohorts is larger than those for partially treated cohorts. Nevertheless, I find that even for partially treated cohorts, individuals are less likely to obtain an advanced degree and more likely to work for smaller companies. I do not find a statistically significant effect on employment in a high-technology industry for partially treated cohorts. In Figure 3.3, I repeat the analysis for outcomes five years post-graduation. The estimated effect of partial exposure on employment in a high-technology industry is positive, and the effect on future employer size is negative, as was the case for first-job outcomes, but the effects are not statistically significant at the 5% level. Effects for fully treated cohorts remain non-zero and statistically significant.

One striking result emerges when looking at impacts on geographic mobility—students who are exposed to CS and Engineering expansions are less likely to leave the CZ of their schools upon graduation (Table 3.2). Relative to their older, unexposed counterparts, these students are on average 1.1 percentage points less likely to move (p -value < 0.01). In the context of a broad literature that highlights the potential inefficacy of expanding the local production of college graduates in order to retain a larger local college-educated workforce, this finding suggests that boosting the production of this subset of STEM graduates can help retain them, at least in the short term. After a few years, however, the estimated effect of the expansions on individuals' probabilities of moving drops to near-zero. This result is perhaps not too surprising given that college-educated workers are quite willing to move to capitalize on favorable job opportunities (Wozniak 2010), which may be more common as these students build their careers over the span of a few years. The lack of persistence in reduced out-mobility does not necessarily mean that the local labor markets in which expanding schools are located do not benefit, via human capital externalities, from an immediate increase larger local college-educated workforce. However, agglomeration economies may be less likely to form if workers do not remain at local firms beyond a few years (Glaeser 2010).

Finally, I show that my results are not sensitive to the inclusion of Great Recession cohorts (graduating in 2007-2009) or to using stricter criteria to define CS and Engineering expansions. The Great Recession represented a short period during the 21st century in which macroeconomic conditions significantly worsened suddenly. Thus even within a narrow range of adjacent cohorts, there is a concern that treated and untreated cohorts around the same expansion event faced significantly different job search environments. I exclude Great

Recession cohorts and re-estimate my baseline specification in Equation 3.1, reporting results in Table C.2 for first-job outcomes and Table C.3 for outcomes five years post-graduation. I repeat this sample restriction for the exposure specification in Equation 3.2, reporting results in Table C.4 and Table C.5. For both models, I find little change in my estimated effects.

I also consider an even stricter definition for a CS or Engineering expansion. While my baseline definition isolates single-year spikes in completed degrees rather than multi-year, continuous increases, I identify a subset of events with even larger jumps in completions. Concretely, I require that the single-year increase in completed degrees in CS or Engineering exceeds 50% and 30 additional degrees. After restricting my sample to students in cohorts at schools with these larger expansions around the time of expansion, I re-estimate my baseline specification and report results in Table C.6 and Table C.7. I also re-estimate my exposure specification and report results in Table C.8 and Table C.9. I once again find little change in my estimated effects.

3.5 Conclusion

Throughout the 21st century, many colleges and universities in the US have made active efforts to boost their production of graduates in STEM fields. In Computer Science and Engineering alone, several hundred schools experienced dramatic single-year surges in their respective numbers of new graduates. I show that these expansions had meaningful labor market impacts for students at these institutions both immediately upon graduation and five years afterwards. Relative to students who completed major declarations just prior to a CS or Engineering expansion at their alma maters, students who entered the same college during the expansion were less likely to obtain advanced degrees but more likely to complete a CS or Engineering degree and find employment in a high-paying, high-technology industry. This estimated impact on employment in high-technology sectors persists five years after graduation. Students impacted by the expansions also were employed in smaller companies, on average. Given a lack of data on individual earnings, I cannot directly measure the expansions' effects on post-graduation earnings, but this analysis would be an interesting and important direction for future research.


I also find that widespread CS and Engineering expansions significantly impacted geographic mobility in the short run. Students who entered college during a CS or Engineering expansion at their alma maters were more likely to remain in the local area immediately after graduation. In light of research on the positive externalities from a larger local college-educated workforce and local policy interests across regions in retaining college graduates, these positive effects suggest that these institution-level policies could benefit entire local labor markets in the short run. Several years after graduation, however, students in cohorts who were exposed to CS or Engineering expansions were no longer significantly more likely to still be working in the same area in which they went to school. Thus any long run effects of these expansions on the local retention of high-skill labor are likely negligible.

Looking ahead, labor demand for STEM talent projects to continue increasing over the next decade.⁵ As schools continue to expand their capacity to produce graduates in STEM fields, even with the creation of new majors in recent years such as Data Science, the evidence from previous expansions in the 21st century suggest that these policies may strongly affect labor market outcomes of students both immediately upon graduation and for several years afterwards. For local policymakers concerned with combating brain drain and creating a larger local STEM workforce, there is insufficient evidence to point to any long run impacts of such policies. And in the present day where remote work has risen substantially and allowed individuals to untie their place of residence from their place of work, the ensuing reduction in barriers to geographic mobility will likely continue to reinforce that conclusion.

⁵Source: <https://nces.nsf.gov/pubs/nsb20261>


3.6 Figures and Tables

Figure 3.1: Impact of Partial and Full Exposure to CS and Engineering Expansions on Completion of CS and Engineering Degrees




Notes: Estimates from Equation 3.2. Dependent variable is an indicator for whether an individual completed a computer science (CS) major (panel a), an engineering major (panel b), or any of the two majors (panel c). Partial exposure is an indicator for whether the individual's cohort of entry to their bachelor's institution is five years prior to an observed sharp increase in completed degrees in CS or Engineering at the institution. Full exposure is an indicator for being in the cohort of entry three or four years prior to this same sharp increase. A sharp increase is defined as both a 30% or more increase and a numerical increase of 20 such degrees from the year prior. Panel a includes only expansions in CS, whereas panel b includes only expansions in Engineering. Panel c includes only concurrent expansions in both fields at the same institution. All models include a dummy variable for female, and fixed effects for school-expansion pairs. I plot 95% confidence intervals, with standard errors clustered at the school level.

Figure 3.2: Impact of Partial and Full Exposure to CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation



Notes: Estimates from Equation 3.2. Dependent variables are: an indicator for whether an individual went on to complete a Masters or Doctorate degree (panel a); in panel b, an indicator for whether their first job was in a high-tech industry (NAICS 51XX– Information; NAICS 334X– Computer and Electronic Manufacturing; NAICS 5413– Engineering; NAICS 5415– Computer Systems Design); log number of employees at first employer post-graduation (panel c); and an indicator for whether, for their first job post-graduation, the individual relocated to a different commuting zone (CZ) from the CZ of their school (panel d). Partial exposure is an indicator for whether the individual’s cohort of entry to their bachelor’s institution is five years prior to an observed sharp increase in completed degrees in CS or Engineering at the institution. Full exposure is an indicator for being in the cohort of entry three or four years prior to this same sharp increase. A sharp increase is defined as both a 30% or more increase and a numerical increase of 20 such degrees from the year prior. All models include a dummy variable for female, and fixed effects for school-expansion pairs. I plot 95% confidence intervals, with standard errors clustered at the school level.

Figure 3.3: Impact of Partial and Full Exposure to CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation



Notes: Estimates from Equation 3.2. Table notes are the same as for Figure 3.2 except that outcomes are based on an individual's job five years post-graduation instead of their first job post-graduation.

Table 3.1: Impact of CS and Engineering Expansions on Major Completions

	Outcome: 1(CS/Engineering Major)		
	CS	Engineering	Pooled
	(1)	(2)	(3)
Post-Expansion Cohort	0.007*** (0.001)	0.018*** (0.002)	0.022*** (0.006)
Observations	708,285	221,176	30,919
Individual Controls	X	X	X
School x Event FE	X	X	X

Notes: Estimates from Equation 3.1. Dependent variable is an indicator for whether an individual completed a computer science (CS) major (column 1), an engineering major (column 2), or any of the two majors (column 3). Post-Expansion Cohort is an indicator for whether the individual's cohort of entry to their bachelor's institution is four years prior to an observed sharp increase in completed degrees in CS or Engineering at the institution. A sharp increase is defined as both a 30% or more increase and a numerical increase of 20 such degrees from the year prior. Column 1 includes only expansions in CS, whereas column 2 includes only expansions in Engineering. Column 3 includes only concurrent expansions in both fields at the same institution. All models include a dummy variable for female, and fixed effects for school-expansion pairs. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table 3.2: Impact of CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation

	Outcome			
	Advanced Degree	High-Tech Industry	Employer Size	Moved
	(1)	(2)	(3)	(4)
Post-Expansion Cohort	-0.014*** (0.001)	0.001* (0.001)	-0.044*** (0.007)	-0.011*** (0.002)
Observations	917,631	917,631	758,610	500,844
Individual Controls	X	X	X	X
School x Event FE	X	X	X	X

Notes: Estimates from Equation 3.1. Dependent variables are: an indicator for whether an individual went on to complete a Masters or Doctorate degree (column 1); in column 2, an indicator for whether their first job was in a high-tech industry (NAICS 51XX– Information; NAICS 334X– Computer and Electronic Manufacturing; NAICS 5413– Engineering; NAICS 5415– Computer Systems Design); log number of employees at first employer post-graduation (column 3); and an indicator for whether, for their first job post-graduation, the individual relocated to a different commuting zone (CZ) from the CZ of their school (column 4). Post-Expansion Cohort is an indicator for whether the individual’s cohort of entry to their bachelor’s institution is four years prior to an observed sharp increase in completed degrees in CS or Engineering at the institution. A sharp increase is defined as both a 30% or more increase and a numerical increase of 20 such degrees from the year prior. All models include a dummy variable for female, and fixed effects for school-expansion pairs. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table 3.3: Impact of CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation

	Outcome		
	High-Tech Industry	Employer Size	Moved
	(1)	(2)	(3)
Post-Expansion Cohort	0.002*** (0.001)	-0.013** (0.006)	-0.002 (0.001)
Observations	917,631	781,103	532,116
Individual Controls	X	X	X
School x Event FE	X	X	X

Notes: Estimates from Equation 3.1. Table notes are the same as for Table 3.2, with the difference that outcomes are now based on an individual's job five years post-graduation instead of their first job post-graduation. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.

Bibliography

- Abel, Jaison R., and Richard Deitz. 2012. “Do Colleges and Universities Increase Their Region’s Human Capital?” *Journal of Economic Geography* 12 (3): 667–691.
- Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. “High Wage Workers and High Wage Firms.” *Econometrica* 67 (2): 251–333.
- Altonji, J. G., and C. R. Pierret. 2001. “Employer learning and statistical discrimination.” *The quarterly journal of economics* 116 (1): 313–350.
- Altonji, Joseph G., Erica Blom, and Costas Meghir. 2012. “Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers.” *Annual Review of Economics* 4 (1): 185–223.
- Altonji, Joseph G., John Eric Humphries, and Ling Zhong. 2023. “The Effects of Advanced Degrees on the Wage Rates, Hours, Earnings, and Job Satisfaction of Women and Men.” In *50th Celebratory Volume*, 50:25–81. Research in Labor Economics. Emerald Publishing Limited.
- Angrist, J. D., and J. S. Pischke. 2009. *Mostly harmless econometrics: An empiricist’s companion*. Princeton university press.
- Autor, David, and Neil Thompson. 2025. “Expertise.” *Journal of the European Economic Association* 23 (4): 1203–1271.
- Autor, David H. 2015. “Why Are There Still So Many Jobs? The History and Future of Workplace Automation.” *Journal of Economic Perspectives* 29 (3): 3–30.
- Autor, David H., David Dorn, and Gordon H. Hanson. 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103 (6): 2121–2168.
- Avery, C., J. S. Howell, and L. Page. 2014. *A Review of the Role of College Applications on Students’ Postsecondary Outcomes*. Research Brief. College Board.
- Bayer, P., S. L. Ross, and G. Topa. 2008. “Place of work and place of residence: Informal hiring networks and labor market outcomes.” *Journal of Political Economy* 116 (6): 1150–1196.

- Beaman, L. A. 2012. "Social networks and the dynamics of labour market outcomes: Evidence from refugees resettled in the US." *The Review of Economic Studies* 79 (1): 128–161.
- Blackwell, M., S. Iacus, G. King, and G. Porro. 2009. "cem: Coarsened exact matching in Stata." *The Stata Journal* 9 (4): 524–546.
- Blanchard, Olivier Jean, and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity* 1992 (1): 1–75.
- Bleemer, Zachary, and Aashish Mehta. 2022. "Will Studying Economics Make You Rich? A Regression Discontinuity Analysis of the Returns to College Major." *American Economic Journal: Applied Economics* 14 (2): 1–22.
- Blom, Erica, Brian C. Cadena, and Benjamin J. Keys. 2021. "Investment over the Business Cycle: Insights from College Major Choice." *Journal of Labor Economics* 39 (4): 1043–1082.
- Blumenstock, J. E., G. Chi, and X. Tan. 2025. "Migration and the value of social networks." *Review of Economic Studies* 92 (1): 97–128.
- Bound, John, Jeffrey A. Groen, Gabor Kezdi, and Sarah Turner. 2004. "Trade in University Training: Cross-State Variation in the Production and Stock of College-Educated Labor." *Journal of Econometrics* 121 (1-2): 143–173.
- Büchel, K., M. V. Ehrlich, D. Puga, and E. Viladecans-Marsal. 2020. "Calling from the outside: The role of networks in residential mobility." *Journal of urban economics* 119:103277.
- Buggle, J., T. Mayer, S. O. Sakalli, and M. Thoenig. 2023. "The refugee's dilemma: Evidence from jewish migration out of nazi germany." *The Quarterly Journal of Economics* 138 (2): 1273–1345.
- Burdett, Kenneth, and Dale T. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39 (2): 257–273.
- Calvo-Armengól, A., and M. O. Jackson. 2004. "The effects of social networks on employment and inequality." *American Economic Review* 94 (3): 426–454.
- Card, D., J. Heining, and P. Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128 (3): 967–1015.
- Card, D., J. Rothstein, and M. Yi. 2024. "Industry Wage Differentials: A Firm-Based Approach." *Journal of Labor Economics* 42 (S1): S11–S59.
- . 2025. "Location, location, location." *American Economic Journal: Applied Economics* 17 (1): 297–336.

- Chetty, R., J. N. Friedman, E. Saez, N. Turner, and D. Yagan. 2020. "Income segregation and intergenerational mobility across colleges in the United States." *The Quarterly Journal of Economics* 135 (3): 1567–1633.
- Cingano, F., and A. Rosolia. 2012. "People I know: job search and social networks." *Journal of Labor Economics* 30 (2): 291–332.
- Combes, P. P., G. Duranton, and L. Gobillon. 2008. "Spatial wage disparities: Sorting matters!" *Journal of urban economics* 63 (2): 723–742.
- Conzelmann, J. G., S. W. Hemelt, B. Hershbein, S. Martin, A. Simon, and K. M. Stange. 2025. "Grads on the go: Measuring college-specific labor markets for graduates." *Journal of Policy Analysis and Management* 44 (3): 741–763.
- Deming, David J., and Kadeem Noray. 2020. "Earnings Dynamics, Changing Job Skills, and STEM Careers." *Quarterly Journal of Economics* 135 (4): 1965–2005.
- Diamond, R. 2016. "The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000." *American economic review* 106 (3): 479–524.
- Diamond, R., and C. Gaubert. 2022. "Spatial sorting and inequality." *Annual Review of Economics* 14:795–819.
- Dustmann, C., A. Glitz, U. Schönberg, and H. Brücker. 2016. "Referral-based job search networks." *The Review of Economic Studies* 83 (2): 514–546.
- Eliason, M., L. Hensvik, F. Kramarz, and O. N. Skans. 2023. "Social connections and the sorting of workers to firms." *Journal of Econometrics* 233 (2): 468–506.
- Federal Reserve Bank of St. Louis. 2025. *Civilian Labor Force Level - Bachelor's Degree and Higher, 25 Yrs. and over*. FRED, Federal Reserve Bank of St. Louis. Retrieved August 10, 2025.
- Ferhat, Samia. 2022. "The Impact of University Openings on Labor Market Outcomes." Working Paper.
- Forsythe, E., and R. Weinstein. 2025. "Recruiting intensity, hires, and vacancies: Evidence from firm-level data." Forthcoming, *The Economic Journal*.
- Gee, L. K., J. Jones, and M. Burke. 2017. "Social networks and labor markets: How strong ties relate to job finding on Facebook's social network." *Journal of Labor Economics* 35 (2): 485–518.
- Glaeser, E. L., and J. D. Gottlieb. 2009. "The wealth of cities: Agglomeration economies and spatial equilibrium in the United States." *Journal of economic literature* 47 (4): 983–1028.
- Glaeser, Edward L. 2010. *Agglomeration Economics*. University of Chicago Press.

- Gottlieb, Paul D. 2004. *Labor Supply Pressures and the Brain Drain: Signs from Census 2000*. Brookings Institution, Living Cities Census Series.
- Granovetter, M. S. 1973. "The strength of weak ties." *American journal of sociology* 78 (6): 1360–1380.
- Green, A. 2024. "Networks and Geographic Mobility: Evidence from World War II Navy Ships." Working paper.
- Guimaraes, P., O. Figueirdo, and D. Woodward. 2003. "A tractable approach to the firm location decision problem." *Review of Economics and Statistics* 85 (1): 201–204.
- Hellerstein, J. K., M. J. Kutzbach, and D. Neumark. 2019. "Labor market networks and recovery from mass layoffs: Evidence from the Great Recession period." *Journal of Urban Economics* 113:103192.
- Hensvik, L., and O. N. Skans. 2016. "Social networks, employee selection, and labor market outcomes." *Journal of Labor Economics* 34 (4): 825–867.
- Hombert, Johan, and Adrien Matray. 2019. "The Long-Term Consequences of the Tech Bubble on Skilled Workers' Earnings." Working Paper.
- Huttunen, K., J. Møen, and K. G. Salvanes. 2018. "Job loss and regional mobility." *Journal of Labor Economics* 36 (2): 479–509.
- Iammarino, S., A. Rodriguez-Pose, and M. Storper. 2019. "Regional inequality in Europe: evidence, theory and policy implications." *Journal of economic geography* 19 (2): 273–298.
- Idaho 66th Legislature. 2022. *House Bill No. 718*. HB 718.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review* 83 (4): 685–709.
- Jia, N., R. Molloy, C. Smith, and A. Wozniak. 2023. "The economics of internal migration: Advances and policy questions." *Journal of Economic Literature* 61 (1): 144–180.
- Kahn, Lisa B. 2010. "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy." *Labour Economics* 17 (2): 303–316.
- Kennan, J., and J. R. Walker. 2011. "The effect of expected income on individual migration decisions." *Econometrica* 79 (1): 211–251.
- Kirkeboen, Lars J., Edwin Leuven, and Magne Mogstad. 2016. "Field of Study, Earnings, and Self-Selection." *Quarterly Journal of Economics* 131 (3): 1057–1111.
- Kitagawa, Fumi, Chiara Marzocchi, Mabel Sanchez-Barrioluengo, and Elvira Uyarra. 2022. "Anchoring Talent to Regions: The Role of Universities in Graduate Retention through Employment and Entrepreneurship." *Regional Studies* 56 (6): 1001–1014.

- Kleven, H. J., C. Landais, E. Saez, and E. Schultz. 2014. "Migration and wage effects of taxing top earners: Evidence from the foreigners' tax scheme in Denmark." *The Quarterly Journal of Economics* 129 (1): 333–378.
- Koenen, M., and D. Johnston. 2024. "Social Ties and Residential Choice: Micro Evidence and Equilibrium Implications." Working Paper.
- Kramarz, F., and O. N. Skans. 2014. "When strong ties are strong: Networks and youth labour market entry." *Review of Economic Studies* 81 (3): 1164–1200.
- Kramarz, F., and D. Thesmar. 2013. "Social networks in the boardroom." *Journal of the European Economic Association* 11 (4): 780–807.
- Krueger, Alan B., and Lawrence H. Summers. 1988. "Efficiency Wages and the Inter-Industry Wage Structure." *Econometrica* 56 (2): 259–293.
- Kuhn, P., and K. Shen. 2013. "Gender discrimination in job ads: Evidence from China." *The Quarterly Journal of Economics* 128 (1): 287–336.
- Lange, F. 2007. "The speed of employer learning." *Journal of Labor Economics* 25 (1): 1–35.
- Manning, A., and B. Petrongolo. 2017. "How local are labor markets? Evidence from a spatial job search model." *American Economic Review* 107 (10): 2877–2907.
- Manski, C. F. 1993. "Identification of endogenous social effects: The reflection problem." *The review of economic studies* 60 (3): 531–542.
- McFadden, D. 1978. "Modeling the choice of residential location." In *Spatial Interaction Theory and Planning Models*, 75–96.
- Molloy, R., C. L. Smith, and A. Wozniak. 2011. "Internal migration in the United States." *Journal of Economic Perspectives* 25 (3): 173–196.
- Montgomery, J. D. 1991. "Social networks and labor-market outcomes: Toward an economic analysis." *American Economic Review* 81 (5): 1408–1418.
- Moretti, E. 2004a. "Estimating the social return to higher education: evidence from longitudinal and repeated cross-sectional data." *Journal of econometrics* 121 (1-2): 175–212.
- . 2004b. "Workers' education, spillovers, and productivity: evidence from plant-level production functions." *American Economic Review* 94 (3): 656–690.
- . 2010. "Local multipliers." *American Economic Review* 100 (2): 373–377.
- . 2012. *The New Geography of Jobs*. Houghton Mifflin Harcourt.
- . 2013. "Real Wage Inequality." *American Economic Journal: Applied Economics* 5 (1): 65–103.

- Moretti, E., and D. J. Wilson. 2017. "The effect of state taxes on the geographical location of top earners: Evidence from star scientists." *American economic review* 107 (7): 1858–1903.
- Moretti, E., and M. Yi. 2025. "Size Matters: Matching Externalities and the Advantages of Large Labor Markets." *NBER working paper No. w32250*.
- Mountjoy, J., and B. R. Hickman. 2021. "The returns to college (s): Relative value-added and match effects in higher education." Working paper.
- National Bureau of Economic Research. 2001. *Business Cycle Dating Committee Announcement, November 26, 2001*. National Bureau of Economic Research.
- Neal, D. 1995. "Industry-specific human capital: Evidence from displaced workers." *Journal of Labor Economics* 13 (4): 653–677.
- Nimier-David, Elio. 2023. "Local human capital and firm creation: Evidence from the massification of higher education in France." Working paper.
- 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.
- Ost, B., W. Pan, and D. Webber. 2024. "College Networks: The Importance of Employer Connections." Forthcoming, *Journal of Labor Economics*.
- Parent, D. 2000. "Industry-specific capital and the wage profile: Evidence from the national longitudinal survey of youth and the panel study of income dynamics." *Journal of Labor Economics* 18 (2): 306–323.
- Rothstein, Jesse. 2023. "The Lost Generation? Labor Market Outcomes for Post Great Recession Entrants." *Journal of Human Resources* 58 (5): 1452–1479.
- Schmieder, J. F., T. Von Wachter, and J. Heining. 2023. "The costs of job displacement over the business cycle and its sources: evidence from Germany." *American Economic Review* 113 (5): 1208–1254.
- Schultheiss, Tobias, Uschi Backes-Gellner, and Curdin Pfister. 2023. "Education Expansion and High-Skill Job Opportunities for Workers: Does a Rising Tide Lift All Boats?" *Labour Economics* 82:102354.
- Schwandt, Hannes, and Till von Wachter. 2019. "Unlucky Cohorts: Estimating the Long-Term Effects of Entering the Labor Market in a Recession in Large Cross-Sectional Data Sets." *Journal of Labor Economics* 37 (S1): S161–S198.
- Shapiro, J. M. 2006. "Smart cities: quality of life, productivity, and the growth effects of human capital." *The review of economics and statistics* 88 (2): 324–335.


- Silva, J. S., and S. Tenreyro. 2006. "The log of gravity." *The Review of Economics and statistics*, 641–658.
- Stuart, B. A., and E. J. Taylor. 2021. "Migration networks and location decisions: Evidence from US mass migration." *American Economic Journal: Applied Economics* 13 (3): 134175.
- Terza, J. V., A. Basu, and P. J. Rathouz. 2008. "Two-stage residual inclusion estimation: addressing endogeneity in health econometric modeling." *Journal of health economics* 27 (3): 531–543.
- Thompson, Derek. 2025. "Something Alarming Is Happening to the Job Market." The Atlantic. <https://www.theatlantic.com/economy/archive/2025/04/job-market-youth/682641/>.
- Wachter, Till von, Joyce Manchester, and Jae Song. 2009. "Long-Term Earnings Losses Due to Mass Layoffs During the 1982 Recession: An Analysis Using U.S. Administrative Data from 1974 to 2004." Working paper.
- Weinstein, R. 2022a. "Firm decisions and variation across universities in access to highwage jobs: Evidence from employer recruiting." *Journal of Labor Economics* 40 (1): 1–46.
- . 2022b. "Local labor markets and human capital investments." *Journal of Human Resources* 57 (5): 1498–1525.
- West, Darrell M. 2025. "Ways to Help Workers Suffering from AI-Related Job Losses." Brookings Institution. <https://www.brookings.edu/articles/ways-to-help-workers-suffering-from-ai-related-job-losses/>.
- Winters, John V. 2014. "STEM Graduates, Human Capital Externalities, and Wages in the U.S." *Regional Science and Urban Economics* 48:190–198.
- Wozniak, A. 2010. "Are college graduates more responsive to distant labor market opportunities?" *Journal of Human Resources* 45 (4): 944–970.
- Yagan, Danny. 2019. "Employment Hysteresis from the Great Recession." *Journal of Political Economy* 127 (5): 2505–2558.
- Zhu, M. 2025. "Beyond the 'old boys' network': Social networks and job finding at community colleges." *Journal of Human Resources* 60 (4): 1405–1434.
- Zimmerman, S. D. 2019. "Elite colleges and upward mobility to top jobs and top incomes." *American Economic Review* 109 (1): 1–47.

Appendix A

Additional Material for College Alumni Networks and Mobility Across Local Labor Markets


A.1 Appendix Figures and Tables

Figure A.1: Pooled 2010-2019 Industry Shares for Prime-Age College-Educated Workers, Revelio vs. ACS




Notes: Plots of industry mix in Revelio vs. ACS, nationwide (top panel) and Washington DC only (bottom panel). The vertical axis reports the pooled share of jobs held by college-educated workers aged 25-54 in a given 2-digit NAICS industry observed in 2010-2019 in the Revelio Labs data (including all workers, not just those displaced in mass layoffs or firm closures). Each job is only counted once, even if active for multiple consecutive years in the 2010-2019 timeframe. The horizontal axis reports the same share in the ACS, pooling the 2010-2019 cross sections. Each point represents a unique 2-digit NAICS code.

Figure A.2: Average Yearly Degrees Awarded from 2010-2019, Revelio vs. IPEDS




Notes: The vertical axis reports the average yearly number of degrees (all degrees) from 2010-2019 observed for each 4-year institution in our data on self-reported educational histories in LinkedIn. The horizontal axis reports the average yearly number of degrees (all degrees) for each 4-year institution as reported in the IPEDS 2010-2019 completions data files.

Figure A.3: A College-Educated Job Seeker on LinkedIn Can Access Many Pieces of Information about a Local Alumni Network




(a) Job Search Page



(b) List of Co-Alumni at Highlighted Job


Notes: The top panel displays a present-day snapshot of LinkedIn's job search tool from the perspective of a user. The user can highlight a specific local labor market and see not only a list of relevant jobs in that local labor market, but also the number of co-alumni who work there. In the bottom panel, we highlight the page that shows the full list of co-alumni after a user has clicked on the number of co-alumni button in the top panel. LinkedIn's job search tool and related tools have been in place since the early 2010s, so this feature is directly relevant for our sample of job seekers.

Figure A.4: Job Seeker's Location Choice Decision in Work History Timeline




Notes: This diagram traces a worker's post-graduation work history, excluding any job spells that happened before the worker's receipt of highest degree. School choice is taken as given. Otherwise, we focus on the location and job choice in the red outline, immediately following exogenous separation from their previous job due to a mass layoff or firm closure.

Figure A.5: Job Seeker's Location Decision Tree



Notes: This diagram traces a displaced job seeker's location decision tree in our model. The job seeker decides whether or not to stay in the city of displacement, and if they leave, then they choose among all remaining alternative destinations.

Figure A.6: Predictive Accuracy of School Strata in Explaining School Major Share



Notes: Data from IPEDS 2010 completions, at the school level. We plot histograms of standardized residuals from a regression of the share of Bachelor's+ degree recipients in 2010 from each major (2-digit CIP code) at a school on school strata fixed effects, where school strata are defined as in the Data Appendix Section A.2. The eleven majors displayed here correspond to the observed majors, with appropriate aggregation for consistency with 2-digit CIP codes, in the Revelio Labs data. From top left to bottom right, we report the majors in descending order of their frequency in Revelio Labs, with Business-related (which is an aggregate of Business, Finance, Accounting and Marketing) as the most common major, and Information Technology as the least common. We have 1,840 schools and 431 school strata. Regressions are weighted by the number of degrees a school awarded in 2010.

Table A.1: Variation in School Major Share Explained by School Strata

Major	R-squared
Business-related	0.40
Engineering	0.41
Education	0.48
Health Sciences	0.48
Social Sciences	0.57
Biology	0.49
Law	0.32
Physical Sciences	0.52
Math and Statistics	0.48
Information Technology	0.33
Architecture	0.41

Notes: Data from IPEDS 2010 completions, at the school level. We report the R-squared from a regression of the share of Bachelor's+ degree recipients in 2010 from each major (2-digit CIP code) at a school on school strata fixed effects, where school strata are defined as in the Data Appendix Section A.2. The eleven majors displayed here correspond to the observed majors, with appropriate aggregation for consistency with 2-digit CIP codes, in the Revelio Labs data. We report the majors in descending order of their frequency in Revelio Labs, with Business-related (which is an aggregate of Business, Finance, Accounting and Marketing) as the most common major, and Information Technology as the least common. We have 1,840 schools and 431 school strata. Regressions are weighted by the number of degrees a school awarded in 2010.

Table A.2: Effect of Alumni Network Size in Destination City on Location Choice from Multinomial Logit Framework

	Outcome: Number of o to d movers				
	(1)	(2)	(3)	(4)	(5)
Log Alumni Network Size	1.013*** (0.020)	0.976*** (0.011)	0.975*** (0.026)	0.461*** (0.013)	0.561*** (0.036)
Observations	890,960	872,900	163,089	135,067	56,872
Origin x School FE	X	X	X	X	X
Destination FE		X			
Destination x Strata FE			X		X
Origin x Destination FE				X	X

Notes: Estimates of the effect of alumni network size in any city on city choice using a multinomial logit alternative to our nested logit model. Estimates from gravity equation (Equation 1.10) via PPML, run on full counts for both movers and stayers. Dependent variable is the number of displaced alumni from school s in city o who choose city d , with $d = o$ possible, 1 year post-displacement. Primary independent variable is the log alumni network size (for school s) in d . Column 1 includes origin $\times$ school fixed effects. Column 2 adds destination fixed effects. Column 3 replaces destination fixed effects with destination $\times$ school strata fixed effects, where strata are defined as in Data Appendix A.2, whereas column 4 instead includes origin $\times$ destination fixed effects. Column 5 includes both. Standard errors are clustered three-ways at the origin $\times$ destination, origin $\times$ school, and destination $\times$ school levels. Significance levels: * 10%, ** 5%, and *** 1%.

Table A.3: Effect of Destination-Origin Difference in Alumni Network Size on Directed Migration

	Outcome: $\log\left(\frac{\text{Number of } o \text{ to } d \text{ movers}}{\text{Number of stayers in } o}\right)$			
	(1)	(2)	(3)	(4)
Difference in Alumni Network Size	0.405*** (0.020)	0.402*** (0.017)	0.330*** (0.023)	0.449*** (0.026)
Observations	6,333	6,317	5,622	1,064
Origin FE	X	X		
Destination FE	X	X		
School FE		X	X	
Origin x Destination FE			X	X
Destination x School FE				X

Notes: Estimates from log-odds gravity equation, mirroring Moretti and Wilson (2017), via OLS. Dependent variable is the log ratio of the following: in the numerator, the number of displaced alumni from school s in city o who move to city d 1 year post-displacement; in the denominator, the number of displaced alumni from school s in city o who stay in o . Primary independent variable is the difference in the log alumni network size (for school s) between d and o . Column 1 includes origin and destination fixed effects. Column 2 adds school fixed effects. Columns 3-4 replace separate origin and destination fixed effects with origin $\times$ destination pair fixed effects. Column 4 replaces school fixed effects with destination $\times$ school fixed effects. Standard errors are clustered three-ways at the origin $\times$ destination, origin $\times$ school, and destination $\times$ school levels. Significance levels: * 10%, ** 5%, and *** 1%.

Table A.4: Effect of Alumni Network Size (of Bachelor's Institution) in City of Displacement on Probability of Staying

	Outcome: 1(Stay metro 1-year post-displacement)		
	(1)	(2)	(3)
log Number of Co-Alumni	0.325*** (0.014)	0.416*** (0.064)	0.348*** (0.075)
Inclusive Value	-0.492*** (0.045)	-0.348*** (0.083)	-0.264** (0.127)
Observations	17,470	13,014	10,820
Individual controls	X	X	X
Institutional controls	X	X	
Layoff year FE	X	X	X
Origin FE	X		
Origin x school metro FE		X	
Origin x school strata FE			X

Notes: Notes: Table notes are the same as those for Table 1.2. The difference here is that network and school definitions are now based on the institution of Bachelor's degree, and not of highest completion.

Table A.5: Effect of Alumni Network Size (of Bachelor's Institution) in Destination City on Directed Migration

	Outcome: Number of <i>o</i> to <i>d</i> movers				
	(1)	(2)	(3)	(4)	(5)
Log Alumni Network Size	0.818*** (0.010)	0.834*** (0.031)	0.609*** (0.015)	0.799*** (0.043)	0.910*** (0.061)
Observations	458,065	50,886	87,107	31,329	17,232
Origin x School FE	X	X	X	X	X
Destination FE	X				
Destination x School Metro FE				X	
Destination x Strata FE		X			X
Origin x Destination FE			X	X	X

Notes: Table notes are the same as those for Table 1.6. The difference here is that network and school definitions are now based on the institution of Bachelor's degree, and not of highest completion.

Table A.6: Effect of Alumni Network Size in City of Displacement on Probability of Staying for Workers Displaced Outside of Home City

	Outcome: 1(Stay metro 1-year post-displacement)		
	(1)	(2)	(3)
log Number of Co-Alumni	0.350*** (0.021)	0.347*** (0.057)	0.354*** (0.074)
Inclusive Value	-0.521*** (0.054)	-0.187** (0.086)	-0.267*** (0.100)
Observations	11,842	7,512	5,895
Individual controls	X	X	X
Institutional controls	X	X	
Layoff year FE	X	X	X
Origin FE	X		
Origin x school metro FE		X	
Origin x school strata FE			X

Notes: Table notes are the same as those for Table 1.2. Model estimated on subset of job seekers who are displaced in a metro that is different from the metro in which they obtained their highest degree. Standard errors not yet adjusted for sampling variability in the inclusive value. Bootstrapped standard errors are roughly twice as large, based on our other models.

Table A.7: Effect of Destination Alumni Network Size on Directed Migration for Workers Displaced Outside of Home City

	Outcome: Number of <i>o</i> to <i>d</i> movers				
	(1)	(2)	(3)	(4)	(5)
Log Alumni Network Size	0.775*** (0.011)	0.819*** (0.038)	0.629*** (0.015)	0.847*** (0.066)	0.880*** (0.078)
Observations	398,978	41,555	69,021	21,119	12,051
Origin x School FE	X	X	X	X	X
Destination FE	X				
Destination x School Metro FE				X	
Destination x Strata FE		X			X
Origin x Destination FE			X	X	X

Notes: Table notes are the same as those for Table 1.6. Model estimated on subset of job seekers who are displaced in a metro that is different from the metro in which they obtained their highest degree.

Table A.8: Effect of Alumni Network Size in Destination City on Directed Migration for Workers who are Displaced Outside of Home City and Do Not Return to Home City

	Outcome: Number of <i>o</i> to <i>d</i> movers			
	(1)	(2)	(3)	(4)
Log Alumni Network Size	0.595*** (0.013)	0.530*** (0.043)	0.411*** (0.019)	0.371*** (0.105)
Observations	341,948	35,640	55,740	9,129
Origin x School FE	X	X	X	X
Destination FE	X			
Destination x Strata FE		X		X
Origin x Destination FE			X	X

Notes: Table notes are the same as those for Table 1.6. Model estimated on subset of job seekers who are displaced in a metro that is different from the metro in which they obtained their highest degree, additionally excluding workers who move back to the metro of their school 1 year post-displacement.

Table A.9: Estimated Impact of Alumni Network Size in City of Displacement on Probability of Staying is Robust to Including Location-Specific Demographic Heterogeneity

	Outcome: 1(Stay metro 1-year post-displacement)	
	(1)	(2)
log Number of Co-Alumni	0.446*** (0.065)	0.755*** (0.117)
Outside Option Value	-0.056* (0.032)	-0.116 (0.128)
Observations	5,788	2,625
Individual controls	X	X
Layoff year FE	X	X
Origin x school strata FE	X	
Origin x dem group FE	X	
Origin x school strata x dem group FE		X

Notes: Table notes are the same as those for Table 1.2, other than differences in the fixed effects included across specifications. Individual controls are age and age-squared — ethnicity and gender are absorbed by our definition of demographic groups below. Column 1 adds origin $\times$ demographic group FE to our baseline model in column 3 of Table 1.2. Demographic groups are defined as one of 12 groups based on the interaction of: gender (male or female), age bracket (22-30, 31-39, 40+) and race (white or non-white). Column 2 replaces origin $\times$ demographic group fixed effects with fully interacted origin $\times$ demographic group $\times$ strata fixed effects. Standard errors are clustered at the origin $\times$ school level but not yet adjusted for sampling variability in the inclusive value. Bootstrapped standard errors are roughly twice as large, based on our other models. Significance levels: * 10%, ** 5%, and *** 1%.

Table A.10: Inclusion of Location-Specific Demographic Heterogeneity Does Not Impact Estimated Effect of Alumni Network Size in Destination City on Directed Migration

	Outcome: Number of o to d movers in School s and Demographic group g			
	(1)	(2)	(3)	(4)
Log Alumni Network Size	0.901*** (0.058)	1.116*** (0.077)	1.044*** (0.140)	0.996*** (0.172)
Observations	27,966	24,605	5,663	2,928
Origin x School x Dem Group FE	X	X	X	X
Destination x Strata FE	X	X	X	
Origin x Destination FE	X	X		X
Destination x Dem Group FE		X		
Origin x Destination x Dem Group FE			X	
Destination x Strata x Dem Group FE				X

Estimates from adapted version of gravity equation (Equation 1.10) via PPML. Dependent variable is now counts of o to d movers (1-year post-displacement) in a origin-school-demographic group (g) combination, where $d \neq o$ and data is collapsed to the $o \times d \times s \times g$ level. Demographic groups are defined as one of 12 groups based on the interaction of: gender (male or female), age bracket (22-30, 31-39, 40+) and race (white or non-white). Primary independent variable is the log alumni network size (for school s) in d . All columns include origin $\times$ school $\times$ demographic group fixed effects. Columns 1 and 2 include origin $\times$ destination fixed effects and destination $\times$ strata fixed effects, where strata are defined as in Data Appendix A.2. Column 2 adds destination $\times$ demographic group fixed effects to column 1. Column 3 replaces origin $\times$ destination fixed effects and destination $\times$ demographic group fixed effects with fully interacted origin $\times$ destination $\times$ demographic group fixed effects. Column 4 returns to including origin $\times$ destination fixed effects but replaces destination $\times$ strata fixed effects with fully interacted destination $\times$ strata $\times$ demographic group fixed effects. Standard errors are clustered three-ways at the origin $\times$ destination, origin $\times$ school, and destination $\times$ school levels. Significance levels: * 10%, ** 5%, and *** 1%.

Table A.11: Estimated Impact of Alumni Network Size in City of Displacement on Probability of Staying is Robust to Exploiting Narrower Variation Within Firm-Strata

	Outcome: 1(Stay metro 1-year post-displacement)	
	(1)	(2)
log Number of Co-Alumni	0.389*** (0.072)	0.537*** (0.105)
Inclusive Value	-0.367*** (0.110)	-0.368** (0.146)
Observations	10,743	2,143
Individual controls	X	X
Layoff year FE	X	X
Origin x school strata FE	X	
Firm x school strata FE		X

Notes: Table notes are the same as those for Table 1.2, other than differences in the fixed effects included across specifications. Column 1 replicates the specification from column 3 of Table 1.2. Column 2 replaces origin $\times$ strata fixed effects with interacted firm $\times$ strata fixed effects, e.g. Microsoft Seattle $\times$ strata, where a firm nests both origin and company. Standard errors not yet adjusted for sampling variability in the inclusive value. Bootstrapped standard errors are roughly twice as large, based on our other models. Significance levels: * 10%, ** 5%, and *** 1%.

Table A.12: Inclusion of Firm-Specific Heterogeneity Does Not Impact Estimated Effect of Alumni Network Size in Destination City on Directed Migration

	Outcome: Number of o to d movers in School s and Firm f		
	(1)	(2)	(3)
Log Alumni Network Size	0.903*** (0.059)	0.907*** (0.147)	0.855*** (0.146)
Observations	27,700	3,590	3,013
School x Firm FE	X	X	X
Destination x Strata FE	X	X	X
Origin x Destination FE	X	X	
Destination x Company FE		X	
Destination x Firm FE			X

Estimates from adapted version of gravity equation (Equation 1.10) via PPML. Dependent variable is now counts of o to d movers (1-year post-displacement) in a school-firm (f) combination, where $d \neq o$ and data is collapsed to the $d \times s \times f$ level. Firms nest the city of displacement o (company $\times o$). Primary independent variable is the log alumni network size (for school s) in d . All columns include school $\times$ firm fixed effects. All columns also include destination $\times$ school strata fixed effects, where strata are defined as in Data Appendix A.2. Columns 1 and 2 include origin $\times$ destination fixed effects. Column 2 adds destination $\times$ company fixed effects to column 1. Column 3 replaces destination $\times$ company fixed effects with destination $\times$ firm fixed effects, which absorbs all factors varying at the origin $\times$ destination level. Standard errors are clustered three-ways at the origin $\times$ destination, origin $\times$ school, and destination $\times$ school levels. Significance levels: * 10%, ** 5%, and *** 1%.

Table A.13: Controlling for Major Does Not Affect Estimated Effect of Alumni Network Size in City of Displacement on Probability of Staying

	Outcome: 1(Stay metro 1-year post-displacement)		
	Logit	LPM	
	(1)	(2)	(3)
log Number of Co-Alumni	0.442*** (0.057)	0.069*** (0.011)	0.089*** (0.013)
Inclusive Value	-0.200*** (0.077)	-0.029** (0.013)	-0.011 (0.033)
Average Marginal Effect of log Number of Co-Alumni	0.071*** (0.009)	0.069*** (0.011)	0.089*** (0.013)
Observations	7,976	8,430	8,167
Individual controls	X	X	X
Layoff year FE	X	X	X
School FE			X
Origin x School Strata FE	X	X	
Origin x Major FE	X	X	
Origin x Major x school strata FE			X

Notes: Estimates from second-step logit and LPM approximations in Equation 1.9, with additional major-specific heterogeneity. Dependent variable is an indicator for whether a displaced worker stays in the same metro 1 year post-displacement. Estimation proceeds in two steps. In a first step, we estimate the parameters in our choice probability model (Equation 1.8b) using our gravity equation for destination choice (conditional on moving) in Equation 1.10. We add destination $\times$ major fixed effects to the gravity equation. We use these estimates to construct the inclusive value. In our second step, we estimate logits and LPMs at the worker level. All models include fixed effects for the calendar year of displacement and individual controls: gender, race, age and age-squared. In columns 1 and 2 we include origin $\times$ school strata fixed effects and origin $\times$ major fixed effects, with column 1 estimating a logit and column 2 the LPM. School strata are defined as in Data Appendix A.2. We define major based on the primary major at institution of highest degree, and include a separate category for missing majors. In column 3, we estimate an LPM equivalent to column 2 but add school fixed effects and fully interacted origin $\times$ major $\times$ strata fixed effects. Standard errors are clustered at the origin $\times$ school level but not yet adjusted for sampling variability in the inclusive value. Bootstrapped standard errors are roughly twice as large, based on our other models. Significance levels: * 10%, ** 5%, and *** 1%.

Table A.14: Inclusion of Major-Specific Heterogeneity Does Not Impact Estimated Effect of Alumni Network Size in Destination City on Directed Migration

	Outcome: Number of o to d movers in School s and Major m			
	(1)	(2)	(3)	(4)
Log Alumni Network Size	0.808*** (0.036)	0.606*** (0.016)	0.887*** (0.063)	0.964*** (0.073)
Observations	60,872	78,947	22,130	19,867
Origin x School x Major FE	X	X	X	X
Destination x Strata FE	X		X	X
Origin x Destination FE		X	X	X
Destination x Major FE				X

Estimates from adapted version of gravity equation (Equation 1.10) via PPML. Dependent variable is now counts of o to d movers (1-year post-displacement) in not just a school s but school-major (m) combination, where data is collapsed to the $o \times d \times s \times m$ level. Primary independent variable is the log alumni network size (for school s) in d . Column 1 includes origin $\times$ school $\times$ major fixed effects and destination $\times$ school strata fixed effects, where strata are defined as in Data Appendix A.2. Column 2 replaces destination $\times$ strata fixed effects with origin $\times$ destination fixed effects. Column 3 includes both. Column 4 adds destination $\times$ major fixed effects. Standard errors are clustered three-ways at the origin $\times$ destination, origin $\times$ school, and destination $\times$ school levels. Significance levels: * 10%, ** 5%, and *** 1%.

Table A.15: Effect of Alumni Network Size (2010-2019 Yearly Average) in City of Displacement on Probability of Staying

	Outcome: 1(Stay metro 1-year post-displacement)		
	(1)	(2)	(3)
log Number of Co-Alumni	0.324*** (0.018)	0.415*** (0.081)	0.378*** (0.097)
Inclusive Value	-0.491*** (0.050)	-0.375*** (0.093)	-0.365*** (0.112)
Observations	17,435	12,943	10,743
Individual controls	X	X	X
Institutional controls	X	X	
Layoff year FE	X	X	X
Origin FE	X		
Origin x school metro FE		X	
Origin x school strata FE			X

Notes: Table notes are the same as those for Table 1.2. We use a different measure of alumni network size relative to our baseline, namely the log yearly average (from 2010-2019) number of alumni in the city of displacement.

Table A.16: Effect of Alumni Network Size (2010-2019 Yearly Average) in Destination City on Directed Migration

	Outcome: Number of <i>o</i> to <i>d</i> movers				
	(1)	(2)	(3)	(4)	(5)
Log Alumni Network Size	0.821*** (0.010)	0.846*** (0.032)	0.614*** (0.015)	0.795*** (0.045)	0.905*** (0.060)
Observations	455,780	51,357	86,973	31,396	17,431
Origin x School FE	X	X	X	X	X
Destination FE	X				
Destination x School Metro FE				X	
Destination x Strata FE		X			X
Origin x Destination FE			X	X	X

Notes: Table notes are the same as those for Table 1.6. We use a different measure of alumni network size relative to our baseline, namely the log yearly average (from 2010-2019) number of alumni in the potential destination city.

A.2 Data Appendix

Mass Layoffs and Firm Closures Definition

From Revelio Labs’ yearly Workforce Dynamics files from 2010-2018, we use monthly changes in employment counts at the firm $\times$ month level to define mass layoff and firm closure events. Our approach closely follows Schmieder, Von Wachter, and Heining (2023). For mass layoff events, we impose the following set of criteria, all of which must be satisfied. First, the firm experienced a 30% year-to-year drop in employment, a 30% 1-month drop, or a 1-month drop of 500+ employees. Second, the firm had at least 25 employees in the prior month or year. Third, employment at the firm did not increase by more than 30% in the 2 years preceding the employment drop, and did not increase by 30% in the 2 ensuing years. Finally, no more than 20% of the laid-off employees were re-employed at the same firm within 1 year after the layoff event.

Our criteria to define firm closure events is as follows. First, the firm must have experienced a 80% year-to-year drop in employment, or 80% 1-month drop in employment. Second, the firm must have had at least 3 employees in the prior month or year. Finally, no more than 20% of the displaced employees at the firm are re-employed at the same firm within 1 year after the closure event.

Data on Individual Characteristics

For all workers in the individual job histories file observed in the US between 2010 and 2019, we also have imputed demographic characteristics from user names reported on their LinkedIn profiles. Revelio Labs estimates the probability that a user is male or female based on their first name and its share of each gender in social security administration data, from which we impute the user’s gender. We identify male users as those with a male probability exceeding 0.5, and female users as those with a female probability exceeding 0.5. We also impute each user’s ethnicity from estimated probabilities that the user is Black, Hispanic, White, Asian-Pacific Islander, Native American, or multiple races. These estimated probabilities come from Revelio Labs’ model predicting ethnicity using first name, last name, and location in US Census data. For each user, we assign the ethnicity with the highest estimated probability. Additional variables that reveal more information on a given user’s activity on the LinkedIn platform are the date that the profile was last refreshed as well as the user’s number of LinkedIn connections on this date. We do not, however, observe to whom a user is connected and thus cannot identify information about the composition of users’ networks of personal connections.

Coarsened Exact Matching of Schools

We compare the relative local alumni network sizes of observably similar schools located in the same geographical area. To create groups of observably similar schools, which we call

our school strata, we use coarsened exact matching (CEM) on several school characteristics observed in 2010 IPEDS data (Blackwell et al. 2009).

We start by defining location zones to ensure that all schools in the same school strata are in the same location zone. Our location zones are defined either as populous cities with many schools, e.g. Boston and Chicago, populous states, e.g. Texas, or small groups of adjacent non-populous states, e.g. North Dakota and South Dakota combined. Within each location zone, we use CEM to create strata based on binned versions of the following school characteristics in 2010: school size (4 bins), admissions rate (3 bins), 75th percentile composite SAT score (4 bins), and share of enrolled freshmen that are in-state (2 bins).

For a minority of colleges and universities, an exact match on the four aforementioned characteristics does not exist within their location zone. In these instances, we standardize these characteristics and assign these schools to the strata to which their nearest neighbor in the same location zone, based on minimum Mahalanobis distance, belongs. Or in the common event two schools are each other's nearest neighbor and no one else's, we create a new strata for those two schools.

We end up with 431 strata for 1,840 4-year colleges and universities (402 strata for 1,399 schools in our estimation sample). Some examples are as follows. In New England, we have a strata that consists solely of highly selective liberal arts colleges located outside of Boston: Middlebury, Bowdoin, Williams, Amherst, and Wesleyan. We have another strata that consists solely of large public universities in Arizona: University of Arizona and Arizona State University. And as a final example, we have a strata that consists solely of universities located in Northern California that are part of the California State University system: Cal State Chico, Cal State Fresno, and Cal State Sacramento.

A.3 Model Derivations

Impact of Origin Network Size on Probability of Staying

We show that an increase in the number of co-alumni in the city of displacement, σ_{os} , increases the probability of staying if $\beta^{Stay} > 0$. We start with a restatement of Equation 1.7b, the choice probability for staying in the origin:

$$Pr(Stay_{ios}) = \frac{\exp \left[\overbrace{(\omega^{Stay} + \rho^{Stay})}^{\equiv \beta^{Stay}} \log(\sigma_{os}) + \overbrace{(A_{os} + Z_{os}) + (B_o^{City} - r_o) + X'_i \eta}^{\equiv H_{ios}} \right]}{\exp \left[\underbrace{(\omega^{Stay} + \rho^{Stay})}_{\equiv \beta^{Stay}} \log(\sigma_{os}) + \underbrace{(A_{os} + Z_{os}) + (B_o^{City} - r_o) + X'_i \eta}_{\equiv H_{ios}} \right] + \exp(\lambda IV_{os})} \quad (A.1)$$

where we use H_{ios} to simplify notation by aggregating terms that do not include σ_{os} . If we use the quotient rule and take the partial derivative of $Pr(Stay_{ios})$ with respect to σ_{os} , we obtain:

$$\frac{\partial Pr(Stay_{ios})}{\partial \sigma_{os}} = \frac{\overbrace{\exp(\lambda IV_{os})}^{+} \left(\overbrace{\beta^{Stay} \sigma_{os}^{\beta^{Stay}-1} \exp(H_{ios})}^{+ \text{ if } \sigma_{os} > 0} \right)}{\underbrace{\left(\exp \left[\beta^{Stay} \log(\sigma_{os}) + H_{ios} \right] + \exp(\lambda IV_{os}) \right)^2}_{+}} \quad (A.2)$$

Recall that $IV_{os} \equiv \log(\sum_{m=1}^{J \setminus \{o\}} \exp(\frac{V_{ioms}}{\lambda}))$. Because the inner summation is over all destinations *excluding* o , the origin network size σ_{os} never appears in V_{ioms} for any $m \neq o$. Thus when taking the partial derivative of $Pr(Stay_{ios})$ with respect to σ_{os} , the inclusive value of the *Leave* nest IV_{os} can be treated as a constant. Because the denominator of the partial derivative in Equation A.2 is a squared term, it is always positive. In the numerator, we note that $\exp(H_{ios})$ and $\exp(\lambda IV_{os})$ are always positive. For any $\sigma_{os} > 0$, we also have that $\sigma_{os}^{\beta^{Stay}-1} > 0$ for any value of β^{Stay} . Thus with only β^{Stay} in the numerator still unaccounted for, the sign of the partial derivative of $Pr(Stay_{ios})$ with respect to origin network size, $\frac{\partial Pr(Stay_{ios})}{\partial \sigma_{os}}$, has the same sign as β^{Stay} . Thus if β^{Stay} is positive (negative), an increase in origin network size increases (decreases) a job seeker's probability of choosing to stay.

Impact of Destination Network Size on Conditional Destination Choice Probability

Now we show that if $\tilde{\beta}^{Leave}$ is positive, then an increase in any destination's network size will increase the job seeker's probability of choosing that destination, conditional on moving. We restate, with algebraic simplifications, the destination choice probability, conditional on moving, from Equation 1.8b:

$$Pr(h_{iods}|Leave) = \frac{\exp \left[\tilde{\beta}^{Leave} \log(\sigma_{ds}) + \frac{1}{\lambda} \overbrace{\left((A_{ds} + Z_{ds}) + (B_d^{City} - r_d) - \tilde{C}_{od} \right)}^{\equiv H_{ods}} \right]}{\sum_{m=1}^{J \setminus \{o\}} \exp \left[\tilde{\beta}^{Leave} \log(\sigma_{ms}) + \frac{1}{\lambda} \underbrace{\left((A_{ms} + Z_{ms}) + (B_m^{City} - r_m) - \tilde{C}_{om} \right)}_{\equiv H_{oms}} \right]} \quad (A.3)$$

where we use H_{ods} and H_{oms} to simplify notation and collect terms that do not feature destination network size σ_{ms} for any $m \neq o$. Using the quotient rule to take the partial derivative of $Pr(h_{iods}|Leave)$ with respect to σ_{ds} for some destination d , we have:

$$\frac{\partial Pr(h_{iods}|Leave)}{\partial \sigma_{ds}} = \frac{\overbrace{\left(\sum_{m=1}^{J \setminus \{o, d\}} \exp \left[\tilde{\beta}^{Leave} \log(\sigma_{ms}) + H_{oms} \right] \right)}^{+} \overbrace{\exp \left[H_{ods} \right]}^{+} \tilde{\beta}^{Leave} \overbrace{\sigma_{ds}^{\tilde{\beta}^{Leave} - 1}}^{+ \text{ if } \sigma_{ds} > 0}}{\underbrace{\left(\sum_{m=1}^{J \setminus \{o\}} \exp \left[\tilde{\beta}^{Leave} \log(\sigma_{ms}) + H_{oms} \right] \right)^2}_{+}} \quad (A.4)$$

Because the denominator in Equation A.4 is a squared term, it is always positive. In the numerator, we have a sum of exponentiated terms, which is always positive since exponentiated terms are themselves positive. Likewise, $\exp[H_{ods}]$ is always positive. If $\sigma_{ds} > 0$, then we also have $\sigma_{ds}^{\tilde{\beta}^{Leave} - 1} > 0$ regardless of the sign of $\tilde{\beta}^{Leave}$. Thus the sign of $\frac{\partial Pr(h_{iods}|Leave)}{\partial \sigma_{ds}}$ depends entirely on the sign of $\tilde{\beta}^{Leave}$. If $\tilde{\beta}^{Leave} > 0$, then we will also have that $\frac{\partial Pr(h_{iods}|Leave)}{\partial \sigma_{ds}} > 0$. Thus an increase in the network size of a destination increases the probability of choosing that destination.

A.4 Additional Notes for Empirical Exercises

Control Function Approach with Historical Instrumental Variables

In Section 1.4, we present two instrumental variables that isolate arguably exogenous variation in alumni network size that comes from historical differences in in-state enrollment share between schools in the same strata. We employ these instruments in a control function approach to estimate β^{Stay} and $\tilde{\beta}^{Leave}$. Our two instruments, one for alumni network size in the origin and the other for alumni network size in a destination, are as follows:

$$W_{os} \equiv \frac{Share_{s,1972}}{1 + dist_{os}} \quad (\text{A.5a})$$

$$W_{ds} \equiv \frac{Share_{s,1972}}{1 + dist_{ds}} \quad (\text{A.5b})$$

where $Share_{s,1972}$ is the share of enrolled students at school s in 1972 who are from the same state in which s is located. This share is interacted with inverse distance, $\frac{1}{1+dist_{ms}}$ for either $m = o$ or $m = d$, where $dist_{ms}$ is the great circle distance (in miles) between school s and city m . To estimate our two parameters, we estimate linear first stage equations and specify linear control functions in a two-stage residual inclusion (2SRI) procedure (Terza, Basu, and Rathouz 2008). We first describe the 2SRI estimate of $\tilde{\beta}^{Leave}$, which we use to then estimate β^{Stay} via 2SRI. We treat $\log(\sigma_{ds(o)})$, the log number of co-alumni from s in a destination d (averaged across the years of displacement of job-seeking alumni from s who are displaced in o), as our potentially endogenous regressor.¹ Our first stage regresses $\log(\sigma_{ds(o)})$ on the instrument W_{ds} and the same fixed effects that are included in our baseline gravity equation in Equation 1.10. We also add $dist_{ds}$ as a regressor. Thus the linear first stage for our directed migration analysis is:

$$\log(\sigma_{ds(o)}) = \alpha_1 + \pi_1 W_{ds} + \zeta_1 dist_{ds} + \phi_{od} + \phi_{os} + \phi_{dK(s)} + \xi_{ods} \quad (\text{A.6})$$

where $o \neq d$.² We use the residuals from this first stage, $\hat{\xi}_{ods}$, as an additional regressor in our gravity equation estimated via PPML:

$$P_{ods} = \exp(\tilde{\beta}^{Leave} \log(\sigma_{ds(o)}) + \zeta_2 dist_{ds} + \tau_2 \hat{\xi}_{ods} + \phi_{od} + \phi_{os} + \phi_{dK(s)}) \varepsilon_{ods} \quad (\text{A.7})$$

The estimated $\tilde{\beta}^{Leave}$ from Equation A.7 is the 2SRI estimate. Compared to our baseline gravity equation of directed migration, the only difference is that we include $dist_{ds}$ as a regressor and the estimated first stage residual $\hat{\xi}_{ods}$. Following our model's nested logit

¹We have added a subscript o to denote that we take the average destination network size across years of displacement within an origin-school pair. We also defined network size this way in our baseline directed migration analysis in 1.3, but did not include an o subscript to avoid superfluous notation.

²Although the dependent variable in the first stage is in logs, the first stage is still linear and estimated via OLS. Thus we avoid the “forbidden regression” in Angrist and Pischke (2009).

structure, we use these estimated parameters from Equation A.7 to construct the inclusive value of the *Leave* nest, which we first defined in Equation 1.7a. Denoting this estimated inclusive value as $\widehat{IV}_{os}$, which includes the estimated first stage residual, we can now obtain a 2SRI estimate of β^{Stay} . We specify a linear first stage using the origin-specific instrument W_{os} :

$$\log(\sigma_{ose_i}) = \alpha_3 + \pi_3 W_{os} + \zeta_3 dist_{os} + \bar{X}'_{ose_i} \eta_3 + \phi_{oK(s)} + \delta_{e_i} + \lambda_3 \widehat{IV}_{os} + \xi_{ose_i} \quad (\text{A.8})$$

where $\log(\sigma_{ose_i})$, the log number of co-alumni from school s in origin o in year e_i , is our potentially endogenous regressor. We include the same set of fixed effects as in our baseline out-mobility analysis. We collapse the individual controls included in our baseline out-mobility analysis to the same level as $\log(\sigma_{ose_i})$, which we denote as $\bar{X}'_{ose_i}$.³ We estimate the first stage equation via OLS at the origin $\times$ school $\times$ year level, weighting observations by the number of individuals in each cell. We use the estimated first stage residual, $\widehat{\xi}_{ose_i}$, as an additional regressor in our binary logit from Equation 1.9, which modeled a job seeker's decision to stay or leave. We also include $dist_{os}$ as a regressor. We obtain our 2SRI estimate of β^{Stay} by re-estimating our binary logit:

$$\begin{aligned} \log\left(\frac{Pr(stay_i)}{1 - Pr(stay_i)}\right) = & \alpha_4 + \beta^{Stay} \log(\sigma_{ose_i}) + \zeta_4 dist_{os} + \tau_4 \widehat{\xi}_{ose_i} \\ & + X'_i \eta_4 + \phi_{oK(s)} + \delta_{e_i} - \lambda_4 \widehat{IV}_{os} + \epsilon_{ios} \end{aligned} \quad (\text{A.9})$$

Because we estimate both β^{Stay} and $\widetilde{\beta}^{Leave}$ in multiple steps, we bootstrap standard errors (cluster re-sampling at the origin-school level, with 400 repetitions).

Multinomial Logit Model of Location Choice

Here, we consider a multinomial logit alternative to our nested logit model of location choice. Rather than first choose whether to stay or leave, and choose among non-origin cities conditional on leaving, we allow job seekers to choose among all cities in one step. This modification simply requires enforcing λ , the GEV parameter, to be equal to 1. We can then estimate a single β that governs both the stay vs. leave decision and directed choice over all cities. Accordingly, we re-estimate Equation 1.10 but on both movers and stayers, not just movers, thus allowing for $d = o$.

Under the multinomial logit structure, yet another approach to estimate a single β would be to adapt Moretti and Wilson (2017). Given our expression in Equation 1.8b for the choice probability of some destination city d , we can express, for job seekers who graduated from school s and are displaced in origin city o , the relative likelihood of choosing d over staying

³We avoid estimating the first stage equation at the individual level and including non-averaged individual controls, since these controls affect the first stage residual yet our potentially endogenous regressor does not vary across individuals within the same origin $\times$ school $\times$ year cell.

in o as a log-odds ratio that is linear in the difference between utility levels in the two cities:

$$\log\left(\frac{P_{ods}}{P_{oos}}\right) = \beta[\log(\sigma_{ds}) - \log(\sigma_{os})] + (A_{ds} - A_{os}) + (B_d^{City} - B_o^{City}) - (r_d - r_o) + (Z_{ds} - Z_{os}) - \tilde{C}_{od} \quad (\text{A.10})$$

where P_{ods} is the number of job-seeking alumni from school s who are displaced in a mass layoff or firm closure in city o and move to city d 1-year post-displacement, and P_{oos} is the number of job-seeking alumni from school s who are displaced in city o but stay in o . We can absorb destination-origin differences in all of the utility components outside of network size with fixed effects, and obtain the following estimating equation, estimated via OLS:

$$\log\left(\frac{P_{ods}}{P_{oos}}\right) = \beta[\log(\sigma_{ds}) - \log(\sigma_{os})] + \phi_{od} + \phi_{ds} + \varepsilon_{ods} \quad (\text{A.11})$$

Here, we include a set of origin metro $\times$ destination metro fixed effects ϕ_{od} , which absorbs differences in citywide amenities $B_d^{City} - B_o^{City}$; and differences in cost of living $r_d - r_o$. Furthermore, ϕ_{od} captures the moving cost from o to d , $\tilde{C}_{od}$. We can directly absorb Z_{ds} , which captures correlated preferences for a potential destination metro among alumni of the same school, as well as A_{ds} , which captures city-school productivity matches across all destinations, by including destination metro $\times$ school fixed effects ϕ_{ds} . We cannot also include origin metro $\times$ school fixed effects and identify β . We could include origin $\times$ school strata fixed effects by invoking our identifying assumptions that $Z_{os} = Z_{oK(s)}$ and an additional assumption that A_{os} can be decomposed into some combination of $A_{oK(s)}$ plus additively separable city-specific and school-specific work utility components, for all schools within a sufficiently narrowly defined strata K . However, in practice, since estimation requires non-zero flows of both movers and stayers for a given origin $\times$ destination $\times$ school and origin $\times$ school cell, respectively, we are constrained by the number of non-zero flows in our sample of job seekers. Thus we decide to prioritize including destination $\times$ school fixed effects, which more directly addresses any immediate, potentially lingering identification concerns.

Incorporation of Same-Industry Networks into our Model of Location Choice

In Section 1.6, we extend our model of location choice, specifically by allowing the work utility in a location to depend on the number of co-alumni in a given city who work in the same industry from which a job seeker is displaced. This extension allows us to quantify the impact of same-industry local alumni network size on location choice.

With no adjustments to work utility y_{ds} , we can now identify β^{Stay} and $\tilde{\beta}^{Leave}$ separately from the fixed city-school component of work utility A_{ds} . We estimate these parameters from our baseline model in both Equation 1.13, our binary logit relating the probability of staying to the same-industry network size, and a simpler version of Equation 1.14, our gravity equation of destination choice conditional on moving.

We also consider a more complex adjustment to the determination of a job seeker's work utility in a local labor market in Equation 1.4 to more flexibly allow for productivity matches between cities, industries and schools. For a worker from school s who is displaced from industry j , suppose the work utility for some job h in city d , y_{dsjh} , is additively separable into a school $\times$ industry component and city $\times$ industry component for the industry from which the job seeker was displaced, as well as a possible city $\times$ school component. The industry-specific terms capture the value of industry-specific human capital, while the city $\times$ school term provides more flexibility in allowing for both city and school to impact work utility. As in Equation 1.3, there is also an idiosyncratic match component ν_h that is i.i.d. with CDF $F(\nu_h)$. Thus our expression for the wage of job h is:

$$y_{dsjh} = \mu_{dj} + \mu_{sj} + \mu_{ds} + \nu_h \quad (\text{A.12})$$

We note that job h does not necessarily have to be in the same industry of displacement j . Thus the industry-specific human capital that is tied to j has a broad impact on earnings, and we do not model potential cross-industry productivity matches, which would require a model of both location and industry choice that is beyond the scope of this paper.

Next, as in our baseline model, we assume that a job seeker has a finite number of job draws in each potential destination d . However, to account for the role of industry-specific human capital accrued in industry j , we now assume that this number is proportional to the number of co-alumni from school s working in j in that city, σ_{dsj} , by a factor of $a < 1$. As before, this assumption does not require that all draws come entirely from jobs where same-industry co-alumni are present, and instead simply reflects the idea that job seekers are more likely to sample more jobs in a city with many co-alumni working in the same industry from which they were just displaced and had accrued more industry-specific human capital. If we continue to assume that $F(\nu_h)$ is Type-1 Extreme Value (McFadden 1978) with scale parameter ω^{Leave} , then the expected value of the highest work utility job y_{dsjh} in a sample of $a\sigma_{dsj}$ draws, which we denote as y_{dsj} , is:

$$y_{dsj} = \mu_{dj} + \mu_{sj} + \mu_{ds} + \omega^{Leave} \gamma + \omega^{Leave} \log(a\sigma_{dsj}) \quad (\text{A.13a})$$

$$\equiv A_{dj} + A_{sj} + A_{ds} + \omega^{Leave} \log(\sigma_{dsj}) \quad (\text{A.13b})$$


where A_{dj} , A_{sj} , and A_{ds} capture the various match-specific components of work utility between cities, industries, and schools. The rest of our model from Section 1.3 remains unchanged. As such, in our directed mobility analysis in Section 1.6, our inclusion of destination $\times$ industry, origin $\times$ school $\times$ industry, and destination $\times$ school fixed effects, absorbs the aforementioned matches. Furthermore, destination $\times$ school fixed effects absorb both the impact of total co-alumni network size (which enters on the amenities side) and correlated location preferences within schools Z_{ds} . Thus our estimating equations in Section 1.6 are theoretically consistent with the model of location choice presented here, extended to incorporate industry-specific human capital.

Appendix B

Additional Material for The Dynamics of High-Skill Worker Reallocation: Evidence from the Dot-Com Bubble


B.1 Appendix Figures

Figure B.1: Impact of Dot-Com Local Shocks on Employment, Longer Time Horizon




Notes: See notes for Figure 2.2. Here, I extend the time frame to 15 years post-recession.

Figure B.2: Impact of Dot-Com Local Shocks on Employment, Within-School Variation




Notes: See notes for Figure 2.2. Here, I use a different group definition for inclusion of group fixed effects. Here, a group is a set of individuals in the same industry $\times$ sex $\times$ degree level $\times$ graduating cohort $\times$ school in 2000. Relative to the baseline definition used to obtain estimates in Figure 2.2, I use school, i.e. institution of highest degree, instead of state of residence in 2000.

Figure B.3: Impact of Dot-Com Local Shocks on Employment, No Pre-Recession Adjustment



Notes: See notes for Figure 2.2. Here, I use a simpler employment measure as my outcome. Rather than define employment as relative to an individual's pre-recession average employment status, I simply use an indicator for employment of at least three months in a given year as the outcome variable.

Figure B.4: Impact of Dot-Com Local Shocks on Geographic Mobility, Within-School Variation



Notes: See notes for Figure 2.5. Here, I extend the time frame to 15 years post-recession.

Appendix C

Additional Material for The Labor Market Impacts of 21st Century STEM Expansions in US Higher Education

C.1 Appendix Tables

Table C.1: Top 50 Computer Science and Engineering Expansions

School	Field	Year	Number of Degrees	Change (Number)	Change (Percent)
University of Maryland-College	CS	2004	553	379	217.82
New York University	Engineering	2014	276	234	549.41
Ohio State University-Main Campus	Engineering	2010	889	229	34.70
School of Visual Arts	CS	2018	289	220	315.83
Baker College	CS	2015	216	194	904.65
Western Michigan University	CS	2003	226	190	527.78
Stony Brook University	CS	2003	297	189	175.00
Kaplan University	CS	2007	214	186	678.18
Michigan State University	Engineering	2016	691	182	35.89
Duke University	Engineering	2004	333	164	97.04
Rutgers University-New Brunswick	Engineering	2012	683	164	31.47
University of Illinois at Chicago	CS	2003	223	162	265.57
University of Maryland-College Park	CS	2004	447	160	55.48
University of Minnesota-Twin Cities	CS	2014	326	156	91.76
Pennsylvania State University-Main Campus	CS	2003	261	156	147.39
Purdue University Global	CS	2019	151	150	15000.00

APPENDIX C. ADDITIONAL MATERIAL FOR THE LABOR MARKET IMPACTS
OF 21ST CENTURY STEM EXPANSIONS IN US HIGHER EDUCATION

137

School	Field	Year	Number of Degrees	Change (Number)	Change (Percent)
University of Toledo	Engineering	2003	356	150	72.82
Kent State	CS	2003	179	148	477.42
University at Kent					
The University of	CS	2015	376	142	60.68
Texas at Austin					
University of	Engineering	2003	487	140	40.55
California-San					
Diego					
Florida State	CS	2002	266	140	110.28
University					
University at	Engineering	2011	528	137	35.04
Buffalo					
Indiana University-	CS	2005	191	126	193.85
Bloomington					
Kennesaw State	Engineering	2018	459	124	37.22
University					
Pace	CS	2002	319	122	61.93
University-New					
York					
University of	Engineering	2017	417	120	40.64
Dayton					
University of	Engineering	2008	475	119	33.43
Pittsburgh-					
Pittsburgh Campus					
University of	CS	2003	331	117	54.67
California-Irvine					
Brigham Young	CS	2019	375	116	45.07
University-Provo					
The University of	CS	2003	417	113	37.17
Texas at Austin					
United States Air	Engineering	2003	345	111	47.44
Force Academy					
Kansas State	CS	2003	151	104	224.73
University					
Colorado State	CS	2014	162	104	179.31
University-Global					
Campus					
Case Western	Engineering	2016	406	104	34.44
Reserve University					
New York	CS	2014	162	104	176.92
University					

School	Field	Year	Number of Degrees	Change (Number)	Change (Percent)
Massachusetts Institute of Technology	CS	2004	265	104	64.09
Florida Polytechnic University	Engineering	2018	106	102	2928.57
Georgia Institute of Technology-Main Campus	CS	2016	365	98	36.70
Purdue University-Main Campus	CS	2012	297	96	47.39
University of Miami	Engineering	2015	278	96	52.33
St John's University-New York	CS	2002	253	92	57.63
University of Utah	CS	2014	174	90	107.14
Boston University	Engineering	2014	368	90	32.37
Oklahoma State University-Main Campus	CS	2012	117	89	317.86
Eastern Michigan University	CS	2003	140	88	169.23
University of Michigan-Ann Arbor	CS	2003	221	88	66.17
Colorado Christian University	CS	2002	160	86	116.22
Full Sail Real World Education	CS	2006	224	82	57.75
University of California-Santa Cruz	CS	2002	162	82	102.50
Montana State University	Engineering	2014	304	81	36.32

Notes: Table showing the 50 largest expansion events in computer science and engineering, as recorded in the IPEDS data. All events are required to have involved an increase in 20 degrees completed in the specified year, with a percent change of at least 30%.

Table C.2: Impact of CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation, No Great Recession Cohorts

	Outcome			
	Advanced Degree	High-Tech Industry	Employer Size	Moved
	(1)	(2)	(3)	(4)
Post-Expansion Cohort	-0.012*** (0.002)	0.001 (0.001)	-0.042*** (0.007)	-0.008*** (0.002)
Observations	817,535	817,535	676,896	451,048
Individual Controls	X	X	X	X
School x Event FE	X	X	X	X

Notes: Estimates from Equation 3.1. Table notes are the same as for Table 3.2, with the difference that Great Recession cohorts (graduating 2007-2009) are excluded. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table C.3: Impact of CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation, No Great Recession Cohorts

	Outcome		
	High-Tech Industry	Employer Size	Moved
	(1)	(2)	(3)
Post-Expansion Cohort	0.002** (0.001)	-0.012* (0.007)	-0.002 (0.002)
Observations	817,535	696,555	473,310
Individual Controls	X	X	X
School x Event FE	X	X	X

Notes: Estimates from Equation 3.1. Table notes are the same as for Table 3.3, with the difference that Great Recession cohorts (graduating 2007-2009) are excluded. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table C.4: Impact of Partial and Full Exposure to CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation, No Great Recession Cohorts

	Outcome			
	Advanced Degree	High-Tech Industry	Employer Size	Moved
	(1)	(2)	(3)	(4)
Partial Post-Expansion Cohort	-0.006*** (0.001)	0.0002 (0.001)	-0.023*** (0.007)	-0.004** (0.002)
Full Post-Expansion Cohort	-0.017*** (0.002)	0.001 (0.001)	-0.053*** (0.006)	-0.010*** (0.002)
Observations	1,683,908	1,683,908	1,398,265	936,156
Individual Controls	X	X	X	X
School x Event FE	X	X	X	X

Notes: Estimates from Equation 3.2. Table notes are the same as for Figure 3.2, with the difference that Great Recession cohorts (graduating 2007-2009) are excluded. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table C.5: Impact of Partial and Full Exposure to CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation, No Great Recession Cohorts

	Outcome		
	High-Tech Industry	Employer Size	Moved
	(1)	(2)	(3)
Partial Post-Expansion Cohort	0.0004 (0.001)	-0.015** (0.007)	-0.002 (0.001)
Full Post-Expansion Cohort	0.001* (0.001)	-0.024*** (0.006)	-0.003* (0.001)
Observations	1,683,908	1,438,883	977,269
Individual Controls	X	X	X
School x Event FE	X	X	X

Notes: Estimates from Equation 3.2. Table notes are the same as for Figure 3.3, with the difference that Great Recession cohorts (graduating 2007-2009) are excluded. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table C.6: Impact of Largest CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation

	Outcome			
	Advanced Degree	High-Tech Industry	Employer Size	Moved
	(1)	(2)	(3)	(4)
Post-Expansion Cohort	-0.014*** (0.002)	0.001 (0.001)	-0.039*** (0.010)	-0.013*** (0.003)
Observations	337,324	337,324	277,544	179,325
Individual Controls	X	X	X	X
School x Event FE	X	X	X	X

Notes: Estimates from Equation 3.1. Table notes are the same as for Table 3.2, with the difference that I tighten my restriction on an expansion event definition. Here, a CS or Engineering expansion event must yield at least a 50% increase in completed degrees and 30 additional completed degrees in the related field relative to the previous year. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table C.7: Impact of Largest CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation

	Outcome		
	High-Tech Industry	Employer Size	Moved
	(1)	(2)	(3)
Post-Expansion Cohort	0.002** (0.001)	-0.015 (0.011)	-0.004* (0.002)
Observations	337,324	285,850	191,379
Individual Controls	X	X	X
School x Event FE	X	X	X

Notes: Estimates from Equation 3.1. Table notes are the same as for Table 3.3, with the difference that I tighten my restriction on an expansion event definition. Here, a CS or Engineering expansion event must yield at least a 50% increase in completed degrees and 30 additional completed degrees in the related field relative to the previous year. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table C.8: Impact of Partial and Full Exposure to Largest CS and Engineering Expansions on Individual Outcomes, First Job Post-Graduation

	Outcome			
	Advanced Degree	High-Tech Industry	Employer Size	Moved
	(1)	(2)	(3)	(4)
Partial Post-Expansion Cohort	-0.006*** (0.002)	0.002 (0.001)	-0.018 (0.012)	-0.006** (0.002)
Full Post-Expansion Cohort	-0.019*** (0.003)	0.002* (0.001)	-0.051*** (0.009)	-0.012*** (0.002)
Observations	687,005	687,005	567,044	367,884
Individual Controls	X	X	X	X
School x Event FE	X	X	X	X

Notes: Estimates from Equation 3.2. Table notes are the same as for Figure 3.2, with the difference that I tighten my restriction on an expansion event definition. Here, a CS or Engineering expansion event must yield at least a 50% increase in completed degrees and 30 additional completed degrees in the related field relative to the previous year. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.

Table C.9: Impact of Partial and Full Exposure to Largest CS and Engineering Expansions on Individual Outcomes, Five Years Post-Graduation

	Outcome		
	High-Tech Industry	Employer Size	Moved
	(1)	(2)	(3)
Partial Post-Expansion Cohort	0.002* (0.001)	-0.00001 (0.011)	-0.002 (0.002)
Full Post-Expansion Cohort	0.003*** (0.001)	-0.023** (0.009)	-0.003 (0.002)
Observations	687,005	583,617	390,670
Individual Controls	X	X	X
School x Event FE	X	X	X

Notes: Estimates from Equation 3.2. Table notes are the same as for Figure 3.3, with the difference that I tighten my restriction on an expansion event definition. Here, a CS or Engineering expansion event must yield at least a 50% increase in completed degrees and 30 additional completed degrees in the related field relative to the previous year. Standard errors are clustered at the school level. Significance levels: * 10%, ** 5%, and *** 1%.