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

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

Doctoral Dissertations, 2011-2026 Economics

2024

Essays on Human Capital: Re-skilling, College Major Choice, and the Incidence of Local Labor Market Shocks

Schulze, Karl

Princeton, NJ : Princeton University

http://arks.princeton.edu/ark:/88435/dsp01z316q498m

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

Essays on Human Capital: Re-skilling, College Major Choice, and the Incidence of Local Labor Market Shocks

Karl Schulze

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Economics
Advisers: Adam Kapor, Leah Boustan, Bo Honoré

September 2024

© Copyright by Karl Schulze, 2024.

All Rights Reserved

Abstract

This dissertation consists of three chapters that study individual investment in human capital and the incidence of economic shocks. I use administrative data to answer these questions, developing and applying structural methods to analyze individual choices in combination with a careful consideration of research design.

Chapter 1 studies the efficacy of using late-in-life educational investments to adapt to economic change. I focus on the Great Recession, an important shock that was accompanied by large increases in adult community college enrollment. I use administrative panel data on earnings and enrollment histories to estimate a dynamic model that captures the tradeoffs workers face in pursuing education later in life. A key contribution of the framework is to flexibly quantify worker skills using a data-driven approach. I find substantial heterogeneity in the earnings benefits of returning to school, but that these effects are limited or negative for those at the margin of enrollment.

Chapter 2, co-authored with Justine S. Hastings, Christopher A. Neilson, and Seth D. Zimmerman, studies how the choice of college and major are related to earnings in a context where selection bias can potentially confound direct measurement. We leverage administrative student data and the centralized structure of the Chilean higher education system to estimate a generalized Roy model of selection into colleges and majors. We find an important role of match effects and sorting on causal gains in generating earnings dispersion, indicating a large role for sorting into fields of study based on comparative advantage.

Chapter 3, co-authored with Sharada Dharmasankar, studies how economic shocks propagate through a local labor market by studying the effects of mass layoff events on an entire labor market. By leveraging matched employer-employee data, we can identify sharp layoff events that affect a specific labor market, cleanly identify which workers are directly and indirectly exposed to a layoff event. We find that mass layoff events permanently scar local labor markets with some evidence of spillover effects.

Acknowledgements

First, I am deeply indebted to my committee – Adam Kapor, Leah Boustan, and Bo Honoré – for their patience, encouragement, and support. I feel extremely lucky to have such an amazing team as my committee. Throughout my years with Adam as my advisor, he has been extraordinarily generous with his time, and I have learned a tremendous amount from him, not only in technical skills but also in shaping my approach to research. Special thanks also go to Chris Neilson, for inspiring me to do my best work and for providing opportunities to prove myself. My advisors are role models, not just for their research output, but also for the generous way they interact with students and colleagues.

I am also thankful to those who have enriched my experience at Princeton. Many faculty members at Princeton have generously provided feedback at various stages, including Zach Bleemer, Nick Buchholz, Ellora Derenoncourt, Henry Farber, John Grigsby, Kate Ho, Jakub Kastl, Alessandro Lizzeri, Steve Redding, Richard Rogerson, and others. My PhD cohort has been amazing and I could not imagine getting through the difficult parts of graduate school without them. In particular, Eduard Boehm and Carolyn Tsao have been generous with their friendship as well as their feedback. Princeton has hosted several postdoctoral scholars that have become friends as well as colleagues: Chris Campos, Ben Scuderi, and Emilio Borghesan. I am also grateful to the Industrial Relations Section, Valerie Ching, and other Princeton University staff for their support of my endeavours.

I am also thankful for the colleagues and teachers who have helped me along the way and whose guidance was instrumental in my decision to attend graduate school. In particular, I want to thank the economists I worked for at the Chicago Federal Reserve – Dan Aaronson, Mariacristna De Nardi, and Bhash Mazumder – for putting me on the path to success and for showing me the joy of the research process. Various teachers and researchers that I encountered during my undergraduate degree at the University of

Chicago were also instrumental in piquing my interest in research. In particular, Pietro Biroli and John Eric Humphries were key for educating me about the graduate school process.

On a personal level, I want to acknowledge my parents, Dean and Patricia; my sister, Helene; and my friends, Richard Chickey, Jason Nelson, Tommy Crawford, Dake Kang, and Mike Stralka; and my family-in-law, Ganapathy, Sowjanya, and Shivani Dharmasankar. You have all been my biggest supporters.

Last but not least, I want to thank my wife (and coauthor), Sharada Dharmasankar. Thank you for standing by my side and for picking me up off the floor when I am feeling down. You are my best friend and have made my life and this journey so much better.

To my family – Dean, Patricia, and Helene.

Contents

Abstract 3

Acknowledgements 4

1 Re-skilling, Up-skilling, and the Role of Education in the Adjustment to Economic Shocks **11**

1.1 Introduction 11

1.2 Data and Motivating Facts 23

1.2.1 Texas Schools Project Data 23

1.2.2 Basic Enrollment Patterns 24

1.2.3 Descriptive Patterns of Selection by Industry and Field of Study . . 27

1.2.4 Sources of Selection and Heterogeneity in Enrollee Earnings Trajectories 30

1.3 The Impact of Exposure to the Great Recession on the Schooling Choices of Adult Workers 33

1.3.1 Defining Great Recession Exposure 34

1.3.2 Baseline Event Analysis 36

1.3.3 Heterogeneity by Industry and Field of Study 39

1.4 An Empirical Model of Adult Labor Supply and Schooling Choices 40

1.4.1 Earnings and Human Capital Dynamics 43

1.4.2 Utility, Choices, and Beliefs 49

1.5 Identification 55

1.5.1 Key Assumptions . . . 55
1.5.2 Margins of Selection and Implied Restrictions on Worker Mobility . 57
1.5.3 Identification of Key Parameters . . . 61
1.6 Estimation . . . 65
1.6.1 Estimation Strategy . . . 65
1.6.2 Estimation Procedure - Step 1: Earnings and Human Capital Parameters . . . 67
1.6.3 Estimation Procedure - Step 2: Transition Parameters . . . 68
1.6.4 Standard Errors . . . 70
1.7 Estimation Results . . . 70
1.7.1 Parameter Estimates . . . 70
1.7.2 Model Fit and Validation . . . 80
1.8 Decompositions and Counterfactual Exercises . . . 83
1.8.1 Effects of Price Cuts and Field-of-Study Specific Subsidies . . . 84
1.8.2 Gender Differences in Enrollment and Earnings Trajectories . . . 91
1.9 Conclusion . . . 93
1.10 Appendix A: Additional Figures and Tables . . . 101
1.11 Appendix B: Data Appendix . . . 113
1.12 Appendix C: Mathematical Appendix . . . 119
1.13 Appendix D: Estimation Appendix . . . 126

2 Student Choices and the Earnings Associated with College Major and Selectivity 132

2.1 Introduction . . . 132
2.2 Institutional Context and Data . . . 140
2.2.1 College Admissions in Chile . . . 140
2.2.2 Data Sources . . . 141
2.2.3 Reports used in the Earnings Analysis . . . 143

2.2.4 Estimation Samples . . . 143
2.3 Model of Student College and Major Choice . . . 146
2.3.1 Overview . . . 146
2.3.2 Choices and Outcomes . . . 147
2.3.3 Specification of Unobserved Heterogeneity . . . 149
2.3.4 Reduced Form Choice Model . . . 151
2.3.5 Likelihood of an Application . . . 153
2.3.6 Decomposing Causal Effects . . . 154
2.4 Research Design and Estimation Approach . . . 156
2.4.1 Research Design and Identification . . . 156
2.4.2 Estimation Procedure . . . 161
2.5 Estimation Results . . . 165
2.5.1 Parameter Estimates and Model Fit . . . 165
2.5.2 What are the Latent Types? . . . 171
2.5.3 Variation in Preferences and Earnings: Do Students Value College-Major Earnings Gains? . . . 172
2.6 Decomposition of Treatment Effects . . . 176
2.7 Counterfactual Experiments . . . 178
2.7.1 Alternate Admissions Policy . . . 178
2.7.2 Seat Reallocation Policy . . . 180
2.8 Conclusion . . . 184
2.9 Appendix A: Additional Figures and Tables . . . 189

3 The Incidence of Mass Layoff Events: Evidence from Brazil 194
3.1 Introduction . . . 194
3.2 Data and Context . . . 200
3.2.1 RAIS Data . . . 200
3.2.2 The Brazilian Economy . . . 201

3.2.3 Employment Regulations in Brazil 202
3.3 Mass Layoffs in Brazil 203
3.3.1 Definition of Flagship Mass Layoff Events 203
3.3.2 Characteristics of Mass Layoff Events 205
3.3.3 Characteristics of Workers at the Layoff Firm 209
3.4 Empirical Strategy 213
3.4.1 Matching Approach 214
3.4.2 Empirical Specification: Local Labor Market-Level Analysis 218
3.4.3 Understanding Mechanisms: Worker-Level Analysis 220
3.5 Local Labor Market-Level Analysis 222
3.5.1 Baseline Results 222
3.5.2 Heterogeneity 229
3.6 Worker-Level Analysis 230
3.7 Conclusion 236
3.8 Appendix A: Additional Figures and Tables 241

Chapter 1

Re-skilling, Up-skilling, and the Role of Education in the Adjustment to Economic Shocks

1.1 Introduction

Many workers experience adverse shocks after having made their initial career and educational choices. For example, between 2000 and 2019, employment in manufacturing shrank by 25 percent, with over half of this decline occurring over the course of the Great Recession.[1] These adverse shocks can have effects that persist throughout a worker's career. Studies of displaced workers (Jacobson et al. [1993]), recessions (Yagan [2019]), and exposure to import competition (Autor et al. [2014]) all find that short-term hardship has long-run consequences. Workers face frictions in finding new jobs, and old skills do not always transfer to new kinds of work, meaning the specificity of human capital is a key determinant of the incidence of such shocks (Neal [1995], Gathmann and Schoenberg [2010], Traiberman [2019], Huckfeldt [2022], Deming and Noray [2022]).

[1] See https://fred.stlouisfed.org/series/MANEMP

Returning to school later in life is a common adjustment strategy that workers use, whether in response to job loss, business cycles, or trade exposure. In 2018, 3.1 percent of U.S. adults between the ages of 30 and 55 were enrolled in some kind of public higher education.[2] Policymakers often suggest investments in training programs - much of which takes place at community colleges - as a means by which workers may adapt to economic change. However, while this policy may increase earnings, it also involves substantial opportunity costs in terms of foregone earnings and lost accumulated work experience. Because adult workers have already made substantial investments in their prior career, new human capital investment may also require foregoing valuable existing skills. As this trade-off is heterogeneous, varying by both industry of employment and field of study, the value of these investments likely depends on who chooses to take up these programs and the nature of their future work transitions.

This paper studies the tradeoffs that low-skill workers face when pursuing later-life human capital investment. I focus on workers in Texas, for whom I have comprehensive administrative data on earnings and enrollment histories, and on one particular shock that was accompanied by substantial increases in community college enrollment, the Great Recession. I ask - which workers take up and benefit from training through the public post-secondary education system? To what degree do these patterns vary among different groups of workers? From a policy perspective, what impact would lowering community college tuition have on the earnings of those induced to enroll and how can we target this policy to maximize earnings gains?

To quantify the heterogeneous earnings benefits of later-life education and identify which workers are most sensitive to counterfactual tuition policies, this paper develops and estimates a dynamic discrete choice model of a worker's decision to return to school and switch industries in response to various kinds of shocks. I begin the paper

[2]This number comes from my calculations in the American Community Survey. While this enrollment rate declines with age, even among the 40 to 55-aged population the enrollment rate stands at 2.0 percent. When restricting the sample to Texas, the 30 to 55 enrollment rate increases to 3.5 percent when restricting the ACS to Texas.

by documenting patterns in the pathways taken between school and the labor market for adult workers and examine how these relationships translate into which workers return to school in response to granular, industry-level shocks from the Great Recession. Building off these results, the empirical model takes into account the imperfect transferability of prior work experience as a key determinant of enrollment choices. I develop a novel, data-driven approach to estimating the transferability of worker skills, which allows me to incorporate types of training and comparative advantage into the empirical model. By incorporating dynamics, the framework endogenizes the full pathway of work and schooling options taken by workers, capturing key complementarities between past work experience, fields of study while in school, and post-schooling work transitions. The model is quite rich, being tailored to match patterns of selection into schooling that have been documented as important in the training literature (Heckman et al. [1999], McCall et al. [2016]). I use the model to assess the distributional impact on earnings from reducing the price of community college and examine margins on which these subsidies could be better targeted to increase earnings.

The main finding is that workers are responsive to changes in tuition but that training is not effective in raising earnings for the average worker induced to enroll by a tuition subsidy. However, these average effects mask substantial underlying heterogeneity. Specifically, I find that lowering the price of community college results in an increase in enrollment of 2.9 percentage points but average earnings losses of 3.8 percent for those induced to enroll, with only 22.9 percent seeing earnings gains in the counterfactual. The model indicates that those who gain would have experienced counterfactually longer unemployment spells in the absence of the subsidy. By contrast, those who experience earnings losses tend to be characterized by failed re-skilling - these workers are induced by the subsidy to move out of their baseline sectors but fail to see earnings gains from training. This is driven by a combination of foregoing old specific skills, losing out on work experience during enrollment, and a failure to find work in a relevant

job after training. The results also indicate an important role for fields of study, where subsidizing technically-oriented fields, which have more general returns across sectors, fares relatively better than subsidizing health care, which sees potentially high returns that are very specific. This specific nature of human capital investment - that some training is only relevant in certain industries - magnifies the negative effects from failing to secure a job in a target industry. Taken together, these results indicate that *training is no substitute for on-the-job work experience* for adult workers and emphasizes the important interaction between education and work transitions. However, while re-skilling can be risky, education can be effective for those who face large frictions in returning to work after a job loss.[3]

I begin by outlining evidence that students select into particular education pathways based on their prior work experience and demographic types. For example, workers from health care make up only 3.4 percent of enrollment in vocational technical degrees (e.g. welding) while making up 44.6 percent in health care related fields. Women are also much more likely to pursue later-life schooling, enrolling at nearly twice the rate of men. I examine how this heterogeneity drives worker responses to shocks using an event study design that compares changes in enrollment between observably similar industries that experience different levels of employment losses from the Great Recession. This analysis shows that workers in more negatively shocked industries respond by pursuing further education - a one standard deviation increase in my measure of an industry's exposure to the Great Recession increases the probability of enrollment by

[3]A natural question is why forward-looking individuals return to school if they experience negative returns. From the perspective of the model, this will be driven by workers' preferences, which will adjust to rationalize why workers return to school at high observed rates in the face of low returns. The model conducts a positive analysis of the distributional earnings effects of training and does not take a stand on a specific market inefficiency justifying government involvement in post-secondary education. In particular, I do not take a stand on whether the tuition subsidy is "good" from a welfare perspective (lowering the cost of college can only make people "better off" in this framework) and focus on effects on earnings. For example, it could be the case that people are misinformed about the earnings value of college and would not go if they had better information (Jensen [2010], Hastings et al. [2015]). At the same time, it could also be the case that schooling has a return not captured by earnings, such as an amenity value of studying. Both of these mechanisms are picked up in the estimated preference for schooling and I cannot distinguish between them.

0.37 percentage points, a 7.1 percent increase relative to the 2007 baseline. However, these effects are heterogeneous by field of study and a worker's baseline industry. For example, I find that about 40 percent of the overall enrollment effect is driven by enrollment in health-related fields and that workers from professional industries (e.g. finance and real estate) are most responsive to the shock. These facts indicate that a worker's decision to enroll in school varies based on their prior labor market experience as well as their chosen field of study.

To understand how these selection patterns matter for earnings and identify the workers most sensitive to counterfactual policies, I develop and estimate a rich empirical framework that quantifies key mechanisms affecting the tradeoffs a worker faces in deciding to return to school. The framework is composed of a dynamic model of enrollment and work decisions combined with a rich statistical decomposition of life cycle earnings, building off work in the dynamic discrete choice literature (Keane et al. [2011]). In each period, workers make the joint choice of sector of employment and field of study, with the goal of maximizing the present discounted value of lifetime utility. Flow utility consists of both earnings and non-pecuniary factors, capturing the preferences of different worker types across sectors and fields of study. A key feature of the framework is that I model earnings as a weighted sum of latent skills in the spirit of Lazear [2009], where an individual's time-varying skills receive industry-specific weights and accumulate with work and schooling choices. This feature captures that workers may have comparative advantage across sectors based observed work histories. For example, a worker with experience in construction may have higher potential earnings in manufacturing than in health care. It also enables a tractable incorporation of fields of study, as certain types of training will be more or less relevant across sectors. The framework allows individuals to respond to various kinds of adverse shocks, including job loss, persistent shocks to earnings on the job, aggregate changes in labor market frictions, and aggregate changes in industry conditions. These create rich persistence in the un-

explained portion of earnings, substantially weakening identification assumptions, as well as enables the model to match transition and enrollment patterns in the presence of possibly persistent adverse shocks.

Estimation of the earnings parameters leverages the panel structure of the data to identify how a chosen sequence of industries and school enrollment decisions maps into earnings, measuring heterogeneous returns to work experience and fields of study. The crucial identifying assumption is that individuals cannot select on within-period changes in the error term in the earnings equation. However, due to the flexibility with which I model skills and the persistent process on the errors, individuals can select on their entire work and schooling history and a prior period's earnings. This amounts to a timing assumption. The intuition is similar to research designs based on movers (Card et al. [2013], Bonhomme et al. [2019]) or matching estimators using lagged earnings and work histories (Heckman et al. [1998a], Leung and Pei [2020]).[4] Specifically, given individuals with identical work and schooling histories up to a point, the model imputes average earnings changes upon an industry switch (or into enrollment) as the counterfactual earnings for stayers (non-enrollees). In practice, the estimation routine makes many such comparisons — occurring across different histories, types of transitions, and over multiple future periods — to inform parameters governing how skills are augmented by different paths of choices. The role of the Great Recession in identification is to

[4]My model allows for match effects based on observed work histories as well as prior earnings to affect selection into education and mobility across employers. In this spirit, my framework is most related to the dynamic model considered in Bonhomme et al. [2019] that relaxes assumptions in Abowd et al. [1999] (AKM)-style analysis. My framework differs in that it allows the entire past history of choices to affect earnings, incorporates education choices, and treats industry as the relevant level of firm classification.

The richness of the model also makes the identification of returns similar to arguments invoked in studies that use non-experimental methods to evaluate training programs. These non-experimental estimators have been shown to reasonably replicate experimental results when using rich lagged earnings histories (Heckman et al. [1998a], Dehejia and Wahba [1999], Card et al. [2018]). My framework is most similar to the strategy in Ashenfelter and Card [1985] in this literature, in that I estimate persistence parameters on the idiosyncratic component of earnings that give forecasts of how quickly shocks to earnings die away. The model differs in that I jointly estimate these persistence parameters at the same time as parameters governing heterogeneous returns, while the dynamic choice model microfounds the selection rule in the Ashenfelter and Card [1985] context.

generate exogenous variation in such transitions, disciplining model parameters through the increased mobility across industries and into education.[5]

A concern with this approach is that individuals who make certain kinds of transitions are selected on earnings in a way that is not accounted for in the framework. I address this concern in several ways. First, the model directly incorporates mechanisms that can generate Ashenfelter [1978]-dip dynamics, such job loss and persistent shocks to earnings, the inclusion of which also substantially weakens identification requirements. Second, I directly check implications of the identification assumptions, such as that future choices should not affect current earnings differences conditional on a state and a choice. This exercise indicates that the potential for residual bias to drive the results to be limited.[6] Third, I assess the fit of the model around job separation and enrollment events, which poses a difficult test for the model to match. I show that the model replicates earnings trajectories around these events, underscoring the credibility of the estimates. Finally, I directly implement matching estimators to estimate enrollment effects based on a richer set of covariates than the model can tractably incorporate. While this exercise misses the relevant population for policy counterfactuals and does not characterize mechanisms, I measure average earnings effects that are close to zero and often negative, indicating that results from the model are plausible when focusing on the policy-relevant population.

The estimates hold several important lessons about the interaction between labor market trajectories and later-life education. First, all else equal, returns from studying health have the potential to be high but are only relevant in the health care sector. In contrast, technically-oriented fields have returns that are lower but rewarded more consistently across sectors. Other fields of study consistently show limited returns relative

[5]I provide proofs of the semi-parametric identification of the skill process in a simple case.

[6]Future versions of this paper can enrich the framework and level of heterogeneity considered, which would further mitigate these concerns. For example, a simple extension would be to include time-invariant unobserved types as in Traiberman [2019]. However, given the limited earnings gains I estimate using matching estimators, I see it as unlikely this additional heterogeneity will substantially change the results.

to working. Second, there are large returns from continuous employment, indicating a substantial cost of job loss and an important role for job-ladder dynamics. Third, preferences are heterogeneous by gender and baseline education across fields and sectors. Men tend to have relatively high preferences for working in construction and manufacturing and lower preferences for working in health care. Fourth, I estimate large costs of switching sectors that are in line with the literature, where I additionally find that schooling plays a role in reducing these switching frictions. This echoes work which argues that part of what education does is to facilitate access to jobs that differ in on-the-job learning (Deming [2023]). Finally, workers are quite sensitive to tuition even at the low levels of tuition seen in the data.

The model is used to evaluate the take up and distributional effects on earnings of tuition subsidies. By modeling heterogeneity in choices and outcomes, the model sheds light on who could be targeted to maximize earnings gains. I focus on a hypothetical surprise $1,000 reduction in the cost of schooling that coincides with the onset of the Great Recession, representing more than 50 percent reduction in price. By comparing simulation results under the baseline and counterfactual data, the model identifies those who comply with the policy and counterfactual changes in earnings for this group. I find that 2.9 percent of all workers are induced to enroll between 2009 and 2015 but that only 22.9 percent of this group see earnings increases. Those who benefit from the policy face difficulty in returning to their old industry after a layoff, being characterized in the baseline scenario by long non-employment spells during the Recession followed by transitions into the low-paying service sector. For these workers, community college provides a pathway back onto the job ladder while also building new skills at a time when they lack on-the-job work experience. In contrast, those who see earnings losses are associated with failed re-skilling - the subsidy induces them to switch out of their baseline industry but does not translate into higher earnings in the target sector. Finally, examining fields of study shows that some of this risky re-skilling is amplified by the

specificity of human capital investment. Despite having the potential to have the highest earnings growth, health care fares the worst in terms of earnings gains due to the irrelevance of the degree for those who do not find employment in that sector.

The results of this paper provide valuable takeaways for policymakers, showing that training programs can be most effective when they combine skills learned in the classroom with straightforward pathways back into the labor market. First and foremost, the findings indicate that later life education is no substitute for work experience once relatively low-skill workers have entered the labor market. Not only do workers face a substantial opportunity cost from foregone earnings, but they also incur lost on-the-job experience. At the same time, this means that the population most likely to benefit from training are those who face substantial difficulty returning to work to gain such work experience. Re-skilling is also risky, as it may be detrimental if it does not translate into a relevant job at the end of education. This is especially true if the human capital investment from schooling is highly specific to narrowly defined sectors. The importance of post-schooling work transitions can rationalize some disparate results from the training literature even when experimental evidence is available, both reinforcing findings from evaluations of sectoral training programs (Katz et al. [2022]), which combine training with assistance transitioning into a targeted industry, and explaining disappointing results from other randomized evaluations (Fortson et al. [2021]).

Related Literature: This paper contributes to our understanding of the ways in which workers may respond to economic shocks. Studies have shown long-lasting negative effects of job loss (Jacobson et al. [1993], Couch and Placzek [2010]) that worsen during recessions (Davis and von Wachter [2011]). Echoing these results, trade shocks (Autor et al. [2013], Autor et al. [2014]) and recessions (Yagan [2019]) show scarring effects on worker and labor market outcomes. A natural question is why these effects are so persistent - one might expect workers to simply transition towards industries, occupations, and labor markets that are less affected by sudden declines in labor demand. However,

it has long been a source of concern that low-skill workers typically do not move across labor markets in response to demand shocks (Bound and Holzer [2000], Yagan [2019]). Moreover, workers face substantial earnings penalties when transitioning to types of work that are less related to their prior work experience (Neal [1995], Parent [2000], Gathmann and Schoenberg [2010]).[7]

Given the evident difficulty that low-skill workers face in adjusting to shocks along the margins discussed above, education holds potential as an alternate means to help workers adapt to sudden economic change. Prior work has found that the propensity to return to schooling is associated with adverse economic conditions (Betts and McFarland [1995], Barrow and Davis [2012]). Most related to the focus of this paper, Foote and Grosz [2019] and Minaya et al. [2023] identify causal increases in enrollment in response to mass layoff events, Barr and Turner [2015] show that more generous unemployment benefits increase enrollment in two-year colleges, and Boustan et al. [2022] find increases in enrollment in response to industry-specific exposure to technological change.[8] A related literature focuses on the impact of the business cycle on field of study choices for first-time students as well as how different college majors fare over the business cycle.[9] What this literature does not tell us is whether returning to school is the best course of action for adult workers in response to the particular shocks they face. My empirical framework advances our understanding of these issues by linking the decision to enroll

[7]The trade literature, for example, has argued that slow adjustment to trade shocks is driven by difficulty in adjusting skills and switching industries or occupations (Artuc et al. [2010], Dix-Carneiro [2014], Traiberman [2019]). Finally, job loss often entails the permanent destruction of valuable worker-employer matches (Lachowska et al. [2020]), which may entail occupational downgrading (Huckfeldt [2022]) or loss of job stability (Jarosch [2023]) due to frictions in finding jobs that substitute for prior employment.

[8]Conversely, Charles et al. [2018] and Schazenbach et al. [2023] find that enrollment declines in response to positive shocks to outside options, which is especially salient for community colleges.

[9]The use of education as an adjustment mechanism extends beyond the extensive margin of enrollment. Blom et al. [2021] finds that exposure to high unemployment causes individuals to select into different fields of study at the time of their first enrollment. Liu et al. [2019] and Esroy [2020] use cohort and state-based variation respectively to show that first-time students substitute towards majors less affected by the Great Recession. Acton [2021] shows that students enrolling in community college substitute towards less-affected majors with similar skill requirements in response to local mass layoff events. Altonji et al. [2016b] finds that the penalty for graduating in a Recession differs by college major. Deming and Noray [2022] show that changes in demand for skills taught in initial college major affect later life-cycle wage growth.

in response to adverse shocks with mechanisms that drive earnings outcomes.

This paper relates to the substantial literature on the return to education for adults, both in terms of training programs and community colleges. The literature on active labor market and training programs (Heckman et al. [1999], McCall et al. [2016]) tends to find mixed evidence on average returns with substantial differences by type of training and worker demographics.[10] The literature that uses panel data methods to estimate the returns to community college tends to find more significantly positive effects when exploiting variation in the number of credits or using non-completers as a comparison group.[11] These papers also tend to find that effects vary by field of study and are higher for women, findings that are also common in the training literature. A subset of these papers seek to estimate returns to community college for workers who have recently experienced an adverse shock, either focusing on displaced (Jacobson et al. [2005]) or unemployed (Leung and Pei [2020]) workers. These papers tend to find positive effects that are especially driven by health and technical degrees but with some types of training having zero effect. In particular, Leung and Pei [2020] point out that positive returns tend to be related to post-enrollment industry switching, which is an endogenous outcome of treatment in this literature. I build on this literature by providing a framework to understand mechanisms driving effects and worker selection into endogenous pathways. For example, my findings suggest that, by focusing on recently laid off workers, this literature may be restricting to a population with heterogeneously positive treatment effects, even if they have internally valid research designs. Policymakers should be

[10] See Card et al. [2018] for a meta-analysis of the literature. Recently, Andersson et al. [2013] evaluate the Workforce Investment Act (WIA) training programs, finding modest positive to negative impacts depending on the specific stream of services. Hyman [2018] finds positive effects for those participating in Trade Adjustment Assistance (TAA) based on quasi-random assignment to case workers. Humlum et al. [2023] find large positive effects for those retraining after suffering physical injuries. In terms of experimental evidence, Fortson et al. [2021] find no impact of WIA training programs on earnings and employment. In sharp contrast, experimental evaluations of sectoral training programs, which target in-demand sectors and provide assistance with job development and placement, show large positive effects (Katz et al. [2022]).

[11] See Kane and Rouse [1995], Jepsen et al. [2014], Belfield and Bailey [2017], and Stevens et al. [2019], among others

careful when extrapolating the results from this literature to those directly affected by policies.

Finally, this paper contributes to an extensive literature that models selection into education and life-cycle skill accumulation. Work on selection into schools for K-12 students (Walters [2018], Abdulkadiroglu et al. [2020]) and college majors (Arcidiacono [2005]) finds that the type of person enrolling is an important driver of observed returns. Some of this literature has considered whether returns to college major are general or specific (Kinsler and Pavan [2015], Eckardt [2022]),[12] finding that high-return education choices are sometimes specific to certain types of work. This paper extends this literature to consider these issues in later-life retraining decisions. This paper also relates to an extensive literature focusing on life-cycle human capital formation (see Keane and Wolpin [1997], Lee and Wolpin [2006], Roys and Taber [2019], among others). Applications of these models have shown that specific human capital determines who is affected the most by shocks, especially in the context of trade exposure (Dix-Carneiro [2014], Traiberman [2019]) or business cycles (Adda et al. [2013], Huckfeldt [2022], Grigsby [2022]).[13] I contribute to these literatures by combining a life-cycle model of worker mobility with a dynamic model of educational investment, tailoring the framework to match patterns of selection that are important in the training literature as well as my specific application.[14]

Organization: The structure of the paper is as follows. Section 1.2 outlines data and motivating facts concerning enrollment in my sample and over the Great Recession; Section 1.3 examines how the Great Recession shock impacts enrollment and field choices

[12] See also Lemieux [2014] or Leighton and Speer [2020] for additional empirical evidence.

[13] This paper is also related to the literature on dynamic treatment effects in sequential discrete choice models (Heckman and Navarro [2007], Heckman et al. [2016], Heckman et al. [2018]). Closely related to this paper, Rodriguez et al. [2022] adapt the empirical framework of Heckman et al. [2018] to understand the dynamic effects of multi-period job training programs. Relative to this literature, my paper fully specifies the earnings and choice process of forward-looking agents over many periods.

[14] A related literature considers the process by which workers learn about skills and match effects. See Miller [1985], Neal [1999], Papageorgiou [2014], Sanders [2014], Guvenen et al. [2020], and Arcidiacono et al. [2024], among others. I abstract from these considerations as the focus of the paper is how experienced workers, who likely have an idea of their skills, may use education to adjust these skills in response to changes in their labor market prospects. While potentially important, incorporation of this mechanism would greatly complicate my framework and make the current richness infeasible.

for adult workers; Section 1.4 outlines the model of work and schooling choices; Section 1.5 and 1.6 discuss identification and estimation, respectively; Section 1.7 presents parameters estimates and initial results from the model; Section 1.8 presents counterfactuals and decomposition exercises; and Section 1.9 concludes.

1.2 Data and Motivating Facts

1.2.1 Texas Schools Project Data

I use individual-level high school and college enrollment records linked with longitudinal earnings data from the state unemployment insurance (UI) system. These data are provided by the Texas Schools Project at the University of Texas at Dallas Education Research Center (ERC). High-school and college enrollment records come from the Texas Education Agency (TEA) and the Texas Higher Education Coordinating Board (THECB), respectively. The TEA records comprise the universe of Texas public high school students. The THECB enrollment data comprise all Texas public colleges (from 1992 onward) and private non-profit colleges (from 2003 onward). I use data on financial aid records for those receiving aid at any college in Texas (from 2001 onward), which includes all sources of federal, state, and categorical aid disbursed during the relevant academic year. When defining enrollment, I include all sectors of college, although my focus on relatively lower-skill workers means that more than 80 percent of enrollment in my baseline sample occurs at public 2-year institutions with almost the entire remainder at other public colleges. For each enrollment spell, I identify fields of study by combining information on the primary Classification of Instructional Programs (CIP) codes of study across semesters and CIP codes for degrees received. Details are provided in Appendix Section B.2.

These data are linked with quarterly earnings records from the Texas Workforce Commission (TWC), capturing the universe of Texas workers subject to the UI sys-

tem.[15] Earnings are deflated to 2019 dollars and winsorized at the 99.9th percentile. The data contain granular measures of industry of employment, which I aggregate into 5-digit North American Industry Classification System (NAICS) codes after some harmonization across years. Additional details on data cleaning are provided in Appendix Section B.2.

1.2.2 Basic Enrollment Patterns

Figure 1.1 gives a first indication of the importance that the business cycle, worker heterogeneity, and fields of study play in driving enrollment decisions. This figure plots age 25-plus enrollment counts in the community college sector as observed in the raw THECB data, broken out by gender and field of study. Appendix Table A1 reports the complete counts over various subgroups.[16]

Role of the Great Recession: *Figure 1.1a shows that the business cycle is a key driver of enrollment decisions, with a 34.7 percent increase in adult enrollment during the Great Recession.*[17]

[15] Missing earnings in UI data could indicate out-migration, working in jobs not covered by UI, or that the individual is out of the labor force. Fortunately, Texas is a state with the lowest out-migration of any state (Aisch et al. [2014]). Moreover, my focus on relatively low-skill workers means that out-migration is less of a concern, as these workers tend not to migrate in response to shocks (Bound and Holzer [2000]). Moreover, mobility responses across locations to the Great Recession are limited (Yagan [2019]). Additionally, I later make several sample restrictions to mitigate concern over these issues within sample, such as allowing longer periods of non-employment than other studies using state UI data (e.g. Lachowska et al. [2020]). Nevertheless, my result should be interpreted as conditional on the sample of workers with some minimal attachment to the Texas labor force over the period of analysis. I compare the representatives of my baseline sample to the same cohorts born in Texas in the American Community Survey in Appendix Figure B1.

[16] The focus of this paper is on older workers who already have some work experience. These age 25-plus enrollees account account for 34 percent of the total in 2007-2008 (317,958 out of a total of 935,455). However, younger enrollees also see increased enrollment during the Great Recession, increasing 20 percent (from 617,497 in 2007-2008 to 740,485 in 2010-2011). This indicates that the age-25 plus enrollees account for about half of the overall enrollment increase, which - given the other adjustment mechanisms available to older workers - merits a focused study. It is likely that the increase in enrollment for younger enrollees reflects the same shocks to outside options and policy changes that driver older enrollment. See Barr and Turner [2013] for an excellent summary of the impacts of the Great Recession on American higher education more generally.

[17] The official dates of the Great Recession are between December 2007 and June 2009. However, this dating does not map cleanly into the slack experienced in the labor market. First, there was a significant delay for the financial crisis to hit the labor market. Unemployment only rose above 5 percent in April 2008 and did not peak at 10 percent until October 2009. Second, it took quite a while for the unemployment rate to return to normal levels — it only returned to below 9 percent in October 2011 and did not return

In the 2007-2008 academic year, total enrollment across both genders was 317,958 for age 25-plus students, spiking to 428,394 in the 2010-2011 academic year. This rapid increase followed a period of minimal changes in enrollment prior to the Recession, where aggregate enrollment actually declined by 1.9 percent between 2005-2006 and 2007-2008. In 2017-2018, enrollment remains 8.1 percent above the pre-Recession level.[18]

Differences by Gender: *Figure 1.1a also shows that women are much more likely to enroll in community college, representing about two-thirds of overall enrollment*. In the 2007-2008 academic year, women make up 64 percent of overall enrollment (204,076 women versus 113,882 men). Both men and women increase enrollment during the Great Recession, with women increasing enrollment at a slightly faster rate than men during the downturn. Female enrollment increases by 32 percent (to 269,558), accounting for 59 percent of the overall increase in enrollment. This is striking as industries dominated by men were most affected in the Great Recession.

Differences by Fields of Study: *Figure 1.1b examines differences by fields of study, underscoring that specialization through different fields matters for these adult enrollees as well as how they respond to the business cycle.* This figure displays enrollment counts for the three largest areas of study (excluding liberal arts and general studies),[19] which make up about half of total enrollment in 2007-2008. Health degrees account for a large share of total enrollment - 22.8 percent of the total in 2007-2008 (72,548) - and are highly cyclical, increasing 46.5 percent to 106,289 in 2010-2011 before declining to a slightly lower enrollment level in 2017-2018 (71,112). Business and STEM fields also see large enrollment increases coinciding with the onset of the Great Recession, increasing by 32.9 percent for business (from 43,751 to 58,152) and 41 percent for STEM (38,876 to 54,819) between

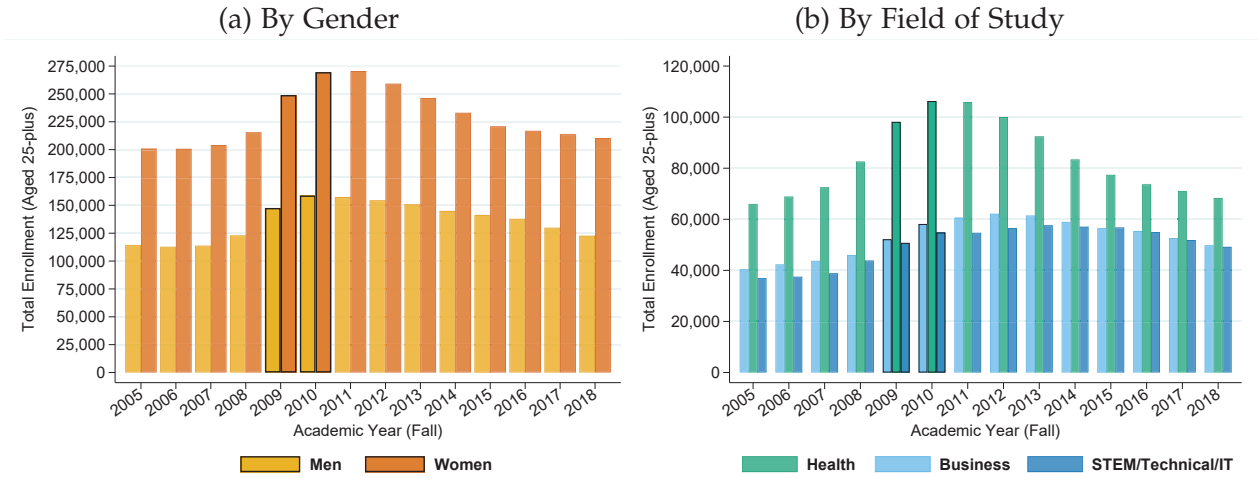
to below 6 percent until September 2014. For this reason, I would not expect the labor market downturn to have its full impact on enrollment until the 2009-2010 academic year and for any transitory effects to have recovered by 2014-2015.

[18]Since these are aggregate numbers, one might expect enrollment to increase over time due to in-migration into Texas. However, I show below that being exposed to employment losses of the Great Recession has a permanent effect on one's propensity to enroll.

[19]Field shares are broken out into more detail on my analysis sample in Figure 1.2.

2007-2008 and 2010-2011, respectively. In contrast to health, the increases in enrollment for these fields persist after the Great Recession, remaining 20.3 percent for business (at 52,642) and 33.5 percent for STEM (at 51,885) above pre-Recession enrollment levels in 2017-2018.[20]

Figure 1.1: Aggregate Adult Enrollment in the Public Two-Year Sector



Notes: Raw annual age 25-plus enrollment counts in the community college sector from THECB data prior to any sample selection based on work histories, including both full and part-time enrollment counts. Classification of enrollments into fields of study is outlined in Appendix Section B.2. Aggregation of CIP codes to broad areas of study is exhibited in Appendix Section B.4. The outlined bars mark school years that occur when the unemployment rate was over 9 percent during the Great Recession: April 2009 until September 2011. The 2009-2010 and 2010-2011 school years fall in this range. See Appendix Table A1 for complete results.

Sample Definitions

I now investigate how these factors relate to an individual's prior work experience and future transitions, which requires some sample selection to incorporate information on labor market dynamics for a consistent panel of workers. I focus my analysis on two panels of adult workers' employment and education histories. The first sample is the *full sample* of workers in the UI data with some attachment to the Texas labor market. The second panel is the *baseline sample*, a subset of the full sample on which I present most of my main analysis and conduct the structural estimation. This second sample

[20]The fields excluded from this figure also see large but less pronounced increases - increasing from 162,783 in 2007-2008 to 209,134 in 2010-2011 for a 28.5 percent growth. These other enrollments are not persistent, remaining 3.3 percent above pre-Recession-levels in 2017-2018 (at 168,143).

is restricted to the 1994 to 1998 cohorts of Texas high school graduates who do not attain advanced education degrees. These individuals are aged 29 to 33 in 2009 and 39 to 43 at the end of the observed panel period in 2019. The advantage of the baseline sample is the ability to observe demographic information for non-enrollees as well as the full history of work and education choices.[21] Additional details on data cleaning, sample definitions, and aggregation of industries and fields of study are discussed in Appendix B. Summary statistics for this sample are shown in Table B1 and comparisons with the American Community Survey are shown in Figure B1.

1.2.3 Descriptive Patterns of Selection by Industry and Field of Study

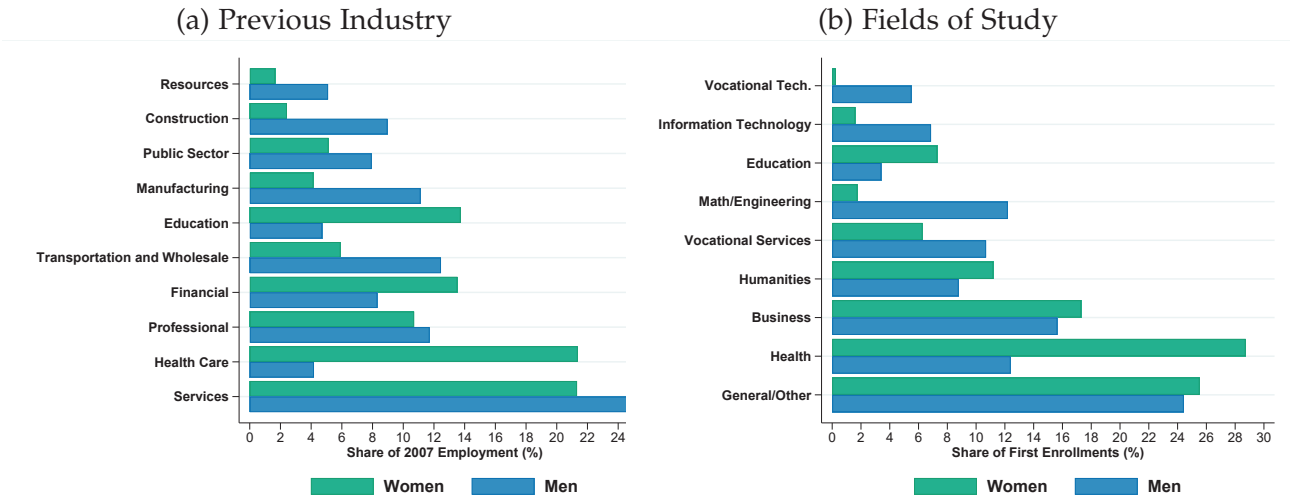
In Figure 1.2, I summarize the distributions of fields of study and industries of employment in the baseline sample, broken out by gender. A person's field of study is the field associated with their first enrollment spell after age 25. Industry is based on 2007 employment.

Differences in Selection by Gender: *Figure 1.2 shows that women tend to work in different industries than men and, when they go to school, they make different enrollment choices.* The largest industry category across both genders is services with 23.4 percent of employment in 2007. Other industries show larger differences in employment, with women more than 5.1 times more likely to work in health care and men more than 3.7 times more likely to work in construction. Turning to fields of study, health, business, and general/other make up a majority of study in the sample - totaling 63.7 percent of enrollments in the sample. Again, there are large gender differences, with women being more than 2.3 times likely to study health care than men and men are 6.9 times more likely to study math/engineering.

Selection into Field of Study by Prior Industry: *Table 1.1 shows that pathways from*

[21] The nature of the data is that I do not observe demographic information for individuals who neither enroll in the college or K-12 system.

Figure 1.2: Field of study and industry shares of enrollment



Notes: Both figures are computed on the baseline sample of $N = 247,973$. See Appendix Table B1 for additional variables. Panel (a) shows industry of employment in 2007 for the industry with highest earnings. Panel (b) shows fields of study associated with first enrollment spell after age 25, preceded by three years of no enrollment. See Appendix Section B.4 for details on how fields of study and industry codes were aggregated. Appendix Figure A1 shows this figure for the full sample, but not separated by gender.

work to school depend on prior industry, suggesting comparative advantage across courses of study based on prior work experience. This table shows the share of enrollment in a field coming from the largest industry of enrollment, excluding services.[22] While health care workers comprise 44.6 percent of enrollment in health fields, only 3.4 percent of enrollment in vocational technical fields and 5 percent in math/engineering comes from health care workers. In the latter two fields, workers from manufacturing make up 19.3 and 17.6 percent of enrollment, respectively. These numbers rise to 48.4 and 41.8 percent when combining manufacturing, construction, wholesale trade, and resource extraction industries.

At the same time that the patterns in Table 1.1 suggest predictable pathways between work and education, the relationships are not deterministic, indicating both re-skilling and up-skilling motivations in driving enrollment decisions. The majority studying health did not previously work in the health care sector, indicating many workers may be transitioning into health care. Individuals coming from service industries are a large share of enrollment

[22] As enrollment from the service sector is uniformly large across fields, focusing on the largest non-service sector is more informative for these relationships.

and are represented fairly equally across all fields of study, indicating that these workers may be up-skilling without regard to comparative advantage.

Table 1.1: Shares of Field Enrollment and Five-Year Switching Rates by Last Industry

	Primary Source Industry		*Future Outcomes*		
	Industry	Share (%)	Switch (%)	Work in Primary (%)	Graduate (%)
A. By Field of Study					
General/Other	Health Care	17.4	36.6	18.6	6.5
Business	Financial	19.1	37.2	18.5	27.1
Humanities	Financial	12.6	43.2	9.6	41.4
Math/Engineering	Manufacturing	17.0	46.9	22.9	31.2
Health	Health Care	41.3	41.4	54.6	36.2
Vocational Tech.	Manufacturing	18.2	52.0	19.0	37.8
Vocational Services	Public Sector	31.8	41.2	37.9	32.4
Education	Education	40.5	33.8	50.6	13.8
Information Technology	Professional	17.8	42.6	19.4	25.9
B. By Overall Enrollment					
Non-Enrollee			30.5		
Enrollee			39.8		25.4

Notes: Sample is workers with 4 quarters of consecutive employment in their baseline industry. Baseline industry is most common 5-digit industry of employment in quarters 5 to 12 prior to enrollment. Fields for enrollees comes from enrollments spells where the individual was not enrolled in the previous three years. Switching industries is defined as switching 1-digit industries defined relative to the baseline industry. Appendix Table A2 shows these results for the full sample.

Future Transitions of Enrollees: *The remaining columns of Table 1.1 show that enrollment is often, but not always, used as a means to switch industries, with differing use across fields of study.* Overall, enrollees tend to switch industries at a higher rate than non-enrollees. 39.8 percent of enrollees will switch 1-digit industries within five years while 30.5 percent of non-enrollees will do the same. These switches are at the 1-digit industry code level, indicating substantial changes in a person's career.[23] Some fields tend to be associated with industry switching more than others. For example, for those studying business, 37.2 percent switch 1-digit industries in five years and 59.1 percent will switch 5-digit industries. In contrast, for those studying vocational technical degrees, 52 percent will switch 1-digit industries while nearly 70 percent will switch 5-digit industries.[24]

[23] Comparing differences in these switching rates indicates that most of this switching is due to these large 1-digit industry changes. The comparable numbers for 5-digit switching is 60 percent for enrollees and 48.5 for non-enrollees, indicating that 80 percent of the gap in switching between enrollees and non-enrollees is largely due to these large 1-digit industry changes.

[24] Additionally, I investigate the how the relationship of industry switching and fields of study varies over the business cycle. These results are displayed in Appendix Figure A2 for both samples. Enrollees have a

Decomposition of Field-Industry Pathways by Gender: *Appendix Table A3 evaluates whether gender differences in enrollment choices are driven by differences in the baseline industry selected by men and women, consistently showing that most of the patterns are driven by differing choices by gender within an industry.* This table decomposes gender differences in enrollment into within- and across- industry differences for each field of study.[25] The results consistently show that about two-thirds of differences in field shares between men and women are due to differences in gender-specific enrollment within industries, whereas one-third is due to industry-specific enrollment rates across gender.[26] This within-industry variation in enrollment by gender will later be captured in gender-specific preferences in the model.

1.2.4 Sources of Selection and Heterogeneity in Enrollee Earnings Trajectories

I now briefly characterize how these patterns relate to earnings, underscoring important mechanisms driving enrollment decisions of adult workers - dynamics and fields of study. Both the decision of *when* to return to schooling and *what* to study are related to an individual's experience in the labor market. I capture these relationships in the model outlined in Section 1.4, where the analysis incorporates dynamic selection, dynamic treatment effects, and multiple unordered intensive margin treatments in a single empirical framework.

Earnings and Employment Dynamics around Enrollment: *Figure 1.3 plots employ-*

higher rate of industry switching rates across all years but there is actually little change in the probability of switching conditional on enrollment over the business cycle. The notable exception is health degrees, which exhibit about a 5pp increase in the rate of switching 1-digit industries conditional on enrollment across both samples.

[25] Table A4 shows the results for the full sample.

[26] For example, conditional on enrollment, women are 16pp more likely to enroll in health care than men. After equalizing the industry employment shares by gender, women still are 10.4pp more likely to enroll in health care. In other words, men tend to enroll less than women in health fields, even if they are working in health care. This indicates that most of the difference is not accounted for by the industries in which men and women work.

ment rates and average earnings around enrollment events, showing key patterns suggestive of complex dynamics that drive the decision to return to school. The typical pattern is a steep drop in earnings and employment during enrollment followed by robust post-enrollment earnings growth, amounting to about a 25 percent increase in earnings over the 10 years considered. First, the decline in earnings while enrolled (i.e., the "lock-in" effect) is preceded by a decline in employment rates prior to enrollment, which is worse for those enrolling in the Recession. The pre-enrollment drop in earnings is a well-known phenomenon in the training literature, reflecting selection based on transitory changes in labor market prospects (Ashenfelter [1978], Heckman and Smith [1999]). Second, we see that pre-enrollment earnings differ based on the timing of enrollment - those enrolling during the Recession have higher pre-enrollment earnings at the same time that they have a sharper drop in pre-enrollment earnings. This shows that accounting for differences in the level of pre-enrollment earnings is important as individuals may be differentially selected on the permanent components of their earnings. Finally, the robust post-enrollment earnings growth is faster for those enrolling in the Recession than other years. This may partially reflect the fact that Recession-era enrollees return to work in a thawing labor market, but may also reflect that these enrollees may differ in their potential for faster earnings growth, which may be related to higher initial earnings. These selection patterns will be incorporated into the modeling framework.

Heterogeneity in Fields of Study: *Figure 1.4 shows that different types of training are associated with different rates of post-enrollment earnings growth.* I sidestep selection bias stemming from extensive margin enrollment decisions in this analysis by focusing on intensive margin differences in earnings growth, demeaned relative to average earnings growth in the sample. More sophisticated analyses are exhibited in Appendix A.2, including results from matching estimators using previous histories. There is a clear ordering to estimated returns, with math/engineering having the highest relative returns, resulting in earnings growth 15 to 20 percent higher than the average field. Education

Figure 1.3: Employment and Earnings Trajectories around Enrollment Event



Notes: Earnings trajectories around target enrollment events (see Appendix B.2 for more information) for the baseline sample. Sample is balanced around the enrollment event. See Figure A3 for the full sample results.

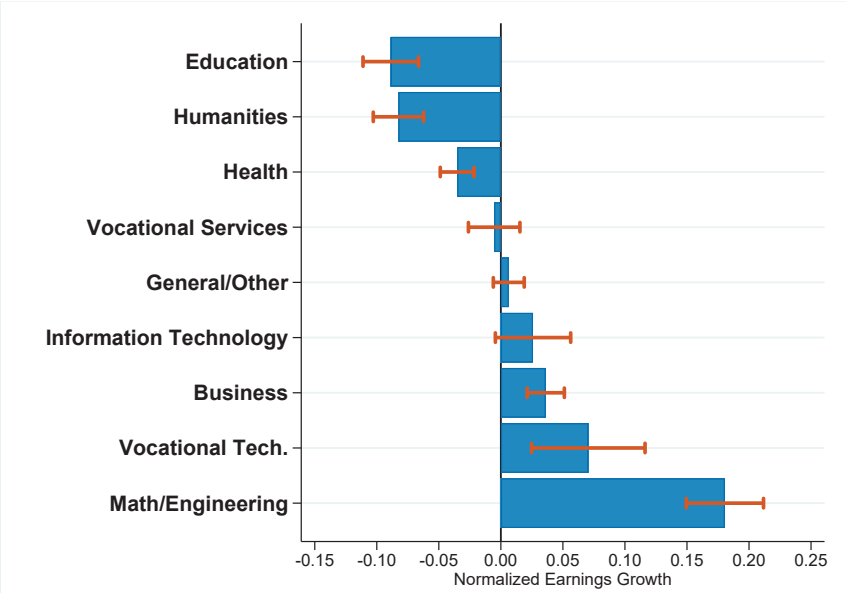
and humanities yield earnings growth that is significantly lower than the average course of study. Appendix Figure A6 shows these patterns survive adjusting for demographics and pre-enrollment employment dynamics. These results show that majors play an important role in my context, as has been shown in other education settings (e.g. Altonji et al. [2016a]).

Comparison with Non-enrollees: *Figure 1.5 constructs counterfactual non-enrollment paths for enrollees using matched controls, showing that there appear to be limited earnings gains from enrollment for enrollees on average.* Matches are constructed in the baseline sample by finding individuals of the same age, gender, pre-enrollment industry, and have similar earnings and employment dynamics prior to enrollment.[27] Employment and earnings prior to enrollment shows that the matching is done well.

However, the results show only an insignificant $14 (53.1) increase in earnings, not accounting for the lock-in effect. We see that part of this is driven is by lower rates of employment in Figure 1.5b, where enrollees are 2 percentage points less likely to be employed. The implication is that enrollees earn more than matched comparisons when

[27] Details are discussed in the notes to Figure 1.5. Results from placebo matches are shown in Appendix Figure A7, which indicates that results are not driven by a mechanical relationship driven by the matching procedure.

Figure 1.4: Relative Post-Enrollment Earnings Growth by Field of Study



Notes: Earnings growth for 41,857 target enrollments (see Section B.2) that occur between 2002 and 2013 in the baseline sample. Additionally, I restrict to enrollments where there is at least one period of work in periods 4 to 6 after enrollment. I use average non-zero earnings in quarters 9 to 12 prior to enrollment as baseline earnings. Using observations 12 quarters after enrollment, I run a regression of percent earnings relative to baseline on dummies for fields of study (excluding a constant) and control for a quadratic in earnings growth centered at 20 quarters post-enrollment and report clustered standard errors at the individual-level.

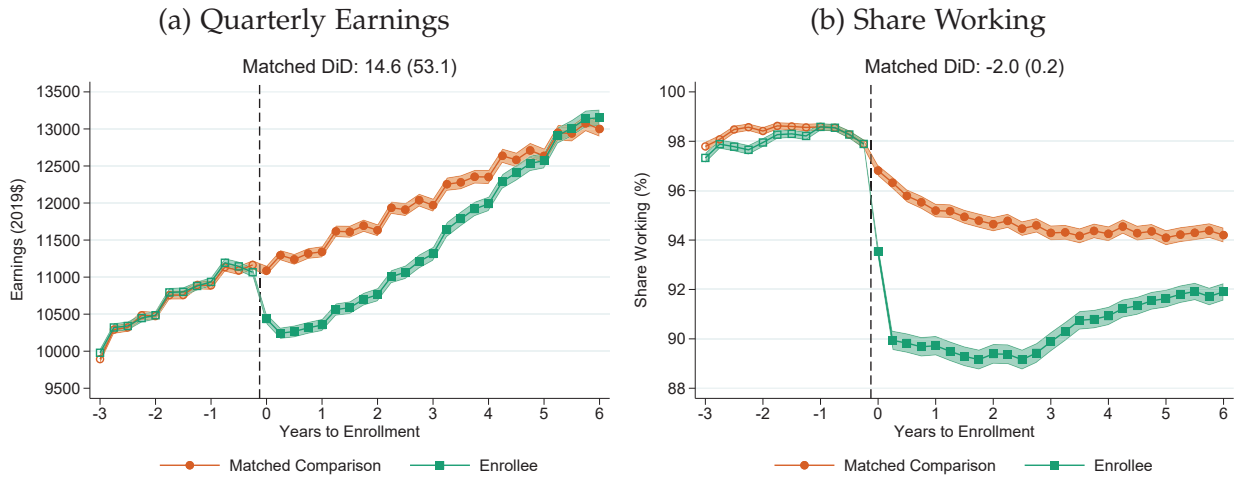
working, but have lower employment rates overall. This is consistent with other research that finds a key role for the extensive margin.[28] An advantage of the structural model is to separate these extensive versus intensive margin effects and aid in identifying the policy-relevant population by modeling choice.

1.3 The Impact of Exposure to the Great Recession on the Schooling Choices of Adult Workers

In this section, I show that exposure to the Great Recession across industries has a causal impact on a worker's decision to return to school. Using an exposure design in the spirit of Autor et al. [2014] and Yagan [2019], I show that *a 1 standard deviation increase in expo-*

[28]For example, Leung and Pei [2020] find that most of their *positive* effects are driven by increases in employment for unemployed enrollees.

Figure 1.5: Matching Estimates of Enrollment



Notes: For an individual enrolling in quarter t_0, I construct matches by first blocking on (i) gender, (ii) graduation cohort (iii) 3-digit NAICS of employment in $t_0 - 8$, (iv) indicators for positive earnings in $t_0 - 1$ to $t_0 - 4$, (v) earnings in each of the three years prior to be within \$5000. Then, I do nearest-neighbor Mahalanobis matching using quarterly earnings in each quarter $t_0 - 1$ through $t_0 - 12$. I use Mahalanobis matching rather than propensity score matching to account for different choices of time of enrollment and fields of study, which could be thought of as different treatments. Results are reported for the 22,594 enrollees and their matched comparisons for whom a control could be found. I check the quality of this matching procedure using placebo matches in Appendix Figure A7. Figure (a) plots total average earnings and Figure (b) plots any employment.

sure to the Great Recession causes a 7.1 percent (0.37pp) increase in the probability of enrollment for adult workers. This research contributes to the extant literature on economic shocks and schooling decisions (e.g., Betts and McFarland [1995], Barrow and Davis [2012], Foote and Grosz [2019], Boustan et al. [2022], Minaya et al. [2023]), where my contribution is to leverage the administrative data to dig into heterogeneity by industries, fields of study, and other outcomes. In addition to providing direct evidence that workers use education as a means to adapt to shocks, this variation will be directly incorporated into the later modeling framework.

1.3.1 Defining Great Recession Exposure

I begin by defining my measure of exposure to the Great Recession and describing its relationship to enrollment decisions. I measure exposure to the Great Recession as the negative of the percent change in quarterly employment between 2007 and 2009 in a 5-digit industry. To get baseline exposure, I link individuals to their 5-digit industry of

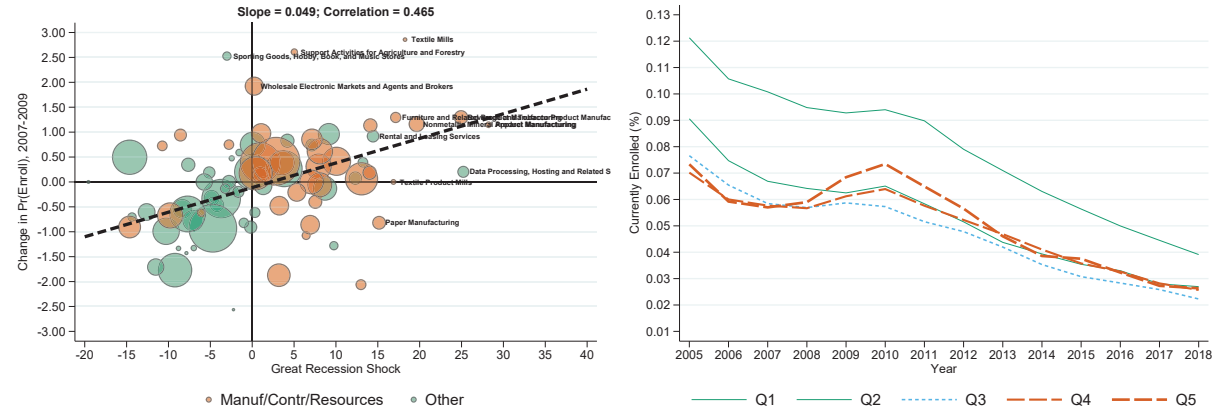
employment in 2007:Q4.[29]

Figure 1.6 shows that industries having a greater measured exposure to the Great Recession see increases in the rate of enrollment. This figure plots the raw variation of the exposure measure, aggregating up to the 3-digit NAICS-level. In particular, a 10 percentage point increase in exposure to the Great Recession is associated with about a 0.49 percentage point increase in the probability of enrollment. Notably, sectors associated with manufacturing, construction, and resource extraction are especially hard-hit, shown in orange.

Figure 1.6: Variation from Industry-Level Employment Shocks from the Great Recession: Baseline Sample

(a) Change in Enrollment by Shock Measure

(b) Enrollment Rate over Time by Quintile of Shock



Notes: Panel (a) shows a scatter-plot of the change in the probability of enrollment against the raw Great Recession exposure measure. Probability of enrollment and the Great Recession exposure measure are aggregated to the 3-digit NAICS level using baseline employment weights. Manufacturing, construction, and resource extraction industries are plotted in orange. Panel (b) shows the probability of being enrolled by year, separately for quintiles of the Great Recession exposure measure. Q4 and Q5 measure the industries most exposed to the Great Recession.

One concern with drawing causal inferences from this variation is that industries that are more exposed to the Great Recession have higher pre-Recession enrollment rates and may be on different enrollment trends. For example, health-oriented sectors tend to have higher baseline enrollment and are among the least exposed to the Great Recession. This is confirmed in Panel (b) of Figure 1.6, which plots enrollment rates by quintiles of the shock measure, where we see that the least shocked industries do

[29]To ensure attachment to this industry and a sufficient earnings history to define control variables, I restrict to the sample of workers who have at least one year of tenure in the baseline industry and are employed at least one quarter in 2006 and 2005. The panel is balanced through 2019, as per the sample restrictions discussed in Appendix B.2.

have higher baseline enrollment and see faster declines. Nevertheless, severely shocked industries also have marked increases in enrollment that represent a stark break in trend. For example, the most severely shocked quintile (Q5) increases its enrollment probability by 1.5pp between 2007 and 2010, from a baseline rate of 5.2 percent, around a 30 percent increase. To account for the potential effect of this aspect of the data, I control for baseline industry fixed effects and include pre-determined individual and industry-level controls interacted with time effects in my regression analysis. This ensures comparisons between individuals most likely to be on similar trends in the absence of the Recession. I also conduct several additional analyses in the appendix, showing my findings are robust and that parallel trends likely hold.[30]

1.3.2 Baseline Event Analysis

Let j_0 be the individual's baseline 5-digit industry in 2007:Q4. I model individual i's outcome in quarter t as:

$$Y_{it} = \alpha_{j_0(i)} + \alpha_t + \sum_{y(t)\neq 2007} \delta_y Z(SHOCK)_{j_0(i)} + \beta_t X_i^{2007} + \gamma_t W_{j_0(i)}^{2007} + \epsilon_{it} \tag{1.3.1}$$

where X_i^{2007} and $W_{j_0(i)}^{2007}$ are individual- and industry-level covariates defined prior to the Great Recession.[31] These coefficients are allowed to vary flexibly with time, weakening the parallel trends assumption to need only hold conditional on these controls. I also control for industry fixed effects $\alpha_{j_0(i)}$ and time fixed effects α_t. I cluster at the j_0-level to

[30] These results are shown in Table A6. Specifically, I investigate the performance of the controls and robustness to a sample defined relative to 1997 baseline industry, showing that parallel trends across industries are likely to hold conditional on the controls. I also use an exposure measure using data on sectoral employment outside Texas from the Statistics of U.S. Businesses. I also use the main sample to assess impacts on three-year industry transition probabilities, displayed in Appendix Figure A4, where the transition probability is redefined relative to an industry in every year.

[31] The individual-level covariates are worker earnings in 2007, worker earnings growth between 2005 and 2007, total experience in the labor market as measured in the TWC data, tenure in baseline industry, as well as age, gender, race, and observed education prior to 2007 for the baseline sample. The industry-level controls are 2007 average industry earnings, industry enrollment rates, and average industry earnings growth.

reflect that the exposure measure varies at the baseline industry level.[32]

Under this specification, δ_y measures the association between the degree to which a baseline industry is shocked and changes in Y_{it} relative to 2007, after removing common trends in the time-invariant controls. Interpreting δ_y as a causal effect of exposure on affected industries requires assuming that individuals with similar X_i^{2007} working in industries with similar $W_{j_0(i)}^{2007}$ would be on similar trends in Y_{it} in the absence of the Recession shock. This assumption is invalidated if workers with similar observables in low-shock industries are on a different paths of the outcome than observably similar workers in high-shock industries. This assumption can be checked by looking at pre-trends prior to 2009, where we should not see a change in the association with the shock measure for pre-Recession years.

Enrollment Effects: *Figure 1.7a shows that individuals working in shocked industries look similar in terms of enrollment rates before the Great Recession and then increase enrollment after the start of the downturn.* A 1 standard deviation increase in exposure causes a 0.37pp increase in the probability of enrollment in 2009, representing an 7.1 percent increase off these shocked industries' baseline enrollment rates (measured at a 1 standard deviation).[33] Interestingly, these enrollment effects persist on those impacted after the Recession has ended. In 2018, individuals shocked by the Recession are still 0.26pp more likely to be enrolled, although this effect is only marginally significant. Point estimates are shown in Appendix Table A5.[34]

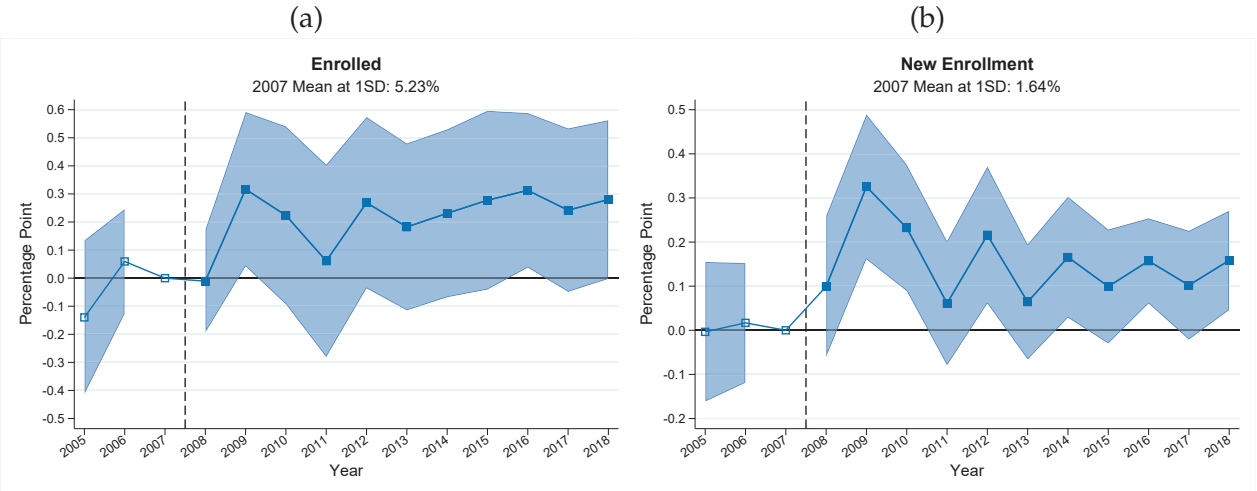
New Enrollment Effects: *Figure 1.7b shows strong effects on new enrollment in exposed*

[32]Since I run this regression on a balanced panel and the exposure varies at the 5-Digit baseline industry level, the inclusion of individual-level fixed effects will not affect the results.

[33]To get baseline averages for these industries, I regress the 2007 outcome on the Z-score of the exposure measure and compute the predicted value at 1 SD. The average 5-digit industries experiences a 1.7 percent decline in employment with a standard deviation of 15.2 percent.

[34]Appendix Table A5 shows that the exposure measure captures deteriorating labor market prospects for workers from shocked industries more broadly, despite being constructed using only employment declines. A 1 standard deviation increase in the shock from the Great Recession leads to a persistent decline in earnings, beginning with a negative $2,654 effect in 2009 (a 4.2 percent decline) and persisting to a negative $1,796 in 2018 (a 2.8 percent decline). While there is a slight pre-trend in earnings, it is small relative to the magnitude of the post-Recession effect.

Figure 1.7: Industry-Level Exposure Shock Event Study Results



Notes: Plots coefficient estimates of δ_y in Equation 1.3.1, which measures changes in the association between the outcome and the exposure measure of the Great Recession relative to 2007 after accounting for trends based on pre-determined controls. The exposure measure in this regression is set to be mean zero and standard deviation of one; the average 5-digit industries experiences a 1.7 percent decline in employment with a standard deviation of 15.2 percent. Panel (a) shows results for the outcome as a dummy for being currently enrolled. Panel (b) shows the results for a new enrollment, defined as being currently enrolled but not enrolled in the previous three years. Standard errors clustered based on an individual's pre-Recession 5-digit industry code. "2007 Mean at 1SD" is based on predicted values from a linear probability model of 2007 enrollment on the exposure measure, evaluated at 1 SD of the exposure measure. See Appendix Table A5 for point estimates and additional results.

industries, indicating that that almost all of the baseline enrollment effect is driven by workers enrolling for the first time. New enrollment is defined as a current enrollment without any enrollment in the previous three years. These new enrollment effects are much larger relative to the baseline, representing a 0.33pp increase or a 20 percent increase. Comparing to the overall enrollment effect, these results indicate that 90 percent of the increase in overall enrollment is made by new enrollees (0.33/0.37) rather than extending enrollment. Interestingly, these effects persist after the Recession ends as well.

Other Results: Looking at other outcomes reported in Appendix Table A5, I find significant effects on graduation but only in 2010, indicating that some enrolling individuals complete degrees but not in large numbers. These patterns are consistent in the full sample, where I estimate enrollment impacts of 0.17 percentage point (5.9 percent increase) on any enrollment and 0.11 percentage point on new enrollment (12.2 percent increase). I also conduct an exhaustive set of robustness exercises which lead to similar results, including a placebo analysis using 1997 as the baseline industry to assess the scope for mean reversion in driving the results. I also compare the results against al-

ternate controls and when using an exposure measure defined using employment losses outside of Texas. These results are shown in Appendix Table A6.

1.3.3 Heterogeneity by Industry and Field of Study

Table 1.2 shows heterogeneity in responsiveness to shocks from the Great Recession by outcome and within broad industry groups. To gain precision, I focus on the full sample and combine 2009 and 2010 coefficients from the event study. Results for the baseline sample, shown in Table A7, exhibit similar magnitudes although are often imprecise.

Heterogeneity by Industry: *Panel (a) of Table 1.2 shows that different types of industries differ in the responsiveness of enrollment to the shocks of the Great Recession, with professionally-oriented being the most responsive.* This panel reports results that subset the sample to compare industries in the same category. This means that each regression in Panel (a) is implicitly comparing changes in the association between enrollment and the Great Recession shock within groups. Professional sectors (which includes financial services and real estate) increases enrollment rates by 0.42 percentage points (or 10.5 percent off of the baseline for this group). Manufacturing and production-oriented sectors (which include construction, resource extraction, and wholesale trade), see increases in enrollment that are significant but less striking. Other industries see noisily estimated or limited increases in enrollment.

Heterogeneity by Field of Study: *Panel (b) of Table 1.2 examines heterogeneity by field of study, showing that by far the biggest enrollment increase comes from enrollment in health-related fields.* In this panel, enrolling in a field of study changes the outcome considered. Enrollment in health increases by 0.09pp in response to the shock, representing a 73.1 percent increase off the baseline. This large percent change is due to low enrollment rates in health prior to the Recession for those working in high-shock industries such as manufacturing and construction. Looking to the overall enrollment effect of 0.25pp (which combines 2009 and 2010) indicates that enrollment in health studies accounts for

38 percent of the total enrollment effect. While imprecise, the comparable number for the baseline sample is a 0.14pp increase in health (a 48 percent effect) that accounts for 40 percent of the total enrollment effect. When looking at other fields of study, we see that all fields somewhat increase enrollment in response to the shock.

Heterogeneity by Field-Industry Pairs: *Looking at the interaction between field and industry shows that many workers use this as an opportunity to up-skill while others use this as an opportunity to re-skill, with the re-skilling effect driven by moves into health care.* Appendix Table A8 reports a full breakdown of industry and field of study combinations for the full sample.[35] The results indicate field-by-industry patterns in enrollment responses to the shock. The response of production-oriented industries to studying health is large, representing 22 percent of the overall effect for these workers. This indicates that many manufacturing workers change tracks in response to the shock. At the same time, the remainder of the response is dominated by enrollment in more technically-oriented fields, where vocational technical degrees and other engineering degrees account for 75 percent of the effect for these workers. We see a similar set of patterns when looking at what drives the effects for professional workers: 22 percent of the effect is driven by health and 32 percent is driven by business enrollment.

1.4 An Empirical Model of Adult Labor Supply and Schooling Choices

The results thus far show systematic patterns of selection by past work experience and fields of study that determine how workers respond to the shocks of the Great Recession. While valuable pieces of evidence, these results do not speak to whether returning to school was the best course of action for these workers or how policy could be better

[35]These crosstabs are almost always statistically insignificant for the baseline sample but typically report similar magnitudes.

Table 1.2: Heterogeneity by Field and Industry: Full Sample

A. By Industry:				
	Production/Manufacturing	Services	Professional/Finance	Other
Enroll (pp)	0.159**	0.053	0.420***	0.129
SE	(0.073)	(0.058)	(0.127)	(0.120)
Effect (%)	[6.4%]	[1.6%]	[10.5%]	[5.3%]
N (Industries)	142	130	79	210
N (Individuals)	411,385	685,468	594,704	1,748,053

B. By Fields of Study:					
	Study Health	Study Business	Study Information Technology	Study Vocational Services	Study Education
Enroll (pp)	0.094***	0.031*	0.014***	0.029***	0.018***
SE	(0.025)	(0.016)	(0.005)	(0.010)	(0.006)
Effect (%)	[73.1%]	[3.9%]	[8.2%]	[22.7%]	[5.8%]
N (Industries)	561	561	561	561	561
N (Individuals)	3,439,610	3,439,610	3,439,610	3,439,610	3,439,610

Notes: Coefficients pool the event study parameters δ_y for 2009 and 2010 from the event study specification in Equation 1.3.1. Standard errors, shown in parentheses, are clustered based on an individual's pre-Recession 5-digit industry code. Standard errors clustered based on an individual's pre-Recession 5-digit industry code. Percent effects, shown in brackets, are based on predicted values from a linear probability model of 2007 enrollment on the exposure measure, evaluated at 1 SD of the exposure measure. Results for the baseline sample are shown in Appendix Table A7 and a full breakdown of outcomes by industry groups are reported in Appendix Table A8. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

designed to aid transitions and increase worker earnings. In order to shed light on these issues, I now outline a structural model of a worker's choice of industry of employment, field of study, and resulting earnings. At a high level, the empirical framework has two components: a dynamic model of work and schooling choices and a rich statistical model of life-cycle earnings. By incorporating dynamics, the model endogenizes the full pathway taken by workers, accounting for selection into education based on past work experience as well as the indirect effect of schooling on future work transitions.

Preliminaries: Time is discrete. In each year t, individual i makes a choice $k \in \{1, ..., K\}$ in order to maximize their present discounted value of lifetime utility,[36] which includes both earnings and non-pecuniary factors. Each action the individual takes corresponds to a joint choice of working in industry j and studying field of study f: $k = (j, f)$ where $j \in \{0, 1, ..., J\}$ and $f \in \{0, 1, ..., F\}$. Individuals can both work in industry j and study in field f, which I call *part-time study* (i.e., $j(k)$>0 and $f(k)$>0) or be entirely non-employed (in which case j=0 and f=0).[37] Let $j(k)$ and $f(k)$ to denote the

[36] Workers enter the labor market at age 25 and retire at 55.

[37] Workers may be fully employed in sector j (in which case $j > 0$ and $f = 0$), study full-time in field f (in which case $j = 0$ and $f > 0$), or be voluntarily non-employed (in which case $f = 0$ and $j = 0$). The $J = 7$ industries are: A General Sector, Construction, Manufacturing, Health Care, Services, Financial Services, and a Public Sector (including Education). The $F = 5$ fields of study are: General Studies, Technical Fields, Health Care, Humanities, and Business.

corresponding sector and field of study resulting from action k.[38] As $J = 7$ and $F = 5$, there are $K = 1 + J + F + F \cdot J = 48$ actions.[39]

Time-invariant worker heterogeneity comes through $q \in \{1, 2, 3, 4\}$ discrete worker types.[40] These are observable demographic groups - defined as gender and observed college attendance before age 25. Each type q differs in terms of their preferences and comparative advantage over work and schooling options.[41] Let $\mathscr{H}_i^{-t}$ denote the *history* of individual i's work and schooling choices prior to period t.

Timing: The timing is summarized in Figure 1.8 below. A worker begins period t in their previous industry $j_{i,t-1}$, previous idiosyncratic productivity $\eta_{i,j_{i,t-1},t-1}$ in that industry, prior schooling state $s_{i,t-1}$ ($s_{i,t-1} = 1$ if in school), and other continuous states Ω_{it} to be discussed in detail below. At the start of the period, individuals experience exogenous separations with probabilities depending on the their starting industry, $j_{i,t-1}$. Upon a separation, workers lose some job-specific skills and get shifted into non-employment. Workers then observe the idiosyncratic component of flow utility, ϵ_{it}^k, and make a choice k, taking into account expectations over realized idiosyncratic productivity η_{ijt} given current values of $\eta_{i,j_{i,t-1},t-1}$. After making a choice, workers realize the shock determining η_{ijt} and consume flow utility u_{qt}^k, which depends on their worker type q, time t, action k, as well as current states. At the end of the period, the state evolves, which includes skill accumulation as outlined in the following section.

[38] There are no restrictions on the choice set an agent has, however, agents will face switching costs to capture labor market frictions.

[39] I do not directly model intensive margin of labor supply. In my setting, *part-time* simply indicates whether the person is engaging simultaneous work and study. These intensive margin labor supply decisions are partly captured in the persistent η shocks, discussed below.

[40] Future work will seek enrich the degree of time-invariant worker heterogeneity by identifying latent worker types in a first-step clustering procedure using the EM algorithm. See Traiberman [2019] and Arcidiacono et al. [2024] for recent applications. These procedures pool information across periods to identify latent types, using persistent patterns in earnings and choices to identify the covariance between comparative advantage and preferences in an analogous manner to a fixed effects argument.

[41] Agents are also differentiated by their graduation cohort. I leave this as implicit, but in practice each of the five cohorts is a different age in each year t. This necessitates solving the model and evaluating choice probabilities separately for each cohort. For tractability, I only estimate the choice model for the 1994 and 1998 cohorts, but use all cohorts in earnings estimation and counterfactuals.

1.4.1 Earnings and Human Capital Dynamics

Earnings: If individual i in period t makes the choice k to work in industry j (i.e, $j = j(k_{it}) > 0$), their log-earnings are decomposed into the following components:

$$y_{ijt}^k = \underbrace{r_{jt}}_{\text{Skill Prices}} - \underbrace{\alpha^{PartTime} D\left[f(k) > 0\right]}_{\text{Part-Time Penalty}} + \underbrace{\sum_{s=1}^{S} \pi_j^{(s)} \omega_{it}^{(s)}}_{\text{Transferable Skills}} + \underbrace{\alpha_j^{\tau} \tau_{ijt}}_{\text{Job-Specific Skill}} + \underbrace{\eta_{ijt}}_{\text{Idiosyncratic Productivity}} \tag{1.4.1}$$

Skill prices r_{jt} are industry-time fixed effects in log-earnings that measure the return to human capital in industry j at time t, capturing time-varying industry conditions that affect wages.[42] These are equilibrium objects that may change due to aggregate or industry-wide productivity shocks or changes in the availability of inputs. I do not directly model changes in the non-stationary wage-setting equilibrium.[43] I discuss the information agents have on aggregate states below. Agents who both work and study face a part-time penalty $\alpha^{PartTime}$, capturing that workers may be less able to supply human capital to the labor market while enrolled.

Transferable Skills: My approach to modeling human capital is quite general: worker skills are neither completely general nor completely specific, with the degree of transferability informed by the data through observed transitions. I adopt a skill weights approach to human capital in the spirit of Lazear [2009]. This framework nests multiple

[42]Restricting skill weights to be positive treats human capital as a Cobb-Douglas production function across model objects. Exponentiating yields:

$$Y_{ijt}^k = R_{jt} \left(\exp\left\{\alpha^{PartTime} D\left[f(k) > 0\right]\right\} \prod_{s=1}^{S} \left(\exp\left\{\omega_{it}^{(s)}\right\}\right)^{\pi_j^{(s)}} \left(\exp\left\{\tau_{ijt}\right\}\right)^{\alpha_j^{\tau}} \exp\left\{\eta_{ijt}\right\} \right) = R_{jt} H_{ijt}^k$$

where R_{jt} is the skill price such that $r_{jt} = \log\left(R_{jt}\right)$ and H_{ijt}^k is human capital supplied to industry j.

[43]Skill prices and separation rates, objects that vary at the industry-time level, should be interpreted as reduced-form for the non-stationary equilibrium that a general equilibrium model would micro-found.

models of human capital.[44] Relative to prior work that adopts related approaches,[45] I do not use measures of skill directly. Using a fully data-driven approach that is based on observed transitions allows me to capture components of human capital not measured in such data. Since skill measures are typically only available for occupations, this approach also enables me to overcome limitations of industry-level data as well as incorporate fields of study.[46]

I model earnings as an industry-specific weighting of S individual-specific latent skills.[47]. Individual i begins a period with latent skill $\omega_{it}^{(s)}$ which gets weight $\pi_j^{(s)}$ in industry j. I choose $S = 3$. Skill weights are restricted to be positive, $\pi_j^{(s)} > 0$, and time-invariant.[48] Based on their current vector of skills $\vec{\omega}_{it}$, a worker has comparative advantage over sectors as skills are portable across industries. All else equal, workers prefer working in the sector that gives the highest weight to their skills.

In addition to these industry-specific returns, each skill is also differentiated in its accumulation process. Current choice k impacts the future values of skills, which change

[44] For example, if $S = J$ and $\Pi = I$, then skills are perfectly non-transferable. In contrast, if $\Pi = \iota_J$ and $S = 1$, then there is only general human capital. My approach quantifies the intermediate cases, under the restriction that $S = 3$. This number is motivated by the fact that the prior work that uses skill measures typically considers two or three dimensions of skills. For example. Deming [2017] considers the returns to cognitive, non-cognitive, and social skills. The choice of $S = 3$ is quite flexible, but is still a substantive restriction.

[45] Some examples are Yamaguchi [2012], Roys and Taber [2019], and Lise and Postel-Vinay [2020].

[46] A restriction of this model is that skill weights are time-invariant. This means that the process by which past experience gets translated into human capital in each industry is constant, with time-varying changes in how this experience is rewarded being captured in the skill price. A literature has documented that skill demands may have increased after the Great Recession (e.g. Hershbein and Kahn [2018], Blair and Deming [2020], Modestino et al. [2020]). These patterns will be captured in systematic changes in the skill prices r_{jt}. This abstraction can be made more realistic in future work by increasing the number of industries workers can choose from.

[47] I use the term latent in that the econometrician does not observe skills prior to estimation. In particular, given a guess of parameters, skills in my context are a deterministic functions of past work histories. This approach is distinct from those that estimate factor models using skill data. For example, in Yamaguchi [2012], a stochastic process governs how skills change over time, so individual skills are not known even given a guess of parameters. In his context, the accumulation process must be disciplined using observed occupation-level measures of skill and an individual's time-varying skills are backed out using the Kalman filter.

[48] One limitation is that the model does not account for occupation or task information. These are partially accounted for through the $\pi_j^{(s)}$. One reason industries may have similar weights is that they have similar occupation distributions, making work experience more transferable between them.

with schooling and work experience and depreciate with non-employment. Each skill has a worker type-specific technology of accumulation, $\Gamma_q^{(s)}()$. Each worker type q and latent skill s follows its own recursive law of motion:

$$\omega_{i,t+1}^{(s)} = \Gamma_q^{(s)}(k|\omega_{it}^{(s)})\omega_{it}^{(s)}$$

where $\Gamma_q^{(s)}$ captures the percent growth in skills and is flexibly parameterized as:

$$\log \Gamma_q^{(s)}(k|\omega_{it}^{(s)}) = \begin{cases} \log\left(\underbrace{1-\gamma_{Depr}^{(s)}}_{\textbf{Depreciation}}\right) & \text{if } j(k)=0 \text{ and } f(k)=0 \text{ (non-employment)} \\ \underbrace{\exp\left\{-\gamma_{Dim:School}^{(s)}(\omega_{it}^{(s)}-1)\right\}\left(\gamma_{School:q}^{(s)}+\gamma_{School:f(k)}^{(s)}\right)}_{\textbf{Education Investments}} & \text{if } j(k)=0 \text{ and } f(k)>0 \text{ (full-time study)} \\ \underbrace{\exp\left\{-\gamma_{Dim:Work}^{(s)}(\omega_{it}^{(s)}-1)\right\}\left(\gamma_{Work:q}^{(s)}+\gamma_{Work:j(k)}^{(s)}\right)}_{\textbf{Learning-by-doing}} & \text{if } j(k)>0 \text{ and } f(k)=0 \text{ (full-time work)} \\ \underbrace{\gamma^{PartTime}\exp\left\{-\gamma_{Dim:School}^{(s)}(\omega_{it}^{(s)}-1)\right\}\left(\gamma_{School:q}^{(s)}+\gamma_{School:f(k)}^{(s)}\right)}_{\textbf{Part-Time Education Investments}} & \text{if } j(k)>0 \text{ and } f(k)>0 \text{ (part-time study)} \\ +\underbrace{\left[1-\gamma^{PartTime}\right]\exp\left\{-\gamma_{Dim:Work}^{(s)}(\omega_{it}^{(s)}-1)\right\}\left(\gamma_{Work:q}^{(s)}+\gamma_{Work:j(k)}^{(s)}\right)}_{\textbf{Part-Time Learning-by-doing}} & \end{cases} \tag{1.4.2}$$

This flexible skill accumulation function captures several sources of differentiated human capital investment. First, $\gamma_{Work:j(k)}^{(s)}$ and $\gamma_{School:f(k)}^{(s)}$ capture how work experience in sector j and studying field f modifies each latent skill. Some fields of study may be useful in some industries but not in others. For example, it may be the case that studying health increases one's earnings as a nurse, but these skills will not be very useful in construction. This would be captured in a high $\gamma_{School:f}^{(s)}$ for nursing and a high $\pi_j^{(s)}$ for health, but a low $\pi_{j'}^{(s)}$ for construction. Second, $\gamma_{School:q}^{(s)}$ and $\gamma_{Work:q}^{(s)}$ capture comparative advantage to investing in human capital through school or work experience by worker type q on each skill margin. For example, college-educated workers may have higher returns to experience and this pattern may be more salient in finance than construction. Third, $\gamma^{PartTime} \in (0,1)$ captures how much skill accumulation is from work compared

to school when studying part-time. If workers are only working or only studying, they receive the full benefits of learning through either work or study but otherwise receive a convex combination. Last, $\gamma^{(s)}_{Dim:School} > 0$ and $\gamma^{(s)}_{Dim:Work} > 0$ capture diminishing benefits to further investing in human capital as skills increase, which may differ by type of skill as well as by schooling or work.[49]

Specific Skill: Job-specific skill τ_{ijt} captures the return to remaining continuously employed. Importantly, this object is entirely lost upon a job separation, grows identically in each industry, and is not modified through schooling. Like transferable skills, this skill receives industry-specific weights and accumulates over time. Unlike transferable skills, the specific skill is memoryless in that the specific sequence of sectors does not matter for its level. This mechanism is meant to capture match effects with current employer, firm-specific human capital, and job-shopping dynamics in a reduced form relationship that is held fixed in the counterfactual (Bagger et al. [2014]).

To capture these job-shopping and job-ladder dynamics, the job specific skill is allowed to grow both within and across periods, depending on if the worker makes a voluntary transition across industries or an involuntary transition into non-employment. If an individual experiences an involuntary separation from their previous job, $sep_{it} = 1$, τ is completely lost: $\tau_{ijt} = 0$ and subsequently grows depending on whether the individual decides to return to work. If an individual is not separated and makes a voluntary switch, τ grows within period $\tau_{ijt} = \gamma^{(\tau)}_{Vol}\tau_{i,j_{i,t-1},t}$. Between periods, τ grows according to the accumulation function $\Psi(k|\tau_{ijt})$, where growth is normalized to unity if an individual starts from $\tau = 0$ due to a voluntary or involuntary switch out of non-employment.

[49] Note that $\gamma^{(s)}_{School:f(k)}$, $\gamma^{(s)}_{Work:j(k)}$, $\gamma^{(s)}_{Work:q}$, and $\gamma^{(s)}_{School:q}$ are not sign-restricted, so individuals can lose skills from taking certain actions.

To summarize, the within-period growth due to transitions is:

$$\tau_{ijt} = \begin{cases} 0 & \text{if } sep_{it} = 1 \quad \textbf{(involuntary separation)} \\ \gamma^{(\tau)}_{Vol}\tau_{i,j_{i,t-1},t} & \text{if } sep_{it} = 0 \text{ and } j(k) \neq j_{i,t-1} \quad \textbf{(voluntary switch)} \\ \tau_{i,j_{i,t-1},t} & \text{if } sep_{it} = 0 \text{ and } j(k) = j_{i,t-1} \quad \textbf{(remain)} \end{cases} \tag{1.4.3}$$

and the across-period growth follows:

$$\tau_{ij,t+1} = \begin{cases} 0 & \text{if } j(k) = 0 \quad \textbf{(non-employment)} \\ 1 & \text{if } \tau_{ijt} = 0 \text{ and } j(k) > 0 \quad \textbf{(entry-level employment)} \\ \Psi(k|\tau_{ijt})\tau_{ijt} & \text{if } \tau_{ijt} > 0 \text{ and } j(k) > 0 \quad \textbf{(experienced employment)} \end{cases} \tag{1.4.4}$$

where the function $\Psi(k|\tau)$ uses a functional form that mirrors the law of motion for transferable skills:

$$\log \Psi(k|\tau) = \gamma^{(\tau)}_{Accum} \exp\{-\gamma^{(\tau)}_{Dim}(\tau - 1)\} \tag{1.4.5}$$

Taken together, these equations mean that specific skill τ_{ijt} is a function of the duration of continuous employment and timing of voluntary transitions.[50]

Persistent Idiosyncratic Productivity: The structure on idiosyncratic productivity η_{ijt} allows for persistent heterogeneity in earnings, both weakening identification requirements and capturing patterns of selection documented as important in the training literature. Idiosyncratic productivity η_{ijt} follows the AR(1) process:

$$\eta_{ijt} = \rho^{j}_{j_{i,t-1}} \eta_{i,j_{i,t-1},t-1} + \varphi_{i,j,j_{i,t-1},t} \tag{1.4.6}$$

[50] A key identifying assumption will be that I can distinguish involuntary and voluntary employment-to-employment transitions across industries (in a way that matters for earnings), which is accomplished by using quarterly data. The normalization of $\tau_{ij,t+1} = 1$ if $\tau_{ijt} = 0$ will separate the level of τ_{ijt} from the levels of α^{τ}_{j}, the industry-specific return to τ_{ijt}.

where $\varphi_{i,j,j_{i,t-1},t} \underset{\text{i.i.d.}}{\sim} N(0, \sigma^{j}_{j_{i,t-1}})$.[51] The persistence parameters, ρ, and variance of the innovations, σ, both depend on chosen in industry j as well as last activity $j_{i,t-1}$ through the type of work transition.[52] Specifically, parameters governing the variance in the innovations differ by the target industry j and whether a worker remains in their current industry (*Remain*), switches industry voluntarily (EE), or switches out of non-employment (EN): σ_j^{Remain}, σ_j^{EE}, σ_j^{EN}. For tractability, the persistence parameters ρ vary by only the type of transition: $\rho^{Remain}, \rho^{EE}, \rho^{EN}$.

The agent knows previous productivity $\eta_{i,j_{i,t-1},t-1}$ at the time of choosing k but $\varphi_{i,j,j_{i,t-1},t}$ is realized after the individual's choice is made. This timing assumption aids greatly in the tractability of the model while still allowing selection in response to shocks, which is needed in order to capture the potential dips earnings prior to enrollment (Ashenfelter [1978]). Specifically, individuals with negative η_{ijt} will respond by taking low-ρ actions while those with high η_{ijt} will take high-ρ actions, resulting in positive values to realized to η_{ijt} over time. Specifically, at estimated parameters, this will mean that individuals will be more likely to switch in response to negative draws of $\varphi_{i,j,j_{i,t-1},t}$.

Initial Conditions: Observables prior to age 25 are used to parameterize initial conditions. This captures that workers may be different in their pre-sample work experiences which may affect their payoff to working in certain industries. These are modeled as:

$$\log\left(\omega_{i1}^{(s)}\right) = z_i' \gamma_{Z:1}^{(s)} + \gamma_{q:1}^{(s)}$$

where z_i is a vector of these initial observables with $\gamma_{Z:1}^{(s)}$ a vector of coefficients for each skill s and dummies for worker type $\gamma_{q:1}^{(s)}$. I observe full work history prior to age 25 for high school graduates. Elements of z_i are indicator variables for most recent

[51] In practice, all that is needed for estimation of the earnings parameters is the mean-independence of $\varphi_{i,j,j_{i,t-1},t}$. Full independence and the normal distributional assumption will be used in choice estimation and counterfactual simulations.

[52] These work transitions are determined by both target j and previous activity $j_{i,t-1}$, necessitating the need for cumbersome notation.

industry of employment in age 23 or 24 (including none). I impose the normalization that $\mathbb{E}\left[\omega_{i1}^{(s)}\right] = 1$ in the population, which helps fix the location and scale of $\omega_{it}^{(s)}$. Normalizations are discussed more in Section 1.5.

For other initial conditions, $j_{i,t-1}$ in $t = 1$ is set to the the same recent industry of employment that determines z_i. Initial specific skill $\tau_{i,j_{i,t-1},t}$ in this industry is determined by Equation 1.4.5 based on observed tenure. The initial persistent shock is drawn independently with zero mean and a variance parameter specific to the first period: $\eta_{ij1} \underset{\text{i.i.d.}}{\sim} N(0, \sigma_j^0)$.

1.4.2 Utility, Choices, and Beliefs

I now describe the behavioral model underlying the earnings dynamics outlined above. This allows us to predict counterfactual worker choices and relate them to changes in earnings.

A worker's choice is determined by the state in which they start the period. In addition to previous activities $j_{i,t-1}$ and $s_{i,t-1}$, these include continuous states Ω_{it}. These continuous states include current transferable skills $\omega_{it}^{(s)}$ $\forall s$, specific skill $\tau_{i,j_{i,t-1},t}$ in the industry at the start of the period, previous period's income $Y_{i,t-1}$, and age Age_{it}: $\Omega_{it} = \{\omega_{it}^{(1)}, \omega_{it}^{(2)}, \omega_{it}^{(3)}, \tau_{i,j_{i,t-1},t}, Y_{i,t-1}, Age_{it}\}$.

Flow Utility: An agent of type q receives the following flow utility from taking action k in year t:

$$u_{qt}^k(j_{i,t-1}, s_{i,t-1}, \Omega_{it}; \vec{\epsilon}_t, \eta_{ijt}) = \underbrace{\mu_q^k(Age_{it})}_{\textbf{Preferences for Actions}} + \underbrace{\beta_Y y_t^k(\Omega_{it}, j_{i,t-1}; \eta_{ijt})}_{\textbf{Preference for Log-Income}} \tag{1.4.7}$$

$$- \underbrace{\beta_P Price_t^k(Y_{i,t-1})}_{\textbf{Sensitivity to Tuition Costs}} - \underbrace{C_{qt}^k\left(j_{i,t-1}, s_{i,t-1}\right)}_{\textbf{Switching Cost}} + \underbrace{\epsilon_t^k}_{\textbf{EV1}} \tag{1.4.8}$$

$$= \bar{u}_{qt}^k(j_{i,t-1}, s_{i,t-1}, \Omega_{it}; \eta_{ijt}) + \epsilon_t^k \tag{1.4.9}$$

which depends upon type-specific non-pecuniary benefits from taking action k, $\mu_q^k()$;[53] preference for labor market earnings, y_t^k; disutility from net tuition costs that vary by full- or part-time study; switching costs $C_{qt}^k()$; and the vector of extreme-value shocks, $\vec{\epsilon}_t$.[54]

Tuition and Financial Aid Functions: I take sticker prices for attending college from the Integrated Post-Secondary Education Data System (IPEDS), accounting for differences in full and part-time tuition.[55] I incorporate administrative financial aid data from THECB in order to get realistic measures of price sensitivity needed for the main counterfactuals, modeling grant aid as a flexible function of the prior year's income $Y_{i,t-1}$.[56] I treat negative values of tuition net of aid as zeros, meaning that individuals with more aid than tuition will not be sensitive to price although still face a cost of switching into school.

Switching Costs: Frictions in transitioning across industries and out of non-employment are captured in switching costs. Switching into non-employment or remaining in the same industry is costless. Let $E(j_{i,t-1}) = D[j_{i,t-1} > 0]$ denote whether a worker starts in an employed state and $U(j_{i,t-1})$ whether they start in a non-employed state, such that $E(j_{i,t-1}) = 1 - U(j_{i,t-1})$. If a worker makes choice k to switch industry (i.e., $j(k) \neq j_{i,t-1}$),

[53] In practice, the function $\mu_k^q(Age_{it})$ consists of indicators for each industry, indicators field of study, an indicator for part-time study, and quadratic slopes in age for being employed full-time. All coefficients are specific to worker type q

[54] I normalize the variance of ϵ to 1 to fix scale. To fix the location of utility, choosing non-employment has utility of 0. As agents prefer log-income, they will prefer to have predictable streams of income over time.

[55] I total reported full-time full-year tuition, fees, and books/supply costs for each two-year college in Texas and average these using full-time equivalent enrollment for that academic year. To create part-time prices, I convert these into a per-credit cost by dividing this average cost by 30 credits, which is equivalent to attending full-year and full-time enrollment. Then, I assume that students enrolled part-time enroll in 18 credits which is close to the observed average for those who work and enroll at the same time. I do this rather than using cost measures from the financial aid data as the prices from IPEDS lined up with reported per-credit tuition costs reported on community college websites.

[56] I model aid as splines with kinks at \$10,000 and \$30,000 and allow estimated aid schedules to vary across groups of years associated with significant policy changes: 2001-2008, 2009-2012, 2013-2019. Agents prior to the Recession expect the 2001-2008 functions to last indefinitely and are surprised by changes in aid in 2009-2012. For tractability, agents know the 2013-2019 aid schedules and expect it to continue indefinitely. I treat these schedules as a calibration and do not adjust standard errors for them. The estimated aid functions are shown in Figure D3.

they incur the following disutility:

$$C_{qt}^{k}\left(j_{i,t-1}, s_{i,t-1}\right) = \exp\{\underbrace{\kappa^{U}U(j_{i,t-1}) + \kappa^{E}E(j_{i,t-1})}_{\textbf{Employment Status}} + \underbrace{\kappa_{q}}_{\textbf{Worker Type}} + \underbrace{\kappa_{j(k)}}_{\textbf{Target Industry}} + \underbrace{\kappa_{FromSchool}D\left[s_{i,t-1}=1\right]}_{\textbf{From School}}$$
$$+ \underbrace{\kappa_{GR}^{U}U(j_{i,t-1})D\left[t \geq 2009, t \leq 2011\right] + \kappa_{GR}^{E}E(j_{i,t-1})D\left[t \geq 2009, t \leq 2011\right]}_{\textbf{Business Cycle}}\} \quad (1.4.10)$$

The parameters $\kappa^{U(j_{-1})}$ and $\kappa^{E(j_{-1})}$ allow switching costs vary by employment status at the start of the period, capturing that search frictions may vary in unemployment. These are further allowed to vary over the business cycle through $\kappa_{GR(t)}^{E(j_{-1})}$ and $\kappa_{GR(t)}^{U(j_{-1})}$, increasing temporarily and unexpectedly between 2009 to 2011. The parameters κ_j allows this friction to vary by target industry, measured in differences relative to the general sector. I also allow worker types to differ in their switching rate κ_q, relative to non-college men. Finally, I allow starting in a schooling state to potentially decrease these switching costs through $\kappa_{FromSchool}$, permitting returning to school to act as a form of job search as has been suggested to be important in the training literature (Heckman and Smith [1999]).

If workers do not start in school and switch into schooling, they incur disutility:

$$C_{qt}^{k}\left(j_{i,t-1}, s_{i,t-1}\right) = \exp\{\kappa_{IntoSchool}\} \qquad \text{if} \qquad s(k) \neq s_{i,t-1} \text{ and } s_{i,t-1} = 0$$

This feature captures difficulties individuals may face in returning to schooling that are not captured by direct tuition costs. For example, registering for courses may be especially costly when one first enrolls.

Exogenous Separations: Individuals starting in an employed state ($j_{i,t-1} > 0$) may be exogenously separated with a probability that varies by industry and year: δ_{jt}. Two important things happen upon separation. First, while they retain their vector of transferable skills ω_{it}, they lose job-specific skills $\tau_{i,j_{i,t-1},t} = 0$. Second, they face the switching costs of non-employed workers. Both exogenous separations δ_{jt} and negative shocks

to productivity η_{ijt} will capture dynamics generating the pre-program/Ashenfelter dip in earnings (Ashenfelter [1978]). Combined with unexpected changes in skill prices r_{jt}, these forces will drive training and transition decisions in response to both individual-level and aggregate shocks.

Beliefs on Aggregate States and Policy Regimes: I treat the Great Recession as a one-time, unanticipated shock. Agents are not myopic and use a forecasting rule for the aggregate states prior to the Recession. Practically, I linearly interpolate skill prices and separation rates for the years 2009 through 2015, assuming agents use these interpolated states to make their decisions prior to 2009.[57] By doing so, I treat agents as having biased beliefs about aggregate skill prices and separation rates in the sense that they are completely surprised by these changes. This is shown below for select industries in Figure 1.11. Similarly, policy changes to financial aid rules as well as increases in switching costs are modeled as unanticipated in the pre-period. It should be emphasized that, due to my two-step approach to estimation, any assumptions about beliefs do not affect estimates of earnings parameters.

Dynamic Problem: Agents make choices taking into account both current and future payoffs. Given current beliefs about aggregate states, this involves taking expectations over future separations δ_{jt}, preference shocks $\vec{\epsilon}_{it}$, and persistent shocks to idiosyncratic productivity η_{ijt}.

At the time of making a decision, workers have experienced exogenous separations based on $\delta_{j_{i,t-1},t}$ and know the value of the preference shocks $\vec{\epsilon}_t$, but do not know the realization of idiosyncratic productivity η_{ijt} in the period. Let $\hat{\eta}$ denote a potential realization drawn from the distribution $F^k(\hat{\eta}|\eta_{i,j_{i,t-1},t-1}, j_{i,t-1})$, depending on k through Equation 1.4.6. Similarly, let $\hat{\Omega}^k_q(\Omega_{it})$ be the future continuous states resulting from tak-

[57] An alternative would be to estimate an autoregressive process on skill prices in each industry and use this as a rational expectations forecasting rule (e.g. Lee and Wolpin [2006]). I do not pursue this approach as the number of sectors is large and the deviations from the trend would enter the state space separately. In addition, the length of the panel is short meaning that there are few observations to inform the process. I consider my approach as the best choice available.

ing action k for worker type q.[58] Net of logit shocks $\vec{\epsilon}_t$,[59] the value from taking action k at t for a worker of type q with $\hat{\eta}$ is:

$$v_{qt}^{k}(j_{i,t-1}, s_{i,t-1}, \Omega_{it}; \hat{\eta}) = \bar{u}_{qt}^{k}(j_{i,t-1}, s_{i,t-1}, \Omega_{it}; \hat{\eta}) + \underbrace{\delta_{j(k),t}\beta V_{q,t+1}\left(0, s(k), \hat{\eta}, \hat{\Omega}_{q}^{k}(\Omega_{it})\right) + \left(1 - \delta_{j(k),t}\right)\beta V_{q,t+1}\left(j(k), s(k), \hat{\eta}, \hat{\Omega}_{q}^{k}(\Omega_{it})\right)}_{\textbf{Expected Continuation Value}} \tag{1.4.11}$$

where $\beta = 0.96$ is the discount factor and $V_{q,t+1}()$ denotes the expected value of starting the next period in a state. The ex-ante value of action k involves integrating over potential innovations in earnings productivity $\hat{\eta}$:[60]

$$\bar{v}_{qt}^{k}(j_{i,t-1}, s_{i,t-1}, \eta_{i,j_{i,t-1},t-1}, \Omega_{it}) = \int_{\hat{\eta}} v_{qt}^{k}(j_{i,t-1}, s_{i,t-1}, \Omega_{it}; \hat{\eta}) dF^{k}(\hat{\eta} | \eta_{i,j_{i,t-1},t-1}, j_{i,t-1}) \tag{1.4.12}$$

Integrating over the preference shocks $\vec{\epsilon}$ gives the expected value of being in a state at time t:

$$V_{qt}(j, s, \eta, \Omega) = \gamma_{Euler} + \log\left(\sum_{k} \exp\{\bar{v}_{qt}^{k}(j, s, \eta, \Omega)\}\right) \tag{1.4.13}$$

The full timing of the model is illustrated in Figure 1.8.

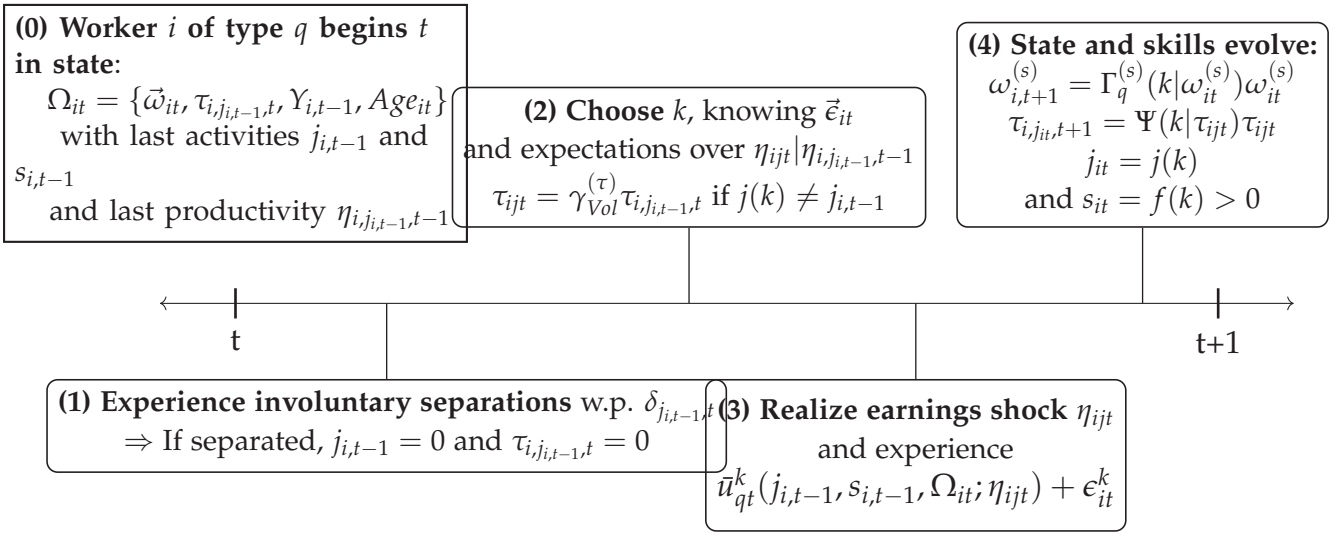
Approximate Solution: Solving this model is computationally difficult. The practical implication resulting from treating the Great Recession an unanticipated shock is that I need to solve the model under two belief regimes: the pre-Recession decision rules (through the whole life-cycle) and the post-Recession decision rules (from 2009 on-wards, when agents are surprised). In addition, I need to solve the model for each worker type q and graduation cohort. Moreover, the state space is large and has many continuous states. Due to the history dependence in skills, there is no renewal action or

[58] This is worker-type q specific as the skill accumulation process $\Gamma_q()$ is q-specific.

[59] The logit shocks are removed from $\bar{v}_{qt}^{k}$ as they are unaffected by realizations of $\hat{\eta}$. They re-enter when computing the expected value of starting in a state in Equation 1.4.13.

[60] I use Gaussian-Quadrature with 5 nodes to evaluate this expectation.

Figure 1.8: Timing of one period of the model



finite dependence argument that allows for estimation without solving the model (e.g. Arcidiacono and Miller [2011], Arcidiacono and Miller [2020]).

To make this problem computationally feasible, I solve the model through backwards recursion on a subset of possible states and approximate the value functions using regression interpolation.[61] Expected value functions $V_{q,t+1}(j, s, \eta, \Omega)$ are approximated using the regression:

$$\tilde{V}_{q,t+1}(j, s, \eta, \Omega) = x(\Omega, \eta, s)'\hat{\pi}^j_{t+1}$$

Thus, a solution to the agent's problem consists of all $\hat{\pi}^j_t$ governing decisions. For the design matrix $x(\Omega, \eta, s)$, I use all linear effects, pairwise interactions, and own-quadratic effects for all continuous states in Ω, excluding Age_{it}.[62] I include an additional fixed effect and linear interactions if the individual was in school in the previous period, $s_{i,t-1} = 1$. Letting $\hat{\pi}^j_{t+1}$ depend on j also implicitly interacts all of these covariates with the discrete employment state j at the start of the period. This creates 2,160 linear regressions that must be run to solve the model.

In choosing the number of states on which to evaluate the value function, I leverage

[61] Similar interpolation methods to quickly compute value functions through regression interpolation have been used in Keane and Wolpin [1997], Lee and Wolpin [2006], Sweeting [2013], and Dix-Carneiro [2014].

[62] Since I interpolate each cohort and time period separately, age is implicitly interacted with all coefficients as well.

the fact that the empirical distribution of states is observed after separately estimating the skill accumulation functions in a first-step.[63] I randomly sample 500 states, sampling from each continuous state's marginal distribution.[64] This approximation performs very well in practice, as shown in in Appendix Figure D1, where I vary the number of evaluation states on which to interpolate.

1.5 Identification

The key endogeneity problem the structure of the model helps address is the dynamic relationship between past industry, current enrollment choices, and future transitions in determining earnings. Identification relies on having rich panel data on earnings and sufficient variation in work and school transitions. The main assumption is the mutual independence of time-varying unobservables, which breaks estimation into two steps. These assumptions are common in the dynamic discrete choice literature dating back to Rust [1987]. The persistent process on η_{ijt} substantially weakens the implications of these assumptions by allowing persistence in η_{ijt} that is correlated with skills while preserving tractability.

1.5.1 Key Assumptions

Assumption 1.5.1. *(Information Sets) The agent's information at the time of choosing k consists of their current state, $\{\Omega_{it}, j_{i,t-1}, s_{i,t-1}, \eta_{i,j_{i,t-1},t-1}\}$, including whether the previously employed state $j_{i,t-1}$ measures realized separations; the currently expected paths of policy regimes, switching costs, separation rates δ_{jt}, and skill prices, r_{jt}; and the unobserved preference shock, $\vec{\epsilon}_t$. In particular, the shock to productivity η_{ijt} is not known at the time of making a choice and agents*

[63] I divide each continuous state into quintiles and add additional states that are higher than the maximum observed value. I draw states at which to evaluate the value function by first sampling from these bins and then sampling uniformly within bin. I repeat this for the marginal distribution of each state.

[64] To account for states outside of this observed distribution (which agents may take into account when making choices), I make sure to sample slightly outside the observed range.

take this into account by taking expectations over potential realizations, conditioning on previous $\eta_{i,j_{i,t-1},t-1}$.

Assumption 1.5.2. *(Unobserved Shock to Utility) Conditional on the state, the preference shocks* $\vec{\epsilon}_t$ *are identically distributed extreme value type 1 and independent over both individuals and time periods.*

These assumptions create separability of the earnings and choice models, which buys tractability while also permitting the model a richness that would otherwise be difficult. In particular, the earnings equation is semi-parametrically identified, meaning that estimated earnings parameters are not sensitive to the structure of the choice model. Since choices are made without knowing independently drawn innovations in productivity η_{ijt}, the full distribution of η_{ijt} is revealed *ex post* in the microdata.[65]

Another issue is whether transitions are voluntary or involuntary. To account for this in a tractable way, I add the following assumption:

Assumption 1.5.3. *(Observed Involuntary Employment-to-Employment Transitions) An involuntary separation is observed,* $sep_{it}^{Obs} = 1$*, if the employed individual experiences one or two quarters of non-employment before switching back to an employed state. If the worker immediately re-enters an employed state, involuntary separations only affect earnings through the specific skills* τ*, as the worker does not experience skill depreciation.*

[65] It is worth comparing these assumptions to alternative setups of the model. First, one could allow idiosyncratic productivity to be drawn in each industry in every time period and be known to the agent at the time they made decisions. In this case, one could in principle allow η and ϵ to be correlated by drawing them from a jointly multivariate normal distribution. This is the approach taken in papers such as Lee and Wolpin [2006] and Dix-Carneiro [2014]. I do not pursue this route for two reasons. First, my inclusion of part-time study blows up the choice space, making the dimensionality of ϵ too large for this procedure to be practical. Second, it would also mean estimation would need to occur in a single step, making computation more difficult.

A different potential setup would be to estimate a model of search, matching, and bargaining (e.g. Postel-Vinay and Robin [2002], Taber and Vejlin [2020], Huckfeldt [2022]). Bargaining with current employers in response to unobserved outside wage offers would also generate correlation between unobserved utility shocks (e.g.,through realized offer probabilities) and unobserved productivity (e.g., through bargaining outcomes). I do not pursue this setup as bargaining is not the main focus of the paper and would necessitate a loss of richness in other dimensions to remain traceable.

While the model is estimated at the annual level, Assumption 1.5.3 leverages the quarterly data to enrich unemployment dynamics. An involuntary switch into non-employment where the individual does not return to work affects earnings in the same way as a voluntary switch into non-employment, as τ_{ijt} is always lost whenever individuals enter non-employment. This assumption abstracts from the earnings impact of short unemployment spells of less than a quarter. One may think of any permanent effects on earnings resulting from short unemployment spells to be captured by the process governing idiosyncratic productivity η_{ijt}.[66]

1.5.2 Margins of Selection and Implied Restrictions on Worker Mobility

There are several well-known identification challenges to estimating the returns to education that are captured in the framework. The two most pressing concerns in estimating returns to education are (1) ability bias, that individuals may be positively selected on earnings in the absence of enrollment, and (2) the pre-program/Ashenfelter dip in earnings, the degree to which individuals are selected on a possibly persistent negative shock. The second source of bias makes comparisons using lagged earnings challenging, as the degree of persistence in the adverse shock affects the appropriate comparison group.[67]

From the perspective of the model, levels of skills ω_{it} and τ_{ijt} and idiosyncratic productivity η_{ijt} proxy for worker ability. Selection on the transitory component of earnings is captured by shocks to productivity η_{ijt}, changes in aggregate skill prices r_{jt}, and re-

[66] In Figure 1.13, I show below that the model reasonably matches the earnings patterns of workers separated during the Great Recession.

[67] For example, if shocks are perfectly persistent, the appropriate comparison group are individuals who have the same pre-enrollment dips in earnings as the treatment group. Failure to make this correct comparison will result in negative estimated returns. If shocks are entirely transient, the appropriate comparison group are individuals with similar earnings prior to the observed dip for enrollees. Failure to make this correct comparison will result in a positive bias, as individuals will be erroneously compared to those with low permanent earnings. See Table 2 in Ashenfelter and Card [1985] for an excellent illustration of how differing base periods affect estimates of training. See also Ashenfelter [1978], Heckman and Robb [1985], Heckman and Smith [1999].

alized separations δ_{jt}. Each of these objects create forces that may drive an individual to enroll in response to these shocks when enrollment may not have been the optimal choice in prior periods. By estimating the parameters of this autoregressive process, I am quantifying the degree of persistence of these shocks directly from the data.[68]

As a result of the structure on the unobservable terms, the crucial identifying assumption becomes a timing assumption: individuals cannot select on within-period changes in the error term in the earnings equation. However, due to the rich structure of the earnings model, individuals can select on their entire work and schooling history and a prior period's earnings. This means that identification comes from movers (or enrollees), with past earnings and work histories proxying for selection in a relatively flexible way that allows for match effects and persistent shocks to earnings.

Intuitively, given individuals with identical work and schooling histories up to a point, the model imputes average earnings changes upon an industry switch (or into enrollment) as the counterfactual earnings for stayers (non-enrollees). The estimation routine outlined in Section 1.6 makes many such comparisons in practice, finding persistent patterns in earnings changes associated with different transitions occurring at different levels of work experience. I prove in a simple case in Appendix Section C.1 how to identify latent skills from work transitions.

The key moment condition justifying my approach to estimation of the earnings and skill parameters is the mean independence of innovations in the error conditional on

[68] Ashenfelter and Card [1985] estimate the training effect by modeling the autocorrelation structure of comparison worker earnings and using this to form predictions about counterfactual trainee earnings. The key insight is that trainees are different on this transitory shock but the researcher can use the covariance structure of earnings and a selection rule in terms of pre-training earnings to forward predict the contribution of this transient shock to post-training earnings. We can think of identification in my model as related in that I am jointly estimating the effect of schooling on skills at the same time as the autocorrelation parameter. The autocorrelation parameters ρ in my model gives expected earnings in subsequent periods in terms of this transient component. Moreover, the selection rule in my context is micro-founded by the dynamic discrete choice model.

other attributes of the model:

$$\mathbb{E}[\varphi_{ij,j_{i,t-1},t}|k_{it}, X_{it}, t] = 0 \tag{1.5.1}$$

where X_{it} is conditioning set consisting of the entire sequence of past employment histories $\mathscr{H}_i^{-t}$, prior period earnings $Y_{i,t-1}$, determinants of initial conditions z_i, and the worker types q_i are allowed to drive selection. This is implied by Assumptions 1.5.1 and 1.5.2 as well as the following restriction on worker mobility patterns: $P_t(k_{it} = k|\varphi_{ij,j_{i,t-1},t}, X_{it}) = P_t(k_{it} = k|X_{it})$. Importantly, φ is allowed to drive selection, but with a lag. This is key to incorporating Ashenfelter [1978] dynamics.

Formally, the assumptions of the structural model imply that $\mathbb{E}[Y_{it}^k|k, X_{it}] = \mathbb{E}[Y_{it}^k|X_{it}]$. Similar assumptions is also invoked in the voluminous literature that uses fixed effects and matching approaches to estimate returns to education in similar panel data contexts.[69] It is also related to the literature on worker and firm decompositions of earnings (Abowd et al. [1999], Card et al. [2013]), specifically those that allow for match effects and persistent heterogeneity (e.g. Bonhomme et al. [2019]). Once we condition on X_{it}, individuals are in the same state and have the same expected earnings from making choice k. Conditional on X_{it}, differences in choices are only driven by the preference shocks, $\vec{\epsilon}_t$, which are independent from all other objects in the model by Assumption 1.5.2.[70]

An implication of the timing assumption is directly assessed in the data in Figure 1.9,

[69] See Jepsen et al. [2014], Stevens et al. [2019], and Altonji and Zhong [2021] for recent examples. The training literature has extensively evaluated this assumption in the context of matching estimators that used lagged earnings histories, finding that these non-experimental approaches replicate experimental estimates (Heckman et al. [1998a], Dehejia and Wahba [1999], Card et al. [2018]) as noting that adding additional variables to the conditioning set matters little beyond lagged earnings histories (Andersson et al. [2013], Caliendo et al. [2017]).

What the structure of the model adds relative to these research designs is a framework for characterizing mechanisms and predicting how workers change their industry and schooling choices in counterfactual environments.

[70] An extension to allow for time-invariant unobserved heterogeneity as in an Arcidiacono and Miller [2011] procedure would these latent worker types to X_{it}. This implicitly adds *future* choices to the conditioning set under the assumption of a finite number of types. I plan to extend my estimation procedure to incorporate this in future work.

which supports the notion that the estimated parameters will likely purged of the most important sources of selection bias. Observables from outside the conditioning set X_{it} should not add information about the potential outcome Y_{it}^k. For example, this would be the case if enrollees were an unobserved type that is not accounted for in X_{it}, in which case future choices would reveal information about the current type. I compare two individuals who should have identical earnings from the perspective of the model based on X_{it} prior to $t = -2$ and make the same choice in $t = -2$, but differ in their future enrollment choices at $t = 0$. For each new enrollee, I construct a matched comparison based on this exercise.[71] Specifically, the model predicts that enrollees and matched comparisons should have identical earnings in the gray-shaded area of Figure 1.9. There is a small difference in earnings in this period on the order of \$964 in non-Recession years and \$802 in Recession years. This difference is statistically significant in non-Recession years. This indicates that there may be a small amount of bias in estimated parameters resulting from negative selection of enrollees into schooling, but overall supports that the structure of the model is supported by the data.[72][73]

I also implement matching estimators based on a much richer set of covariates than what is able to be taken account of in the model. While these results abstract from the policy-relevant population by not modeling choices, they indicate whether the results of negative returns for marginal enrollee are plausible. Mirroring the results from the model, I estimate average treatment effects on the treated that are close to zero and exhibit substantial heterogeneity. These results are shown in in Appendix Section A.2.

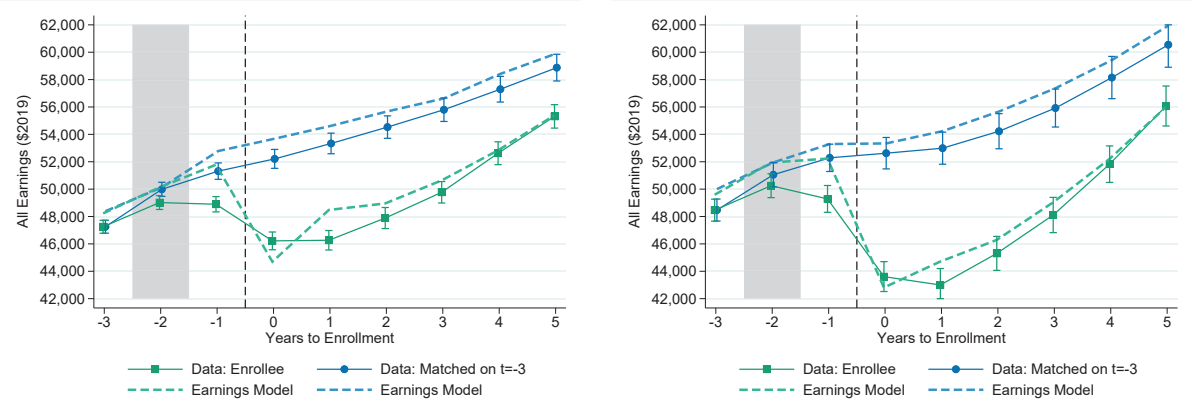
[71] To account for the role of auto-correlated productivity, I find the individual who has the closest earnings within a common history, restricting to comparisons within a \$1000 difference of the enrollee earnings in $t = -3$

[72] This figure also shows the fit of the earnings model on this sample, holding choices fixed and forecasting future earnings based on $\eta_{i,j_{i,t-1},t-1}$.

[73] As these are forecasts based on lagged states, the apparent lack of model fit in $t = -1$ is not surprising as this gap reflects that individuals have negative realizations of the innovation in productivity prior to enrollment $\varphi_{i,j,j_{i,t-1},t}$.

Figure 1.9: Check of Selection Assumptions

(a) Non-Recession (b) Enroll in 2009-2010



Notes: This figure uses observed earnings histories to check the implication of the independence of changes in time-varying unobservables, as implied by Assumptions 1.5.1 and 1.5.2. I match individuals based on information used in the model prior to $t = -2$ and stratify earnings histories based on enrollment choices made in $t = 0$. An implication of the timing assumptions is that future enrollment decisions should not affect current earnings, as individuals are in identical states based on past histories. For each enrolling individual, I match them to a non-enrollee that has identical work and schooling histories up until $t = -3$ (at the coarseness of the model), makes the same labor supply decision in $t = -2$, and is restricted to the closest comparison with earnings within \$1,000 bandwidth of earnings in $t = -3$. Conditional on these observables, individuals have the same earnings in $t = 2$ from the perspective of the model. Dashed lines show predicted earnings from the model at estimated parameter values, predicting earnings using $\eta_{i,j_{it-1},t-1}$ and current observed choice $j(k_{it})$, ignoring the contribution of current innovations.

1.5.3 Identification of Key Parameters

I now turn to a more focused discussion concerning identification of specific parameters. These arguments are summarized in Table 1.3. As is well known, dynamic discrete choice models cannot generically distinguish between flow utility and continuation values (Magnac and Thesmar [2002]). I fix the discount factor to $\beta = 0.96$, normalize flow utility to zero for non-employment at age 30, and fix the variance of idiosyncratic preference shocks to 1.

Latent Skills: Skills are unobserved quantities prior to estimation. Just like utility, we need to fix their location and scale to recover them from the data. Observably high earnings may be rationalized by high skill prices r_{jt}, by high skills $\omega_{it}^{(s)}$, or by high returns to skills $\pi_j^{(s)}$. I discuss the identification arguments behind these components at a high-level here and provide semi-parametric identification proofs for a simple version of the model in Appendix Section C.1. The bottom line of these proofs is that Γ and Ψ are in principle directly recoverable from the data given the mean-independence assumptions on φ, additive separability, and a sufficient number of informative moves.

Scale is fixed by assuming that there are fewer skills than industries, that skills evolve deterministically, and that skills weights are invariant over time. This allows cross-industry but within-work history comparisons of average earnings to identify the skill accumulation process. For example, comparing earnings differences across industries for individuals with the same initial z_i identifies skill weights, since workers with the same z_i have the same first-period skills by assumption. Then, location is fixed by normalizing each initial skill to be one in the population, $\mathbb{E}[\omega_{i1}^{(s)}] = 1$ for each s, separating the level of skills from the level of skill prices.[74] The match component τ has similar scale and location issues, which are fixed through its growth process as exhibited in Equation 1.4.5. Specifically, τ goes to zero upon non-employment, which fixes location, and one period of tenure in an industry sets it to unity, which fixes scale.

Implicit in these arguments and those formally outlined in Appendix Section C.1 is the assumption that there are enough transitions to identify the parameters. At the extreme, we cannot identify the transferability of work experience if individuals never move.

Separation Rates, Switching Costs, and Flow Payoffs: The assumptions that guarantee separability of the model means that transition probabilities can be non-parametrically identified after identifying the earnings parameters, since states are known given the earnings parameters. If separation rates were known, this would permit identification of conditional choice probabilities (CCPs). By inverting the CCPs and fixing the discount factor, we can recover differences in value of taking an action at each state relative to some reference value through standard arguments (Hotz and Miller [1993], Magnac and Thesmar [2002]). Changing relative values of work versus non-employment is captured through a work-type specific quadratic relationship with age relative to age 30. I cali-

[74] I choose this normalization with the goal of creating concave age-skill profiles. Alternate normalizations one could have considered are normalizing one skill weight for each skill (e.g. forcing the weight for skill 2 to be 1 in construction) or normalizing one of the growth rate parameters. At a minimum, we need at least S normalizations since $\pi^{(s)} \cdot \omega^{(s)}$ could equivalently be rationalized by $\tilde{\pi}^{(s)} \cdot \tilde{\omega}^{(s)}$, where $\tilde{\pi}^{(s)} = \alpha\pi^{(s)}$ and $\tilde{\omega}^{(s)} = \frac{\omega^{(s)}}{\alpha}$. I plan to assess the sensitivity of my estimates to these alternate normalizations in future work.

Table 1.3: Identification

Structural Parameters	Variation in Data
Earnings Parameters	
Skill Transferability: $\pi_j^{(s)}$	Cross-industry, within-history differences in earnings
Skill Accumulation: $\Gamma_q^{(s)}$	Within-industry, cross-history differences in earnings
Return to Specific Skill: α_j^{τ}	Within-industry differences in earnings loss after switches into non-employment
Specific Skill Accumulation: Ψ	Earnings growth for industry stayers versus switchers
Specific Skill Voluntary Switch: γ^{Vol}	Earnings growth across voluntary industry transitions
Initial Skills: $\omega_1^{(s)}$	Cross-industry differences in earnings conditional on z_i
Skill Prices: r_{jt}	Average industry-year earnings
Part-Time Penalty: $\alpha^{PartTime}$	Within-history earnings loss for part-time school
Idiosyncratic Shocks Persistence: ρ_j	Within-industry, within-history covariance across lagged earnings
Idiosyncratic Shocks Variance: σ_j	Within-industry, within-history in earnings changes
Preference Parameters:	
Preferences: μ_k^q	Common directions of transitions not explained by earnings gains
Income Coefficient: β_Y	r_{jt} and ω are excluded from utility directly. Unexpected shifts in r_{jt} and how different skill levels respond traces out responsiveness to income maximization.
Tuition Coefficient: β_P	Sensitivity to education take-up over predicted financial aid schedule
Friction Parameters:	
Separation Rates: δ_{jt}	Calibration from the ACS
Switching Cost - Non-employed: κ^U	Duration of non-employment spells
Switching Cost - Employed: κ^E	Frequency of industry switching in data
Switching Cost - From School: $\kappa_{S_{-1}}$	Frequency of sectoral switching out of enrollment
Switching Cost - Into School: $\kappa_{S(k)}$	Frequency of transitions into school not explained by price/aid

brate separation rates using annual transition rates from employment to unemployment by industry in the American Community Survey.[75]

Once separation rates and the relative payoffs of actions in different states are identified, an additional issue is separating the time-invariant preferences for actions from switching costs. Here, the probability of remaining in the same state pins down the mean utility while the probability of switching identifies mean utilities plus the switching costs.[76]

Persistence and Shock Parameters: Shock parameters σ are identified by within-

[75]These are displayed in Figure 1.11(b). Specifically, I compute the unemployment rate for aged 25 to 55 adults in Texas who were employed at least 26 weeks in the prior year. I aggregate the variable INDNAICS to match my industry definitions. For individuals who are currently unemployed, this variable defines the last industry of employment. I ignore sampling error resulting from these rates which I believe is reasonable due to the size of the ACS/Census. Similar to skill prices, the values for 2009 to 2015 are interpolated in the beliefs prior to the Recession.

[76]Consider an illustrative example with only two industries $\{j_1, j_2\}$ and non-employment. We wish to separate mean utilities μ_{j_1} and μ_{j_2} and switching costs $C(j_1|j_2)$ and $C(j_2|j_1)$. By the previous arguments,

history, across-person variation in earnings. The persistence parameters ρ are identified off of the degree to which this variation dies away within individuals. More practically, the assumption that the initial distribution of the productivity shock is independent reveals its full distribution. Then, the timing assumptions on the innovations in individual productivity η allows the econometrician to observe how this object is related across periods, given a guess of parameters. In particular, the full distribution of η is identified in each period under the semi-parametric assumptions on the innovation. This identification argument is implicit in the estimation procedure outlined below.

Income Coefficient: The income coefficient β_Y is a key parameter that captures the degree to which individuals maximize income versus non-pecuniary factors, which affects the incentive to invest in human capital. This parameter is identified from exclusion restrictions that provide independent variation in the earnings equation. Specifically, skill prices and skills levels are excluded from utility directly, only affecting an individual's decision through their valuation of income. Since the Great Recession is a surprise, workers cannot be selecting on industry conditions during the downturn prior to the Recession. The degree to which they respond to changes in industry conditions helps identify β_Y.[77]

Tuition Coefficient: The sensitivity of individual choices to net tuition prices β_P also affects human capital investments. This is identified from estimated financial aid

once we identify separation rates, we can invert choice probabilities to recover:

$$\begin{aligned} u(j_1|j_1) &= \mu_{j_1} & \qquad u(j_1|j_2) &= \mu_{j_1} \\ u(j_2|j_1) &= \mu_{j_2} - C(j_2|j_1) & \qquad u(j_1|j_2) &= \mu_{j_1} - C(j_1|j_2) \\ u(0|j_1) &= 0 & \qquad u(0|j_2) &= 0 \end{aligned}$$

which has 4 equations in 4 unknowns. Generalizing to J industries, we get $J \times J$ equations and $J + J \cdot (J-1) = J \times J$ unknowns. This argument can be repeated for every possible state and worker type.

[77] For example, in the context of the original Roy [1951] model, where individuals choose between hunting and fishing, changes in skill prices in my model are analogous to changes in the market price of rabbits or fishes. This shifts the monetary return to hunting or fishing holding all else equal. As both comparative advantage and preferences are held fixed, changes in prices help reveal maximization on pecuniary versus non-pecuniary factors. See French and Taber [2011] for a thorough discussion on the non-parametric identification of the generalized Roy model.

schedules as well as assumptions on beliefs. Financial aid schedules shift the cost of schooling, so the elasticity of schooling as individuals receive additional aid through higher lagged income identifies this parameter. Moreover, changes in this schedule are modeled as unexpected prior to 2009 - how individuals of different potential financial aid shift their decisions before and after the onset of the Great Recession aids in identification of this parameter.

1.6 Estimation

Estimation is through a maximum likelihood procedure that is broken into two steps for tractability.[78] In the first step, parameters of earnings and the skill process are estimated. In the second step, parameters governing transitions are estimated by maximizing the likelihood of observed choices, conditioning on earnings parameters. The separability of these two steps depends on Assumptions 1.5.2, 1.5.1, and 1.5.3 as discussed in the previous section.

1.6.1 Estimation Strategy

I first clarify the structure of the estimation problem. Let θ_Y denote all parameters affecting the earnings process and θ_C denote the parameters governing transition probabilities. θ_C includes both utility and switching cost parameters. The parameters θ_Y contain several sub-components. First, there are the non-linear parameters α_Y which affect earnings through the skill accumulation functions in Equation 1.4.2 and Equation 1.4.5. Second, there are auxiliary parameters λ_Y, which can be read off the data given a guess of α_Y.

[78] This multi-step approach has three key advantages. First, estimation of the earnings equation is separated from assumptions entering the choice equation, meaning that identification does not depend on beliefs or exact distributional assumptions, only on semi-parametric assumptions on the errors. The procedure is therefore more robust to misspecification than joint estimation. Second, breaking estimation into multiple steps aids in tractability by simplifying the estimation problem in each step. Finally, this approach deals with restrictions that often stem from using administrative data, as the more computationally expensive second step can be implemented outside the secure computing environment.

These auxiliary parameters λ_Y include skill prices r_{jt}, which affect earnings only linearly conditional on α_Y, and parameters governing the variance of the earnings shock process σ_j. As individuals can select on idiosyncratic productivity η, meaning $\rho \in \alpha_Y$ as these are non-linearly estimated.

The observed data in each period are log-earnings conditional on employment, y_{it}; observed work and schooling choices, k_{it}; and observed separations at the beginning of period t, sep_{it}^{Obs}.[79] The econometrician also observes the time-invariant worker type q_i and covariates determining initial skills z_i. Recall that $\mathscr{H}_i^{-t}$ denote the observed history for individual i up to the time of making a choice in period t.

We seek to find values of $\theta = \{\theta_Y, \theta_C\}$ that maximize the likelihood of the observed histories:

$$L(\theta) = \prod_{i=1}^{N} \prod_{t=1}^{T_i} f\left(y_{it} | k_{it}, sep_{it}^{Obs}, \Omega(\mathscr{H}_i^{-t}; \alpha_Y), q_i; \theta_Y\right) P\left(k_{it}, sep_{it}^{Obs} | \Omega(\mathscr{H}_i^{-t}; \alpha_Y), q_i; \theta_Y, \theta_C\right) \tag{1.6.1}$$

Equation 1.6.1 shows key features of the estimation problem. The two pieces of the likelihood are dependent only through worker type q_i and θ_Y - parameters θ_C do not affect the likelihood of y_{it} once we condition on k_{it}, sep_{it}^{Obs}, $\mathscr{H}_i^{-t}$, and type q_i. In addition, the current state Ω is affected by the worker's prior history $\mathscr{H}_i^{-t}$ and non-linear parameters governing earnings α_Y. This is due to the fact that skills are history dependent and different α_Y will result in different levels of skill in period t given the same $\mathscr{H}_i^{-t}$.

Taking logs shows the separability of the estimation procedure when maximizing the likelihood:

$$l(\theta) = \underbrace{\sum_{i=1}^{N} \sum_{t=1}^{T_i} \log f\left(y_{it} | k_{it}, sep_{it}^{Obs}, \Omega(\mathscr{H}_i^{-t}; \alpha_Y), q_i; \theta_Y\right)}_{\textbf{Step 1: Earnings and Human Capital Parameters}} + \underbrace{\sum_{i=1}^{N} \sum_{t=1}^{T_i} \log P\left(k_{it}, sep_{it}^{Obs} | \Omega(\mathscr{H}_i^{-t}; \alpha_Y), q_i; \theta_Y, \theta_C\right)}_{\textbf{Step 2: Transition Parameters}} \tag{1.6.2}$$

[79] This follows from Assumption 1.5.3.

I exploit the fact that one can estimate this likelihood in two steps in estimation, which is made possible by Assumptions 1.5.1 and 1.5.2.[80]

1.6.2 Estimation Procedure - Step 1: Earnings and Human Capital Parameters

The left part of the likelihood in Equation 1.6.2 can be maximized using a non-linear least squares (NLLS) procedure. This estimation routine has the same first-order condition with respect to parameters in α_Y as the left side of the Equation 1.6.2 but only invokes the mean independence assumption on innovations in idiosyncratic productivity φ. Full independence and normality are explicitly invoked to recover σ_j and in the second-step of estimation of the choice model.

Given a guess of non-linear earnings parameters α_Y, skill prices r_{jt} are simply recovered as averages of earnings residuals within each industry-year cell. This involves iterating through a worker's entire sequence of choices to keep track of history-dependent states $\Omega(\mathscr{H}_i^{-t}; \alpha_Y)$, including the contribution of idiosyncratic productivity in the previous period, $\eta_{ij_{i,t-1},t-1}$. After minimizing this NLLS problem, parameters governing the variance of the shock process are computed using the full distribution of residuals.

For the purposes of clarifying the estimation procedure (and abusing notation slightly to ignore the direct dependence of j on k), Equations 1.4.1 and 1.4.6 can be rewritten as:

$$y_{it} = r_{jt} + h\left(k_{it}|\Omega(\mathscr{H}_i^{-t}; \alpha_Y); \alpha_Y, \lambda_Y(\alpha_Y)\right) + \rho_j^{j_{-1}}\eta_{i,j_{-1},t-1} + \varphi_{ij,j_{i,t-1},t} = y(k_{it}|\mathscr{H}_i^{-t}, q_i; \alpha_Y) + \varphi_{ij,j_{i,t-1},t} \quad (1.6.3)$$

where r_{jt} is the skill price resulting from choice k_{it}, $h\left(k_{it}|\Omega(\mathscr{H}_i^{-t}; \alpha_Y); \alpha_Y, \lambda_Y(\alpha_Y)\right)$ is the contribution of supplied human capital, $\rho_j^{j_{-1}}\eta_{i,j_{i,t-1},t-1}$ is the contribution from the persistence in idiosyncratic productivity, and $\varphi_{ij,j_{i,t-1},t}$ is the contribution from shocks to

[80] The separability of estimation also easily permits an extension to the model based on clustering workers into latent types in a manner consistent with the model. This is because the two likelihoods are linked only through worker types. See Traiberman [2019] and Bonhomme et al. [2019] for examples of such applications.

idiosyncratic productivity.

The NLLS objective is defined as:[81]

$$Q^{NLLS}(\alpha_Y) = \frac{1}{2N}\sum_{i=1}^{N}\frac{1}{T_i}\sum_{t}\left(y_{it} - y(k_{it}|\mathscr{H}_i^{-t}, q_i; \alpha_Y)\right)^2 = \frac{1}{2N}\sum_{i=1}^{N}\frac{1}{T_i}\sum_{t}\varphi^2_{ij,j_{i,t-1},t} \quad (1.6.4)$$

where $\varphi_{ij,j_{i,t-1},t}$ is the error defined above. The key moment condition justifying this NLLS procedure is given in Equation 1.5.1 above.

Given a guess of α_Y and prior values of $\eta_{i,j_{i,t-1},t-1}$, skill prices are estimated internally as the average of industry-year residuals:

$$\hat{r}_{jt} = \frac{1}{N_{jt}}\sum_{i \text{ s.t.} j(k_{it})=j}\left(y_{it} - h\left(k_{it}|\Omega(\mathscr{H}_i^{-t};\alpha_Y);\alpha_Y,\lambda_Y(\alpha_Y)\right) + \rho_j^{j-1}\eta_{i,j_{-1},t-1}\right) \quad (1.6.5)$$

Including an individual's productivity in the previous period in Equation 1.6.4 and 1.6.5 presents a computational challenge. Agents select on $\eta_{i,j_{i,t-1},t-1}$ but this quantity also depends on prior period skill prices $r_{j,t-1}$. To overcome this issue, I sequentially update current skill prices r_{jt} given lagged productivity $\eta_{i,j_{i,t-1},t-1}$, followed by updating current productivity η_{ijt} given current skill prices r_{jt} and residuals $\varphi_{ij,j_{i,t-1},t}$. The exact algorithm is outlined in Appendix Section D.1:

1.6.3 Estimation Procedure - Step 2: Transition Parameters

I now turn to estimation of θ_C. Given values of α_Y, we can compute the ex-ante unobservable state $\Omega(\mathscr{H}_i^{-t};\alpha_Y)$ for latent skills in terms of work histories. Thus, the only thing the econometrician does not observe at this point is if a transition into non-employment

[81] I minimize this objective using Gauss-Newton approximations of the Hessian based on the outer-product of each observation's gradient. To deal with sensitivity to starting points, I first explore the parameter space using stochastic gradient descent algorithms with multiple starting points and multiple restarts. Specifically, I use mini-batch samples of 1000 individuals to evaluate the gradient from randomly chosen starting points, restarting the algorithm and moving the current set of points randomly towards the current best point in the full-sample. In simulated data, this procedure reliably gets parameters in the region of the global minimum even when starting from remote starting points. Finally, to assess whether the final parameters are stuck in a local minimum, I do a one direction line search for each parameter.

(or full-time study) is due to an involuntary separation or voluntary transition into non-employment. I calibrate the separation rates from the ACS, both simplifying the estimation problem and enriching the model in terms of industry-time variation.

This second-step is estimated outside of the ERC's secure data environment using a second maximum likelihood procedure that maximizes the likelihood of observed transitions over states:

$$Q^{CCP}(\theta_C|\hat{\theta}_Y) = \sum_{s=1}^{S} w_s \sum_{k=1}^{K} \hat{P}_{k|s}(\hat{\theta}_Y) \log P(k|\Omega_s, \hat{\theta}_Y, \theta_C) \tag{1.6.6}$$

where $\hat{P}_{k|s}$ is the reduced-form probability of an individual in the baseline sample choosing action k at state s, discussed below. Weights w_s are computed based on the share of person-year observations in a cell that defines s. Equation 1.6.6 is thus an approximation of the likelihood governing choices in Equation 1.6.2.

Evaluating Equation 1.6.6 requires solving the dynamic problem on each evaluation. This is because the history-dependence of skills and non-stationarity of problem makes finite-differencing and other forward simulation techniques difficult to implement (Arcidiacono and Miller [2011], Arcidiacono and Miller [2020], Hotz and Miller [1993]). The reduced-form conditional choice probabilities $\hat{P}_{k|s}$ are predicted from K linear probability models that pool across years of the same and condition on θ_Y.[82] There is then the issue of choosing states at which to evaluate the objective function 1.6.6. I leverage the fact that I observe the states once α_Y is estimated, picking states based on the observed

[82]For covariates for this linear probability model, I use (a) a constant, (b) splines in year with a break at 2009: $\{year, year - 2009, D(year \geq 2009\}$, (c) indicators for each industry of prior employment and additional indicators for full-time or part-time study in the previous period, (d) continuous states affecting choices that include age, skills, periods of tenure in prior activity, lagged incomes, predicted financial aid, and previous idiosyncratic productivity, (e) worker types q, and (f) all pairwise and quadratic interactions for the previous. At each state, I bound the predicted probabilities and between 1e-100 and 1-1e-100 to avoid negative or unusually large probabilities.

distribution in the data.[83]

1.6.4 Standard Errors

I compute the asymptotic variance-covariance matrix of the parameters using the non-parametric bootstrap over 100 bootstrap samples. These standard errors account for clustering at the individual level by drawing each individual's panel with replacement when estimating θ_Y and $P_{k|s}$. The second-step of the estimation procedure simply maximizes the likelihood of the (smoothed) probabilities in the data $\hat{P}_{k|s}$, a step which does not use any individual-level data directly. To compute standard errors on estimates of θ_C, I repeat the second-stage estimation conditioning on each bootstrap draw of $\hat{\theta}_Y^b$ and $\hat{P}_{k|s}^b$. This accounts for covariance between $\hat{P}_{k|s}$ resulting from the panel data.[84]

1.7 Estimation Results

This section focuses on a selected overview of parameter estimates, fit and validation of the model, and some key implications. Complete estimates are reported in Appendix Section D.3.

1.7.1 Parameter Estimates

Skill Weights and Skill Process: Estimates of skill weights are shown in panel A of Table D1. These results indicate that skill 1 tends to be general skill while and skills 2

[83] Specifically, I divide the continuous states that define Ω into quartiles based on the observed distribution of states implied by $\hat{\theta}_Y$. Further, I also only evaluate the objective at the years 2003, 2006, 2009, 2012, 2017 and for the 1994 and 1998 cohorts for simplicity. These years, cohorts, quartiles of continuous states, discrete schooling state $s_{i,t-1}$, and discrete employment state $j_{i,t-1}$ define the state cell s. I compute the share of individuals in s choosing k, using the median values of continuous states to define Ω_s on which I evaluate the model. Finally, I only disclose cells with at least 5 observations resulting in a total of $S = 51,430$ cells representing $1,379,330$ person-year observations. I use the linear probability models to predict $\hat{P}_{k|s}$ at each state.

[84] I need to bootstrap since the same individual may contribute to CCPs estimated across multiple states. Note that standard errors do not account for (a) choice of states on which to evaluate the model or (b) calibration of industry separation rates or financial aid parameters.

and 3 are relatively more specific. The weights on skill 1 range from about 0.56 (0.07) for construction to about 0.67 (0.08) for services. In other words, a 1 unit increase in skill 1 increases earnings by 55 percent to 67 percent across these sectors, respectively. A one unit increase in skill 2 increases earnings from about 7 percent for services to 21 percent for manufacturing. Skill 3 has returns ranging from about 23 percent for health care and the public sector to about 49 percent for financial services and the general sector.

Each skill differs substantially in its accumulation process as defined by Equation 1.4.2, as shown in panels B and D of Table D1. These numbers indicate that different work and studying choices show substantial variation on each skill margin, including the fact that some skills may depreciate upon taking an action. These parameters are normalized to the men with no prior college attendance and 1 unit of the target skill. Moreover, these represent earnings gains from the prior period on each skill margin holding all else equal, rather than relative to an alternate choice.

Non-college men have a 0.1 increase in the log-rate of accumulation of skill 1 for each period they work in health care whereas those studying health care have a 0.3 increase in this growth. Combining this information with the skill weights above, this indicates that working in health care boosts earnings through skill 1 by 6.6 percent ($100 \times (\exp\{0.100\} - 1) \times 0.627$) while studying health care boosts earnings by 22.2 percent ($100 \times (\exp\{0.303\} - 1) \times 0.627$) in the same sector. In contrast, the comparable numbers for skill 2 are -0.8 percent ($100 \times (\exp\{-0.040\} - 1) \times 0.203$) for working in health care and 2.8 percent ($100 \times (\exp\{0.127\} - 1) \times 0.203$) for studying health. For skill 3 these are -6.5 percent ($100 \times (\exp\{-0.143\} - 1) \times 0.489$) for working in health care and -30.2 percent ($100 \times (\exp\{-0.963\} - 1) \times 0.489$) for studying health.

The previous discussion shows that the process by which estimated parameters map into earnings is complex and reported only for non-college men normalized to a skill level of 1. To ease interpretation in translating what parameter estimates mean for earnings growth from taking certain actions, Tables 1.4 and 1.5 show implied one-year earn-

ings growth due to skills using the ex-post distribution of skills in simulated data. To make comparisons relative to an action, I compute these numbers relative to working full time in the prior industry. It is important to note that this exercise ignores earnings impacts from skill prices, specific skills, and selection on idiosyncratic productivity over time. Looking at fields of study in Table 1.4, we see three key patterns. First, only technical and health studies have any scope to boost average earnings - the other fields have uniformly negative earnings impacts across sectors. Second, while health has the highest earnings impacts on the health care and public sectors, it has negative earnings impacts in all other sectors. This underscores the importance of the specificity of human capital - some fields are only valuable if individual obtain a job in the target sector. Finally, there is substantial heterogeneity in effects. For example, the average increase in earnings from technical studies is half of the standard deviation in returns.

Table 1.4: Fields of Study and One-Year Sector-Specific Earnings Growth from Latent Skills

	General	Construction	Manufacturing	Health Care	Service	Financial	Public
General	-14.7	-11.9	-14.6	-10.6	-9.6	-14.2	-9.8
SD	[5.0]	[3.6]	[5.2]	[4.8]	[2.8]	[4.7]	[4.4]
Technical	5.7	3.3	6.2	5.8	0.6	5.2	5.3
SD	[11.4]	[7.6]	[11.9]	[10.9]	[4.2]	[10.5]	[9.9]
Health	-3.0	-9.6	-0.7	10.9	-17.1	-5.1	9.5
SD	[22.2]	[14.9]	[23.3]	[21.1]	[8.8]	[20.5]	[19.2]
Humanities	-4.2	-5.1	-3.7	-0.2	-6.5	-4.6	-0.3
SD	[8.5]	[5.7]	[8.9]	[8.2]	[3.5]	[7.8]	[7.5]
Business	-5.5	-4.2	-5.6	-4.4	-2.9	-5.2	-4.0
SD	[5.4]	[3.8]	[5.6]	[5.3]	[3.0]	[5.0]	[4.8]

Notes: Constructed on baseline simulated dataset of 1,000,000 at estimated parameter values. Each (f, j') cell reports average growth in log-earnings from studying field f (rows) and switching to industry j' (columns), computed as $100 \times \sum_{s=1}^{S} \pi_{j'}^{(s)} \left(\Gamma_{q_i}^{(s)}(f|\omega_{it}^{(s)}) - \Gamma_{q_i}^{(s)}(j_{-1}|\omega_{it}^{(s)}) \right)$ and averaged over the realized distribution of ω_{it}. These growth rates do not account for changes in earnings due to other objects impacting earnings. Standard deviations of these individual-level growth rates are reported below.

Similar patterns for the transferability of sector-specific work experience across potential future sectors are shown in Table 1.5. The general sector, manufacturing, and financial services all have the highest growth in earnings from skill. Working in the service sector tends to result in skill depreciation with the least amount of depreciation

occurring in the service sector itself. Working in health care or the public sector tend to result in a small increase in skills in those sectors but result in skill depreciation in other sectors. Last, the overall standard deviation in implied returns to sector-specific experience is much lower than for field of study in Table 1.4.

Table 1.5: Learning-by-Doing and One-Year Sector-Specific Earnings Growth from Latent Skills

	General	Construction	Manufacturing	Health Care	Service	Financial	Public
General	3.7	3.0	3.7	3.1	2.7	3.6	2.9
SD	[3.4]	[2.6]	[3.4]	[3.1]	[2.5]	[3.2]	[2.9]
Construction	2.1	2.0	2.0	1.5	2.2	2.1	1.4
SD	[3.4]	[2.6]	[3.4]	[3.1]	[2.5]	[3.2]	[2.8]
Manufacturing	3.6	3.1	3.6	2.6	2.8	3.6	2.4
SD	[3.3]	[2.6]	[3.4]	[3.1]	[2.4]	[3.2]	[2.8]
Health Care	-1.4	-2.0	-1.2	0.6	-2.5	-1.6	0.5
SD	[3.3]	[2.6]	[3.4]	[3.1]	[2.4]	[3.2]	[2.8]
Service	-4.4	-2.7	-4.6	-4.7	-1.0	-4.0	-4.3
SD	[3.5]	[2.7]	[3.5]	[3.1]	[2.5]	[3.3]	[2.8]
Financial	4.3	3.5	4.3	3.0	2.7	4.1	2.8
SD	[3.4]	[2.7]	[3.4]	[3.1]	[2.4]	[3.2]	[2.8]
Public	-0.9	-1.5	-0.6	1.0	-2.1	-1.1	0.9
SD	[3.3]	[2.6]	[3.4]	[3.1]	[2.4]	[3.2]	[2.8]

Notes: Constructed on baseline simulated dataset of 1,000,000 at estimated parameter values. Each (j, j') cell reports average growth in log-earnings from working in industry j (rows) and working in industry j' (columns), computed as $100 \times \sum_{s=1}^{S} \pi_{j'}^{(s)} \left(\Gamma_{q_i}^{(s)}(j|\omega_{it}^{(s)}) - \Gamma_{q_i}^{(s)}(j_{-1}|\omega_{it}^{(s)}) \right)$ and averaged over the realized distribution of ω_{it}. These growth rates do not account for changes in earnings due to other objects impacting earnings. Standard deviations of these individual-level growth rates are reported below.

The heterogeneous returns to work experience in Panel C of Table D1 show that women tend to accumulate skills 1 and 2 slower than men but have some comparative advantage in accumulating skill 3. For example, women working in health care increase earnings in the general sector through skill 1 by 6.4 percent if non-college ($100 \times (\exp\{0.111 - 0.014\} - 1) \times 0.627$) and 6.5 percent if college ($100 \times (\exp\{0.111 - 0.012\} - 1) \times 0.627$), whereas the same numbers are 7.3 percent and 9 percent for men, respectively. In contrast, men with prior college appear to more easily accumulate all skills through work. The schooling returns in panel E show that men with no college have an easier time accumulating skills via school overall, with the notable exception being that men with prior college have a much easier time accumulating skill 2.

Panel F of Table D1 shows that the earnings gain from human capital investments declines as workers accumulate skills, captured in the diminishing investment parameters $\gamma^{(s)}_{Dim:Work}$ and $\gamma^{(s)}_{Dim:School}$. For skill 1, investments from work are 70 percent ($100 \times (\exp\{-1.189\} - 1)$) less productive for individuals with 2 units of skill 1 than at 1 unit. Turning to skill 2, we see that work investments do not diminish as quickly but that investments from schooling diminish as rapidly as skill 1. In contrast, these parameters are close to zero for skill 3. Since skill 3 tends to diminish over time, this means that the rates at which the skill depreciates is relatively constant over the life-cycle. Panel F also shows parameters governing the effects of part-time study and non-employment on skills. We see that about 57 percent of part-time human capital investments come through work whereas 43 percent comes through schooling. Skills depreciate slightly in non-employment - individuals lose 3.2 percent of their stock of skill 1 while non-employed, little of skill 2, and about 0.7 percent of skill 3. Returning to Panel A shows that workers face about a 12 percent cut in earnings while enrolled relative to full-time (non-enrolled) work.

Finally, Panel G of Table D1 shows how initial conditions are determined, where covariates are defined prior to age 25. We see that skill 1 tends to be uniformly higher for those working, with non-college men having the lowest levels of skill. Skill 2 is more differentiated in terms of sectors - with construction and financial services having higher values. In terms of gender, women with prior college having lower values. In contrast, Skill 3 is highest for non-college men and for those working in manufacturing and the public sector, while being lower for those working in services.

Table D1 shows similar parameters for the specific skill. Panel A shows that this component is weighted relatively equally across sectors, with each sector returning about 30 percent higher earnings for each subsequent unit of skill 1. The specific skill grows by about 6.9 percent at one unit of skill for each period worked ($100 \times (\exp\{0.067\} - 1)$) translating into about 2.25 percent higher earnings in the general sector. The efficacy

investments in specific skill is quickly diminishing - one period worked in the general sector raises earnings by 0.32 percent at 2 units of the specific skill, as shown in Panel F. Finally, specific skill tends to *grow* across voluntary employment transitions by about 2.1 percent. This growth likely captures the fact that individuals switch industries with an outside job offer in hand.[85]

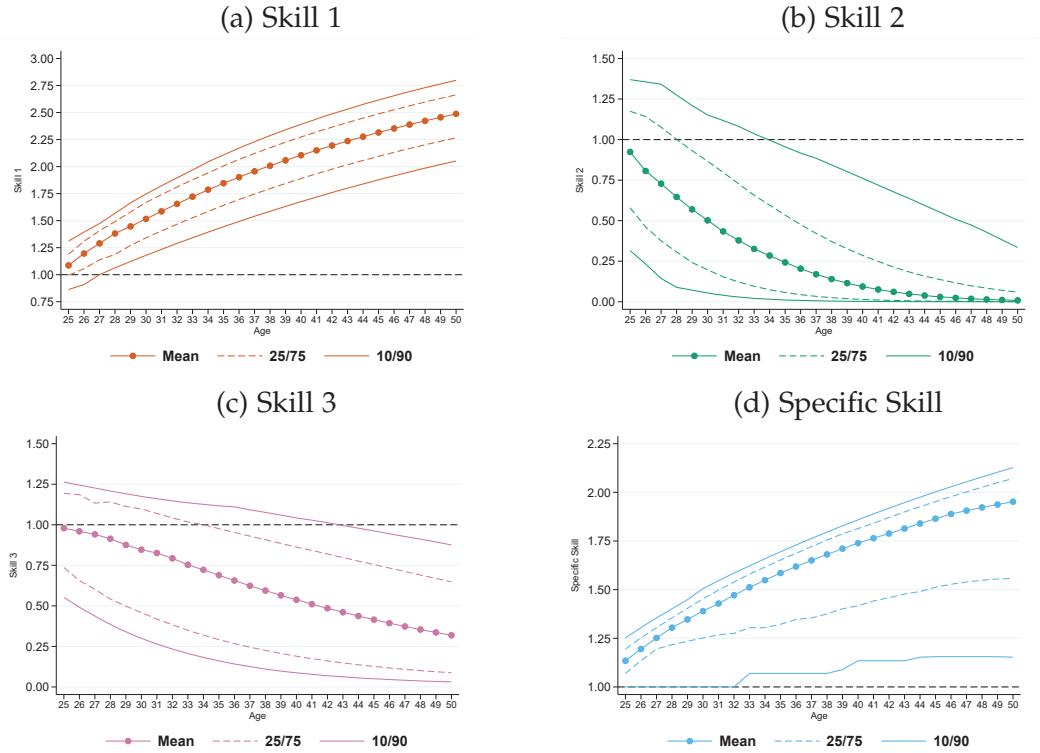
Figure 1.10 and Appendix Figure A5 show the ex-post distribution of skill in simulated data. Figure 1.10 shows how the marginal distributions of skill change with age. First, we see that all skills start on average close to 1, reflecting the normalization of the initial condition. Latent skill 1 grows deterministically over the life cycle with relatively low variation at each age. The specific skill, shown in panel (d), also increases over the life cycle. However, there is a large tail to the distribution, reflecting that individuals can fall off a job ladder in a way that permanently affects their earnings. In contrast, latent skill 2 and latent skill 3 diminish over the life cycle but have a large amount of variation.[86]

Table 1.6 shows an ensemble decomposition of earnings with respect to each component on the realized skill distribution in the administrative data, broken out into groups of ages. First, the earnings model does quite a good job at explaining the variation in the data - the implied R^2 at ages 30 to 34 of the NLLS procedure is 0.744 when including lagged errors and 0.268 when only looking at skill and price components. The persistent component captures about half of the variance in earnings, 47.6 percent. Each skill also explains a non-trivial share of total earnings variance with important changes over time. The importance of skill 1 rises sharply as individuals age and that of skill 2 declines over age. The importance of the specific skill τ is fairly constant at around 9 percent of the total variance in earnings, reflecting that remaining continuously employed is quite

[85] Workers anticipate that if they pay the cost to voluntarily switch they will accumulate this. This stands in contrast to a random search model, where workers passively receive offers and then decide to move.

[86] Appendix Figure A5 does a compatible exercise to examine ex-post sorting behavior by plotting average skill components in each industry over time. Here, we visually see the importance of skill 1 and the specific skill. We also see cross-sector sorting on skill 2 and skill 3. For example, individuals working the health care sector have no skill 3 while those working in the service sector have no skill 2 on average.

Figure 1.10: Marginal Distributions of Latent Skills by Age



Notes: Plots the marginal distribution of each skill at each age in a simulated dataset of N =1,000,000 . At each age, the average, 10th, 25th, 75th, and 90th percentile of each skill is computed across all workers, without regard to labor force status.

important for earnings.

Table 1.6: Ensemble Variance Decomposition of Earnings into Skill Components

	Log-Earnings		Contribution of:					
	Variance	Share Explained	r_{jt}	$\pi_j^{(1)}\omega_{it}^{(1)}$	$\pi_j^{(2)}\omega_{it}^{(2)}$	$\pi_j^{(3)}\omega_{it}^{(3)}$	$\alpha_j^{\tau}\tau_{it}$	$\rho_j^{j-1}\eta_{i,t-1}$
Age 25-29	0.33	60.0%	3.3%	1.9%	4.6%	4.8%	8.9%	36.5%
Age 30-34	0.39	74.4%	2.9%	5.1%	3.8%	6.0%	9.0%	47.6%
Age 35-39	0.44	77.9%	3.1%	7.0%	3.0%	6.8%	9.0%	49.1%

Notes: Conducts an ensemble decomposition of the variance in log-earnings in the baseline sample of N =247,973. For each employed individual, I report $100 \times Cov(y_{it}, w_{it})/Var(w_{it})$ where w_{it} is one of the listed model components. The unexplained portion of earnings is φ.

Examining these key patterns allows for a brief characterization of the estimated skills: skill 1 tends to capture a general skill that cannot be modified later in life. Skill 2 tends to capture more specific investments where health and technical studies have returns, which college educated men tend to be better at accumulating and are less rewarded in services or construction sectors. Skill 3 tends to capture skills that are specifically not rewarded in the health and public sectors. In particular, high-levels of

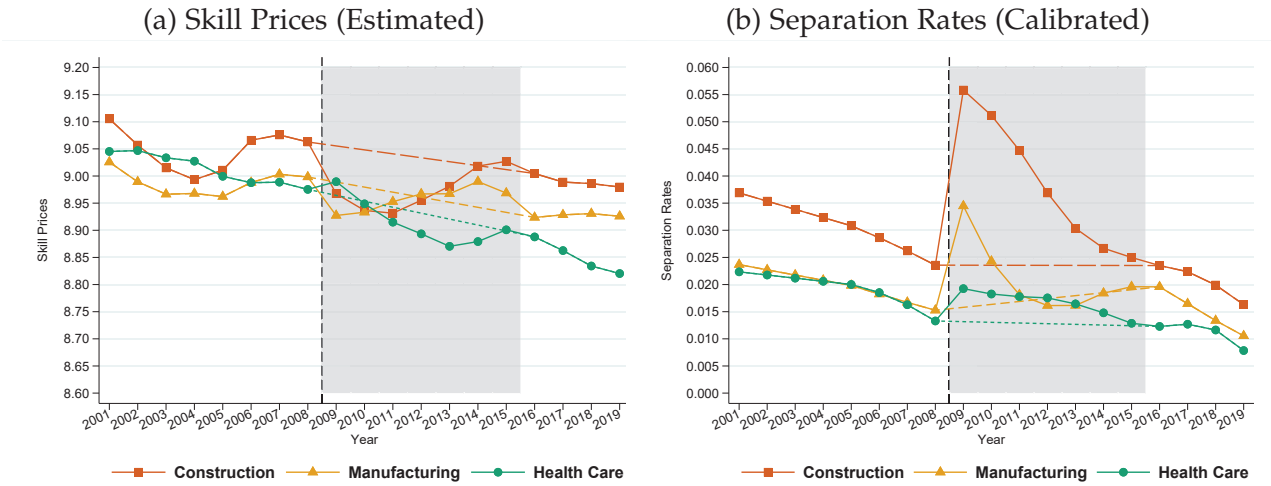
this skill generates an opportunity cost of switching into health and the public sector for those with high levels.

Persistence and Shock Process Parameters: The parameters of the shock process are shown in Table D2. Shocks to productivity are quite persistent for those remaining in their current industry with an autocorrelation parameter of 0.888 (0.001). This is lower for those voluntarily switching industries at 0.719 (0.002) and higher for those switching out of non-employment at 0.350 (0.003). The variance parameters indicate that remaining in your current industry is lower-variance, but switching is more uncertain, especially coming out of non-employment. The public sector is routinely the lowest variance in productivity draws and construction is among the highest.

Aggregate Parameters: Separation rates and skill prices generate industry-by-time variation and affect which workers are most affected by the business cycle. These are shown for key industries in Figure 1.11 and for all industries in Appendix Figure D2. Panel (a) shows skill prices recovered as auxiliary parameters in the estimation procedure. Industries differ in their levels of skill prices as well as their sensitivity to the business cycle. For example, construction has the highest return to skill and is the most affected by declines over the business cycle. In contrast, health care is less sensitive to the business cycle but has lower skill prices. Separation rates, shown in Panel (B), also show differentiation in baseline levels, with construction having the highest baseline rate of layoffs. Industries are differentially affected by the business cycle. Corroborating Figure B1, construction and manufacturing have large increases in separation rates in 2009.[87]

[87] For example, construction sees 5.5 percent of the workforce get laid off, whereas health care sees only a small increase to about 2 percent.

Figure 1.11: Skill Prices and Separation Rates - Selected Industries



Notes: Plots skill estimated skill prices r_{jt} (see Equation 1.4.1) and calibrated separation rates δ_{jt} (see Equation 1.4.11) for selected industries. Dashed lines indicate the biased beliefs agents act on prior to the impact of the Great Recession, which linearly interpolates the years 2009 and 2015. See Appendix Figure D2 for all industries.

Figure 1.11 also includes a graphical depiction of how beliefs about the Great Recession are modeled, previously discussed in section 1.4.2. The lightly shaded area in Figure 1.11 indicates the region over which individuals project the future aggregate state prior to the Recession, during which they have biased beliefs. These biased beliefs are shown in the dashed line. There is a clear dip in skill prices and increases in separation rates during this period, indicating that the approach to beliefs approximate the region over which the economy is pushed off its long-running trend by the Recession well.

Income Coefficient: Table D3 shows the estimated preference and transition parameters. I estimate a preference for log-income of 0.738 (0.024). This is a key parameter as it governs the degree of income maximization in the model, determining the degree to which individuals make forward-looking investments in human capital .[88]

Switching Costs: Panel C of Table D3 shows time-varying switching costs that vary by employment status, target industry, and type of worker. Since utility is linear in log-income, we can translate these into dollar terms. Since switching costs are paid once, they should be compared to a present value increase in an income earnings, which

[88]For a comparable estimate, Ransom [2022] estimates 0.916 on expected log-earnings in a model of geographic mobility.

can be solved by multiplying switching costs by $1-\beta$. This implies switching costs for employed non-college men of 20.7 percent of income. Since the average observed earnings in the sample is \$60,635, this implies moving costs on the order of \$12,127.[89] For comparison, Traiberman [2019] estimates occupational switching costs of 17 percent of income, Ransom [2022] estimates geographic switching costs of 30-40 percent of income, and Kennan and Walker [2011] estimate geographic switching costs that are equivalent to a \$17,897 flow increase.[90]

The non-employed face switching costs that are 17 percent higher ($100 \times \left(\frac{\exp\{1.401\}}{\exp\{1.244\}} - 1\right)$. Sectors also differ in the utility cost of switching into them, reflecting structural industry differences in the job-finding rate. Most substantially, the service sector has switching costs that are 30 percent lower ($100 \times \left(\frac{\exp\{1.244-0.359\}}{\exp\{1.244\}} - 1\right)$ than the general sector, reflecting that these are easier jobs to find. The switching costs increase in the Recession and especially so for the employed workers. This indicates increasing frictions from making voluntary switches but only a slight increase in switching out of non-employment, all else equal.

The estimates indicate an important role for schooling in terms of increasing the rate of job finding. The parameters indicate that, all else equal, being currently in school lowers the switching cost by 26%. This provides an indirect role for schooling in increasing human capital - it can serve a job search function, providing an indirect mechanism for schooling to increase earnings.

Tuition Sensitivity and Switching Costs into School: The estimates indicate that workers face substantial frictions in paying for and returning to school. The parameter governing sensitivity to tuition β_P is estimated to be 0.467 (0.016). Using the same conversion as for switching costs above, this implies that individuals would need a present

[89] The average observed earnings in the sample is \$60,635 or about 11.01 log-points. The equivalent present value percent increase is then: $100 \times \exp\left((1-\beta)\frac{\exp\{\kappa^E\}}{0.738}\right) - 1\% = 20.7$ percent

[90] Due to my inclusion of γ^{Vol}, which I estimate to slightly increase specific skills across voluntary transitions and gives an incentive to switch to increase earnings, I expect my estimates of sectoral switching to be slightly higher than the literature.

value increase in income of 2.6 percent in order to offset the current dis-utility of tuition payments. For the average worker, this translates into an increase in income of $1,552.

Preference Heterogeneity: Finally, Table D3 indicates heterogeneous preferences by gender and prior education. The first column gives the common preferences, normalized to men with no prior college, while the remaining columns report these preferences for other worker types. We see that women tend to have higher preferences for working in health care and the financial sector, while relatively disliking construction and manufacturing. Preferences for working full-time tend to be convex over the life cycle.

1.7.2 Model Fit and Validation

To assess the fit of the model, I simulate a dataset of 1,000,000 individuals at estimated parameters values. I simulate initial conditions for these individuals by grouping the raw observed initial conditions into finely grained cells and drawing from the resulting distribution.[91] I compare the earnings and choice of various cuts of the real data to this simulated dataset. While all moments of the data are implicitly targeted in estimation, the procedure does not force each moment to match mechanically.

Figure 1.12 shows that the model fits basic patterns of the data quite well. In Panel (a), the model fits the life-cycle log-earnings profile of observed worker types, with the possible exception of early in their careers. Panel (b) and (c) shows that the model fits the work and schooling choices of each of these demographic types. Importantly, the model is able to reproduce differences by gender and education in sectoral and enrollment choices. Finally, panel (d) shows the probability of enrollment over time. The model is able to fit the level in enrollment by age but misses some of the steepness of the descent over the life cycle. In addition, the model does not precisely fit the increase in enrollment

[91] Specifically, to simulate initial conditions, I group the raw data into cells based on a dummy for education before age 25, gender, previous industry before age 25 (including non-employment), and years of tenure in this activity and draw from the resulting frequency distribution. To get the initial lag of income, I compute the mean and standard deviation of income within each cell and draw from the resulting normal distribution.

during the Recession for this sample.[92]

Figure 1.12: Model Fit: Basic Patterns



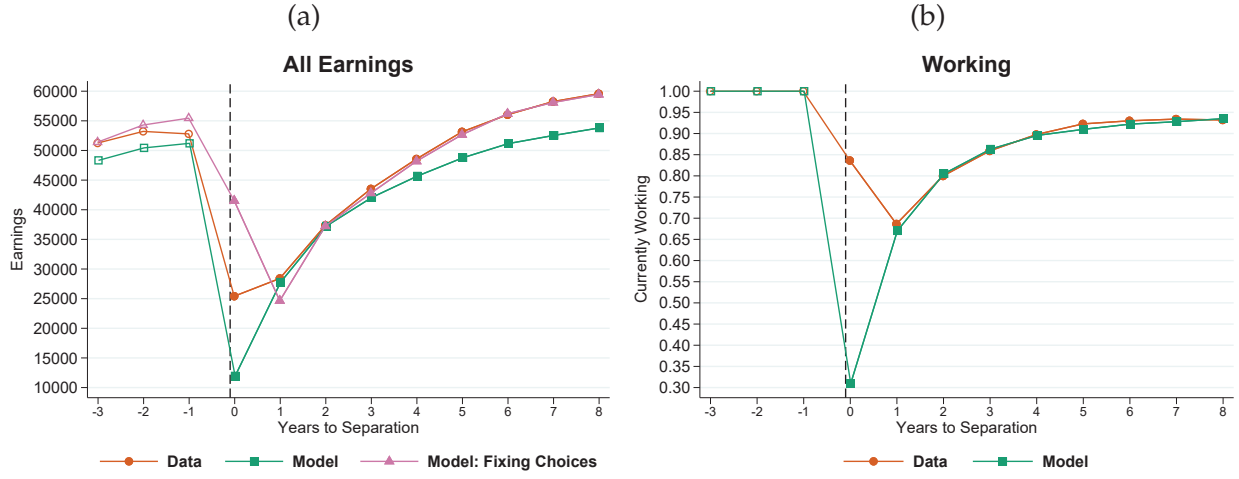
Notes: Compares fit in baseline sample with simulated data of size N=1,000,000 . Panel (a) plots the model predictions of log-earnings in marked-lines with data in dashed lines. Panels (b) and (c) shows sector of employment and field of study for each of these groups, outlining the data in dark lines. Panel (d) shows the probability of enrollment in each year, with data shown in orange and the model in green.

Next, I turn to examining how the model predictions stack-up when compared to quasi-experimental events studies. In Figure 1.13, I show that the model roughly matches employment probabilities and average total earnings for workers with a tenure of three years in one of the $J = 7$ industries and who enter non-employment in 2009. Through the lens of the model, this includes both short non-employment spells (observed exogenous separations) and complete switches into non-employment (endogenous or exogenous). Panel (a) shows that the model does a good job of matching earnings trajectories around

[92]This pattern regarding enrollment over the Recession is hard to see since young workers enroll at a higher rate. Since I focus on five cohorts, the sample mechanically ages over time and life-cycle trends swamp out business cycle trends. Each year the enrollment probability declines except for in 2009 and 2010, which ends around to a 1pp in extra enrollment after age adjusting.

a job separation. Importantly, the model roughly matches the earnings levels prior to displacement, indicating that it captures which kinds of workers get separated at the start of the Recession. The pink (triangle) line shows predicted earnings holding transitions fixed to those in the raw data, which indicates that most differences are due to the choice model not matching transitions precisely rather than poor fit of the earnings model. The model does not match expected earnings at the exact time of separation, which is mostly due to a lack of fit of the earnings parameters. Panel (b) shows why - many separated workers have some employment in the year of impact in the data and the earnings model expects these individuals to have some earnings based on their observed employment status. This is clearly due to workers who have some employment with extremely low earnings in the year of separation, likely measuring part-time work (e.g. only a few weeks the quarter after separation) that the model is not set up to capture (or is not measured in the quarterly data). The model otherwise replicates the share working almost exactly otherwise, which is very encouraging.

Figure 1.13: Model Fit: Employment and Earnings of Workers Separated in 2009



Notes: Share employed and total earnings, including zeros, for $N = 13,487$ workers in the baseline sample who are employed for three years prior to an entry into non-employment in 2009. Orange line (circles) shows the fit of the raw data, green lines (squares) show the fit of data simulated from the full model, pink (triangle) line shows the fit of the earnings model holding choices fixed to the path observed in the data. Share employed for the pink line fits the orange line mechanically. For earnings, the pink line plots expected earnings from the model, accounting for idiosyncratic productivity by adding the observed $\rho\eta_{-1}$ as well as adding $\sigma_\varphi^2/2$ to account for the expected impact of the idiosyncratic innovation on log-earnings converted into levels. Since choices are fixed, the pink line fits share employed mechanically.

Figure 1.14 shows these patterns for new enrollees around an enrollment event that

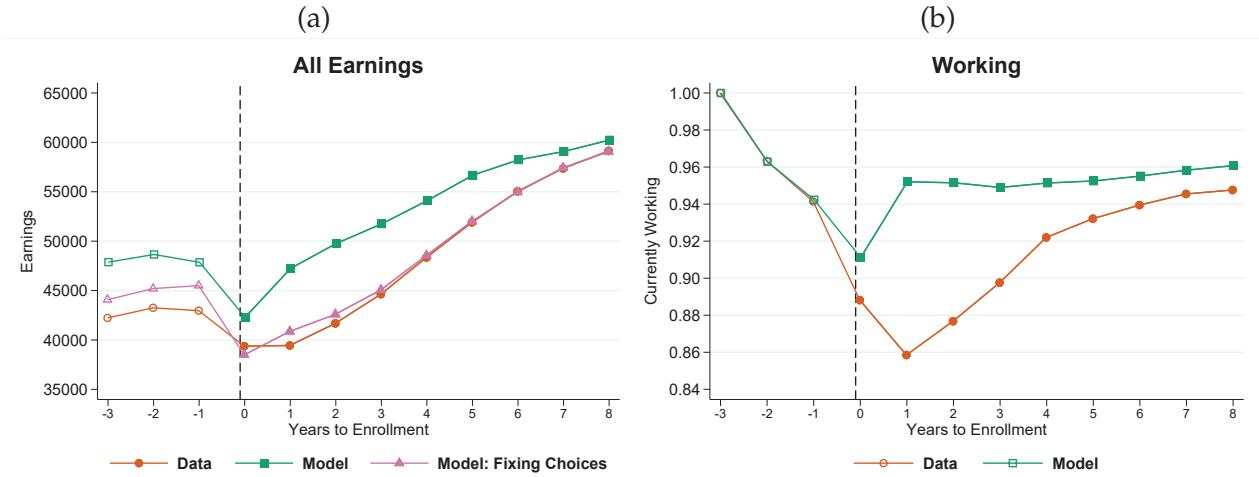
begins in 2009 or 2010. This is defined for individuals who shift into enrollment after not being enrolled in the previous three years and who have some work in year three prior to enrollment. The model roughly matches the earnings trajectories around enrollment, though admittedly worse than it does for worker separations. The most important differences are due to the model over-predicting the employed share of workers after enrollment. Comparing to the pink (triangle) line in Panel (a), which shows the fit of the earnings model, we see that almost all of the lack of fit on earnings comes from the choice model not fitting the sequence of transitions correctly. In the base period, the model exactly reproduces the employed share of workers and roughly matches pre-enrollment earnings, but over-predicts earnings by $5,159. Comparing overall earnings growth, growth in average earnings for enrollees is 40 percent in the data and 27 percent in the model due to this positive selection in the model. Most of these differences disappear when fixing the path of choices to those in the data, as shown in the pink line, where the pre-period gap in earning shrinks to $2,539 and the overall earnings growth is 34 percent. This gives some comfort that the differences in fit are not likely due to biased estimates of returns. Importantly, the fact that the model over-predicts how quickly people return to work likely results in an upward bias in the counterfactual results.

1.8 Decompositions and Counterfactual Exercises

I now use the estimated model to assess the effects of counterfactual policies and decompose mechanisms driving worker choices. I follow the same simulation procedure as in Section 1.7.2 to draw a large simulated dataset, but change primitives affecting workers' choices. To make individuals comparable across simulations, I fix realizations of random outcomes across simulations.[93]

[93]These include the time-varying unobservables, including innovations in idiosyncratic productivity φ, time-varying preferences ϵ, and outcomes of separations δ. I also fix the draws of the initial conditions.

Figure 1.14: Model Fit: Earnings of Workers Enrolling in 2009 or 2010



Notes: Share employed and total earnings, including zeros, for $N = 14,950$ workers in the baseline sample who enroll in 2009 or 2010 after not being enrolled in the prior three years and with some employment in year three prior to enrollment. Orange line (circles) shows the fit of the raw data, green lines (squares) show the fit of data simulated from the full model, pink (triangle) line shows the fit of the earnings model holding choices fixed to the path observed in the data. For earnings, the pink line plots expected earnings from the model, accounting for idiosyncratic productivity by adding the observed $\rho\eta_{-1}$ as well as adding $\sigma^2_\varphi/2$ to account for the expected impact of the idiosyncratic innovation on log-earnings converted into levels. Since choices are fixed, the pink line fits share employed mechanically.

1.8.1 Effects of Price Cuts and Field-of-Study Specific Subsidies

I assess the effect of price changes by making an unexpected and permanent change in sticker prices, occurring in 2009, which will thus not affect worker behavior prior to the Recession. This fixes initial human capital investments and lessens the influence of initial conditions.[94] I focus on the propensity to enroll between 2009 and 2015, before the oldest cohort in the sample turns forty, and on changes in total worker earnings in their forties.[95]

Figure 1.15 shows the effect of these prices changes on enrollment. Mirroring the discussion in Section 1.7.1, workers are quite sensitive to tuition prices. Panel (a) examines the change in the probability a worker enrolls at any point between 2009 and 2015 and Panel (b) shows the change in the annual enrollment rate over the same period. The model indicates that that a $1,000 price decrease raises the propensity of workers to

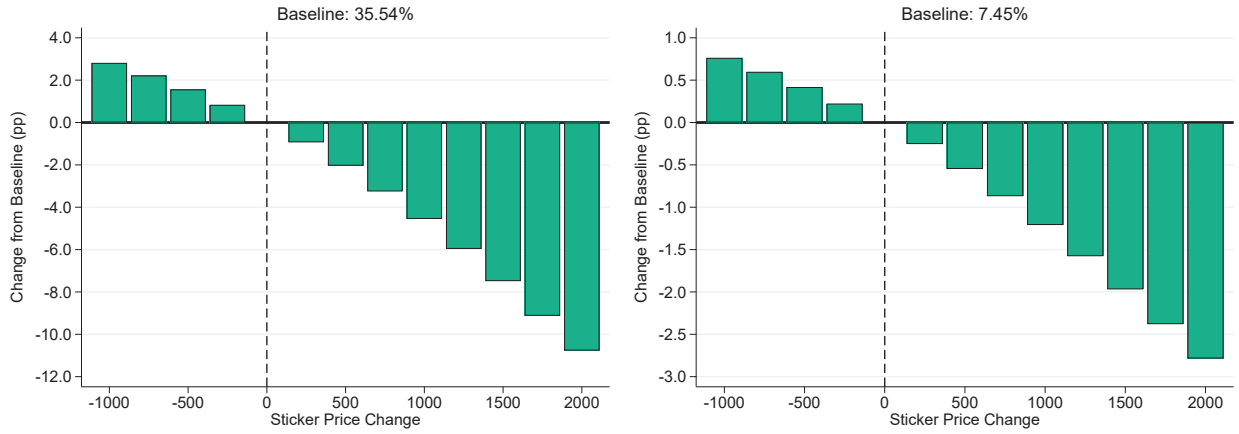
[94]For changes to the part-time prices, I compute a credit-equivalent price reduction and apply this to all enrolling workers.

[95]These numbers do not include the lock-in effect from enrolling, resulting from part-time penalties while enrolled or taking time away from work while studying. Excluding zeros or focusing on the discounted version of the realized earnings streams gives similar results.

enroll in any period between 2009 and 2015 by about 2.9pp and the annual rate by 0.8pp. By comparison, a $1,000 price increase lowers enrollment in 2009 to 2015 by 4.1 percent and the annual rate by 1.3 percent. The seeming asymmetry between a price increase and price decrease is partially due to the fact that individuals with financial aid are less affected by price declines but will be affected by price increases. For reference, 35.5 percent of workers ever enroll and 7.5 percent enroll annually in the baseline simulation, indicating that these are substantial changes.

Figure 1.15: Take-up and Earnings Effects for Varying Tuition Subsidy Levels

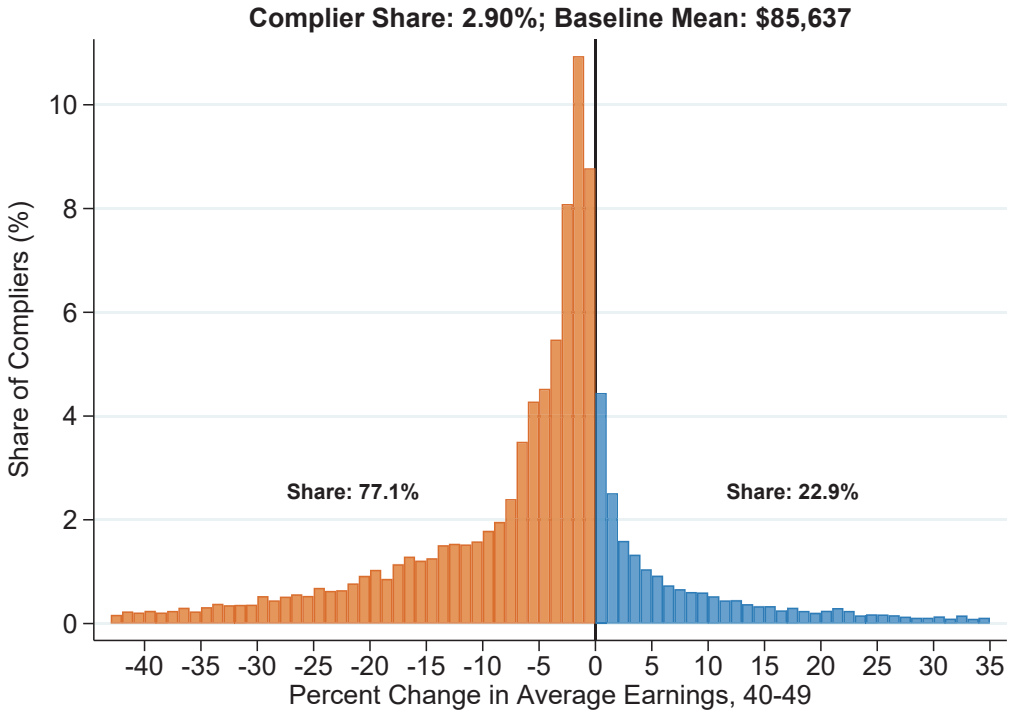
(a) Ever Enroll in 2009 to 2015 (b) Annual Enrollment in 2009 to 2015



Notes: This figure shows the changes in the probability of enrollment in 2009 and 2015 from an unexpected and permanent decrease or decrease in the sticker in $2009. Each bar looks at differences of $250, comparing outcomes for in two simulated datasets of 1,000,000 individuals, holding random realizations fixed across datasets, only changing the price. For changes to the part-time prices, I compute a credit-equivalent price reduction by multiplying the subsidy by 18/30, where 30 credits is full-time and 18 credits is the average for part-time in the data.

Figure 1.16 narrows the focus to those induced to shift behavior by the $1,000 price cut, focusing on earnings changes between the baseline and counterfactual scenario. Most individuals experience earnings losses from enrollment: only 22.9 percent of those enrolling due to the price cut see earnings increases. These changes in earnings are clustered near zero. However, the negative earnings losses have a long left-tail, resulting in an average earnings loss for compliers of -3.8 percent and a median loss of 2.9 percent. While negative, the wide variation in earnings changes also indicates the potential for better targeting.

Figure 1.16: Earnings Effects for those shifted by Tuition Subsidy



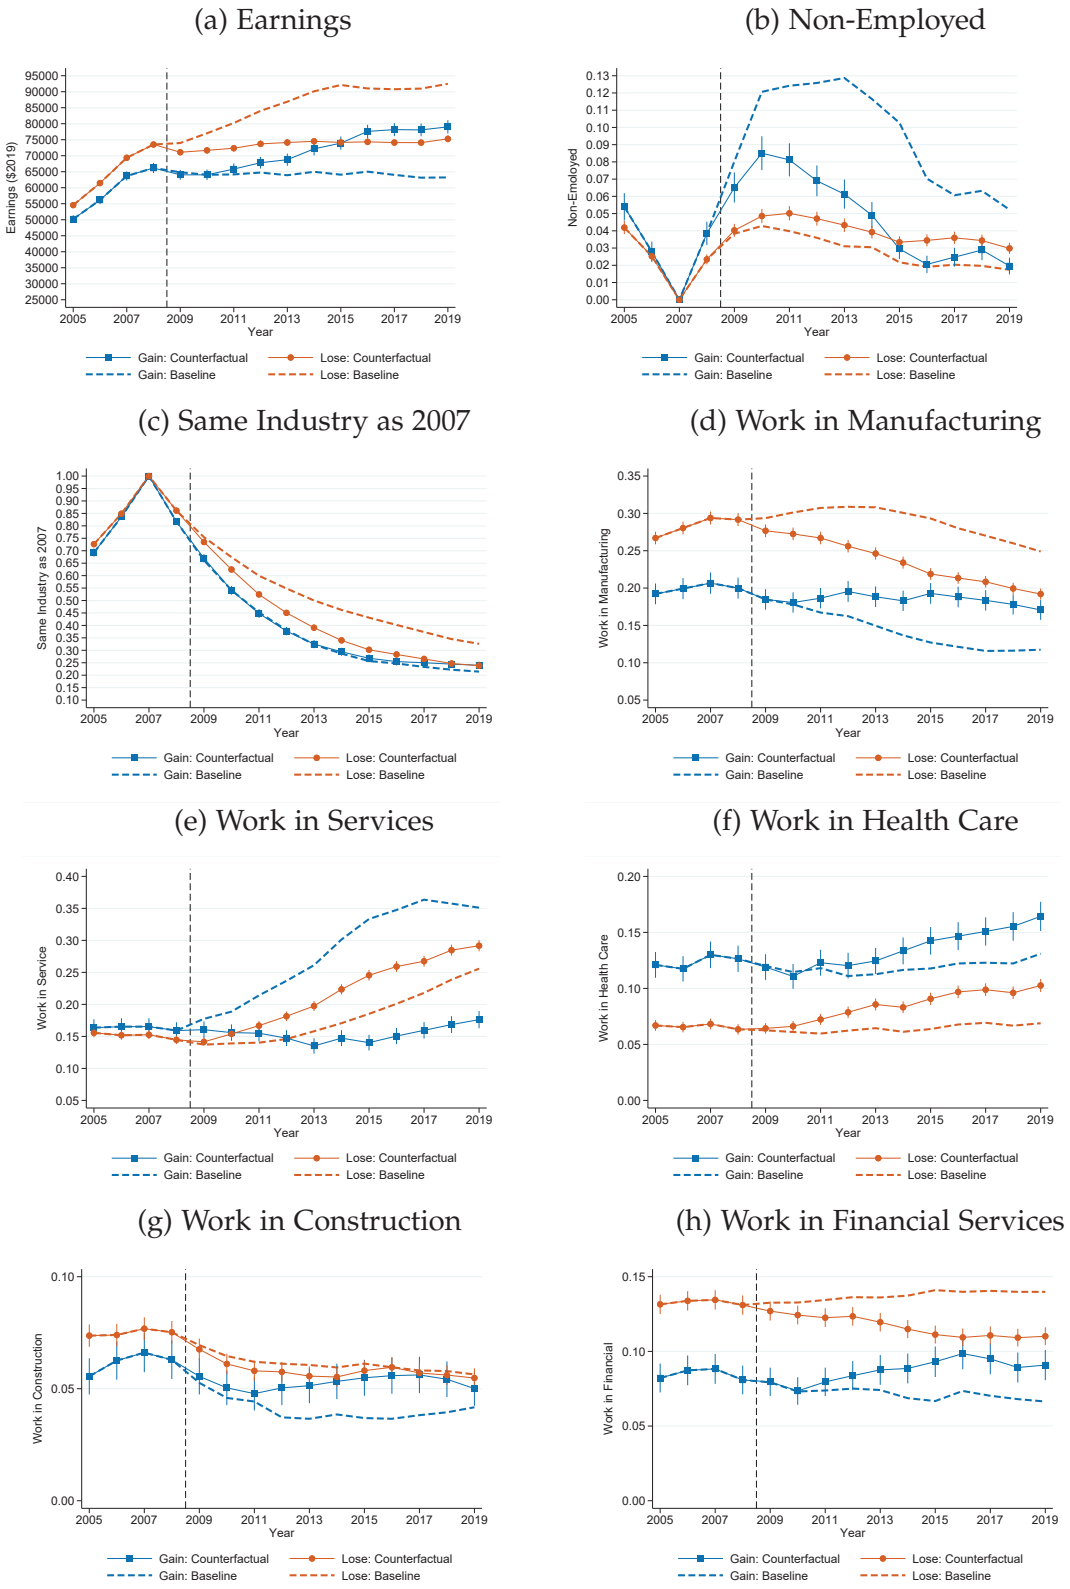
Notes: To construct this figure, I sum all earnings between the ages of 40 and 49 in the baseline and counterfactual scenario, for each simulated individual induced to change behavior as a result of the $1,000 price decrease. This holds fixed the sequence of idiosyncratic draws in both simulated datasets. This figure focuses on earnings between the ages of 40 to 49, but other outcomes give similar results.

Figure 1.17 investigates the drivers of this heterogeneity by summarizing the pathways taken for those who benefit in terms of earnings gains (defined as greater than 5 percent) and those who lose in terms of earnings (defined as more than a 5 percent earnings loss), comparing pathways under both the baseline and counterfactual for these groups.[96] The most striking pattern is in Panel (b) - those who benefit from the policy would have been much more likely to be non-employed between 2009 and 2015 under the baseline scenario. There is a marked reduction of non-employment in 2009 of about 4pp that persists through the rest of the panel for this group. In contrast, those who lose earnings see a slight increase in their rate of non-employment. Panel (c) shows that these individuals are much less likely to be employed in their baseline industry, while those who benefit see little change in this probability. The remaining panels show differences in the sectoral choices for these groups. Those who lose earnings tend to be more likely to be employed in manufacturing and financial services. There is a striking change in the decline in employment in the service sector in the post-2015 era for those who benefit, declining from around 35 percent to around 15 percent. The other sectoral probabilities show that those who gain tend to move out of manufacturing and financial services.

Figure 1.18 takes a closer look at how earnings changes between the benchmark and counterfactual scenarios relate to the probability of remaining in one's pre-Recession industry. Those who see earnings losses were induced to move out of their baseline industry by the subsidy, dropping from 40 percent to 20 percent. In contrast, those who benefit were going to have a lower baseline rate of remaining in their 2007 industry, but this is not affected by the counterfactual policy. If anything, those with very high earnings gains are more likely to remain in their baseline industry under the policy. These results give strong indication that up-skilling may be more effective than re-skilling, at least in this context.

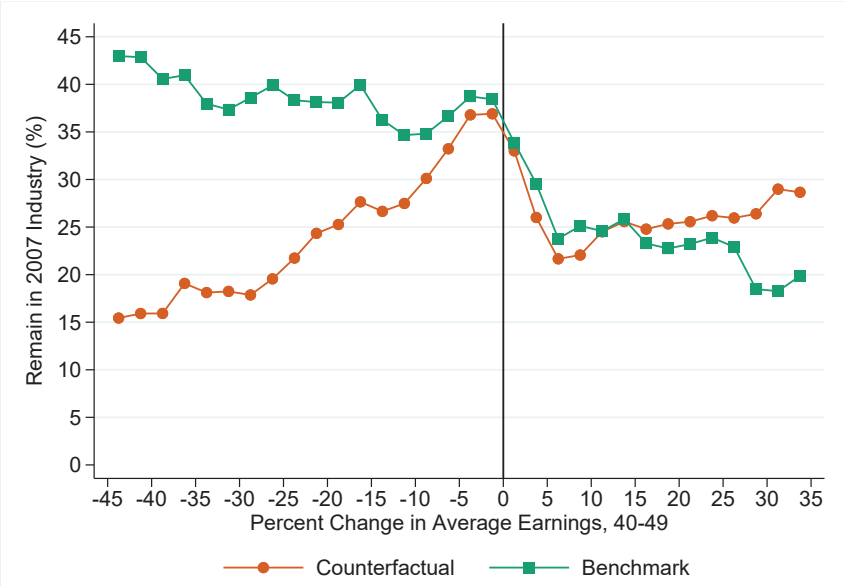
[96] Since the policy is a surprise, these groups have the same path prior to 2009 by construction. Also by construction, Panel (a) shows that those who gain from the policy see large increases in earnings while those that lose see large declines.

Figure 1.17: Sources of Earnings Gains and Losses for those Induced to Enroll



Notes: Plots history of variables for simulated individuals from Figure 1.16. Gain plots for those whose earnings increase by more than 5 percent; while losers are defined as earnings losses of more than 5 percent. Dashed lines plot these series under the baseline environment; solid lines plot for individuals in the counterfactual with price decreases. Error bars indicate simulation error resulting from sampling from simulated data.

Figure 1.18: Probability of Remaining in Baseline Industry by Earnings Change



Notes: Plots the probability of remaining in 2007 industry in 2016, in the benchmark and counterfactual simulated data. Earnings changes are grouped into 2.5pp bins and smoothed over one neighbor in each direction.

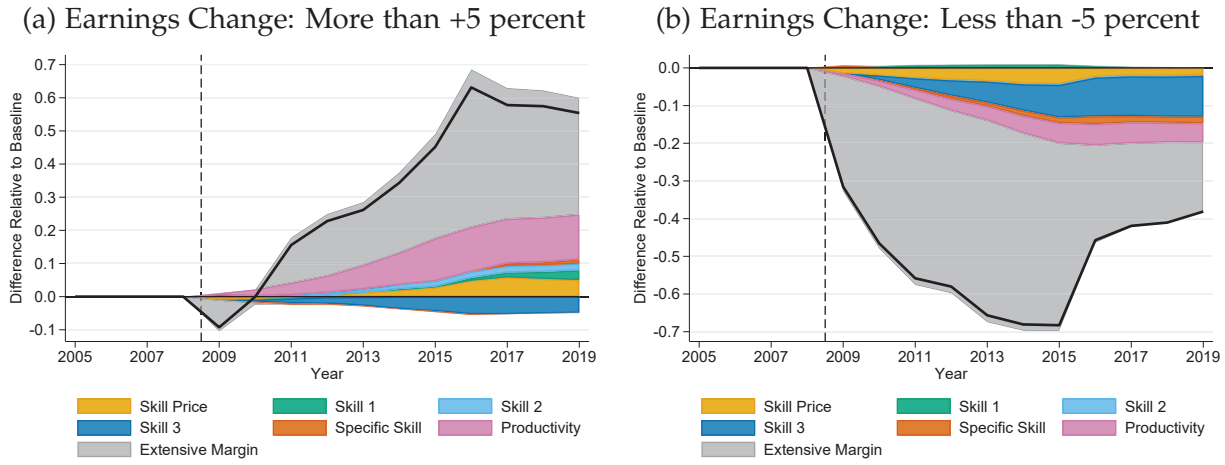
Figure 1.19 decomposes changes in log-earnings into the components outlined in Equation 1.4.1.[97] A large share of the difference is driven by participation: those who benefit end up working more and those who lose earnings tend to work less. However, there are also key patterns on the intensive margin which become more important over time. First, all workers have counterfactually lower values of skill 3, which is the skill least rewarded in health care and the public sector. Second, those who benefit have much higher levels of skill prices and idiosyncratic productivity in later years, indicating that they are able to build these components over time through post-schooling work transitions. For these workers, education has put them on a path to working in higher earning industries and jobs. In contrast, those who lose earnings have lower values of

[97] I focus on log-earnings to keep the decomposition linear. A key issue is dealing with the extensive margin differences in earnings, since Figure 1.17 shows differences in participation is important. The extensive margin component of these decompositions reflect differences in participation holding average log-earnings conditional on participation to the baseline scenario: $\mathbb{E}[y^0|D^0 = 1] \cdot \left(P(D^1) - P(D^0)\right)$, where y^0 and D^0 are earnings and participation under the baseline scenario. This follows from the decomposition:

$$\mathbb{E}[y^1] - \mathbb{E}[y^0] = \mathbb{E}[y^1|D^1 = 1]P(D^1) - \mathbb{E}[y^0|D^0 = 1]P(D^0) = \left(\mathbb{E}[y^1|D^1 = 1] - \mathbb{E}[y^0|D^0 = 1]\right) P(D^1) + \mathbb{E}[y^0|D^0 = 1] \left(P(D^1) - P(D^0)\right)$$

these components, indicating they are knocked off the job ladder after failing to find a job. Importantly, skill 1 and skill 2, which are modified by schooling, increase but are a relatively small share of changes in earnings components. This indicates that much of what schooling does for those who benefit is to ease transitions back into the workforce, although it does increase earnings directly as well.

Figure 1.19: Decomposition of Sources of Earnings Gains and Losses

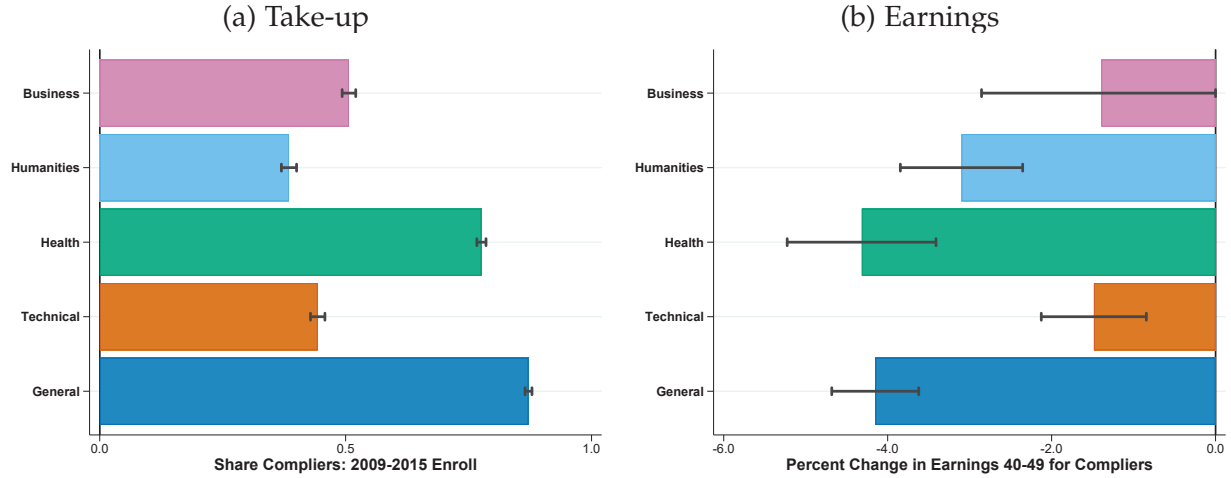


Notes: Decomposes difference in earnings between the baseline scenario and the counterfactual in terms of underlying model components.

Figure 1.20 investigates the potential role for field-specific subsidies to increase earnings. Each row in the figure is a different counterfactual that only subsidies a specific field of study beginning in 2009. Fields are substantially different in both the propensity for individuals to respond as well as in earnings effects. Panel (a) shows the share of the data that are compliers with each field-specific subsidy. Health and general studies both result in higher rates of take-up, mirroring the event study evidence. However, when looking at earnings effects in Panel (b), all fields have negative effects, with some key differences. In particular, health and general studies lead to larger earnings losses, but technical studies yields the least negative earnings losses. Business subsidies have effects on the order of technical studies but with a lot of variation.

Taken together, I interpret these results as indicating the substantial nuance surrounding education for adult workers. First and foremost, education is no substitute for direct

Figure 1.20: Take-up and Earnings Effects for Field-of-Study Specific Subsidies



Notes: Error bars are computed from simulated dataset of 1,000,000 , and thus represent simulation error. The large role for simulation error is due to the a large variance in earnings impacts and also the low rate take-up of the policy. In principle, this could be eliminated by taking more draws, but the number of draws already taken is cumbersome to manage.

work experience, with much of the positive effects of education being driven by its indirect effect on work transitions. On the one hand, this means that it can be effective for the long-term unemployed who face difficulty acquiring this experience, building a pathway back to previous types of work. On the other hand, re-skilling through education can be a risky investment for those who do not face such dire circumstances. Underscoring this risk, fields of study that deliver general returns, such as technical training, seem to be more effective than fields associated with high but specific returns. If individuals do not finish their education with a job lined up, pursuing later-life education may be quite detrimental.

1.8.2 Gender Differences in Enrollment and Earnings Trajectories

The descriptive analysis shows that men and women make quite different labor supply and education choices, with women enrolling at twice the rate of men. The model estimates of preferences and comparative advantage that vary by gender and prior education can be used to provide some insight into drivers of these differences. However, care should be taken in interpreting these numbers as gender-specific preferences cap-

ture any mechanism that drives differences in transitions net of earnings gains between men and women, not just gender-specific amenities.[98] At the same time, differences in earnings parameters only capture the degree to which groups are more or less able to convert their work and schooling history into earnings, which can be due to many factors.[99] Nevertheless, these numbers are useful in understanding the degree to which earnings gaps are due to pathways taken versus earnings differences within the same pathway.

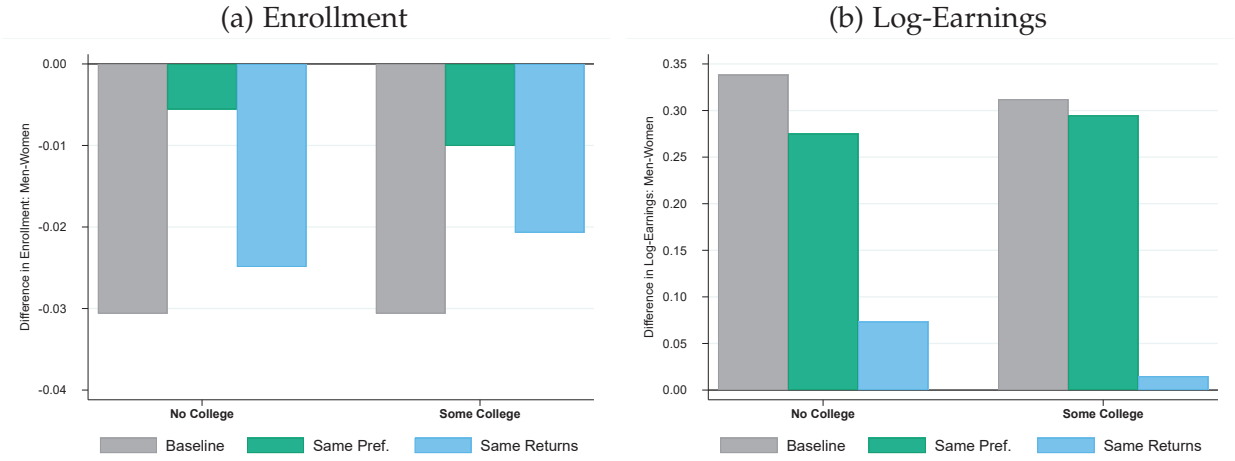
Panel (a) of Figure 1.21 investigates whether comparative advantage or preferences drives differences in enrollment. I sequentially shut down gender differences in preferences and earnings parameters, by assigning women to the same parameter values as men with the same initial education. Most of the difference in enrollment is due to gender-specific preferences rather than gender-specific returns. The gray bar shows the gap in the baseline simulated data, indicating that men enroll 3.1pp less than women of the same education level. When equalizing preferences between men and women, the enrollment gap shrinks to 0.5pp for non-college and 1.0pp for college. This remaining portion is due to differences in gender-specific returns, accounting for 18 percent and 33 percent of the enrollment gap.

Panel (b) of Figure 1.21 conducts a similar exercise to analyze the gender earnings gap. Men earn 0.3 log-points more than women on average at baseline. Equalizing preferences within education groups shrinks this earnings gaps slightly, indicating that women have slight preferences for working in industries for which they have lower earnings growth. However, the majority of the gap is driven by differences in earnings potential between men and women given the same preference parameters, accounting for 78 percent of the earnings gap among the non-college group and 95 percent of the

[98] For example, if men lack information about jobs in health care, this will picked up as a man-specific disutility of working in health care. Similarly, if women are discriminated in hiring into construction, this will show up as a woman-specific disutility of construction.

[99] For example, if women receive within-industry discrimination in promotions, this will be captured in a lower return to experience.

Figure 1.21: Decomposition of Gender Differences in Enrollment and Earnings



Notes: These figures shows differences in enrollment and log-earnings (if working) between men and women of the same education level under the baseline scenario (gray), the scenario setting women's transition parameters to men's within the same pre-determined education level (green), and the scenario setting women's earnings parameters to men's within the same pre-determined education-level (blue). These figures do not adjust for differences in the covariance between gender and initial conditions. Based on 1,000,000 simulated individuals.

earnings gap among the college group.

1.9 Conclusion

In this paper, I study the trade-offs that workers face when pursuing later-life human capital investment. The analysis combines rich administrative on earnings and enrollment records, a large shock that drives adult enrollment in the form of the Great Recession, and a structural model of worker choices to adjust to shocks by switching industries or returning to school. The modeling framework takes seriously that a worker's benefits from returning to school are heterogeneous, depending on their preferences and past work experience, while addressing selection issues documented in the prior literature.

We learn that there is substantial nuance that determines the effectiveness of training and not all workers should use education to adjust to shocks. I find both that workers are responsive to price reductions but that the average worker who responds to the counterfactual sees earnings declines. Re-skilling can be risky if it does not translate into a job in the target sector and returning to work quickly, often in the origin sector, is

more effective for those who are able to do so. However, this means that the population that benefits the most from schooling are workers who would have had long periods out of work in the absence of returning to school. For these workers, education both helps build valuable skills and serves as a path back onto the job ladder. This may explain why studies that condition on displaced or unemployed workers find positive effects (e.g., Jacobson et al. [2005]). The structural model is necessary to uncover this fact as future industry switching is an endogenous outcome that can be affected by education directly.

Some extensions should be kept in mind for future research but are outside the scope of this paper. First, the framework does not deal with equilibrium considerations, analyzing choices for five cohorts of high-school graduates while holding the path of wage-setting equilibria fixed. Accounting for such equilibrium adjustment would require an overlapping generations model of the labor market. To my knowledge, the only paper that deals with this issue in the context of higher education is Heckman et al. [1998b] and there is scope to consider these issues in the context of multi-dimensional human capital. Second, the framework does not incorporate spatial considerations: in reality, workers may adjust to shocks by moving to less-shocked local labor markets. However, low-skill workers do not seem to use this margin to respond to shocks (Bound and Holzer [2000]), a fact that is especially true in the Great Recession (Yagan [2019]). Finally, information frictions may drive some of the results as preferences in my model do not distinguish between information frictions and true amenity values from studying. The effectiveness of training policies could be enhanced by providing workers information about types of training, economic conditions in the target industry, or worker fit for a specific field-industry pathway, topics that should be explored in future work.

Bibliography

Atila Abdulkadiroglu, Parag A. Pathak, Jonathan Schellenberg, and Christopher R. Walters. Do parents value school effectiveness? *American Economic Review*, page 1502–1539, 2020.

John M. Abowd, Francis Kramarz, and David N. Margolis. High wage workers and high wage firms. *Econometrica*, 67(2):251–333, 1999.

Riley K. Acton. Community college program choices in the wake of local job losses. *Journal of Labor Economic*, 39(4):1129–1154, 2021.

Jerome Adda, Christian Dustmann, Costas Meghir, and Jean-Marc Robin. Career progression, economic downturns, and skills. *NBER Working Paper no. 18832*, 2013.

Gregor Aisch, Robert Gebeloff, and Kevin Quealy. Where we came from and where we went, state by state. *The New York Times*, 2014.

Joseph G. Altonji and Ling Zhong. The labor market returns to advanced degrees. *Journal of Labor Economics*, 39(2):303–360, 2021.

Joseph G. Altonji, Peter Arcidiacono, and Arnaud Maurel. The analysis of field choice in college and graduate school: Determinants and wage effects. *Handbook of the Economics of Education*, 5:305–396, 2016a.

Joseph G. Altonji, Lisa B. Kahn, and Jamin D. Speer. Cashier or consultant? entry labor market conditions, field of study, and career success. *Journal of Labor Economics*, 34(1): S361–S401, 2016b.

Fredrik Andersson, Harry J. Holzer, Julia I. Lane, David Rosenblum, and Jeffrey Smith. Does federally-funded job training work? nonexperimental estimates of wia training impacts using longitudinal data on workers and firms. *NBER Working Paper no 19446*, 2013.

Peter Arcidiacono. Affirmative action in higher education: Do admission and financial rules affect future earnings? *Econometrica*, 73(5):1477–1524, 2005.

Peter Arcidiacono and Robert A. Miller. Conditional choice probability esimtation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6): 1823–1867, 2011.

Peter Arcidiacono and Robert A. Miller. Identifying dynamic discrete choice models off short panels. *Journal of Econometrics*, 215:473–485, 2020.

Peter Arcidiacono, Esteban Aucejo, Arnaund Maurel, and Tyler Ransom. College attrition and the dynamics of information relevation. *mimeo*, 2024.

Erhan Artuc, Shubham Chaudhuri, and John McLaren. Trade shocks and labor market adjustment. *American Economic Review*, 100(3):1008–1046, 2010.

Orley Ashenfelter. Estimating the effect of training programs on earnings. review of economics and statistics. *Review of Economics and Statistics*, 60(1):47–57, 1978.

Orley Ashenfelter and David Card. Using the longitudinal structure of earnings to estimate the effect of training programs. *The Review of Economics and Statistics*, 67(4): 648–660, 1985.

David H. Autor, David Dorn, and Gordon H. Hanson. The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–2168, 2013.

David H. Autor, David Dorn, Gordon H. Hanson, and Jae Song. Trade adjustment: Worker-level evidence. *Quarterly Journal of Economics*, 129(4):1799–1860, 2014.

Jesper Bagger, Francois Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics. *American Economic Review*, 104(6):1551–1596, 2014.

Andrew Barr and Sarah Turner. Expanding enrollments and contracting state budgets: The effect of the great recession on higher education. *The ANNALS of the American Academy of Political and Social Science*, 650(1):168–193, 2013.

Andrew Barr and Sarah Turner. Out of work and into school: Labor market policies and college enrollment during the great recession. *Journal of Public Economics*, 124:63–73, 2015.

Lisa Barrow and Jonathan Davis. The upside of down: Postsecondary enrollment in the great recession. *Economic Perspectives*, 2012.

Clive Belfield and Thomas Bailey. The labor market returns to sub-baccalaureate college: A review. *Center for Analysis of Postsecondary Education and Employment*, pages 1–31, 2017.

Julian R. Betts and Laurel L. McFarland. Safe port in a storm: The impact of labor market conditions on community college enrollments. *The Journal of Human Resources*, 30(4): 741–765, 1995.

Peter Q. Blair and David J. Deming. Structural increases in demand for skill after the great recession. *American Economic Association Papers and Proceedings*, 110:362–365, 2020.

Erica Blom, Brian C. Cadena, and Benjamin J. Keys. Investment over the business cycle: Insights from college major choice. *Journal of Labor Economics*, 39(4):1043–1082, 2021.

Stephane Bonhomme, Thibaut Lamadon, and Elena Manresa. A distributional framework for matched employer-employee data. *Econometrica*, 87(3):699–739, 2019.

John Bound and Harry J. Holzer. Demand shifts, population adjustments, and labor market outcomes during the 1980s. *Journal of Labor Economics*, 18(1):20–54, 2000.

Leah Platt Boustan, Jiwon Choi, and David Clingingsmith. Automation after the assembly line: Computerized machine tools, employment, and productivity in the united states. *National Bureau of Economics Working Paper no. 30400*, 2022.

Marco Caliendo, Robert Mahlstadt, and Oscar A. Mitnik. Unobservable, but unimportant? the relevance of usually unobserved variables for the evaluation of labor market policies. *Labour Economics*, 46:14–25, 2017.

David Card, Jörg Heining, and Patrick Kline. Workplace heterogeneity and the rise of

west german wage inequality. *The Quarterly Journal of Economics*, 128(3):967–1015, 2013.

David Card, Jochen Kluve, and Andrea Weber. What works? a meta analysis of recent active labor market program evaluations. *Journal of the European Economic Association*, 16(3):894–931, 2018.

Kerwin Kofi Charles, Erik Hurst, and Matthew J. Notowodigdo. Housing booms and busts, labor market opportunities, and college attendance. *American Economic Review*, 108(10):2947–2994, 2018.

Kenneth A. Couch and Dana W. Placzek. Earnings losses of displaced workers revisited. *American Economic Review*, 100(1):572–589, 2010.

Steven J. Davis and Till M. von Wachter. Recessions and the cost of job loss. *National Bureau of Economic Research Working Paper no. 17638*, 2011.

Rajeev H. Dehejia and Sadek Wahba. Causal effects in nonexperimental settings: Reevaluating the evaluation of training programs. *Journal of the American Statistical Association*, 94(448):1053–1062, 1999.

David J. Deming. The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4):593–640, 2017.

David J. Deming. Why do wages grow faster for educated workers? *National Bureau of Economic Research Working Paper no. 31373*, 2023.

David J. Deming and Kadeem Noray. Earnings dynamics, job skills, and stem careers. *The Quarterly Journal of Economics*, 135(4):1965–2005, 2022.

Rafael Dix-Carneiro. Trade liberalization and labor market dynamics. *Econometrica*, 82 (3):825–885, 2014.

Dita Eckardt. Are chemists good bankers? returns to the match between training and occupation. *Working Paper*, 2022.

Fulya Y. Esroy. The effects of the great recession on college majors. *Economics of Education Review*, 77:1–20, 2020.

Andrew Foote and Michel Grosz. The effect of local labor market downturns on postsecondary enrollment and program choice. *Education and Finance Policy*, pages 593–622, 2019.

Kenneth Fortson, Paul Burkander, Annalisa Mastri, Peter Schochet, Linda Rosenberg, Sheena McConnel, and Ronald D'Amico. Providing public workforce services to job seekers: 30-month impact findings on the wia adult and dislocated worker programs. *Mathematica Policy Research*, 2021.

Eric French and Christopher Taber. Identification of models of the labor market. *Handbook of Labor Economics*, page 537–617, 2011.

Christina Gathmann and Uta Schoenberg. How general is human capita? a task-based approach. *Journal of Labor Economics*, 28(1):1–49, 2010.

John Grigsby. Skill heterogeneity and aggregate labor market dynamics. *Working Paper*, 2022.

Fatih Guvenen, Burhan Kuruscu, Satoshi Tanaka, and David Wiczer. Multidimensional skill mismatch. *American Economic Journal: Macroeconomics*, 12(1):210–244, 2020.

Justine Hastings, Chrisopher Neilson, and Seth Zimmerman. The effects of earnings disclosure on college enrollment decisions. *National Bureau of Economics Working Paper*

no. 21300, 2015.

James J. Heckman and Salvador Navarro. Dynamic discrete choice and dynamic treatment effects. *Journal of Econometrics*, 136:341–396, 2007.

James J. Heckman and Richard Robb. Alternative methods for evaluating the impact of interventions: An overview. *Journal of Econometrics*, 30:239–267, 1985.

James J. Heckman and Jeffrey A. Smith. The pre-programme earnings dip and the determinants of participation in a social programme. implications for simple programme evaluation strategies. *The Economic Journal*, 109:313–348, 1999.

James J. Heckman, Hidehiko Ichimura, Jeffrey Smith, and Petra Todd. Characterizing selection bias using experimental data. *Econometrica*, 66(5):1017–1098, 1998a.

James J. Heckman, Lance Lochner, and Christopher Taber. Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents. *Review of Economic Dynamics*, 1:1–58, 1998b.

James J. Heckman, Robert J. Lalonde, and Jeffrey A. Smith. The economics and econometrics of active labor market programs. *Handbook of Labor Economics*, 3A:1865–2097, 1999.

James J. Heckman, John Eric Humphries, and Gregory Veramendi. Dynamic treatment effects. *Journal of Econometrics*, 191:276–292, 2016.

James J. Heckman, John Eric Humphries, and Gregory Veramendi. Returns to education: The causal effects of education on earnings, health, and smoking. *Journal of Political Economey*, 126(S1):S197–S246, 2018.

Brad Hershbein and Lisa B. Kahn. Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *The American Economic Review*, 108(7):1737–1772, 2018.

V. Joseph Hotz and Robert A. Miller. Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies*, 60:497–529, 1993.

Christopher Huckfeldt. Understanding the scarring effects of recessions. *American Economic Review*, 112(4):1273–1310, 2022.

Anders Humlum, Jakob R. Munch, and Pernille Plato. Changing tracks: Human capital after the loss of ability. *Working paper*, 2023.

Benjamin G. Hyman. Can displaced labor be retrained? evidence from quasi-random assignment to trade adjustment assistance. *Working paper*, 2018.

Louis Jacobson, Robert J. LaLonde, and Daniel Sullivan. Earnings losses of displaced workers. *American Economic Review*, 83(4):685–709, 1993.

Louis Jacobson, Robert J. LaLonde, and Daniel Sullivan. The impact of community college retraining on older displaced workers: Should we teach old dogs new tricks? *Industrial and Labor Relations Review*, 58(3):398–415, 2005.

Gregor Jarosch. Searching for job security and the consequences of job loss. *Econometrica*, 91(3):903–942, 2023.

Robert Jensen. The (perceived) returns to education and the demand for schooling. *The Quarterly Journal of Economics*, 125(2):515–548, 2010.

Christopher Jepsen, Kenneth Troske, and Paul Coomes. The labor-market returns to community college degrees, diplomas, certificates. *Journal of Labor Economics*, 32(1):

95–121, 2014.

Thomas J. Kane and Cecilia E. Rouse. Labor market returns to two and four-year college. *American Economic Review*, 85(3):600–614, 1995.

Lawrence F. Katz, Jonathan Roth, Richard Hendra, and Kelsey Shaberg. Why do sectoral employment programs work? lessons from workadvance. *Journal of Labor Economics*, 40(S1):S249–S291, 2022.

Michael P. Keane and Kenneth I. Wolpin. The career decisions of young men. *Journal of Political Economy*, 105(3):473–522, 1997.

Michael P. Keane, Petra E. Todd, and Kenneth I. Wolpin. The structural estimation of behavioral models: Discrete choice dynamic programming methods and applications. *Handbook of Labor Economics*, 4A:331–461, 2011.

John Kennan and James R. Walker. The effect of expected income on individual migration decisions. *Econometrica*, 79(1):211–251, 2011.

Josh Kinsler and Ronni Pavan. The specificity of general human capital: Evidence from college major choice. *Journal of Labor Economics*, 33(4):933–972, 2015.

Marta Lachowska, Alexandre Mas, and Stephen A. Woodbury. Sources of displaced workers' long-term earnings losses. *American Economic Review*, 110(10):3231–3266, 2020.

Edward Lazear. Firm-specific human capital: A skill-weights approach. *Journal of Political Economy*, 117:914–940, 2009.

Donghoon Lee and Kenneth I. Wolpin. Intersectoral labor mobility and the growth of the service sector. *Econometrica*, 74(1):1–46, 2006.

Margaret Leighton and Jamin D. Speer. Labor market returns to college and major specificity. *European Economic Review*, pages 1–25, 2020.

Thomas Lemieux. Occupations, fields of study and returns to education. *Canadian Journal of Economics*, 47(4):1047–1077, 2014.

Pauline Leung and Zhuan Pei. Further education during unemployment. *Princeton University Industrial Relations Section Working Paper no. 642*, 54(4):986–1036, 2020.

Jeremy Lise and Fabien Postel-Vinay. Multidimensional skills, sorting, and human capital accumulation. *American Economic Review*, 110(8):2328–2376, 2020.

Shimeng Liu, Sun Weizeng, and John V. Winters. Up in stem, down in business: Changing major decisions. *Contemporary Economic Policy*, 37:476–491, 2019.

Thierry Magnac and David Thesmar. Identifying dynamic discrete decision processes. *Econometrica*, 70(2):801–816, 2002.

Brian McCall, Jeffrey Smith, and Conny Wunsch. Government-sponsored vocational education for adults. *Handbook of the Economics of Education*, 5:479–652, 2016.

Robert A. Miller. Job matching and occupational choice. *Journal of Political Economy*, 92 (6):1086–1120, 1985.

Veronica Minaya, Brendan Moore, and Judith Scott-Clayton. The effect of job displacement on public college enrollment: Evidence from ohio. *Education of Economics Review*, 92:1–12, 2023.

Alicia Sasser Modestino, Daniel Shoag, and Joshua Ballance. Upskilling: Do employers demand greater skill when workers are plentiful? *Review of Economics and Statistics*,

102(4):793–805, 2020.

Derek Neal. Industry-specific human capital: Evidence from displaced workers. *Journal of Labor Economics*, 13:653–677, 1995.

Derek Neal. The complexity of job mobility among young men. *Journal of Labor Economics*, 17(2):237–261, 1999.

Theodore Papageorgiou. Learning your comparative advantages. *Review of Economic Studies*, 81:1263–1295, 2014.

Daniel Parent. Industry-specific capital and the wage profile: Evidence from the national longitudinal survey of youth. *Journal of Labor Economics*, 18(2):306–323, 2000.

Fabien Postel-Vinay and Jean-Marc Robin. Equilibrium wage dispersion with worker heterogeneity. *Econometrica*, 70(6):2295–2350, 2002.

Tyler Ransom. Labor market frictions and moving costs of the employed and unemployed. *Journal of Human Resources*, 57(S):S137–S166, 2022.

Jorge Rodriguez, Fernando Saltiel, and Sergio Urzua. Dynamic treatment effects of job training. *Journal of Applied Econometrics*, 37(2):242–269, 2022.

A.D. Roy. Some thoughts on the distribution of earnings. *Oxford Economic Papers*, 3(2): 135—146, 1951.

Nicolas A. Roys and Christopher R. Taber. Skill prices, occupations, and changes in the wage structure for low skilled men. *National Bureau of Economic Research Working Paper no. 26453*, 2019.

John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033, 1987.

Carl Sanders. Skill accumulation, skill uncertainty, and occupational choice. *Working Paper*, 2014.

Diane Whitmore Schazenbach, Julia A. Turner, and Sarah Turner. Raising state minimum wages, lowering community college enrollment. *National Bureau of Economic Research working paper no. 31540*, 2023.

Ann Huff Stevens, Michal Kurlaender, and Michel Grosz. Career technical education and labor market outcomes. *Journal of Human Resources*, 54(4):986–1036, 2019.

Andrew Sweeting. Dynamic product positioning in differentiated product markets: The effects of fees for musical performance rights on the commercial radio industry. *Econometrica*, 81(5):1763–1803, 2013.

Christopher Taber and Rune Vejlin. Estimation of a roy/search/compensating differential model of the labor market. *Econometrica*, 88(3):1031–1069, 2020.

Sharon Traiberman. Occupations and import competition: Evidence from denmark. *American Economic Review*, 109(12):4260–4301, 2019.

Christopher R. Walters. The demand for effective charter schools. *Journal of Political Economy*, 126:2179–2220, 2018.

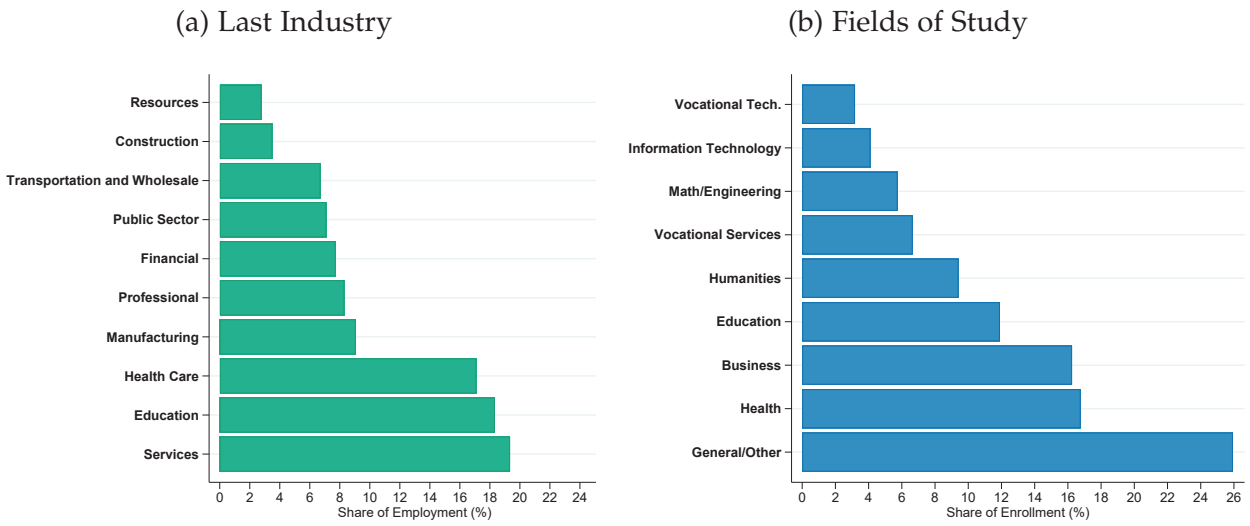
Danny Yagan. Employment hysteresis from the great recession. *Journal of Political Economy*, 127(5):2505–2558, 2019.

Shintaro Yamaguchi. Tasks and heterogeneous human capital. *Journal of Labor Economics*, 30(1):1–53, 2012.

A Additional Figures and Tables

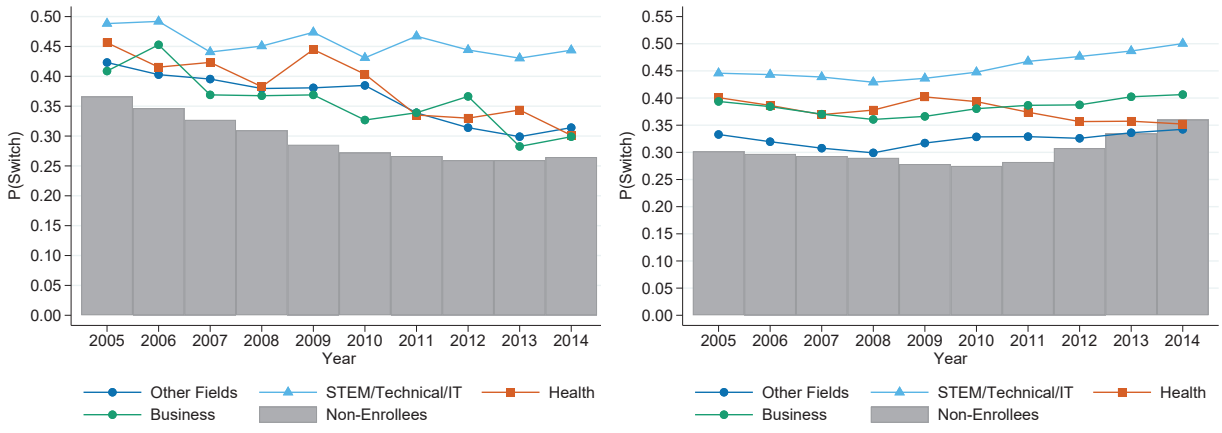
A.1 Additional Results

Figure A1: Field of study and industry shares of enrollment: Full Sample



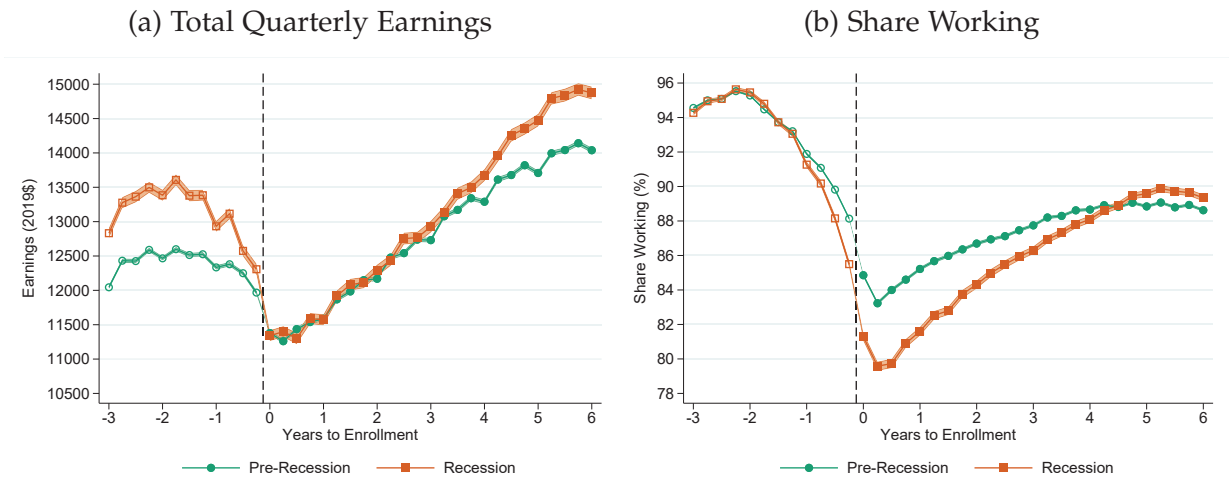
Notes: See notes to Figure 1.2.

Figure A2: Industry Switching and Field of Study over Time



Notes: See notes to Table 1.1.

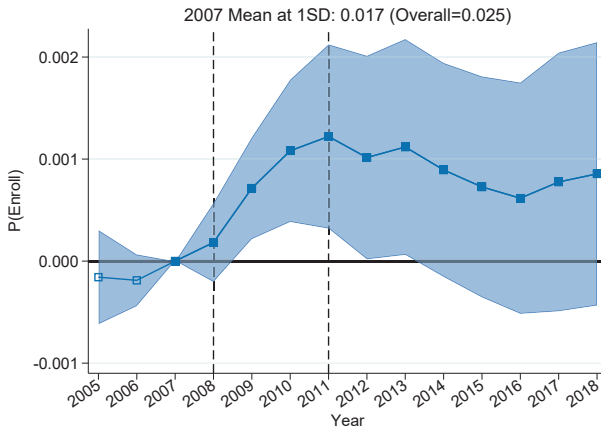
Figure A3: Employment and Earnings Trajectories around Enrollment Event: Full Sample



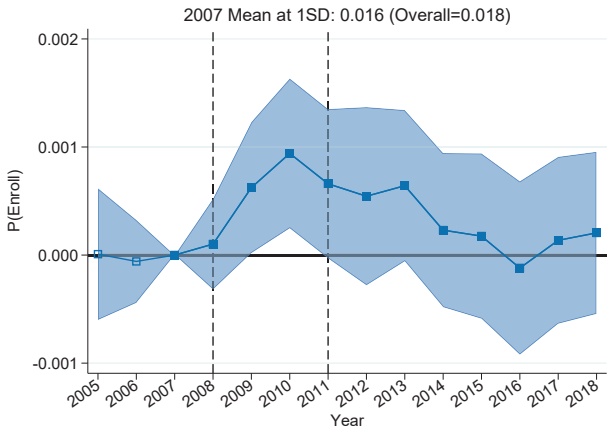
Notes: See notes to Figure 1.3.

Figure A4: Event Study Results using Industry-Level Panel of 3-Year Transition Probabilities

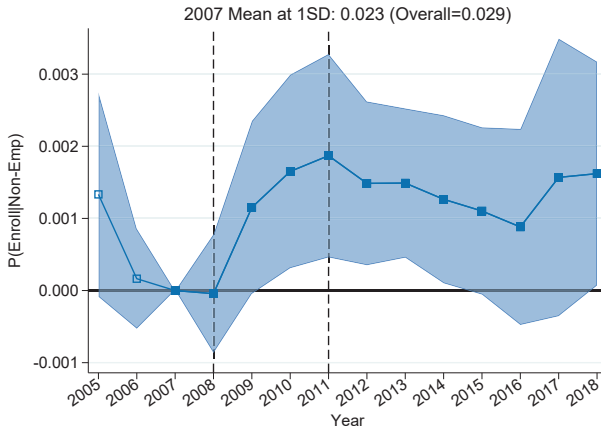
(a) Enrollment



(b) Enrollment - Unweighted



(c) Enroll after Non-employment



(d) Switch 1-Digit Industry after Enrollment

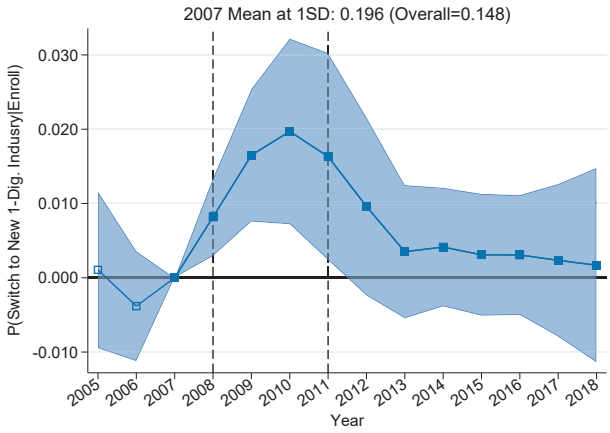


Figure A5: Decomposition of Skills in Each Industry and Time

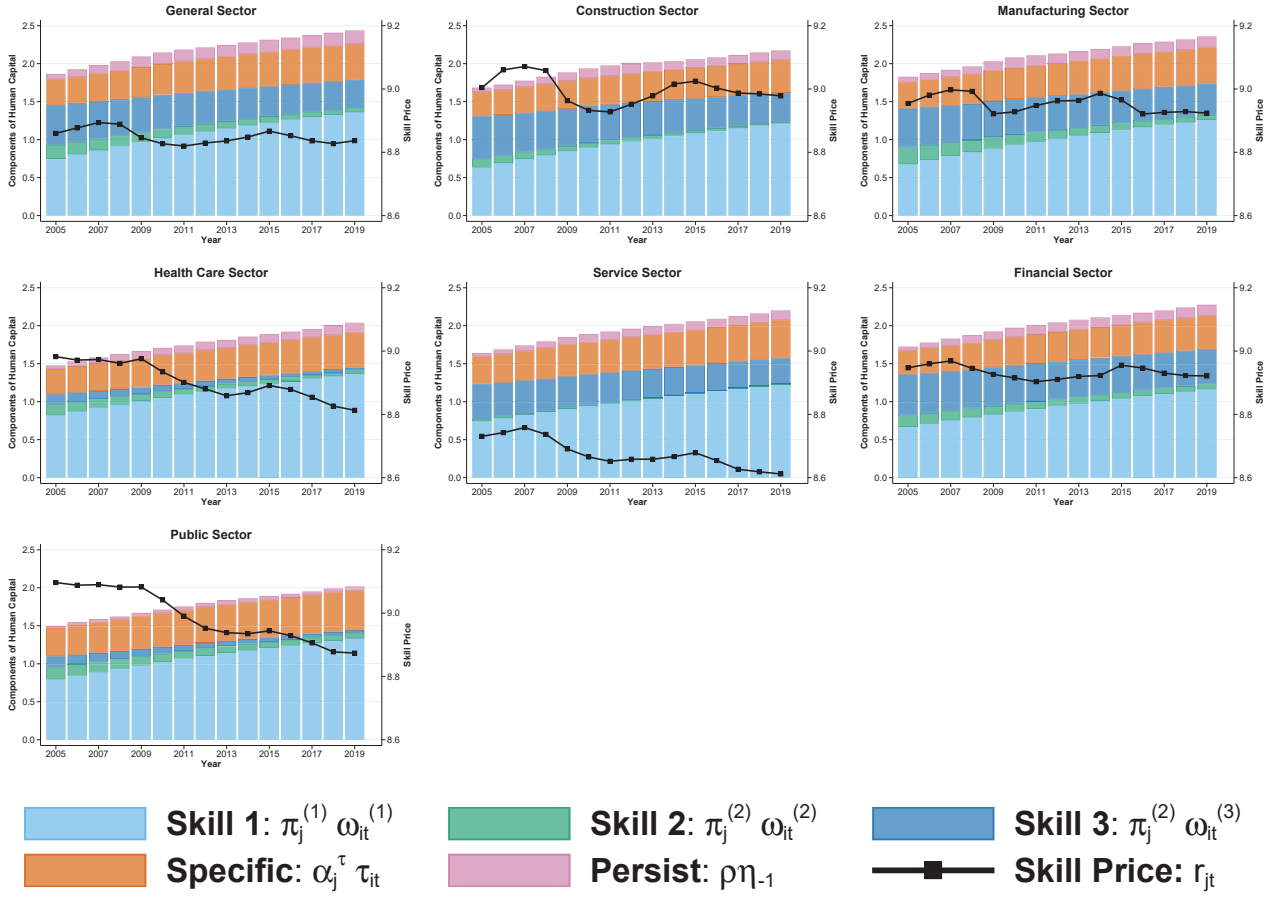


Table A1: Aggregate Enrollment in THECB Data

	Enrollment Count			*Increase from 2007-2008 (%)*		*Group Shares (%)*	
	2007-08	2010-11	2017-18	2007-08 to 2010-11	2007-08 to 2017-18	In 2007-08	Of 2007-08 to 2010-11 Change
A. By Gender							
Male	113,882	158,836	129,972	39.5	14.1	35.8	40.7
Female	204,076	269,558	213,810	32.1	4.8	64.2	59.3
B. By Prior Education							
No Prior College	121,433	155,450	105,091	28.0	-13.5	38.2	30.8
Prior College Attendance	170,057	235,239	196,192	38.3	15.4	53.5	59.0
Prior College Graduation	26,468	37,705	42,499	42.5	60.6	8.3	10.2
C. By Field of Study							
General/Other	104,242	131,289	112,400	25.9	7.8	32.8	24.5
Business	43,751	58,152	52,642	32.9	20.3	13.8	13.0
Humanities	26,048	28,085	20,996	7.8	-19.4	8.2	1.8
Math/Engineering	16,172	22,626	20,924	39.9	29.4	5.1	5.8
Health	72,548	106,289	71,112	46.5	-2.0	22.8	30.6
Vocational Tech.	8,123	12,627	11,325	55.4	39.4	2.6	4.1
Vocational Services	21,577	30,479	21,147	41.3	-2.0	6.8	8.1
Education	10,916	19,281	13,600	76.6	24.6	3.4	7.6
Information Technology	14,581	19,566	19,636	34.2	34.7	4.6	4.5
D. By Education Sector							
Community	317,958	428,394	343,782	34.7	8.1	61.6	76.1
Public 4-Year	190,200	219,439	222,889	15.4	17.2	36.9	20.1
Private	7,974	13,441	6,900	68.6	-13.5	1.5	3.8
E. By Age							
< 25	617,497	740,485	778,267	19.9	26.0	66.0	52.7
$25+$	317,958	428,394	343,782	34.7	8.1	34.0	47.3

Notes: Enrollment counts and changes in the THECB post-secondary education data prior to any sample selection using the TWC employment data. Panels A to C shows total enrollment in the community college sector for those aged 25-plus. Panel D shows enrollment counts across sectors for those aged 25-plus. Panel E shows enrollment counts in the community college sector for all age groups.

Table A2: Shares of Field Enrollment and Five-Year Switching Rates by Last Industry: Full Sample

	Primary Source Industry		*Future Outcomes*		
	Industry	Share (%)	Switch (%)	Work in Primary (%)	Graduate (%)
A. By Field of Study					
General/Other	Health Care	15.7	35.3	16.3	5.2
Business	Financial	14.5	39.0	14.9	31.1
Humanities	Education	21.5	39.7	29.8	39.6
Math/Engineering	Manufacturing	24.3	44.0	25.3	28.4
Health	Health Care	43.0	39.5	54.0	34.9
Vocational Tech.	Manufacturing	26.4	45.7	23.0	27.3
Vocational Services	Public Sector	29.6	39.8	34.3	33.9
Education	Education	73.6	18.0	78.8	48.2
Information Technology	Manufacturing	16.4	42.5	12.9	20.1
B. By Overall Enrollment					
Non-Enrollee			31.0		
Enrollee			36.4		27.1

Notes: See notes to Table 1.1.

Table A3: Decomposition of Gender Differences in Field of Enrollment by Last Industry: Baseline Sample

	Raw Female-Male Difference	*Within-Industry Female-Male Difference*		*Across-Industry Female-Male Difference*	
	Difference (p.p.)	Difference (p.p.)	Share (%)	Difference (p.p.)	Share (%)
General/Other	1.0	0.9	85.3	0.2	14.7
Business	2.0	2.8	138.1	-0.8	38.1
Humanities	2.1	2.3	107.8	-0.2	7.8
Math/Engineering	-10.0	-7.1	71.2	-2.9	28.8
Health	16.0	10.4	64.7	5.7	35.3
Vocational Tech.	-4.9	-3.5	72.1	-1.4	27.9
Vocational Services	-5.1	-3.2	62.1	-1.9	37.9
Education	4.1	2.7	65.5	1.4	34.5
Information Technology	-5.4	-5.3	98.2	-0.1	1.8

Notes: This figure decomposes differences in field of study shares by gender into within and across industry components using the law of total probability: $P(f|g) = \sum_j P(f|g,j)P(j|g)$ where j is industry, f is field, and $g \in \{M, W\}$ is gender. Using this identity, I perform the following decomposition:

$$P(f|F) - P(f|M) = \underbrace{\sum_j \left[P(f|W,j) - P(f|M,j)\right] P(j|W)}_{\textbf{Within Industry}} + \underbrace{\sum_j P(f|M,j)\left[P(j|W) - P(j|M)\right]}_{\textbf{Across Industry}}$$

The within industry share measures differences in field shares between men and women in the same industry, holding industry shares fixed to women. The across industry share measures differences in field shares driven by differences in where men and women work, holding the within-industry propensity to enroll in a field to men.

Table A4: Decomposition of Gender Differences in Field of Enrollment by Last Industry: Full Sample

	Raw Female-Male Difference	*Within-Industry Female-Male Difference*		*Across-Industry Female-Male Difference*	
	Difference (p.p.)	Difference (p.p.)	Share (%)	Difference (p.p.)	Share (%)
General/Other	0.5	1.1	205.1	-0.6	105.1
Business	0.1	1.8	2441.0	-1.7	2341.0
Humanities	2.4	2.0	83.7	0.4	16.3
Math/Engineering	-9.6	-6.4	66.5	-3.2	33.5
Health	12.3	8.2	66.4	4.1	33.6
Vocational Tech.	-7.0	-4.7	67.0	-2.3	33.0
Vocational Services	-2.6	-1.5	60.0	-1.0	40.0
Education	8.3	3.2	39.2	5.0	60.8
Information Technology	-4.5	-3.8	83.7	-0.7	16.3

Notes: This figure decomposes differences in field of study shares by gender into within and across industry components using the law of total probability: $P(f|g) = \sum_j P(f|g,j)P(j|g)$, where j is industry, f is field, and $g \in \{M, W\}$ is gender. Using this identity, I perform the following decomposition:

$$P(f|F) - P(f|M) = \underbrace{\sum_j \left[P(f|W,j) - P(f|M,j)\right] P(j|W)}_{\textbf{Within Industry}} + \underbrace{\sum_j P(f|M,j)\left[P(j|W) - P(j|M)\right]}_{\textbf{Across Industry}}$$

The within industry share measures differences in field shares between men and women in the same industry, holding industry shares fixed to women. The across industry share measures differences in field shares driven by differences in where men and women work, holding the within-industry propensity to enroll in a field to men.

Table A5: Preferred Industry Exposure Estimates

	Baseline Sample				Full Sample			
	Enrolled	New Enrollment	Earnings	Graduated	Enrolled	New Enrollment	Earnings	Graduated
$Z(Shock)_j \times$ D(2005)	-0.10	-0.00	359.99**	-0.19***	-0.09*	0.00	566.15***	-0.12***
	(0.13)	(0.08)	(142.54)	(0.06)	(0.05)	(0.01)	(201.83)	(0.04)
$Z(Shock)_j \times$ D(2006)	0.06	0.02	445.61***	-0.06	-0.04*	-0.03**	616.24***	-0.03***
	(0.09)	(0.07)	(113.38)	(0.04)	(0.02)	(0.01)	(174.29)	(0.01)
$Z(Shock)_j \times$ D(2008)	0.01	0.10	-961.58***	0.01	0.05**	0.03*	-1215.50***	-0.00
	(0.09)	(0.08)	(223.59)	(0.04)	(0.03)	(0.02)	(348.70)	(0.01)
$Z(Shock)_j \times$ D(2009)	0.37***	0.33***	-2653.57***	-0.01	0.17***	0.11***	-3396.18***	0.00
	(0.14)	(0.08)	(444.77)	(0.04)	(0.05)	(0.02)	(570.83)	(0.01)
$Z(Shock)_j \times$ D(2010)	0.28*	0.23***	-2343.58***	0.08**	0.23***	0.09***	-3270.28***	0.03**
	(0.15)	(0.07)	(517.75)	(0.04)	(0.06)	(0.02)	(730.50)	(0.01)
$Z(Shock)_j \times$ D(2011)	0.12	0.06	-1872.04***	0.01	0.21***	0.03	-2364.53***	0.03**
	(0.16)	(0.07)	(581.93)	(0.06)	(0.07)	(0.02)	(748.34)	(0.01)
$Z(Shock)_j \times$ D(2012)	0.30**	0.22***	-1759.58***	0.06	0.21***	0.04**	-2301.22***	0.03*
	(0.15)	(0.08)	(648.53)	(0.06)	(0.08)	(0.02)	(847.10)	(0.01)
$Z(Shock)_j \times$ D(2013)	0.19	0.06	-1381.96**	0.04	0.19**	0.02	-1721.28*	0.02
	(0.14)	(0.07)	(693.78)	(0.05)	(0.09)	(0.02)	(964.04)	(0.02)
$Z(Shock)_j \times$ D(2014)	0.22	0.17**	-1366.22*	-0.04	0.22**	0.04*	-1663.69	0.02
	(0.14)	(0.07)	(770.17)	(0.04)	(0.09)	(0.02)	(1059.63)	(0.02)
$Z(Shock)_j \times$ D(2015)	0.26*	0.10	-1892.98**	0.05	0.23**	0.05**	-1877.45*	0.02
	(0.15)	(0.07)	(801.11)	(0.04)	(0.09)	(0.02)	(984.34)	(0.02)
$Z(Shock)_j \times$ D(2016)	0.29**	0.16***	-1845.12**	0.01	0.23**	0.04	-1441.62*	0.01
	(0.13)	(0.05)	(810.22)	(0.05)	(0.10)	(0.03)	(807.99)	(0.02)
$Z(Shock)_j \times$ D(2017)	0.20	0.10	-1684.71**	0.00	0.23**	0.04*	-586.25	0.02
	(0.14)	(0.06)	(808.36)	(0.04)	(0.11)	(0.02)	(681.64)	(0.02)
$Z(Shock)_j \times$ D(2018)	0.26*	0.16***	-1796.20**	-0.05	0.22**	0.04*	88.95	0.01
	(0.13)	(0.06)	(903.66)	(0.04)	(0.11)	(0.03)	(651.60)	(0.02)
Baseline Controls	X	X	X	X	X	X	X	X
Demographic Controls	X	X	X	X				
2007 Outcome Mean	7.3	2.2	56790	0.7	4.4	1.2	86349	0.5
at 1SD Z(Shock)	5.2	1.6	63456	0.5	2.9	0.9	92709	0.3
N (Individuals)	561	561	561	561	561	561	561	561
N (Industries)	189,073	189,073	189,073	189,073	3,439,610	3,439,610	3,439,610	3,439,610

Notes: Estimates of parameters δ_y based on the event study specification in Equation 1.3.1. Standard errors, shown in parentheses, are clustered based on an individual's pre-Recession 5-digit industry code. Outcome means in 2007 at 1SD of the shock measure are based on predicted values from a linear probability model of 2007 enrollment on the exposure measure, evaluated at 1 SD of the exposure measure. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: Robustness of Industry Exposure Analysis

	External Shock		1997 Baseline		Alternate Controls I		Alternate Controls II	
	Enrolled	Earnings	Enrolled	Earnings	Enrolled	Earnings	Enrolled	Earnings
$Z(Shock)_j \times$ D(2005)	-0.220*	96.864	-0.182**	178.559	-0.140	341.717**	0.030	311.738**
	(0.114)	(144.022)	(0.072)	(187.853)	(0.140)	(144.161)	(0.135)	(151.280)
$Z(Shock)_j \times$ D(2006)	-0.100	334.632**	-0.034	95.139	0.060	465.500***	0.051	322.171***
	(0.094)	(142.855)	(0.032)	(142.962)	(0.095)	(112.433)	(0.091)	(111.925)
$Z(Shock)_j \times$ D(2008)	-0.021	-540.820***	-0.009	1016.685***	-0.011	-822.039***	0.044	-955.561***
	(0.079)	(177.219)	(0.032)	(287.155)	(0.094)	(234.085)	(0.097)	(245.230)
$Z(Shock)_j \times$ D(2009)	0.287***	-1096.239***	0.007	1537.426***	0.317**	-2507.155***	0.303**	-2567.577***
	(0.109)	(340.518)	(0.060)	(494.729)	(0.140)	(465.762)	(0.144)	(487.152)
$Z(Shock)_j \times$ D(2010)	0.216	-620.344*	0.015	2073.452***	0.223	-2118.237***	0.162	-2279.154***
	(0.137)	(371.348)	(0.069)	(649.541)	(0.162)	(546.004)	(0.163)	(583.000)
$Z(Shock)_j \times$ D(2011)	0.206	1.896	0.015	1166.552**	0.061	-1512.203**	0.004	-1963.609***
	(0.166)	(413.105)	(0.064)	(555.696)	(0.175)	(623.679)	(0.167)	(648.115)
$Z(Shock)_j \times$ D(2012)	0.157	486.977	0.041	-172.431	0.269*	-1277.925*	0.199	-2003.834***
	(0.165)	(431.583)	(0.068)	(508.232)	(0.156)	(714.853)	(0.155)	(693.484)
$Z(Shock)_j \times$ D(2013)	0.181	972.859**	0.098	-498.096	0.182	-793.205	0.048	-1817.688**
	(0.145)	(465.266)	(0.075)	(594.491)	(0.152)	(784.282)	(0.152)	(710.321)
$Z(Shock)_j \times$ D(2014)	0.168	1057.021**	0.085	-615.669	0.231	-685.414	0.054	-1805.448**
	(0.174)	(530.254)	(0.079)	(621.329)	(0.152)	(882.636)	(0.147)	(773.164)
$Z(Shock)_j \times$ D(2015)	0.189	1149.714**	0.088	-692.418	0.278*	-1179.636	0.086	-2427.194***
	(0.161)	(566.631)	(0.084)	(734.356)	(0.162)	(927.656)	(0.155)	(813.748)
$Z(Shock)_j \times$ D(2016)	0.114	1378.318**	0.114	-499.890	0.313**	-1165.222	0.160	-2358.616***
	(0.167)	(591.672)	(0.084)	(868.961)	(0.140)	(943.109)	(0.132)	(875.048)
$Z(Shock)_j \times$ D(2017)	0.126	1290.957**	0.114	-1052.173	0.242	-956.394	0.032	-2146.538**
	(0.164)	(629.570)	(0.088)	(938.671)	(0.148)	(957.642)	(0.134)	(851.302)
$Z(Shock)_j \times$ D(2018)	0.179	1568.184**	0.164*	-1946.960*	0.280*	-996.831	0.076	-2382.846**
	(0.176)	(707.415)	(0.088)	(1060.677)	(0.144)	(1075.996)	(0.131)	(929.174)
2007 Outcome Mean	7.2	57196			7.3	56790	7.3	56790
at 1 SD Z(Shock)	5.3	63665			5.2	63456	5.2	63456
Demographic Controls	X	X					X	X
Industry Controls	X	X	X	X	X	X		
1-Digit Controls							X	X
Sample	Baseline	Baseline	Full	Full	Baseline	Baseline	Baseline	Baseline
N (Individuals)	529	529	561	561	561	561	561	561
N (Industries)	174,855	174,855	1,924,805	1,924,805	189,073	189,073	189,073	189,073

Notes: Estimates of parameters δ_y based on the event study specification in Equation 1.3.1. Standard errors, shown in parentheses, are clustered based on an individual's pre-Recession 5-digit industry code. Outcome means in 2007 at 1SD of the shock measure are based on predicted values from a linear probability model of 2007 enrollment on the exposure measure, evaluated at 1 SD of the exposure measure. "External Shock" reports estimates using n exposure measure defined using industry-level employment losses defined outside of Texas using data from the Statistics of U.S. Businesses. "1997 Baseline" reports coefficients relative to 1997:Q4 as the base period rather than 2007:Q4. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: Heterogeneity by Field and Industry: Baseline Sample

A. By Industry:	Production/Manufacturing	Services	Professional/Finance	Other
Enroll (pp)	0.404	0.082	0.708**	-0.029
SE	(0.348)	(0.262)	(0.323)	(0.325)
Effect (%)	[7.6%]	[1.4%]	[10.4%]	[-0.9%]
N (Industries)	142	130	79	210
N (Individuals)	15,282	41,803	41,923	90,065

B. By Fields of Study:	Study Health	Study Business	Study Information Technology	Study Vocational Services	Study Education
Enroll (pp)	0.136*	0.095*	0.020	0.045	0.018
SE	(0.078)	(0.049)	(0.020)	(0.028)	(0.018)
Effect (%)	[48.2%]	[6.8%]	[6.9%]	[14.1%]	[8.2%]
N (Industries)	561	561	561	561	561
N (Individuals)	189,073	189,073	189,073	189,073	189,073

Notes: Coefficients pool the event study parameters δ_y for 2009 and 2010 from the event study specification in Equation 1.3.1. Standard errors, shown in parentheses, are clustered based on an individual's pre-Recession 5-digit industry code. Standard errors clustered based on an individual's pre-Recession 5-digit industry code. Percent effects, shown in brackets, are based on predicted values from a linear probability model of 2007 enrollment on the exposure measure, evaluated at 1 SD of the exposure measure. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

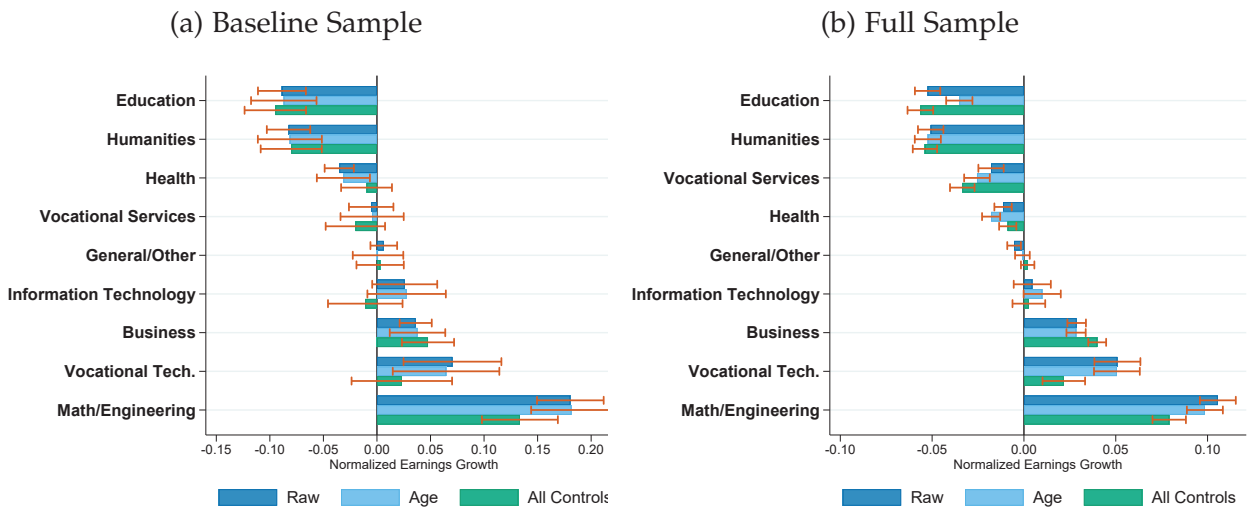
Table A8: Heterogeneity by Field and Industry: Full Sample

	All	Production/Manufacturing	Services	Professional/Finance	Other
A. Main Outcomes					
Earnings	-3727.360***	-1627.413***	-1705.510***	-2627.081***	-4617.686***
	(699.068)	(619.682)	(434.856)	(711.498)	(1656.782)
	[-4.0%]	[-1.7%]	[-2.2%]	[-2.7%]	[-4.9%]
Enrolled (pp)	0.245***	0.159**	0.053	0.420***	0.129
	(0.061)	(0.073)	(0.058)	(0.127)	(0.120)
	[8.3%]	[6.4%]	[1.6%]	[10.5%]	[5.3%]
New Enrollment (pp)	0.107***	0.050	0.084***	0.204***	0.064**
	(0.018)	(0.036)	(0.024)	(0.047)	(0.027)
	[12.5%]	[6.6%]	[8.7%]	[18.1%]	[9.3%]
B. Most Affected Fields					
Study Health (pp)	0.094***	0.036**	-0.000	0.092***	0.109**
	(0.025)	(0.017)	(0.035)	(0.023)	(0.047)
	[73.1%]	[17.1%]	[-0.0%]	[23.2%]	[-28.2%]
Study Business (pp)	0.031*	0.016	0.021	0.135**	0.001
	(0.016)	(0.020)	(0.016)	(0.053)	(0.024)
	[3.9%]	[2.3%]	[2.8%]	[10.9%]	[0.2%]
Study Information Technology (pp)	0.014***	0.016**	0.013	0.021	0.009
	(0.005)	(0.008)	(0.011)	(0.015)	(0.007)
	[8.2%]	[10.3%]	[7.3%]	[8.7%]	[7.2%]
Study Vocational Services (pp)	0.029***	0.011	-0.010	0.007	0.016
	(0.010)	(0.008)	(0.013)	(0.011)	(0.017)
	[22.7%]	[14.6%]	[-5.0%]	[4.4%]	[14.5%]
Study Education (pp)	0.018***	-0.000	0.026***	0.022**	-0.004
	(0.006)	(0.004)	(0.008)	(0.011)	(0.010)
	[5.8%]	[-0.6%]	[24.7%]	[15.4%]	[-0.6%]
C. Other Fields					
Study General/Other (pp)	0.02	-0.03**	0.00	0.06**	0.02
	(0.01)	(0.02)	(0.03)	(0.03)	(0.03)
	[3.2%]	[-7.5%]	[0.5%]	[7.3%]	[4.3%]
Study Humanities (pp)	0.02	-0.00	-0.01	0.02	-0.04
	(0.02)	(0.01)	(0.02)	(0.04)	(0.03)
	[3.8%]	[-0.7%]	[-1.2%]	[4.1%]	[-7.8%]
Study Vocational Technical (pp)	0.01	0.05***	-0.01	0.02**	0.01
	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)
	[9.5%]	[35.0%]	[-12.9%]	[42.3%]	[5.8%]
Study Math/Engineering (pp)	0.02	0.07***	0.02**	0.04	0.01
	(0.01)	(0.02)	(0.01)	(0.04)	(0.02)
	[4.6%]	[14.0%]	[7.2%]	[11.9%]	[2.9%]
N (Industries)	561	142	130	79	210
N (Individuals)	3,439,610	411,385	685,468	594,704	1,748,053

Notes: Coefficients pool the event study parameters δ_y for 2009 and 2010 from the event study specification in Equation 1.3.1. Standard errors, shown in parentheses, are clustered based on an individual's pre-Recession 5-digit industry code. Standard errors clustered based on an individual's pre-Recession 5-digit industry code. Percent effects, shown in brackets, are based on predicted values from a linear probability model of 2007 enrollment on the exposure measure, evaluated at 1 SD of the exposure measure. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

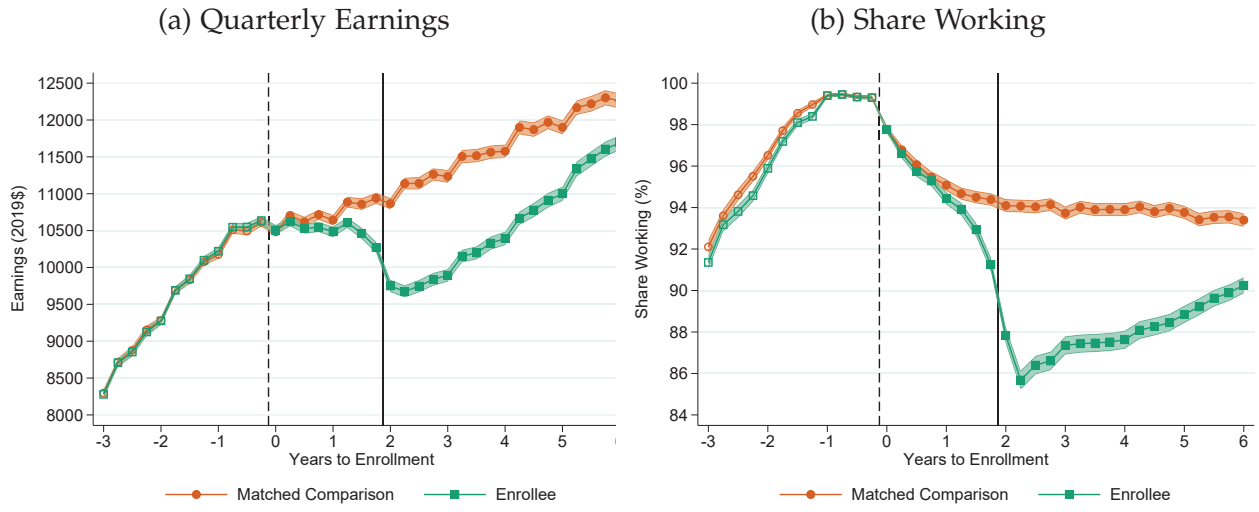
A.2 Additional Returns Estimates

Figure A6: Field of Study: By Sample and Controls



Notes: This figure repeats the empirical exercise conducted in Figure 1.4. The second and third sets of bars ("Age" and "All Controls") add in studentized covariates that are also interacted with the quadratic in post-enrollment trends (centered at 20 quarters). Inclusion of these covariates and interactions with trends adjust for compositional differences in the covariates across field of study while still centering the reported estimates on the average, composition-adjusted earnings growth for that pathway. Controls for "Age" are age and age-squared. "All Controls" includes other demographics (indicators for gender, black, Hispanic, some prior college, and prior college graduation), total earnings in year 3 (quarters 9 to 12) prior to enrollment, total earnings in year 2 (quarters 5 to 8) prior to enrollment, annualized earnings growth between year 1 (quarters 1 to 4) and year 3 prior to enrollment, pairwise quadratic interactions between these, an indicator variable for any non-employment in pre-enrollment quarters 1 to 4, and an indicator variable for only being non-employed in quarters 1 to 4 prior to enrollment.

Figure A7: Matching Estimator: Placebo Check



Notes: This figure reproduces the matching procedure outlined in the notes to Figure 1.5, using a "fake" enrollment date two years prior to the actual enrollment quarter for each enrollee. This fake enrollment event is plotted in the dashed vertical line, and the real enrollment date is plotted in the solid vertical line. Figure (a) plots total average earnings and Figure (b) plots any employment.

Table A9: Heterogeneity in Matching Estimates by Field of Study and Graduation

	All		*Graduate within 4 years*			*Non-graduates*		
	Earnings	Employed"	Earnings	Employed	Switch	Earnings	Employed	Switch
A. Unconditional								
All	14.6	-0.020***	1167.2***	-0.008*	0.142***	-295.5***	-0.023***	0.044***
	(53.1)	(0.002)	(113.9)	(0.004)	(0.008)	(59.8)	(0.002)	(0.004)
B. By Field of Study								
General/Other	-339.8***	-0.022***	557.2	-0.011	0.061*	-389.4***	-0.022***	0.031***
	(105.8)	(0.004)	(476.1)	(0.018)	(0.033)	(108.5)	(0.004)	(0.007)
Business	253.1*	-0.013***	754.0***	-0.009	0.124***	105.5	-0.014***	0.022**
	(131.0)	(0.005)	(253.0)	(0.009)	(0.021)	(152.2)	(0.005)	(0.011)
Humanities	-340.5**	-0.033***	231.9	-0.028***	0.142***	-654.6***	-0.035***	0.015
	(149.8)	(0.006)	(281.0)	(0.011)	(0.020)	(172.9)	(0.007)	(0.013)
Math/Engineering	1667.6***	-0.014*	4280.2***	-0.003	0.258***	772.1**	-0.018**	0.081***
	(333.4)	(0.008)	(599.0)	(0.014)	(0.037)	(393.1)	(0.009)	(0.021)
Health	239.2**	-0.025***	2088.7***	0.004	0.131***	-501.2***	-0.036***	0.066***
	(107.4)	(0.004)	(194.1)	(0.007)	(0.014)	(126.8)	(0.005)	(0.009)
Vocational Tech.	-16.8	-0.005	900.0	0.007	0.280***	-588.5	-0.012	0.092**
	(612.1)	(0.015)	(841.4)	(0.025)	(0.051)	(843.2)	(0.019)	(0.042)
Vocational Services	-132.1	-0.020***	-20.2	-0.025**	0.160***	-178.0	-0.018***	0.073***
	(157.0)	(0.006)	(269.6)	(0.010)	(0.023)	(191.8)	(0.006)	(0.014)
Education	-168.1	-0.008	119.6	0.024	0.203***	-194.8	-0.011	0.060***
	(149.6)	(0.007)	(514.0)	(0.032)	(0.049)	(156.4)	(0.007)	(0.013)
Information Technology	-496.3	-0.000	771.1	0.019	0.039	-861.4**	-0.006	0.028
	(364.8)	(0.009)	(655.0)	(0.019)	(0.046)	(429.4)	(0.010)	(0.023)

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Data Appendix

B.1 Sample Definitions

My analysis focuses on two balanced panels of individuals' work and education decisions:

Full Sample: *This sample represents the universe of workers in the Texas UI data meeting the following sample definitions*:

(a) Have at least four consecutive quarters of full-time equivalent work,[100] during which they are not enrolled. I define the first quarter of the first such period as the quarter of labor market entry for this sample and measure variables such as experience and tenure relative to this base period.

(b) I define a worker as employed if they have quarterly earnings above $2,400 and restrict to individuals for whom the longest period of non-employment is 20 quarters (5 years) and they have some work experience in each of the years 2017, 2018, 2019. This restriction is meant to focus on workers with labor force attachment, while allowing for the inclusion of long-term unemployed, an important feature of the Great Recession (Yagan [2019]) as well as individuals leaving the labor force temporarily to pursue education.

While this data is great in its breadth, I am not able to link individuals with demographic information unless they enter either the Texas K-12 or higher education systems. For this reason, I focus on a second sample of high school graduates for whom I see complete

[100] Full-time equivalent work is defined as quarterly earnings above $4,800, or working 40 hours a week for 12 weeks at $10 an hour.

labor market histories. It is on this sample that I estimate the structural model.

Baseline Sample: *This is sub-sample of the previous panel of workers for which I have demographics and complete work histories, on which I conduct most analysis and estimate the structural model.* These data have the following additional restrictions:

(c) Graduated from a high school in the TEA data in 1994 to 1998.

(d) Do not earn and advanced degree (defined as a masters degree or PhD).

(e) Does not have work or college enrollment before they are aged 14 and are not missing demographic data.

(f) Have initial work experience, defined as at least 4 consecutive quarters of non-enrolled work between ages 21 and 25. This is needed to define initial industry of employment. It also helps to ensure they do not begin their work history in a high education state (e.g. individuals pursuing masters/PhDs who are in education in this period are meant to be excluded).

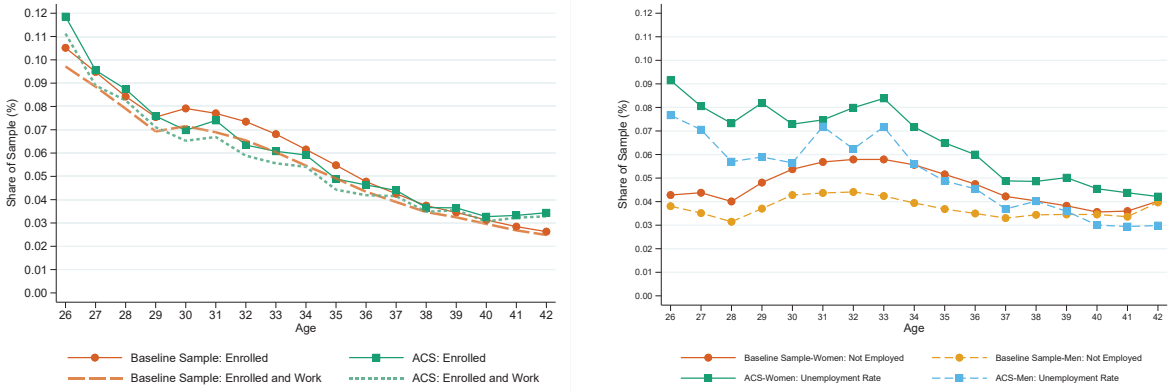
(g) Have one calendar year of 4 quarters of non-enrolled work between the ages of 25 and 30. This ensures that individuals choose an industry in this period.

The ages in this sample will thus range from 29 to 33 at the onset of the Great Recession in 2009 and 39 to 43 at the end of the analysis period in 2019. In Figure B1, I compare the evolution of key variables for the baseline sample panel to cohorts of the same age range in the American Community Surveys.

Figure B1: Comparison of Baseline Sample with Statistics in the ACS

(a) Enrollment

(b) Non-employment (ERC) v. Unemployment Rate (ACS)



Notes: These figures compare key variables from the 2006 to 2018 1-Year American Community Survey (ACS) waves, for the 1976 to 1980 birth cohorts in Texas who are marked as is in the labor force (LABFORCE=2). While not an exact comparison, I use the sample in the labor force from the ACS, since my ERC sample restriction limit the data to those attached to the labor force. Figure (a) compares my administrative marker for current enrollment to the ACS variable for currently attending college/graduate school (GRADEATT=6 or 7) and attending a public institution (SCHLTYPE=2). Figure (b) compares non-employment to the variables for unemployed (EMPSTAT=2 if EMPSTAT=1 or 2). OLF shares in the ACS are much higher than in my data (e.g. 25.2% in 2007).

B.2 Key Variable Definitions

Industry of Employment: I focus the analysis on the industry code related to the job with the highest earnings in a quarter but using earnings summed across all jobs. I use industry NAICS codes defined at the 5-digit level.[101]

Fields of Study: Information on programs of study pursued in college are reported as CIP codes from two sources: (i) the major field of study in each enrollment period or (ii) the CIP code associated with credentials received upon graduation. I focus on 2-digit CIP codes and further aggregate into broad areas of study, as exhibited in Table B2. I do

[101] This level of detail is quite low-level, differentiating between bars and restaurants. I do some harmonization to account for the fact that some industry codes may lose employment due to changes in how the underlying firm is classified. For example, I link industry codes where large numbers of individuals switch to the same new industry at the same time. I also combine industries with very small (< 100) employment. I never allow industries to be linked across 3-Digit codes. Further details are available upon request. This procedure results in 561 unique industries (from a total of 857). Additional details are available upon request.

some imputation to focus on the principal field of study associated with an enrollment spell.[102]

Enrollment Timing: I first map semesters into business quarters to be used with the UI data. Raw enrollment records are at the semester level and I convert these into quarterly data using the following schema: (a) Fall to Q3 and Q4, (b) Spring to Q1 and Q2, (c) Summer I to Q2, (d) Summer II to Q3. This mapping roughly maps to that used in Jepsen et al. [2014]. For the fall and spring semesters, I divide completed hours equally between quarters.

New and Target Enrollment: I define individuals as being *newly enrolled* if they were not enrolled in the prior three calendar years. For each individual, I define a *target enrollment*, as the first quarter when they transitioned from work to school as adults. This is the enrollment timing used in the returns analysis of Section 1.2.4. Specifically, I define this as the first new enrollment after labor market entry where the individual had some work experience in years 2 and 3 prior to enrollment.

B.3 Summary Statistics

[102]Many of the enrollment CIP codes are reported as general studies or as missings, likely not reflecting true courses taken. I focus on the "principal field of study" for an enrollment spell, which imputes the main CIP code pursued in a meriod of enrollment based on CIP code of degrees obtained, the number of periods a CIP code is reported as the main course of study, and the number of credit hours associated with each period of study towards a CIP code. do some imputation based on CIP codes pursued throughout an entire spell and from degrees obtained. Details are available upon request.

Table B1: Summary Statistics for Baseline Sample

	Men	Women
A. Basic Variables		
Black	0.09	0.15
Hispanic	0.29	0.31
Some Prior College	0.21	0.17
B. 2007 Variables		
Work in 2007	0.97	0.95
Services	0.25	0.21
Health Care	0.04	0.21
Manufacturing	0.11	0.04
Earnings in 2007	57504.01	43255.02
Enrollment in 2007	0.06	0.10
C. 2010 Variables		
Work in 2010	0.95	0.94
Earnings in 2010	62302.13	46348.88
Enrollment in 2010	0.06	0.10
D. 2018 Variables		
Work in 2018	0.96	0.96
Earnings in 2018	85641.18	58876.98
Enrollment in 2018	0.02	0.04
E. Post-Work Education Outcomes		
Ever Enroll	0.18	0.27
Ever Graduate with Certificate	0.02	0.02
Ever Graduate with Associates	0.03	0.05
Ever Graduate with B.A.	0.02	0.03
F. At First Post-Work Enrollment		
Study in Business	0.16	0.17
Study in Math/Engineering	0.12	0.02
Study in Health	0.12	0.29
Study in Vocational Services	0.11	0.06
Study in Education	0.03	0.07
Number of Observations	126,431	121,542

B.4 Aggregation of NAICS and CIP Codes

Table B2: Mapping between CIP Codes and Fields of Study

Broad Area of Study	Area of Study in Model	2-Digit CIP
General/Other	General	24. Liberal Arts and Sciences, General Studies and Humanities
		32. Basic Skills
		34. Health-Related Knowledge and Skills
		36. Leisure and Recreational Activities
		37. Personal Awareness and Self Improvement
		99. Not Classified
Business	Business	52. Business, Management, Marketing, and Related Support Services
Humanities	Humanities	5. Area, Ethnic, Cultural, and Gender Studies
		9. Communication, Journalism, and Related Programs
		16. Foreign Languages, Literatures, and Linguistics
		22. Legal Professions and Studies
		23. English Language and Literature/Letters
		30. Multi/Interdisciplinary Studies
		38. Philosophy and Religious Studies
		39. Theology and Religious Studies
		42. Psychology
		45. Social Sciences
		50. Visual and Performing Arts
		54. History
STEM	Technical	4. Architecture and Related Services
		10. Communication Technologies/Technicians and Support Services
		11. Computer Information Sciences and Support Services
		14. Engineering
		15. Engineering Technologies/Technicians
		27. Mathematics and Statistics
		40. Physical Sciences
		41. Science Technologies/Technicians
Health	Health	26. Biological and Biomedical Sciences
		51. Health Professions and Related Clinical Services
Vocational Technical	Technical	1. Agriculture, Agricultural Operations, and Related Sciences
		3. Natural Resources and Conservation
		21. Technology Education
		29. Military Technologies
		46. Construction Trades
		47. Mechanic and Repair Technologies/Technicians
		48. Precision Production
		49. Transpiration and Materials Moving
Vocational Services	General	12. Personal and Culinary Services
		19. Family and Consumer Sciences/Human Sciences
		31. Parks, Recreation, Leisure, and Fitness Studies
		43. Security and Protective Services
		44. Public Administration and Social Service Professions
Education	General	13. Education
		25. Library Science

C Mathematical Appendix

C.1 Identification of Earnings and Human Capital Accumulation Parameters

This section shows semi-parametric identification of key parameters in the earnings equation of the model, given linearity of the earnings equation, time-invariance of earnings coefficients π and α, and the mean independence assumptions on the stochastic process governing η.

For ease of exposition, this section makes several simplifications relative to the baseline model that are without loss of generality. First, only consider identification of the model where there is only one worker type, as the arguments can be repeated conditional on worker type. Parameters that exist across worker types, such as skill prices r_{jt} and skill weights π_j, will have additional cross-worker type restrictions in the more general model, which aids in identification. Second, consider a version of the model model where there is no schooling decision, as the arguments easily extend to that case but notation becomes complicated (e.g. one must consider part-time penalties, etc.). Thus, each period workers make the choice of which industry j to work in, denoted j_{it}. These consist of $J = 7$ industries plus non-employment, $j_{it} = 0$. To ensure that identifying equations are not redundant, we need these industries to be distinct in their parameters, π_j, α_j, and r_{jt}.

Finally, consider a version where the earnings shock η is drawn from a distribution that is mean independent of other components entering the earnings equation. Allow-

ing for dependence across periods means the conditioning arguments used in the proofs will just need to include lagged earnings histories as well. What is necessary for identification regarding η is that the initial η is mean-independent of initial conditions and that innovations in η are mean independent of the state and choices in the state.

The equations in the model can be summarized as:

$$y_{ijt} = r_{jt} + \pi_j' \omega_{it} + \alpha_j \tau_{it} + \eta_{it}$$

$$\omega_{i,t+1}^{(s)} = \Gamma^{(s)}(j_{it}|\omega_{it}^{(s)})\omega_{it}^{(s)}$$

$$\text{Covariates } z_i \text{ determine initial skill } \omega_{i1}^{(s)}(z_i) \text{ with } \mathbb{E}[\omega_{i1}^{(s)}] = 1, \forall s$$

$$\tau_{i,t+1} = \Psi(j_{it}|\tau_{it}, j_{i,t-1})\tau_{it} \text{ with } \Psi(j_{it}|0, j_{i,t-1}) = 1 \text{ and } \Psi(j_{it} = 0 \text{ or } sep_{it} = 1|\tau_{it}, j_{i,t-1}) = 0$$

$$\mathbb{E}[\eta_{it}|j_{it}, \omega_{it}, \tau_{it}] = 0$$

Let F_z denote the distribution of z. For exposition, I consider z to be discrete taking on n_z possible values. The main text makes functional form restrictions on Γ and Ψ, which I do not impose here.

The econometrician observes y_{ijt}, j_{it}, sep_{it}^{Obs}, and their histories, $\mathscr{H}_{-t}$. The identification arguments rest on appeals to the fact that we have a sufficient richness in observed transitions. In particular, this can be thought of as an overlap assumption where at least some workers with work history $\mathscr{H}_{-t}$ that choose each action. This gives both enough distinct identifying equations to recover the parameters of the model.

Skill Weights and Initial Skills: Consider comparisons restricted to individuals who

enter work from non-employment so that $\tau = 0$ can be ignored. The initial earnings of workers in the first period who have initial observables z can be written as:

$$y_{ij1}(z) = r_{j1} + \pi_j' \omega_{i1}(z) + \eta_{it}$$

We can use the assumptions on the earnings innovation η to take expectations so that the η drops out:

$$\bar{y}_{j1}(z) = r_{j1} + \pi_j' \omega_{i1}(z)$$

Fixing industry j, consider differences in earnings for workers with initial covariates z_A from a reference group with z:

$$\bar{y}_{j1}(z_A) - \bar{y}_{j1}(z) = \pi_j' \left(\omega_{i1}(z_A) - \omega_{i1}(z) \right)$$

This differences out skill prices. Now, integrate z over its population distribution:

$$\begin{aligned}
\int_z \bar{y}_{j1}(z_A) - \bar{y}_{j1}(\tilde{z}) dF_Z(\tilde{z}) &= \int_z \pi_j' (\omega_{i1}(z_A) - \omega_{i1}(z)) dF_Z(\tilde{z}) \\
\bar{y}_{j1}(z_A) - \int_z \bar{y}_{j1}(\tilde{z}) dF_Z(\tilde{z}) &= \pi_j' \omega_{i1}(z_A) - \pi_j' \int_z \omega_{i1}(\tilde{z}) dF_Z(\tilde{z}) \\
\bar{y}_{j1}(z_A) - \int_z \bar{y}_{j1}(\tilde{z}) dF_Z(\tilde{z}) &= \pi_j' \left(\omega_{i1}(z_A) - \iota_S \right)
\end{aligned}$$

where ι_S is an S-dimensional vector of ones. The second line follows as the distribution F_z must integrate to one and the last line follows from the normalization on initial skills. The left-hand side of this equation is observed, as $\int_z \bar{y}_{j1}(\tilde{z}) dF_Z(\tilde{z})$ are earnings in j for

difference z, weighted by the population distribution of Z.

For every z_A, we need to recover S initial skills parameters $\omega_{i1}(z_A)$. Thus, there are $n_Z \cdot S$ parameters. For every industry j, we need to recover S skill weights $\pi_j^{(s)}$, so there are $J \cdot S$ of these parameters. This gives a total of $(n_Z + J) \cdot S$ unknowns. In contrast, we have $J \cdot (n_z - 1)$ comparisons given by the equation. This means that we a linear system with more equations than parameters as long as $(n_z + J)S \leq J(n_z - 1)$. For this application, $J = 7$ and $S = 3$, meaning that we have identified the skill weights and initial skills as long as $n_z \geq 4$, which is satisfied.

Skill Accumulation Parameters: I now discuss identification of $\Gamma^{(s)}$, assuming that π and $\omega_0(z)$ are known. These parameters are identified through within-industry but across-work history variation in average earnings.

Consider two individuals, A and B, who have the same initial Z and tenure and thus ω_{i1} and τ_1. Consider a comparison such that $j_{2A} = j_{2B} = l$ but $j_{1A} \neq j_{1B} \neq l$. Then, comparing earnings in period 2 gives:

$$\begin{aligned}\bar{y}_{2A} - \bar{y}_{2B} &= \left(r_{l2} + \pi_l' \Gamma(j_{1A}|\omega_1) \cdot \omega_1 + \alpha_l^\tau \tau_2\right) - \left(r_{l2} + \pi_l' \Gamma(j_{1B}|\omega_1) \cdot \omega_1 + \alpha_l^\tau \tau_2\right) \\ &= \pi_l' \left[\Gamma(j_{1A}|\omega_1) - \Gamma(j_{1B}|\omega_1)\right] \cdot \omega_1\end{aligned}$$

Since A and B-type histories work in the same final industry, skill prices are differenced out. Moreover, since they have the same initial τ and transition each period (but not to l!), they have the same final τ_2, which also gets differenced out.

We have J choices for l and $J-1$ choices for j_{1A} and j_{2B} where ordering doesn't matter, giving $J \cdot \frac{(J-1)(J-2)}{2}$ possible comparisons. We have $J \cdot S$ Γ parameters to identify. Thus, for this single comparison, we have identification of Γ as long as $\frac{(J-1)(J-2)}{2} > S$, which is satisfied for $J = 7$ and $S = 3$.

We can repeat these arguments conditional on work histories $\mathscr{H}_{-t}$ in each period t to recover Γ for every possible ω. In practice, the functional forms on Γ implied by Equation 1.4.2 determines this process in a tractable way.

Specific Skill Parameters: The process governing τ and its industry-specific weight α_j^τ are recovered by exploiting the assumptions regarding involuntary separations. Consider two types of individuals, C and D, who have the same prior work history and same skills at the start of period $t-2$ and work in the same industry in periods j_t, j_{t-1} and j_{t-2}. They will thus have the same ω_t.

Suppose D suffers and involuntary separation at the beginning of period $t-1$ and C suffers an involuntary separation at the beginning of period t. Then, $\tau_{Ct} = 0$ and $\tau_{Dt} = 1$. Comparing average differences in earnings between these two groups will pin down α_j^τ.

Then, Ψ can be recovered through average differences between stayers and switchers who end up in the same industry. Consider individuals E and F who have the same work history $\mathscr{H}_{-t}$. For simplicity, assume they have have $\tau_t = 1$ (so they are separated at the beginning of t or transition from non-employment). E continues to work at $j_{t-1} = l$ in t and $t+1$. F switches to $j_{Ft} \neq j_{t-1}$ before choosing $j_{t-1} = j_{t+1} = l$. Then, average

differences in $t+1$ earnings between these two types of individuals is:

$$\begin{aligned}\bar{y}_{t+1,E} - \bar{y}_{t+1,F} &= \left(r_{l,t+1} + \pi_l'\Gamma(j_{tE}|\omega_t)\cdot\omega_t + \alpha_l^\tau\Psi(j_{Et}|\tau_t,l)\tau_t\right) - \left(r_{l,t+1} + \pi_l'\Gamma(j_{tF}|\omega_t)\cdot\omega_t + \alpha_l^\tau\Psi(j_{Ft}|\tau_t,l)\tau\right. \\ &= \pi_l'\left[\Gamma(j_{tE}|\omega_t) - \Gamma(j_{tF}|\omega_t)\right]\cdot\omega_t + \alpha_l^\tau\left[\Psi(j_{Et}|\tau_t,l) - \Psi(j_{Ft}|\tau_t,l)\right]\tau_t\end{aligned}$$

The terms $\pi_l'\left[\Gamma(j_{tE}|\omega_t) - \Gamma(j_{tF}|\omega_t)\right]\cdot\omega_t$ and α_j^τ are already identified by previous arguments. Average differences between these groups thus identify $\Psi(j_{Et}|\tau_t,l) - \Psi(j_{Ft}|\tau_t,l)$ exactly (one equation with one unknown). There are $\frac{J(J-1)}{2}$ such differences in J unknowns, meaning that Ψ is identified as long as $J \geq 3$. We can then extend the previous argument to $\tau_t \neq 1$.

Skill Prices: The remaining parameters governing average earnings are the skill prices r_{jt}. These are simply residuals of earnings net of already identified quantities.

For initial skill prices r_{j1}, consider comparisons originating from non-employment so that $\tau = 0$. Then, making the appeal to large numbers and exploiting the initial skill normalization, skill prices are just:

$$r_{j1} = \int_z \left(\bar{y}_{j1}(z) - \pi_j'\omega_1(z)\right) dF_z(z) = \int_z \bar{y}_{j1}(z)dF_z(z) - \pi_j'\iota_S \quad \forall j$$

which also shows how the normalization pins down the location of skills relative to skill prices.

Remaining skill prices are computed by integrating over work histories conditional

on working in industry j at time t. Formally:

$$r_{lt} = \bar{y}_{lt} - \pi_l' \mathbb{E}_{\mathscr{H}_{-t}}[\omega_{it} | j_{it} = l] - \alpha_l^\tau \mathbb{E}_{\mathscr{H}_{-t}}[\tau_{it} | j_{it} = l]$$

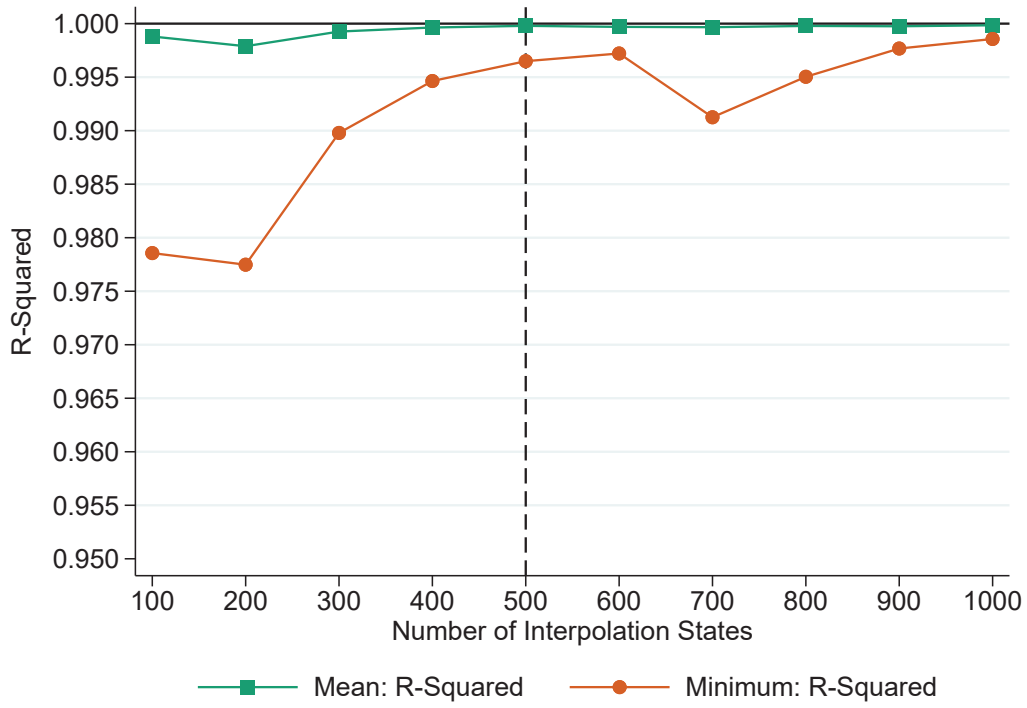
Distribution of η: We have made no assumptions about η other than mean independence. Since we now know all other parameters of the model, we can thus recover its distribution semi-parametrically:

$$\eta_{it} = y_{it} - r_{j_i,t} - \pi_{j_i}' \omega_i(\mathscr{H}_{-t}) - \alpha_{j_i}^\tau \tau_i(\mathscr{H}_{-t})$$

This shows that the model can in principle be much more flexible regarding the distribution of the errors than what is imposed in estimation. In particular, identification relies heavily on linearity, semi-parametric assumptions on the error terms, and having a sufficient overlap between choices and work histories. However, the arguments do not rely on parametric restrictions on the unobservables at all

D Estimation Appendix

Figure D1: Fit of the Value Function Interpolation and Sensitivity to Additional Interpolation States



Notes: This figure shows the fit of the value function interpolation, where the dashed line shows the number of states used for interpolation in estimation. I interpolate the value function separately for each age (30 periods from 25 to 54), graduation cohort (2 for 1994 and 1998 in estimation), last employment state (7 industries, non-employment, and the separation state), and worker type (4) for a total of 2,160 linear regressions. The orange (circles) line shows the minimum R-Squared over all these regressions while the teal (squares) line shows the average.

.

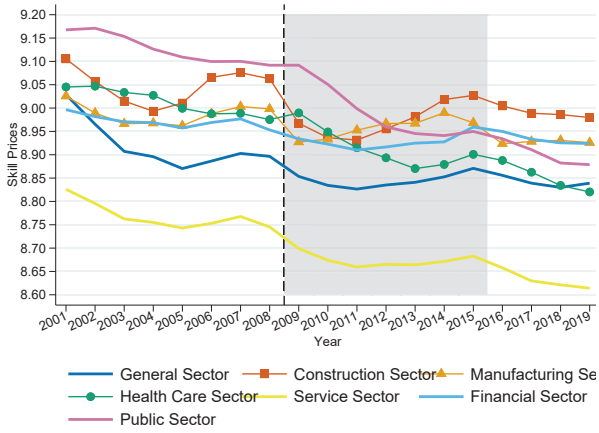
D.1 Algorithm to Evaluate NLLS Objective

An evaluation of Q^{NLLS} for a guess α_Y proceeds according to the following algorithm:

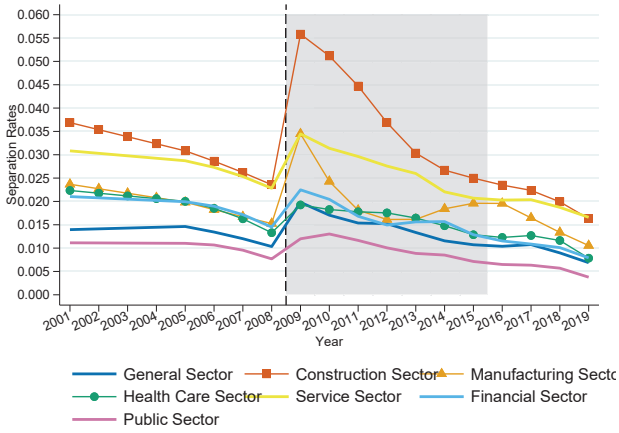
1. *Impose normalization*: Compute unrestricted $\tilde{\omega}_{1i}^{(s)}$ using α_Y for each i and s. Use these to compute average $\bar{\omega}_1^{(s)}$ by averaging in the population and use to force $\mathbb{E}[\omega_{1i}^{(s)}] = 1$ by setting $\omega_{1i}^{(s)} = \tilde{\omega}_{1i}^{(s)} / \bar{\omega}_1^{(s)}$.

Figure D2: Skill Prices and Separation Rates - All Industries

(a) Skill Prices (Estimated)



(b) Separation Rates (Calibrated)



2. *Evaluate human capital*: Compute $h_{it} = h\left(k_{it}|\Omega(\mathscr{H}_i^{-t};\alpha_Y);\alpha_Y,\lambda_Y(\alpha_Y)\right)$ for all i and t

3. *Iterate over years*: Starting in year t with $r_{j,t-1}$ for each j and $\eta_{i,j_{i,t-1},t-1}$ for all i:

   (a) Compute skill prices using Equation 1.6.5

   (b) For each individual i:

   - Compute current idiosyncratic productivity: $\eta_{ijt} = y_{it} - r_{jt} - h_{it}$
   - Compute the current earnings residual: $\varphi_{ij,j_{i,t-1},t} = \eta_{ijt} - \rho^j_{j_{i,t-1}}\eta_{i,j_{i,t-1},t-1}$.

     Note that in the first period an agent works, $\varphi_{ij,j_{i,t-1},t} = \eta_{ij1}$.

4. *Compute objective*: Using Equation 1.6.4 by adding up $\varphi^2_{ij,j_{i,t-1},t}$

Estimates of choice-dependent variances of the shock process in terms of α_Y are:

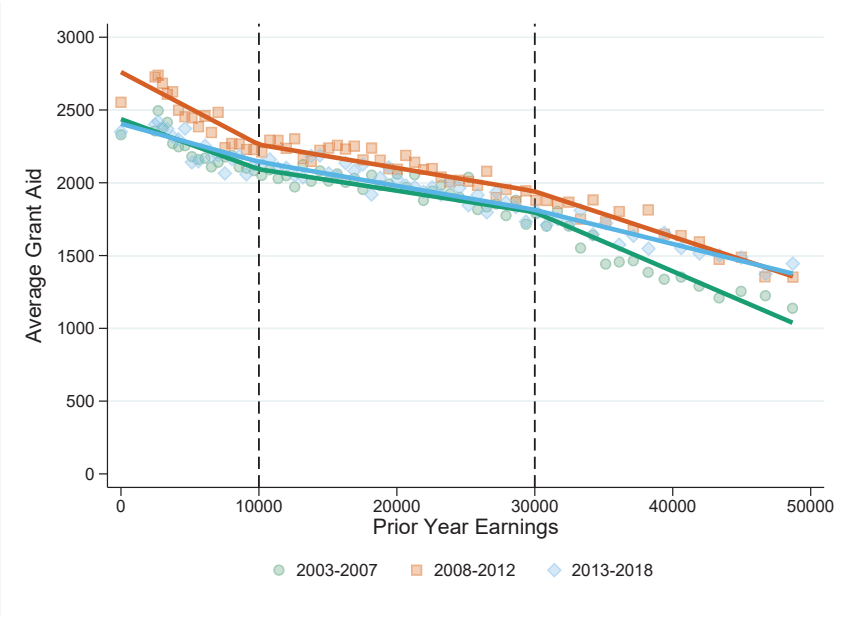
$$\hat{\sigma}_j^{j_{-1}}(\alpha_Y) = \sqrt{\frac{1}{N_{j,j_{-1}}}\sum_{(it)\text{ s.t. }j(k_{it})=j\text{ and }j(k_{i,t-1})=j_{-1}}^{N_{j,j_{-1}}}\left(y_{it} - y(k_{it}|\mathscr{H}_i^{-t},q_i;\alpha_Y)\right)^2}$$

where the sum is over person-years working in industry j and $N_{j,j_{-1}}$ is the total ever

working in j with activity j_{-1} at the start of the period. In practice, this is allowed to vary by the target industry j and by whether the individual remains in their current industry ($j(k_{it}) = j(k_{i,t-1})$), switches industries (voluntarily), or switches from non-employment (including an involuntary separation).

D.2 Estimated Financial Aid Schedules

Figure D3: Calibrated Financial Aid Schedules



D.3 Full Parameter Estimates

Table D1: Skill Weights, Skill Accumulation, and Match Accumulation Parameters: Π, Γ, and Ψ

	Skill 1	Skill 2	Skill 3	Specific Skill	Part-Time
A. Industry-Specific Return: $\pi_j^{(s)}$, $\alpha_j^{(\tau)}$,α_j^{PT}					
General	0.627	0.203	0.489	0.324	0.887
	(0.071)	(0.019)	(0.040)	(0.002)	(0.001)
Construction	0.561	0.136	0.470	0.311	0.887
	(0.065)	(0.020)	(0.040)	(0.003)	(0.001)
Manufacturing	0.601	0.212	0.463	0.320	0.887
	(0.069)	(0.018)	(0.039)	(0.002)	(0.001)
Health Care	0.647	0.190	0.232	0.328	0.887
	(0.069)	(0.018)	(0.034)	(0.002)	(0.001)
Services	0.672	0.069	0.484	0.361	0.887
	(0.072)	(0.014)	(0.041)	(0.002)	(0.001)
Financial	0.619	0.187	0.494	0.289	0.887
	(0.068)	(0.020)	(0.041)	(0.003)	(0.001)
Public Sector	0.608	0.173	0.220	0.344	0.887
	(0.061)	(0.019)	(0.032)	(0.002)	(0.001)
B. Learning-by-doing: $\gamma_{Work:j(k)}^{(s)}$					
Work in General	0.111	-0.051	-0.034		
	(0.010)	(0.009)	(0.005)		
Work in Construction	0.106	-0.148	-0.029		
	(0.010)	(0.025)	(0.007)		
Work in Manufacturing	0.088	-0.073	-0.015		
	(0.007)	(0.008)	(0.003)		
Work in Health Care	0.100	-0.040	-0.143		
	(0.009)	(0.010)	(0.011)		
Work in Services	0.050	-0.503	-0.028		
	(0.006)	(0.075)	(0.005)		
Work in Financial	0.065	-0.015	-0.011		
	(0.007)	(0.012)	(0.005)		
Work in Public Sector	0.105	-0.031	-0.137		
	(0.009)	(0.007)	(0.009)		
Work in Last Sector: $\gamma_{Accum}^{(\tau)}$				0.324	
				(0.002)	
C. Heterogeneity in Learning-by-doing: $\gamma_{Work:q}^{(s)}$					
D(Female-No College)	-0.014	-0.200	0.009		
	(0.003)	(0.034)	(0.004)		
D(Male-College)	0.023	0.021	0.008		
	(0.004)	(0.006)	(0.003)		
D(Female-College)	-0.012	-0.075	0.024		
	(0.003)	(0.012)	(0.004)		
D. Schooling: $\gamma_{School:f(k)}^{(s)}$					
Study in General	0.100	-0.198	-0.162		
	(0.022)	(0.069)	(0.028)		
Study in Technical	0.067	0.037	-0.031		
	(0.021)	(0.047)	(0.028)		
Study in Health Care	0.303	0.127	-0.963		
	(0.030)	(0.062)	(0.079)		
Study in Humanities	0.128	-0.001	-0.196		
	(0.021)	(0.052)	(0.036)		
Study in Business	0.116	-0.106	-0.043		
	(0.019)	(0.058)	(0.027)		
E. Heterogeneity in Schooling Returns: $\gamma_{School:q}^{(s)}$					
D(Female-No College)	-0.028	-0.141	-0.012		
	(0.021)	(0.060)	(0.033)		
D(Male-College)	-0.005	0.091	-0.052		
	(0.018)	(0.047)	(0.032)		
D(Female-College)	-0.009	-0.061	-0.032		
	(0.018)	(0.056)	(0.030)		
F. Other Accumulation Parameters:					
Depreciation: $\gamma_{Depr}^{(s)}$	0.032	0.000	0.007		
	(0.005)	(0.000)	(0.007)		
Diminishing Work Investments: $\gamma_{Dim:School}^{(s)}$, $\gamma_{Dim}^{(\tau)}$	1.189	0.000	0.051	0.311	
	(0.067)	(0.000)	(0.009)	(0.003)	
Diminishing School Investments: $\gamma_{Dim:School}^{(s)}$	2.005	1.983	0.001		
	(0.139)	(0.397)	(0.000)		
Loss from Voluntary Transition: $\gamma_{Vol}^{(\tau)}$				0.320	
				(0.002)	
Part-Time Share From Work: $\gamma^{PartTime}$					0.434
					(0.042)
G. Initial Conditions: $\gamma_{Z:1}^{(s)}$, $\gamma_{q:1}^{(s)}$					
D(General)	0.101	-0.003	0.022		
	(0.038)	(0.075)	(0.034)		
D(Construction)	0.318	0.147	0.052		
	(0.040)	(0.033)	(0.039)		
D(Manufacturing)	0.293	-0.417	0.172		
	(0.040)	(0.050)	(0.024)		
D(Health Care)	0.111	0.030	-0.030		
	(0.023)	(0.070)	(0.160)		
D(Services)	0.214	-0.141	-0.729		
	(0.054)	(0.083)	(0.117)		
D(Financial)	-0.081	0.088	-0.160		
	(0.071)	(0.069)	(0.070)		
D(Public Sector)	0.158	-0.043	0.404		
	(0.040)	(0.048)	(0.042)		
D(Female-No College)	0.478	0.339	-0.197		
	(0.051)	(0.041)	(0.046)		
D(Male-College)	0.154	0.377	-0.182		
	(0.038)	(0.042)	(0.046)		
D(Female-College)	0.088	-0.065	-0.207		
	(0.027)	(0.020)	(0.023)		

Table D2: Shock Process Parameters

	Remain	*EE Switch*	*EN Switch*	*Initial*
σ_j :General	0.275	0.399	0.643	0.487
	(0.001)	(0.002)	(0.002)	(0.003)
Construction	0.316	0.442	0.625	0.494
	(0.002)	(0.004)	(0.003)	(0.004)
Manufacturing	0.272	0.395	0.637	0.499
	(0.001)	(0.002)	(0.002)	(0.002)
Health Care	0.265	0.412	0.581	0.500
	(0.001)	(0.003)	(0.002)	(0.002)
Services	0.291	0.447	0.580	0.512
	(0.001)	(0.002)	(0.001)	(0.001)
Financial	0.255	0.379	0.578	0.438
	(0.001)	(0.003)	(0.003)	(0.003)
Public Sector	0.188	0.378	0.513	0.394
	(0.001)	(0.002)	(0.002)	(0.002)
ρ :	0.888	0.719	0.350	
	(0.001)	(0.002)	(0.003)	

Table D3: Choice Parameters

	Common	Difference: Female, No College	Difference: Male, College	Difference: Female, College
A. Action Utility: μ_q^k				
Work in General	-7.777	-0.047	0.064	0.013
	(0.279)	(0.013)	(0.010)	(0.013)
Work in Construction	-7.496	-0.324	-0.033	-0.329
	(0.276)	(0.013)	(0.011)	(0.013)
Work in Manufacturing	-7.539	-0.163	-0.054	-0.176
	(0.276)	(0.012)	(0.009)	(0.013)
Work in Health Care	-7.710	0.263	0.051	0.180
	(0.270)	(0.014)	(0.009)	(0.013)
Work in Services	-7.351	-0.035	-0.055	-0.137
	(0.265)	(0.010)	(0.009)	(0.012)
Work in Financial	-7.722	0.123	0.047	0.105
	(0.273)	(0.013)	(0.011)	(0.014)
Work in Public Sector	-7.492	0.019	-0.030	-0.020
	(0.272)	(0.014)	(0.010)	(0.013)
Study General	-3.308	-0.241	0.581	0.580
	(0.052)	(0.106)	(0.029)	(0.026)
Study Technical	-3.716	-1.439	0.452	-0.674
	(0.053)	(0.105)	(0.026)	(0.027)
Study Health Care	-3.536	-0.199	0.716	0.819
	(0.068)	(0.138)	(0.038)	(0.038)
Study Humanities	-4.534	-0.297	0.857	1.065
	(0.058)	(0.104)	(0.027)	(0.030)
Study Business	-4.264	-0.134	0.809	0.878
	(0.057)	(0.106)	(0.028)	(0.030)
Part-Time Study	0.547	0.593	-0.262	-0.074
	(0.039)	(0.094)	(0.016)	(0.022)
Full-Time Work $\times(\text{Age}-30)$	-0.006	-0.022	0.010	-0.005
	(0.003)	(0.001)	(0.001)	(0.001)
Full-Time Work $\times(\text{Age}-30)^2$	0.001	0.002	-0.000	0.001
	(0.000)	(0.000)	(0.000)	(0.000)
B. Elasticities:				
Preference for Log-Income: β_Y	0.738			
	(0.024)			
Sensitivity to Net Tuition: β_P	0.467			
	(0.016)			
C. Switching Costs: κ				
Origin: Non-Employed	1.401	0.011	0.053	0.056
	(0.007)	(0.002)	(0.001)	(0.002)
Origin: Employed	1.244			
	(0.004)			
Great Recession: Non-Employed	0.017			
	(0.001)			
Great Recession: Employed	0.058			
	(0.000)			
Target: Construction	0.227			
	(0.004)			
Target: Manufacturing	-0.012			
	(0.004)			
Target: Health Care	-0.064			
	(0.005)			
Target: Services	-0.359			
	(0.005)			
Target: Financial	0.100			
	(0.006)			
Target: Public Sector	0.092			
	(0.004)			
From Schooling	-0.301			
	(0.006)			
Into Schooling	0.396			
	(0.003)			

Chapter 2

Student Choices and the Earnings Associated with College Major and Selectivity

2.1 Introduction[1]

Recent evidence documents that the earnings premia associated with different fields of study can be as large as the extensive margin choice to attend college at all (Altonji et al. [2012]).[2] To what degree can society increase earnings by the way it allocates access to higher education? On the one hand, it could be the case that high average earning majors have large causal effects on earnings, supporting recent attempts by policymakers to increase enrollment in STEM majors. On the other hand, observational differences in earnings premia could simply reflect sorting into majors on the basis of ability – individuals who attend high-earning majors would have had high earnings

[1]This chapter is coauthored with *Justine Hastings*, *Christopher A. Neilson*, and *Seth D. Zimmerman*.

[2]Using majors reported in the American Community Survey, Altonji et al. [2012] estimate a gap in log wages rates between high school and college graduates of 0.577. In contrast, male electrical engineering majors earn a premium of 0.561 relative to general education majors, which only shrinks to 0.264 after controlling current occupation. See Table 3 of their paper for more information.

regardless of what they study. At the same time, high-earning fields may truly generate large increases earnings, but perhaps only for those who seek to attend them or have a sufficient level of academic preparation prior to college. In this case, reallocating seats away from low-earnings majors may actually harm some students with particularly large match effects in these fields. It may also be the case that students do not sort on earnings gains at all. Even if we increased the number of slots available to STEM majors, students may not be willing to sign up for these courses.

The degree to which these different mechanisms explain observational patterns is an empirical question. While a proper accounting of individual heterogeneity has been one of the most important advances in work on the causal return to education (Card [1999], Heckman et al. [2006]), definitive answers on the sources of field-specific premia has remained elusive. This is largely due to the unordered nature of many available options, making interpretation of available evidence difficult. For example, one important set of results looks to evidence from regression discontinuity (RD) designs (Kirkebøen et al. [2016], Hastings et al. [2013]) finding that college majors have impacts for students at the margin of admission. However, these studies generate a profusion of sometimes contradictory evidence based on the margin of interest. Moreover, what these research designs say for students away from the margin of an admission cutoff is less clear. The current literature suffers from a lack of a coherent framework that ties the available variation in the data together.

This paper makes progress on these issues by leveraging the centralized structure of the Chilean higher education system to estimate a generalized Roy model of student selection into colleges and majors. Assignment of college seats in Chile is centralized – students submit preferences through rank order lists (ROL) and scarce slots at each program are allocated via a Deferred-Acceptance (DA) algorithm. This process generates both rich information on student preferences through their applications as well as on the causal effect of programs on earnings through RD designs generated by admission

cutoffs at oversubscribed programs. We use this information to estimate a joint model of student preferences for post-secondary education programs and earnings outcomes, which enables us to extrapolate treatment effects away from admission cutoffs. Our empirical approach allows for rich sources of ability bias and sorting on gains in the presence of unordered and horizontally differentiated options for human capital investment. The empirical model also allows for a flexible structure on preferences without imposing strong assumptions on rationality or information sets, aside from the truth-telling assumptions needed to recover preferences from student reports. We use the empirical model to decompose the earnings premia associated with different programs into average treatment effects, selection bias, and match effects. The model allows us to evaluate which fields of study society would find the most valuable to expand in order to achieve income gains, taking into account feasibility restrictions imposed by student preferences.

Our first main finding is the importance of match effects and sorting on causal gains in generating observed patterns of field-specific premia, which indicates a large role for sorting into fields of study based on comparative advantage. This echoes suggestive evidence in Kirkebøen et al. [2016]. We decompose the income distribution predicted by our model at assigned options, finding that program average treatment effects (ATE) explain relatively little of the variation in earnings. Specifically, we find that program-level ATEs explains only 1.2 percent of total variation in individual earnings. There is a substantial role for classic ability bias, which explains 32 percent of earnings dispersion. In contrast, student-by-program match effects explains 66.4 percent of the variation, most of which is driven by match effects on observables. Looking at program-level effects, we find similar magnitudes with a larger role for sorting on unobserved gains.

While our results on earnings imply a large role for heterogeneous earnings gains, our preference estimates show that students show a modest preference for these heterogeneous earnings gains. In particular, mean program utility has a correlation of about

0.250 with program-level ATEs. Importantly, this result is not driven by the structure of our selection model, as our reduced form model of preferences does not impose rationality on individual choices. However, while programs are horizontally differentiated in terms of earnings gains (due to comparative advantage), reduced form utility has a much more vertical structure. The common component of utility explains 67 percent of the total variation in preferences across unobserved types, whereas only 22 percent is explained by the common component in earnings. This indicates that while students value earnings gains, this is not the only attribute that determines their choices. This leaves the effects of counterfactual allocation policies unclear, as there is no guarantee students will take up additional seats rather than entering the labor market directly or shifting to different programs.

This leads us to examine how changes in counterfactual admission policies and capacity allocations to different programs would change the distribution of earnings, taking into account restrictions placed on the feasibility of policies by student preferences. We consider two sets of counterfactual experiments that change this index. The first set keeps program capacities fixed and alters the order in which programs admit students. Programs in Chile rank students based on an index of program-specific weights on grades and exam scores that are set prior to student applications. Our first counterfactual considers changes to this ranking procedure. We find that the current admissions system increases average earnings relative to randomly deciding the order in which students are admitted – random offers would lead to a 10.7 percent reduction in average earnings and a 0.7pp decline in share of students receiving an on-platform match. However, we find that admission based solely on an average between verbal and quantitative test scores would only marginally increase earnings, leading to a 1.7 percent increase in average earnings and a 0.1pp increase in the share of students receiving an assignment. This indicates that the current system does a reasonably good job in allocating students on the basis of earnings gains.

The second set of counterfactual experiments examines how policymakers might reallocate capacity across programs, taking into account limitations on feasibility imposed by student preferences. We reallocate capacity from low-ATE programs to high-ATE programs and compare this to reallocating from low-average earnings to high-average earning programs. We find that reallocating 10 percent of seats in such a manner increases average earnings by 1 percent for ATE but 2 percent for average program earnings. This slightly larger effect for average program earnings is due to the fact that high earnings programs also have high treatment effects on the treated (TOT), i.e., causal gains for those who attend them. The incidence of these reallocation schemes affect different fields of study, reflecting the differences between ATE and TOT. In general, when reallocating based on ATE, enrollment is increased in the humanities, social studies, and agricultural programs. These fields have returns that generalize across the student population. In contrast, when reallocating based on average earnings, health and technical fields are expanded as these fields have high returns but that are specific to certain students. We hypothesize that the difference is likely due to the necessary student preparation to take advantage of these areas of study.

While these results may be specific to the Chilean context, they provide some caution against policies that uniformly push students into high-earning majors – one-size-fits-all recommendations may be costly. The finding of the importance of sorting on gains is substantially different than what is found in the K-12 literature in the United States that focuses on test score gains (Abdulkadiroglu et al. [2020]). The results also show that there may need to be a significant amount of pre-college human capital investment needed for students to take advantage of high-return majors.

Related Literature: The first literature that we contribute to concerns the returns to college, in particular those focusing on the returns to fields of study. Many of these studies rely on selection on observable assumptions to identify causal effects or only focus on certain sub-populations of students. Most focus only on returns to selectivity and ignore

college major (see Lovenheim and Smith [2023] for a recent review). Hoekstra [2009] and Zimmerman [2014] study RDs generated by admission cutoffs at specific public institutions in the United States, finding large returns even for marginal admits. Similarly, Bleemer [2024] studies an RD generated by a statewide change to college admission policies and finds large returns to college selectivity even for students who perform poorly on standardized exams. A related literature examines the returns to institutional quality using application fixed effects (Dale and Krueger [2002], Dale and Krueger [2014], Mountjoy and Hickman [2021]), which tend to find low returns to selectivity except for perhaps low income students. Our approach to modeling unobservables through preference types is similar in spirit, but is disciplined by targeting additional RD evidence and implied bias in the OLS estimates. A handful of papers have exploited RD designs generated by admission cutoffs from centralized assignment (Kirkebøen et al. [2016], Hastings et al. [2013]) or by identifying hidden admission cutoff policies (Mountjoy [2024]). These papers estimate returns across a multitude of institutions and fields of study but only for students at the margin of admission. While these approaches provide credible treatment effects estimates across a variety of programs and student populations, each estimate is local to a specific cutoff and it is difficult to compare across estimates. By nesting these effects into a single framework and explicitly modeling student choices, our model allows us to extrapolate away from the marginal admit and to predict how students would adjust to a broader set of counterfactual policies.

A recent literature seeks to use additional structure or variation in order to predict treatment effects across the distribution of students. Such approaches are especially challenging with multiple unordered treatments, like the case of college major choice considered in this paper. This is because compliers pushed into one treatment by an instrument can come from multiple origin treatments (Heckman and Urzua [2010]). Solutions typically rely on having multiple continuous instruments (Mountjoy [2022]), directly observable next available options (Kirkebøen et al. [2016]), or the imposition of a

constant treatment effects assumption on a subset of students (Walters and Kline [2016], Abdulkadiroglu et al. [2022]). Our approach is to flexibly model unobserved heterogeneity using variation from student applications and RD estimates generated by admission cutoffs. Closely related to our approach is Barahona et al. [2023] and Abdulkadiroglu et al. [2020], who use application data to estimate joint models of school choice and test score outcomes under centralized assignment. Our paper builds on their work by providing a more flexible structure for selection on unobservables[3] and by incorporating administrative data on earnings reported to a tax agency, rather than focusing on test score gains. Our approach is also closely related to Walters [2018], who combines application data and variation from multiple instruments to estimate a model of selection into charter schools. Instead we focus on sorting into college majors on the basis of earnings. Our focus on higher education is also closely related to existing structural models of sorting into college and fields of study on the basis of earnings gains (Willis and Rosen [1979], Arcidiacono [2004], Heckman et al. [2018], Arcidiacono et al. [2024]). We contribute to this literature by disciplining our models with administrative data and quasi-experimental variation from RD research designs.

Organization: The structure of the paper is as follows. Section 2.2 outlines data and key institutional features of college assignment in Chile; Section 2.3 outlines the model of college-major choice and earnings returns that microfounds this evidence; Section 2.4 discusses identification and estimation; Section 2.5 presents parameters estimates and initial results from the model; Section 2.6 presents decompositions of treatment effects from the model, holding student choices and seat allocations fixed; Section 2.7 considers counterfactual admissions and seat allocation policies, allowing students to adjust

[3]The control function approach of Abdulkadiroglu et al. [2020] estimates an education choice model similar to ours and then, conditioning on each student's preference ranking, controls for the expected value of preference shocks when estimating outcome equations. The functional form assumptions behind our estimation approach is substantially weaker, relying on a discrete mixture of unobserved types that have unrestricted match effects based on program characteristics. Moreover, our empirical approach accommodates data access limitations, as we cannot estimate computationally demanding models from within the tax agency.

their applications and the mechanism to adjust student assignments; and Section 2.8 concludes.

2.2 Institutional Context and Data

2.2.1 College Admissions in Chile

Like many countries around the world, but unlike the United States, Chile uses a centralized system to assign students to higher education options. Students apply to college-by-major combinations by submitting a rank ordered list (ROL) of their preferences for different schooling options. Students can list up to a maximum of 8 or 12 options depending on the year.[4]

Colleges in turn rank students based on an index determined by their entrance exam scores and high school grades, assigning the weights to each component at the time they pick the available programs and corresponding number of available seats. The market clears using a deferred acceptance (DA) algorithm, which generates sharp admissions cutoffs at oversubscribed programs. This means that students are assigned to the highest ranked option for which they are above the admissions cutoff. These assignments are binding; students cannot switch options after assignment and must wait a year and reapply in order to change their enrollment.

Until the late 1990s, almost all students attended one of the Council of Rectors of the Universities of Chile (CRUCH) universities which are the key options allocated through centralized assignment.[5] These not-for-profit institutions include a mix of public, private, and religious colleges. The two most selective CRUCH institutions are the *Universidad de Chile* (the flagship public university) and the *Pontificia Universidad Catolicá*

[4]We focus on the 1981 to 1985 cohorts of applicants. Students can rank 8 options in 1981 and 12 options in 1982 to 1985.

[5]These universities saw expansion in the early 1980s before stabilizing at the 25 traditional universities (the G25) through 2012, when 12 additional universities were added to the platform (Kapor et al. [2024]). Our sample from the early 1980s contains 24 institutions, which differs from the traditional G25 due to a few consolidations and spinoffs. Specifically, the Universidad Católica del Maule, Catholic University of Temuco, and the Universidad Católica de la Santísima Concepción were campuses of Pontificia Universidad Catolicá de Chile prior to 1991. In addition, two universities in our sample, Instituto Profesional de Chillán and Instituto Profesional De Valdivia, were absorbed in 1988 by the Universidad Del Bio-bio and the Universidad Austral de Chile, respectively.

de Chile (a private Catholic university), both of which are world-class universities. We define a degree program as a combination of institution-by-major-by-campus. Within institutions, fields differ in their selectivity and course content. We define a major as a consistent program of study, at which course content measures are constant.

Students must take the *Prueba de Aptitud Académica* (PAA), the national college entrance exam in order to be considered for admission.[6] This exam is administered by *Departamento de Evaluación, Medición y Registro Educacional* (DEMRE), the central testing authority operating under the authority of CRUCH. Scores range from 150 to 850 with a mean and median of 500. All exam takers must take quantitative and verbal exams, with additional optional exams covering specific subjects such as physics, biology, and history. Colleges assign weights to these main exam scores, subject exam scores, and high school GPA to create a deterministic index of priority to a specific degree program.[7]

At the time of submitting applications, students know their exam scores and have some idea of where they might be admitted. As a result of the cap on the number of applications, ROLs may display some strategic considerations (e.g. "reach" and "safety" school behavior). However, the exact cutoffs vary unpredictably across years, still generating unpredictable variation in admissions outcomes. We discuss the behavioral assumptions we place on this application process in detail in Section 2.4.1.

2.2.2 Data Sources

Application and Admission Records: DEMRE provided individual-level records of admissions exams scores from 1968 through 2018 and for college applications from 1977 to 2018. For 2003 and before, we digitize these records manually from hard copies. For many application cycles, we also observe applicants' submitted preference rankings and the admissions rules associated with this information. From this data, we are able to

[6] After 2002, the college entrance exam becomes the *Prueba de Seleción Universitaria* (PSU).

[7] We abstract from the choice to take a subject exam in the first place as well as any consideration of pre-college investment in human capital that may alter scores or GPA.

reconstruct the outcomes of an admissions cycle, including a student's assigned and next-available options. As discussed in Section 2.4.1, we follow Kapor et al. [2024] and restrict to a set of credible applications in which we assume these ranks reveal true preferences.

Test Score Records: We digitize data on PAA/PSU scores from 1967 to 2003 from hard copy records at DEMRE and link these using individual identifiers to the applications and admissions data. We use the quantitative and verbal scores as controls in our analysis, which are observed for everyone.

Information on Available Options: We use administrative information on program cutoffs supplied by DEMRE. We collect information on course content at the level of college major, which will be key to tractably incorporate heterogeneity in preferences and earnings gains for specific courses of study. We also measure program quality or selectivity as the average of math and verbal scores over the years 1977 to 2003.

Secondary School Information: The data from DEMRE also contain demographic information and the high school attended. Students in Chile can either attend public, subsidized voucher, or private high schools, and one key measure of socioeconomic status will be whether the student attended a private secondary school. See Hsieh and Urquiola [2006] for more information on K-12 education in Chile. In addition, we use location of the high school to get a measure of distance to available to each college option.[8]

Tax Records: Application and admission records are linked to individual tax records for the years in compliance with Chilean privacy laws.[9] These records include all labor earnings for the years 2006 to 2016. We include zero values of earnings in the analysis,

[8]Specifically, we identify the communa (or neighborhood) of the high school and use centroids of the communa as the location of the student. We then compute distances to the campus of each program using the great-circle distance formula.

[9]This disclosure is required by the Chilean government. SOURCE: Information contained herein comes from taxpayers' records obtained by the Chilean Internal Revenue Service (*Servicio de Impuestos Internos*), which was collected for tax purposes. Let the record state that the Internal Revenue Service assumes no responsibility or guarantee of any kind from the use or application made of the aforementioned information, especially in regard to the accuracy, validity or integrity.

meaning that our results should be interpreted as combining both intensive and extensive margin effects on earnings. We also top code the top one percent of earnings within cohort and experience bins to limit the influence of outliers. All monetary values are reported in 2011 pesos.

2.2.3 Reports used in the Earnings Analysis

OLS regressions and summary statistics use students' assigned options. These are defined as the first option for which the student i crosses program j's cutoff, $Index_{ij} > \pi_j$. π_j is the program cutoffs, that is, the $Index_{ij}$ for the lowest ranked student who was admitted to j. Threshold crossing regression discontinuity analysis focuses on closely admitted or rejected options. We define an application as *marginal* if $|Index_{ij} - \pi_j| < 15$. We define the *target* option as the first ranked option that is marginal.

We further stratify the RD analysis by the student's *next available* option, which is identifiable as we observe the full ROL. This is the next option they would have been assigned to had the student not been assigned to the target option. Either the target or the next option can be the actual assignment. In the case that the student is marginally rejected from the target option, the next available option is their actual assignment. If the student is marginally admitted to the target option, the next available option is the program to which they would have counterfactually been assigned. The next available option is the outside option in the case there are no alternative assignments.

2.2.4 Estimation Samples

While we use a broader set of students to estimate reduced form returns at the tax agency, we use individuals from the 1981 to 1985 application cohorts to estimate the structural model. We restrict the estimation sample to students having valid distance measures. We also restrict the sample to those submitting applications, meaning that

our results should be interpreted as being for the universe of college applicants. This yields a total sample size of 196,738 over 5 years. We briefly summarize the patterns of student choices present in the data in Table 2.1.

As discussed in Section 2.4.1, we restrict student reports to within a bandwidth of the admissions cutoff, up to the inclusion of a safety school. In Table 2.1 we show summary statistics on the credible application portfolios used in estimation. Note that our procedure allows us to sidestep strategic considerations stemming from maximum list length, as less than 1% of application portfolios in each year reach the maximum list length once we restrict to credible applications.

Table 2.1: Summary Statistics of Applicants by Cycle

	(1) 1981	(2) 1982	(3) 1983	(4) 1984	(5) 1985
A. Demographic Characteristics:					
Male-Private (%)	15.0	14.7	14.2	12.5	13.6
Male-Public (%)	44.8	45.7	45.8	45.6	45.1
Female-Private (%)	9.6	10.1	9.3	9.9	9.9
Female-Public (%)	30.6	29.5	30.7	32.0	31.4
Student in Santiago (%)	48.7	47.6	43.0	41.1	43.3
B. Application Characteristics:					
In RD Sample: Has Next Option (%)	35.4	37.0	35.2	36.7	36.7
In RD Sample: No Next Option (%)	16.9	10.2	19.4	19.3	17.5
Any Assignment (%)	81.0	86.1	69.8	68.1	69.8
Any Rejected Application (%)	60.9	53.2	69.0	71.5	68.5
Any Outside Option Rank (%)	23.7	28.2	42.5	45.3	45.0
C. Characteristics of Assigned Option:					
In Same Region (%)	53.1	54.6	54.3	54.0	54.3
In Santiago (%)	34.4	37.6	31.9	32.2	33.7
STEM Credits (%)	2.2	2.4	2.3	2.2	2.3
Humanities Credits (%)	1.7	1.6	1.7	1.7	1.6
Area: Education (%)	29.6	26.6	30.9	31.7	29.2
Area: Technology (%)	22.8	25.3	23.9	25.5	27.6
Area: Natural Science (%)	11.1	11.4	11.1	10.0	11.3
Area: Business Administration (%)	7.0	10.8	8.6	9.2	8.3
Area: Health (%)	6.5	10.0	7.1	5.0	4.4
Area: Agriculture (%)	5.8	4.9	6.7	6.1	8.6
Area: Art and Architecture (%)	5.3	4.3	3.9	4.5	3.0
Area: Social Science (%)	6.0	2.8	3.9	3.8	3.2
Area: Humanities (%)	4.6	2.4	2.2	2.5	2.6
Area: Law (%)	1.4	1.5	1.7	1.7	1.8
D. Number of Credible Applications:					
0 (%)	3.4	4.1	9.3	10.0	9.7
1 (%)	14.7	11.9	15.6	12.7	11.3
2 (%)	37.9	41.0	30.8	28.6	30.4
3 (%)	19.2	18.1	15.5	15.0	14.6
4 (%)	11.5	11.1	11.0	10.9	11.0
5 (%)	7.1	6.5	7.0	8.1	8.2
6 (%)	3.9	3.7	4.6	5.7	5.5
7 (%)	1.8	1.9	2.9	3.8	3.7
8 (%)	0.5	0.9	1.7	2.5	2.5
9 (%)	0.0	0.5	0.9	1.5	1.5
10 (%)	0.0	0.2	0.5	0.8	0.9
11 (%)	0.0	0.1	0.2	0.4	0.4
12 (%)	0.0	0.0	0.0	0.1	0.2
Observations	33227	37812	42476	42129	41094

Notes: Summary statistics of the estimation sample by application cycle. Applications are restricted to within the bandwidth of credible applications as discussed in Section 2.4.1.

2.3 Model of Student College and Major Choice

2.3.1 Overview

Programs rank students based on an admission index $Index_{ij}$, a known function that weights student exam scores and is set prior to the application process. Combined with ex ante capacities and submitted rank order lists by the student, this index determines program's admissions cutoffs π_j and the the student's resulting feasible set of options, denoted Ω_i^F. As discussed in Section 2.4.1, we allow Ω_i^F to contain any not clearly infeasible option. This allows students to rank programs from which they may be marginally rejected ex-post.

The upshot of using a DA mechanism is that students will be assigned to the most preferred option at which they are above the admissions cutoff. Students i choose among programs $j \in \Omega_i^F$ based on the utility of each option, which consists of both pecuniary (i.e., income maximizing) and non-pecuniary factors. These components are allowed to be correlated in an arbitrary way. The way students shift their application portfolios in response to changing available options helps to identify unobservable preferences among otherwise identical individuals.

Students are characterized by a vector of entrance exam scores which include math and language as well as a series of other exams, $PAA_i = \{Math_i, Lang_i, \dots\}$; an indicator for if they attended a private high school to capture differences in socioeconomic status, $Priv_i$; and an indicator for female $Female_i$. The observable features of the student that affect their payoffs be the vector

$$X_i = \{LowMath_i, LowLang_i, HighMath_i, HighLang_i, Priv_i, Female_i\}$$

captures all observable features of a student that systemically affect utility and earnings payoffs. Low and high scores are defined as the highest and lowest tercile of the verbal

and quantitative exams.[10] In addition, students are characterized by their location loc_i, which determines if if they are in a metropolitan region $Metro_i$ and their distance to each program's campus $Dist_{ij}$.

Programs j are characterized by a broad area-of-study $a(j)$, a college major $m(j)$, an institution $c(j)$, a vector W_j capturing observable program characteristics, and the location of their campus loc_j. The observable characteristics in W_j include major-level measures of course content, as well as the selectivity of the program. We measure selectivity as the average of $Math_i$ and $Lang_i$ scores for admitted students to a program over the years of 1977 to 2003.[11] We also differentiate if a campus is located in one of the large metropolitan regions $Metro_j$ that contains the cities of Santiago (the capital and largest city), Valparaíso, or Concepción.

2.3.2 Choices and Outcomes

The value of attending a program V_{ij} depends on both the log-income the student will earn at that program Y_{ij} but also on non-pecuniary value of attending the program U_{ij}. Students weigh these two factors according to parameter ϕ:

$$V_{ij} = U_{ij} + \phi Y_{ij} \tag{2.3.1}$$

In addition, students have access to the outside option of entering the labor force or taking an off-platform option:

$$V_{i0} = U_{i0} + \phi Y_{i0}$$

The trade-off between Y_{ij} and U_{ij} captures that studying engineering may result in

[10] The additional test scores in PAA_i are optional exams that affect the $Index_{ij}$ of the student at program j. In other words, given a set of admission weighted that varies by program, $Index_{ij} = Index(PAA_i, Rules_j)$. The final cutoffs π_j used by the program is an equilibrium outcome of all indices and submitted reports by all students in entering the mechanism.

[11] We do not allow these program selectivity measures to change in the counterfactual, assuming it captures a fixed quality characteristic of the program. This abstracts from equilibrium adjustments of program characteristics, for example, due to peer effects or supply-side responses.

high earnings but students may find studying engineering undesirable, for example, due to the difficulty of the courses. Conversely, a student may find studying in humanities intellectually satisfying, despite the fact that it may result in low earnings. Our framework allows that this tradeoff to flexibly vary across the student population. For example, despite humanities resulting in lower earnings on average, it may not be the case that these students have high payoffs to studying engineering.

We decompose the pecuniary and non-pecuniary benefits of education into program-specific average effects as well as components that capture selection on comparative and absolute advantage. Expected log-earnings are specified as:

$$Y_{ij} = \alpha_{i1} + \alpha_j + \tau_{ij} \tag{2.3.2}$$

and the non-pecuniary value of a program is:

$$U_{ij} = \delta_j + \psi_{ij} - \lambda_i(loc_i, loc_j) + \epsilon_{ij} \tag{2.3.3}$$

The parameter α_{i1} captures selection on absolute advantage across all on-platform options. This will capture the classic notation of ability bias (Griliches [1977]). Individuals also have outside options earnings Y_{0i} that we discuss below where differences between α_{i1} and Y_{i0} capture extensive margin returns to higher education attendance. The effects α_j and δ_j capture mean earnings gains and amenity values at program j that is common across students. The individual-by-program effects τ_{ij} and ψ_{ij} are the observed and unobserved match components that determine both returns and preferences. We allow τ_{ij} and ψ_{ij} to be arbitrarily correlated across the student population.

In our implementation, two determinants of utility are excluded from log-earnings but are key to predicting counterfactual choices. First, sensitivity to geographic distance as captured in $\lambda_i()$ is excluded from the earnings Equation 2.3.2. This exclusion plays a limited role in identification of earnings gains but is included to better estimate

substitution patterns and more accurately predict how assignments will change under counterfactual environments. Second, ϵ_{ij} is individual i's idiosyncratic preference for option j that is excluded from earnings and known at the time of application. This piece rationalizes residual differences in application behavior not explained by our model of observables and unobservables. Note that we do not allow for any idiosyncratic or unanticipated component of earnings, meaning that students have preferences over forecasts of their expected earnings.[12]

We specify δ_j and α_j in terms of major effects $m(j)$ and institution indicators $c(j)$:

$$\delta_j = \delta_{m(j)} + \delta_{c(j)}$$

$$\alpha_j = \alpha_{m(j)} + \alpha_{c(j)}$$

In our implementation, we exclude major effects for humanities and exclude institution effects for the University of Chile. When computing average utility and earnings effects below, we compute average predicted earnings or utility across the student population and demean within admissions cycles.

2.3.3 Specification of Unobserved Heterogeneity

To allow the non-pecuniary and pecuniary components of utility to be arbitrarily correlated, we classify students into Q unobservable and discrete types. One can interpret these as "preference types", where preferences incorporate utility from program value added as well as non-pecuniary values. Each type q differs in their outside option and on-platform mean earnings, Y_{i0} and α_{i1}; their utility from the outside option, U_{i0}; sensitivity to distance, λ_i; and their match specific components of earnings and utility, ψ_{ij}

[12]This is a practical limitation as the outcome moments we match are based on conditional average earnings. This makes it difficult to capture idiosyncratic error in earnings using the moments we match. We experimented with adding a measurement error component to Y_{ij} but found it difficult to identify from the data. The ex-ante known unobservables do a sufficient job in matching moments based on the variance of earnings for all and admitted students.

and τ_{ij}:

$$U_{i0} = U_0^{q(i)} + X_i'\delta_0^X + \epsilon_{i0}$$

$$Y_{i0} = Y_0^{q(i)} + X_i'\alpha_0^X$$

$$\alpha_{i1} = \alpha_1^{q(i)} + X_i'\alpha_1^X$$

and ψ_{ij} and τ_{ij} vary as:

$$\psi_{ij} = \underbrace{\delta_{a(j)}^{q(i)} + W_j'\delta_W^{q(i)}}_{\psi_{ij}^{Obs}} + \underbrace{\delta_{a(j)}^X(X_i) + W_j'\delta_W^X(X_i)}_{\psi_{ij}^{Unobs}}$$

$$\tau_{ij} = \underbrace{\alpha_{a(j)}^{q(i)} + W_j'\alpha_W^{q(i)}}_{\tau_{ij}^{Obs}} + \underbrace{\alpha_{a(j)}^X(X_i) + W_j'\alpha_W^X(X_i)}_{\tau_{ij}^{Unobs}}$$

where $\alpha_{a(j)}^{q(i)}$ and $\alpha_W^{q(i)}$ capture unobserved match effects in earnings that vary with program j's broad areas of study $a(j)$ and program-level covariates W_j, while $\alpha_{a(j)}^X(X_i)$ and $\delta_W^X(X_i)$ capture observable match effects by area of study $a(j)$ and covariates W_j as well. The preference terms $\delta_{a(j)}^{q(i)}$, $\delta_W^{q(i)}$, $\delta_{a(j)}^X(X_i)$, and $\delta_W^X(X_i)$ are defined similarly.

We do not directly include match with institution for tractability. The implication is that institutions are vertically ranked in terms of preferences and earnings value added beyond compositional differences due to differences in areas of study and other program-level covariates.[13] However, match with institution is indirectly captured to the extent that programs associated with different colleges differ in terms of their majors offered and the selectivity of these programs.

The function governing the sensitivity to log-distance is specified as:

$$\lambda_i(loc_i, loc_j) = \lambda_{In}^{q(i)} \log\Big(Dist_{ij}\Big) InRegion_{ij} + \lambda_{Out}^{q(i)} \log\Big(Dist_{ij}\Big)(1 - InRegion_{ij}) + \lambda_{Metro}^{q(i)} Metro_j$$

[13]To avoid perfect collinearity, we also exclude area-of-study effects for humanities and institution effects for the University of Chile. For stability in our estimation, we set the mean of W_j to zero. For all interactions with latent types, we exclude the main effect for $q = 1$.

where $InRegion_{ij}$ indicates if the student i is in the same region as school j and $Metro_j$ is an indicator for if program j is physically located in regions containing the large cities of Santiago, Valparaíso, or Concepción.

We specify the prior distribution of the unobservable types q according to a multinomial logit in terms of the predetermined covariates, X_i:

$$\Pi_q(X_i) = \frac{\exp\{\gamma_0^q + \gamma_X^q X_i\}}{1 + \sum_{q'=2}^{Q} \exp\{\gamma_0^{q'} + \gamma_X^{q'} X_i\}} \tag{2.3.4}$$

This specification allows the distribution of the latent types to be flexibly related to individual-level covariates, X_i, while still allowing the distribution to depend on unobservables through the intercept terms, γ_0^q. We denote the population shares of types as $\overline{\Pi}_q$, which is $\Pi_q(X_i)$ integrated over the population distribution of X_i.[14]

2.3.4 Reduced Form Choice Model

The structure of the model admits a reduced form structure (in the proper sense) on the preferences governing application behavior. Intuitively, if a group of students has a known comparative advantage at a certain set of options, this should show up systematically in their rank order lists, to the extent that they value earnings gains associated with these characteristics. At the same time, if a group of students has a high non-pecuniary benefit for certain programs (e.g. close to home, less math intensive), this should show up systematically in their reports as well. The reduced form choice model allows us to estimate the distribution of these unobservables in a first step without taking a stand on whether reports are determined by pecuniary or non-pecuniary attributes of a program.

[14] Inclusion of X_i in both the utility and earnings equations and the type shares will not cause an identification problem as long as types are not perfectly correlated with these covariates. For example, we could estimate our model separately for men and women and identify latent type shares conditional on gender. As long as types are not redundant within gender (e.g. all men are type 1), then inclusion of X_i in $\Pi_q()$ simply allows for these latent type shares to shift across gender. In practice, our approach imposes cross-equation restrictions by making the types constant across gender, allowing information to be shared across covariate cells.

In a second step, we relate these unobservables to earnings gains, although we cannot identify ϕ without additional variation or stronger assumptions.

The linearity of the model allows us to rearrange Equation 2.3.1, 2.3.2, and 2.3.3 into:

$$\tilde{V}_{ij} = \tilde{\delta}_j + \tilde{\psi}_{ij} - \lambda_i(loc_i, loc_j) + \epsilon_{ij} \tag{2.3.5}$$

where $\tilde{\delta}_j$ are reduced form program level mean preferences and $\tilde{\psi}_{ij}$ are reduced form match effects in preferences. These are given by:

$$\tilde{\delta}_j = \delta_j + \phi\alpha_j$$

$$\tilde{\psi}_{ij} = \psi_{ij} + \phi\tau_{ij}$$

Due to standard location issues associated with indirect utility models, we can only identify relative differences in utility and cannot distinguish between a intercept terms for on-platform options at the same time as the outside option. In the reduced form choice model, we thus let differences in α_0^q and Y_0^q enter the utility of the outside option which is specified as:

$$\tilde{V}_{i0} = U_0^q + \phi\left(Y_0^q - \alpha_0^q\right)$$

As only relative differences in utility matter, the agents only consider differences between earnings in the outside option relative to any on-platform attendance. To deal with scale, we also normalize the variance of the i.i.d. preference shock to unity and estimate all coefficients without a nurmeraire.

We will not be able to identify the parameter ϕ, which separates true preferences for program attributes from a preference for maximizing income. This does not matter for the counterfactuals we consider, but should be kept in mind for interpretation. This reason is because structural preferences and earnings gains are allowed to be arbitrarily correlated with each other. Associations between earnings and reduced form utility

could be evidence of a either a true preference for earnings gains but could also reflect amenity values that are arbitrary correlated with program value added. For example, if program A has large estimated returns but individuals do not apply for that program, our framework cannot speak to whether students do not care about earnings gains or whether program A has a strong dis-amenity value but students *do* prefer earnings, all else equal. We are however able to analyze the relationship between the reduced form of preferences and program-level earnings gains, as we do below in Section 2.5.

2.3.5 Likelihood of an Application

Assuming that idiosyncratic preferences ϵ_{ij} are distributed independently and identically extreme value, we can write the probability of rank j_r for a type q individual who has previously ranked $j_{1:r-1} = \{j_1, ..., j_{r-1}\}$ as:

$$Pr(j_r|\Omega_i^F, j_{1:r-1}, X_i, q) = \frac{\exp\{\bar{V}_j(X_i, q)\}}{\sum_{k \in \Omega_i^F \setminus j_{1:r-1}} \exp\{\bar{V}_k(X_i, q)\}}$$

where $\bar{V}_j(X_i, q)$ is the deterministic component of full indirect utility for option j as given in Equation 2.3.5, net of the realization of preference shocks ϵ_{ij}. For brevity, we are subsuming loc_i and omitting the dependence of this likelihood on a precise parameterization.

Continuing to condition on type q, the probability of an observed rank order list R_i is given by:

$$Pr(R_i|\Omega_i^F, X_i, q) = \prod_r Pr(j_r|\Omega_i^F, j_{1:r-1}, X_i, q)$$

which includes the probability that an individual ranks the outside option in a short list (i.e., they prefer the outside option to all on-platform options).

Finally, the full likelihood of a student's rank ordered application involves integrating

over the prior distribution of the types q:

$$Pr(R_i|\Omega_i^F, X_i) = \sum_{q=1}^{Q} Pr(R_i|\Omega_i^F, X_i, q)\Pi_q(X_i)$$

The econometrician knows the student's observable characteristics X_i, feasible choice sets Ω_i^F, and observed rank of education options R_i. Given this information and a guess of parameters that determine this likelihood, we can back out the posterior probability of a student being of type q from Bayes' rule:

$$\tilde{\Pi}_{q|i}(R_i, \Omega_i^F, X_i) = \frac{Pr(R_i|\Omega_i^F, X_i, q)\Pi_q(X_i)}{Pr(R_i|\Omega_i^F, X_i)} \tag{2.3.6}$$

which we often just write as $\tilde{\Pi}_{q|i}$ for simplicity and denote $\overline{\Pi}_q$ the population shares of type q: $\overline{\Pi}_q = \frac{1}{N}\sum_{i=1}^{N} \tilde{\Pi}_{q|i}(R_i, \Omega_i^F, x_i)$.

2.3.6 Decomposing Causal Effects

The application likelihood allows us to compute the posterior probability of unobserved types q at estimated parameters using Bayes' rule. This in turn allows us to compute conditional expectations given this information. Let Y_{ij}^q be individual i's predicted earnings at program j, conditioning on them being of type q. Then, given their reports and observable information, the posterior expected earnings is:

$$\mathbb{E}[Y_{ij}|X_i, R_i, \Omega_i^F] = \sum_q Y_{ij}^q \tilde{\Pi}_{q|i}(R_i, \Omega_i^F, X_i) \tag{2.3.7}$$

We will henceforth drop the conditioning on X_i, R_i, Ω_i^F for ease of exposition.

Using Equation 2.3.2, we can decompose these into components due to average effects of the program $\overline{ATE}_j$, on-platform selection bias A_{1i}, and match effects/sorting gains

M_{ij}:

$$\mathbb{E}[Y_{ij}] = \underbrace{\alpha_j}_{\textbf{Average Effects: } \overline{ATE}_j} + \underbrace{\sum_q \alpha_{1i}^q \tilde{\Pi}_{q|i}}_{\textbf{Selection Bias: } A_{1i}} + \underbrace{\sum_q \tau_{ij}^q \tilde{\Pi}_{q|i}}_{\textbf{Match/Sorting Gains: } M_{ij}} \tag{2.3.8}$$

where M_{ij} can be further decomposed into components arising from observable and unobservable contributions. The difference matters to the extent that policymakers could target access towards certain types of students (e.g. based on test scores).

Turning to program-level effects, we can also decompose observable differences in average earnings across programs into similar components. Let j_i^{Assign} be the program to which individual i is assigned. Then, we can condition on the population of students assigned to program j and take expectations over that population in addition to the probability of each student's unobservables:

$$\mathbb{E}[Y_{ij}|D_{ij}=1] = \underbrace{\mathbb{E}[A_{1i}|j_i^{Assign}=j]}_{\textbf{Average Selection Bias}} + \underbrace{\left(\mathbb{E}[Y_{ij}|j_i^{Assign}=j] - \mathbb{E}[A_{1i}|j_i^{Assign}=j]\right)}_{\textbf{Treatment on the Treated}} \tag{2.3.9}$$

$$= \underbrace{\mathbb{E}[A_{1i}|j_i^{Assign}=j]}_{\textbf{Average Selection Bias}} + \underbrace{\overline{ATE}_j}_{\textbf{Average Effects}} + \underbrace{\left(\mathbb{E}[Y_{ij}|j_i^{Assign}=j] - \mathbb{E}[A_{1i}|j_i^{Assign}=j]\right) - \overline{ATE}_j}_{\textbf{Average Match/Sorting Gains}} \tag{2.3.10}$$

where the treatment on the treated (TOT) is the causal effect *for those who attend* program j. This can differ from the average treatment effect (ATE) due to selection on gains.

Our model also admits gains from on-platform attendance relative to an off-platform outside option and abstracts from intensive margin effects. For each individual i, this will simply be the difference between posterior on-platform "ability" ($A_{1i} = \sum_q \alpha_{1i}^q \tilde{\Pi}_{q|i}$) and their posterior off-platform earnings ($A_{0i} = \sum_q Y_{0i}^q \tilde{\Pi}_{q|i}$). The difference between A_{1i} and A_{0i} is the average effect of on platform attendance. We also look to what extent students sort into higher average value-added programs when they attend on-platform (difference between $\overline{ATE}_{j_i^{Assign}} + A_{1i}$ and A_{0i}) and the degree this differs due to posi-

tive/negative sorting on gains.

2.4 Research Design and Estimation Approach

Our empirical approach leverages three key sources of variation to recover the primitives of the model from observational data: student applications, observational earnings differences that include OLS-corrected estimates of program-specific premia, and RD designs which serves as an experiment local to the region around an admissions cutoff. First, we estimate reduced form utility parameters and the distribution of unobservables using student applications. Then, conditioning on the posterior distribution of unobservables, we fit empirical moments based on RD, OLS, and descriptive evidence. This approach mixes variation generated from application behavior, where we have easy access to microdata, with indirect inference earnings moments based on reduced from estimates calculated inside the tax agency. The estimation strategy adapts to limitations to data access and computational resources.

2.4.1 Research Design and Identification

Information in Student Applications and Assumptions on Application Behavior

Student applications recover both the distribution of unobserved types as well how the indirect utility for each program varies with observables X_i and unobservables q. We can recover the distribution of unobservables from applications alone since these rank order lists contain information on both match-specific non-pecuniary benefits ("true" preferences) as well as match-specific returns. Intuitively, if a group of students has a known comparative advantage at a certain set of options, this should show up systematically in their rank order lists to the extent that they value earnings gains associated with these characteristics. At the same time, if a group of students has a high non-pecuniary benefit for certain programs (e.g. close to home, less math intensive), this should show up sys-

tematically in their reports as well. The reduced form choice model allows us to estimate the distribution of these unobservables in a first step without taking a stand on whether reports are determined by pecuniary or non-pecuniary attributes of a program.

The structure we impose on the distribution of unobservables is relatively weak, but recovering preferences from applications is not without assumptions.[15] In this regard, we closely follow Kapor et al. [2024] and infer preferences from reports by assuming that students truthfully report their preferences in a subset of relevant options. These are programs at which a student can be reasonably confident of admission (i.e., not a clear "reach"), but are also not ranked below a clearly available option (i.e., not preferred less than a clear "safety"). The key advantage of this approach is that it allows us to use more information to estimate substitution patterns without making the strong assumption that all reports are truthful.

We take a range around each program's cutoff π_j:

$$\underline{\pi}_j < \pi_j < \overline{\pi}_j$$

For each student i, we define:

- **Ex-ante Clearly Infeasible** if $Index_{ij} < \underline{\pi}_j$ (i.e., programs where the student has no chance)
- **Ex-ante Marginal** if $\underline{\pi}_j < Index_{ij} < \overline{\pi}_j$ (i.e., relevant programs)
- **Ex-ante Clearly Feasible** if $\overline{\pi}_j < Index_{ij}$ (i.e., safe options)

We restrict student reports to the set of ex-ante marginal applications that are reported as preferred to an ex ante clearly feasible option. We use a bandwidth $\overline{\pi}_j - \underline{\pi}_j = 25$ in trimming the reports, the same as in Kapor et al. [2024].[16] We deviate slightly in including the next available option if the assigned option is a clearly feasible option. We

[15] For example, it could be that student reports are strategic. In this case, our estimates would conflate indirect utility with the likelihood a student is admitted.

[16] If an individual reports a complete list of only clearly infeasible options, they are considered as ranking the outside option. By definition, no clearly infeasible options will result in admission nor will the individual be included as a marginal admit/reject.

do this to match the margins of choice at which the regression discontinuity analysis is conducted. Moreover, as the bandwidth on which RD designs are estimated is narrower than this bandwidth, all marginally rejected applications are included as credible applications. Note that this procedure will always include the individual's assigned and next available option. We include the first clearly feasible option in the reported preferences used in estimation. We then compare these trimmed reports to all ex-ante marginal and ex-ante clearly feasible options, only excluding the ex-ante clearly infeasible options.[17]

Based on their $Index_{ij}$ and program cutoffs π_j, a student's feasible set of options is defined by the programs that are not clearly infeasible and are also available in the application cycle:

$$\Omega_i^F = \{j : j \text{ available in } t(i) \text{ and } Index_{ij} + 25 \geq \pi_j\}$$

Let $R_i = \{j_1, ..., j_r, ...\}$ denote the students rank-ordered list, where j_r is the r-th rank. Finally, we infer a student as ranking the *outside option* if: (1) the complete list before any application drops is a short-list and (2) the individuals does not report a clearly feasible option and (3) the student has some truthful options available to report.[18]

Given this structure, arguments that identify utility parameters from applications under centralized assignment are standard. First, within-application correlation in characteristics of listed programs helps to identify unobserved preferences. For example, if there is a subset of students that systematically rank history programs very highly in a way that is not related to observables, the preference model interprets these students

[17] By including all clearly feasible options, if $\overline{\pi}_j = \underline{\pi}_j = \pi_j$, then we make the comparison that the assigned option alone is preferred to all alternative options such that $Index_{ij} > \pi_j$. This is the procedure of Fack et al. [2019] based on stability. If clearly feasible options are excluded, we do not get this nesting of their procedure.

[18] Our approach is similar in spirit to Fack et al. [2019], who use the stability property of deferred acceptance to only look at assigned options within the ex-post feasible set. In particular, they estimate preferences by interpreting the assigned option as the students' most preferred option within the set that is available to them. Our strategy converges to theirs as the window defined by $\overline{\pi}_j$ and $\underline{\pi}_j$ converges to the realized cutoffs π_j. The implication is that a narrower bandwidth places fewer restrictions on student behavior.

as having an unobservably high preference for history. This structure of the reduced form preference model does not take a stand on if this is due to a particularly strong amenity value of studying history or a particularly high earnings gain from studying history. Second, we assume that the index and cutoffs themselves are excluded from student preferences directly. Substitution patterns are inferred from how students shift application behavior as available options change.

Evidence from Regression Discontinuities

Using only variation in applications and relating this to observed earnings would be analogous to research designs that control for application fixed effects (Dale and Krueger [2002], Dale and Krueger [2014], Mountjoy and Hickman [2021]) or use control function approaches based on application portfolios (Abdulkadiroglu et al. [2020]).[19] Incorporating RD evidence provides additional variation to pin down program average effects and sorting gains, exploiting the idea that RD designs act as a randomized experiment around the admission cutoff. However, combining this research design with a fully specified model allows us to extrapolate away from the narrow subset of the population around the cutoff.

The RD estimates are purged of selection bias, but not differences due to the match-specific returns due to sorting on gains around the admission cutoff. The key assumption behind RD designs is that students on both sides of an admission cutoff are identical except for their assignment, even if those at the cutoff differ from the population more broadly. An RD estimate across the threshold for students applying to a marginally admitted or rejected target school, j', compared to a next available option, j'', will identify:

$$\underbrace{\mathbb{E}[Y_{i,j'} - Y_{i,j''} | Marginal\ j', Next\ j'']}_{\Delta_{RD}^{j',j''}} = \underbrace{\alpha_{j'} - \alpha_{j''}}_{\Delta_{\alpha}^{j',j''}} + \underbrace{\mathbb{E}[\tau_{j'} - \tau_{j''} | Marginal\ j', Next\ j'']}_{\Delta_{\tau}^{j',j''}}$$

[19] Arcidiacono [2004] estimates a structural model using student application (and admission) portfolios, major choice, and earnings and graduation outcomes.

where $\Delta_{\tau}^{j',j''}$ is the causal effect of crossing the threshold *for those at the threshold margin* (i.e., differences in TOT for marginal applicants) and $\Delta_{\alpha}^{j',j''}$ is the effect for a random individual in the population (i.e., differences in ATE).

As individuals are assumed to have identical unobservables on both sides of the cut-off, the average on-platform ability terms will cancel out. If $j'' = 0$, the next available option is no attendance, then the above formula must include differences in on- and off-platform ability terms, α_{1i} and Y_{i0}. In addition, we also include RDs based on individuals who are marginally rejected into the outside option, which is key for identifying unobservable earnings gains in the outside option. In the case when we only use RD estimates at target options and do not include information on next available options, these RDs are a weighted average of causal effects at the next available options. We match many RD estimates, across different majors and by different combinations of target and next broad areas of study to inform the model parameters.

Evidence in OLS and Descriptive Earnings Statistics

OLS and descriptive patterns in earnings provide biased estimates of returns, as shown in Equation 2.3.9. This variation is not purged of selection bias and is thus key for estimating parameters governing mean on-platform earnings a student would earn at any program when compared to the evidence. In addition, the greater sample size used to estimate these moments allows us to use more of the data to identify relevant parameters.

In practice, we match regressions based on field of study fixed effects that control for student observables. By simultaneously matching average earnings by field of study and field of study effects from regressions with controls we are able to isolate the component due to observable selection bias and match effects from the unobservable component.

2.4.2 Estimation Procedure

We estimate the model in two steps. In the first step, we estimate the reduced form likelihood of applications using the Expectation-Maximization (EM) algorithm (Dempster et al. [1977], Heckman and Singer [1984]) to uncover the distribution of unobserved preferences. Such methods have been applied in a related, but decentralized, setting by Arcidiacono [2004]. In the second step, we use the posterior distribution of unobservables, conditioning on observed student applications, to match summary statistics, OLS regressions, and RD estimates.

We denote $\theta = \{\alpha, \beta\}$ as the full vector of parameters, with α as the vector governing earnings. β are parameters of the reduced form utility model, which includes parameters γ that determine the prior probability of the unobserved types.[20]

First Step Estimation: Reduced Form Preferences and Unobservables through the EM Algorithm

To obtain estimates of β, we seek to maximize the full log-likelihood of observed (credible) applications:

$$L(\beta) = \sum_{i}^{N} \log \left(\sum_{q=1}^{Q} Pr(R_i | X_i, \Omega_i^F, q; \beta) \Pi_q(X_i; \beta) \right)$$

Maximizing this flexible likelihood based on a mixture of unobserved discrete types is known to be computationally difficult. Arcidiacono and Jones [2003] shows that maximizing this likelihood is equivalent to iteratively maximizing a surrogate likelihood that conditions on the posterior of unobserved types and subsequently updating the posterior. Specifically, at iteration t of the procedure with current parameter guess β^{t-1}:

1. *E-Step*: Compute the posterior likelihood of individual i being type q at $\tilde{\Pi}_{q|i}(R_i, \Omega_i^F, X_i; \beta^{t-1})$,

[20] If we instead denote δ parameters governing "true" preferences net of earnings gains, we can always write $\beta = f(\delta, \alpha, \phi)$ according to Equation 2.3.5. However, as stated above, we cannot disentangle ϕ from correlation between δ and α.

using Equation 2.3.6.

2. *M-Step*: Update β^t by maximizing the surrogate likelihood:[21]

$$\mathscr{E}(\beta|\beta^{t-1}) = \sum_{i=1}^{N}\sum_{q=1}^{Q} \log\left(Pr(R_i|X_i,\Omega_i^F,q;\beta)\Pi_q(X_i;\beta)\right)\tilde{\Pi}_{q|i}(R_i,\Omega_i^F,X_i;\beta^{t-1})$$

which is additively separable once we condition on the posterior. We repeat this procedure until convergence.[22] As is standard in the EM literature to avoid local minima, we use multiple random starting points and random initial posteriors, selecting the parameter vector that results in the highest likelihood.

Second Step Estimation: Earnings Effects via Indirect Inference

With the reduced form preference estimates $\hat{\beta}$ and resulting posterior distribution of unobservable types in hand, we can now condition on observed assignments to estimate outcome parameters α. From the tax agency we are not able to retrieve microdata on earnings outcomes and are limited to extracting aggregate statistics. We choose to estimate OLS regressions of earnings on college assignment and individual characteristics as well as estimating different models with regression discontinuity designs. We extract these estimates and use them to generate indirect inference moments which we can calculate outside the tax agency with the model. We also match some aggregate statistics that we view as particularly important to target.

In total, we target 368 moments in order to estimate 266 parameters:

1. *Major and Institution OLS Regressions (135 moments)*: Coefficients from a regression of log-earnings on institution (23) and major-level fixed effects (105 coefficients) with quadratic controls in quantitative and verbal PAA scores (4 coefficients) as

[21] In principle, the prior parameters could be maximized separately from choice parameters in a second sub-step due to the log-additivity of the surrogate likelihood conditional on the posterior. We find it easier to do this jointly in practice.

[22] To speed up convergence, on each M-Step, we maximize the likelihood using a single Brendt-Hall-Hall-Hausman iteration. We repeat this until $|\beta^t - \beta^{t-1}| < 10^{-6}$.

well as indicator variables for combinations of gender-by-private high school (3 coefficients). The regression we target here also controls for experience and year effects, although we do not target those. This regression controls for all main effects of the major fixed effects and excludes the University of Chile from the institution effects.

2. *Major RD Estimates (96 moments)*: These are RD estimates for students marginally accepted or rejected from a program associated with one of 105 majors.[23] These pool target options that offer the same major and aggregates over the distribution of next available options.
3. *Area-by-Area RD Estimates (83 moments)*: These are broad area of study RD estimates for students marginally accepted or rejected from a program associated with one of 10 broad areas of study. Unlike the major RD estimates, these stratify the estimates based on a marginal applicant's next available option based on their reported preferences, which can include the outside option.[24]
4. *Area Mean Earnings (10 moments)*: For each broad area of study, we target the average earnings of individuals assigned to the area.
5. *Average Earnings by Test Scores and Demographics (40 moments)*: We match the test score gradient in log-earnings by verbal and quantitative scores for different sets of observable groups. To do so, we target average log-earnings within 5 bins of verbal and quantitative test scores for men, women, public high school, and private high school (4 options). These moments are computed on the universe of test takers.
6. *Variance Terms (4 moments)*: As decompositions of predicted earnings are key outputs of the model, we want to do a good job matching these moments related to the dispersion of earnings. At the program-level, we target the variance of program-level mean earnings as well as the mean of program-level mean earnings, weighting

[23] 9 majors are excluded because they fall below the desired sample size of 5000 to be included as a targeted RD moment.

[24] Of 110 possible area-by-area RD estimates, 27 are excluded because they fall below the desired sample size of 5000 to be included as a targeted RD moment

programs equally. At the individual-level, we target the variance in log-earnings for the entire population of test takers as well as the variance in log-earnings conditional on receiving an assignment.

To compute the RD moments in the tax data, we restrict to students whose admission index at a program was within 15 points of an admission cutoff at a program and control for a linear trend above and below the cutoff. To compute the RD moments in estimation, we directly impose the identification assumptions justifying the research design – that individuals are identical on observables and unobservables on both sides of the admission cutoff. Specifically, for those marginally rejected from a program, the model gives predictions of counterfactual earnings for the same individual at both target and next-available options. This approach both takes the research design of the RD seriously in estimation.

Standard Errors and Weighting Matrix

The sequential nature of the estimation poses challenges for inference. Our setting is nonstandard in that the first step uses microdata on student reports, while the second step matches moments computed using restricted tax data with simulated moments based on earnings predictions *for the same population of applicants on which the first step is estimated*. As stated above, it is infeasible to estimate the model using microdata on earnings due to computational and data access limitations at the tax agency.

We adopt a conservative approach to inference based on computing the asymptotic covariance of the moments as implied by the estimation sample of applicants and treating the statistics from the tax agency as known quantities. While in principle one may be concerned about sampling error in the moments we match from the tax agency, we are more concerned about sampling variation driven in our smaller estimation sample.[25] However, this approach still allows for sampling error to affect estimates, which

[25]We also avoid matching moments from the tax agency where we have little data to further avoid concerns with precision in these estimates. For example, we do not match RD moments where there are fewer

is larger for moments where we have fewer assignments to match on. Specifically, we compute each individual observation's contribution to an aggregate moment and use this to compute the asymptotic covariance of these moments in the estimation sample.[26] This permits us to compute the joint asymptotic covariance of the first and second step parameters using standard GMM techniques.

In the earnings estimation, we use a two-step GMM procedure to optimally weight the moments. First, we estimate the model weighting each targeted moment by the inverse squared value of the moment computed at the tax agency. This approach accounts for the scale of different targeted moments but is otherwise agnostic about the appropriate amount of weight to put on each individual moment. Then, after an initial estimation, we update the weight matrix to be the inverse variance of the moment implied by our aforementioned procedure.

2.5 Estimation Results

2.5.1 Parameter Estimates and Model Fit

Parameters governing the distribution of the prior for unobserved types is reported in Table A1 along with an ex-post characterization of each type. Parameters from both the model of college choice and earnings outcomes are reported in Table A2, Table A3, and Table A4.

Choice Model: The choice model fits well, with a detailed set of fit exercises shown in Figure A1. We focus on two moments of the data that are difficult to match in practice and are key for our second step estimation and the counterfactual analysis.

than 5000 observations used in computing the estimate.

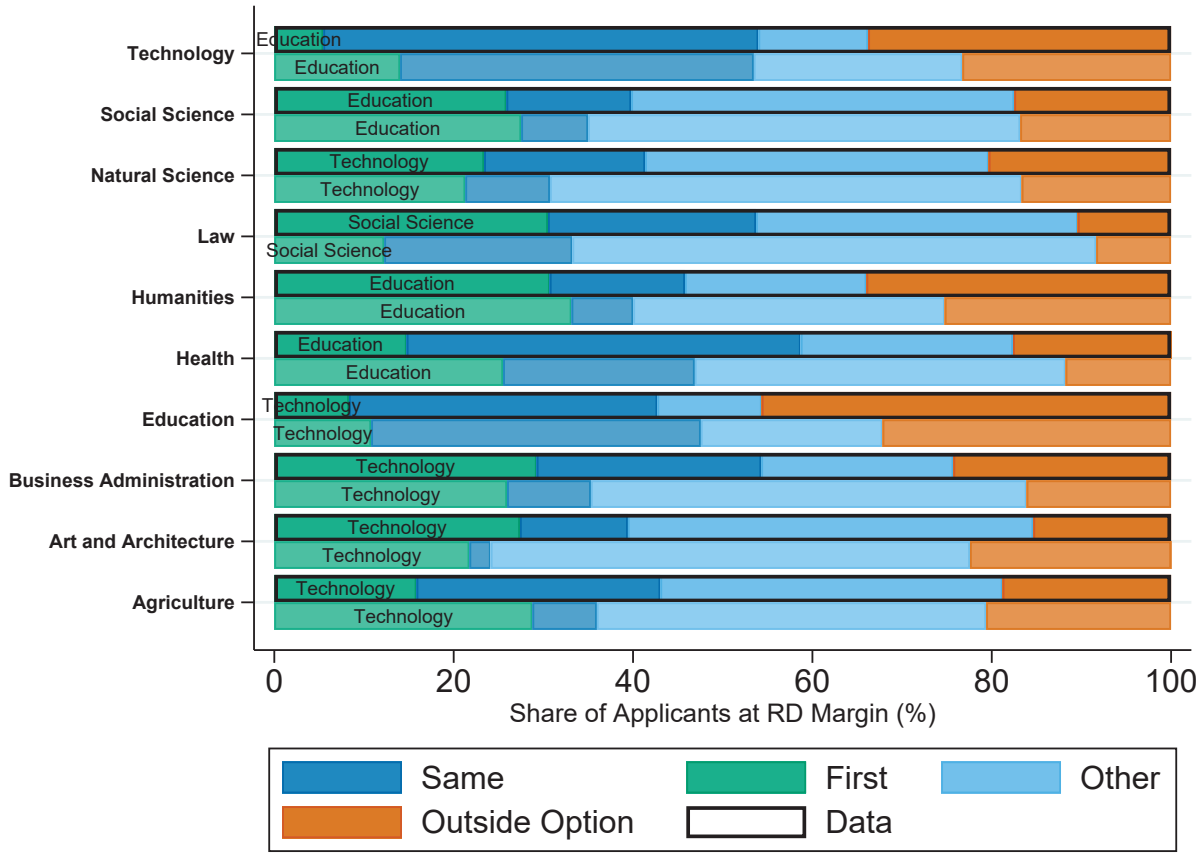
[26] For OLS moments, this involves computing individual regression errors implied by the model and adjusting the variance of these errors by the second moment of the design matrix. Similarly, we make adjustments using the delta method for moments based on the variance of earnings in the population and the variance of program-level average earnings. We also make adjustment to the variance-covariance matrix for the fact that many moments are computed on subsamples of the data (and thus have a higher variance). Details of these procedures are available upon request.

First, Figure 2.1 shows the fit of broad area of study for applications at the RD margin. It is important that this cut of the data fits well, as we will use this population of students to fit the earnings model around the admission cutoffs. The fact that this aspect of the data has a good fit indicates that we are matching persistent within-application patterns for preferences for different fields of study, even at this narrow margin. For these marginal students, we stratify by the target options at which as student is marginally rejected or admitted and show the distribution of next available options. In particular, we categorize the next available options based on if they are in the same broad field of study (dark blue), the most popular next alternative area in the data (green), all other areas of study (light blue), and then the outside options (orange). The statistics from the raw data is outlined. For example, of those marginally rejected/admitted into a Social Science program, about 25 percent of these students had a next ranked option to which they were admitted that was in Education. The share of students whose next available option is the outside option also exhibits good fit.

Second, Figure 2.2 shows the fit of the choice model when we redraw preference shocks and feed predicted ranks into the assignment mechanism. While we exploit truth-telling assumptions in a subset of credible applications to estimate preferences, we do not actually evaluate the mechanism in estimation. However, since we do so when we run our counterfactuals, we hope that the model does a reasonable job predicting the assignments and cutoffs at estimated parameters under the baseline environment. As we are randomly drawing idiosyncratic utility and latent types, there will be some uncertainty when evaluating the mechanism but key aggregate statistics should match well.

We find these moments fit well, as shown in Figure 2.2. Panel (a) shows the predicted share of assignments at each broad area of study and the outside option of no assignment. With perhaps the exception of the outside option (26 percent in data versus 20 percent in the model) and technology degrees (28 percent in data versus 19 percent

Figure 2.1: Fit of the Choice Model at the RD Margin



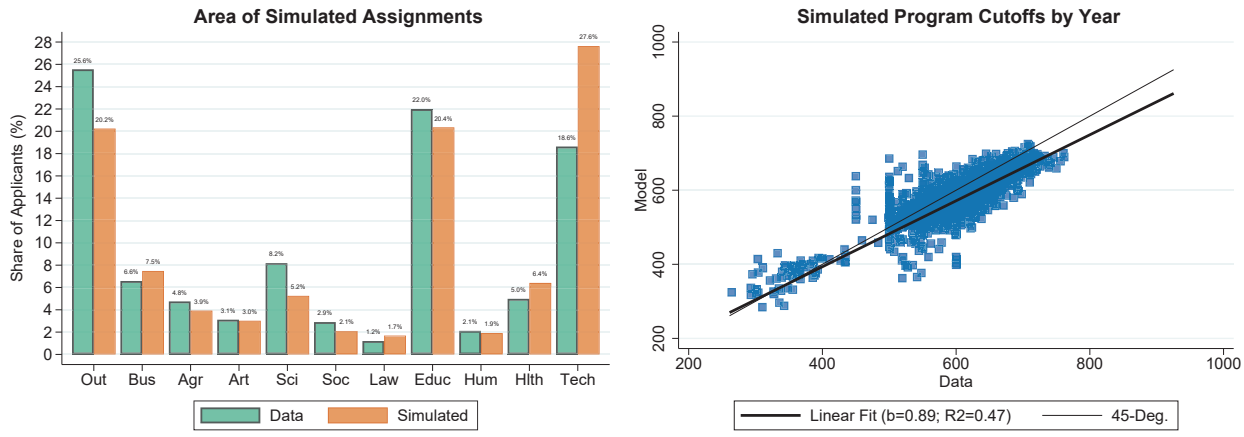
Notes: Data is outlined with a black border. For each application within the RD bandwidth of 15 points relative to the admissions cutoff, this plots the distribution of next available options relative to the the target area of study. Target option is the option to which the individual was marginally admitted/rejected: for those marginally rejected, next available will be the option assigned. For each target area, "First" is the share in the most common next available option for the target area in the data; "Same" is the share of individuals whose next options is a program within the same broad area of study; "Other" combines all other options; and "Outside Option" are those who do not have a next available on-platform option.

in the model), we match these area of study shares extremely well. Panel (b) compares the admissions cutoffs simulated by running the mechanism with new preference ranks compared with the administrative reports of program cutoffs. These also match very well, with a regression coefficient of 0.89 and and R^2 of 0.47.

The fact that this aspect of the choice model matches the data gives us some comfort in using the model to extrapolate away from the observed data.[27] More importantly, this exercise validates our use of administrative reports of the cutoffs and shows a limited role for the sample drops mentioned in Section 2.2.

Figure 2.2: Fit of Mechanism at Choice Parameters

(a) (b)



Notes: Simulation results from redrawing preference shocks at estimated model parameters and feeding into the DA mechanism using capacities and index scores in the data. Figure (a) shows the composition of broad areas of study as well as the outside option. Figure (b) shows admission cutoffs in the simulated dataset (i.e., the index score of the last person admitted to an option).

Earnings Model: We now turn to evaluating the fit and parameters of the earnings model, holding assignments fixed to those in the data as is done in estimation. As we have more moments than parameters, we can use this extra information to do a joint test of the fit of the model. At the estimated earning parameters, we compute a Sargan-Hansen J-statistic of 59.4, indicating that the over-identifying restrictions imposed by the moments do not reject the model at the 102 degrees of freedom.[28]

[27] When comparing counterfactuals to the baseline, we will always simulate new assignments and hold unobservables fixed across counterfactual experiments in order to compare apples to apples.

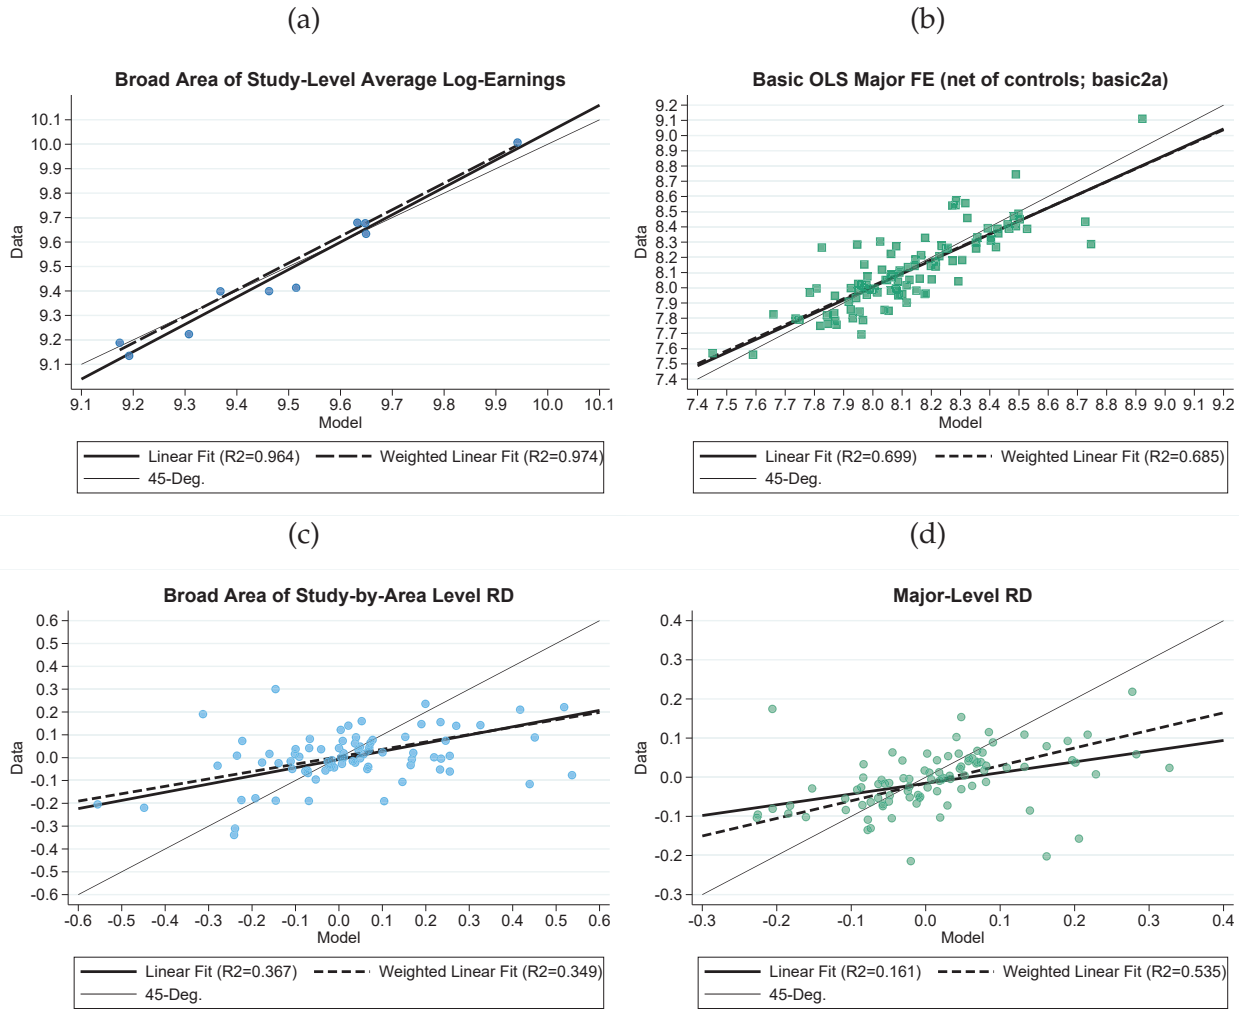
[28] The 10% critical value for this test is 126.

Figure 2.3 shows the fit of the model at moments that we target in estimation. Panel (a) shows the average log-earnings at assignments in each broad area of study. These moments are estimated very precisely and receive a large amount of weight in the second step of the estimation. As a result, we fit these average moments extremely well. Panel (b) shows the fit of major fixed effects in the OLS regression that controls for institution effects, quadratics in verbal and quantitative test scores, and gender-by-high school type effects. Again, these targeted moments exhibit good fit, with an R^2 of 0.699.

Panel (c) and and (d) show how the model fits at the RD moments for the target-next RDs for broad areas as well as for the target RDs for majors. These receive relatively less weight individually as the number of individuals in each RD makes up a small share of the overall sample. However, we match general patterns well, with an R^2 of 0.367 for the area-by-area RDs. For RDs that have relatively larger in absolute value, the model tends to predict more "extreme" values for these effects. We also see that when we weight the major RDs by the precision of the estimates, the fit improves dramatically with an R^2 of 0.535.

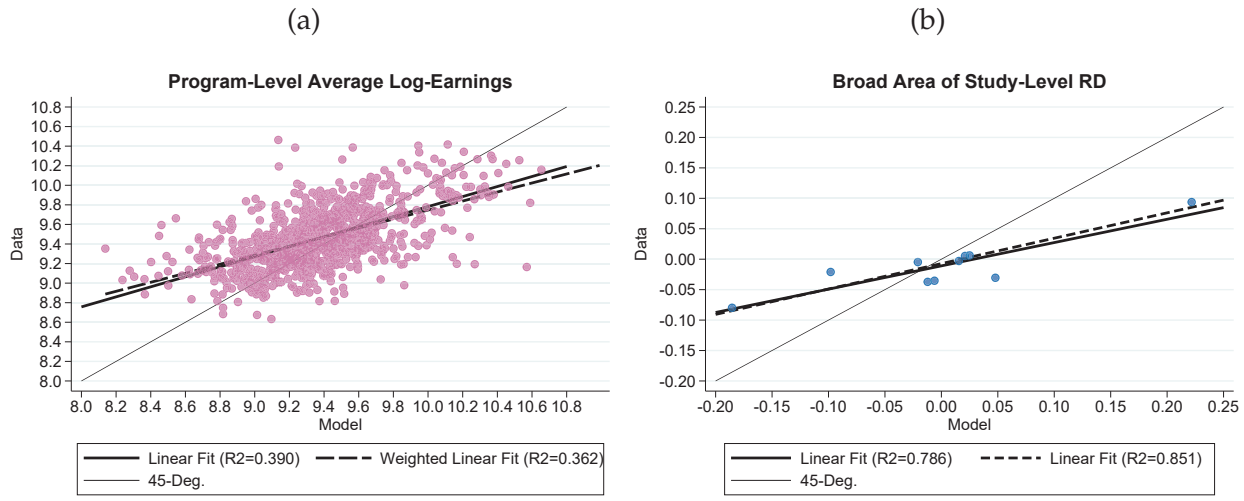
Figure 2.4 assesses the fit of the outcome model at moments that are not explicitly targeted in estimation. Panel (a) looks at the average log-earnings in each DEMRE program, averaging across years and comparing to averages in the data. Even though we do not attempt to match anything at this level of granularity, we still get a strong relationship between average earnings predicted by the model and the raw data. Panel (b) looks at the predicted RD estimates from Figure 2.3 aggregated up to the broad area-level and without stratifying by next best options. We see that we fit these moments quite well. Again, we see that for larger RD estimates in the data, the model tends to predict even larger effects, although it matches the order of the estimates quite well.

Figure 2.3: Fit of Earnings Model: Targeted Moments



Notes: Plots fit of moments targeted in the second-step earnings estimation. Moments are predicted at estimated earnings parameters using the posterior distribution of unobservables predicted by an applicant's observed rank order list. "Weighted Linear Fit" is weighted by the estimated standard error using data from the tax agency.

Figure 2.4: Fit of Earnings Model: Non-Targeted Moments



Notes: Plots fit of moments not targeted in the second-step earnings estimation. Moments are predicted at estimated earnings parameters using the posterior distribution of unobservables given an applicant's observed rank order list. "Weighted Linear Fit" is weighted by the estimated standard error using data from the tax agency.

2.5.2 What are the Latent Types?

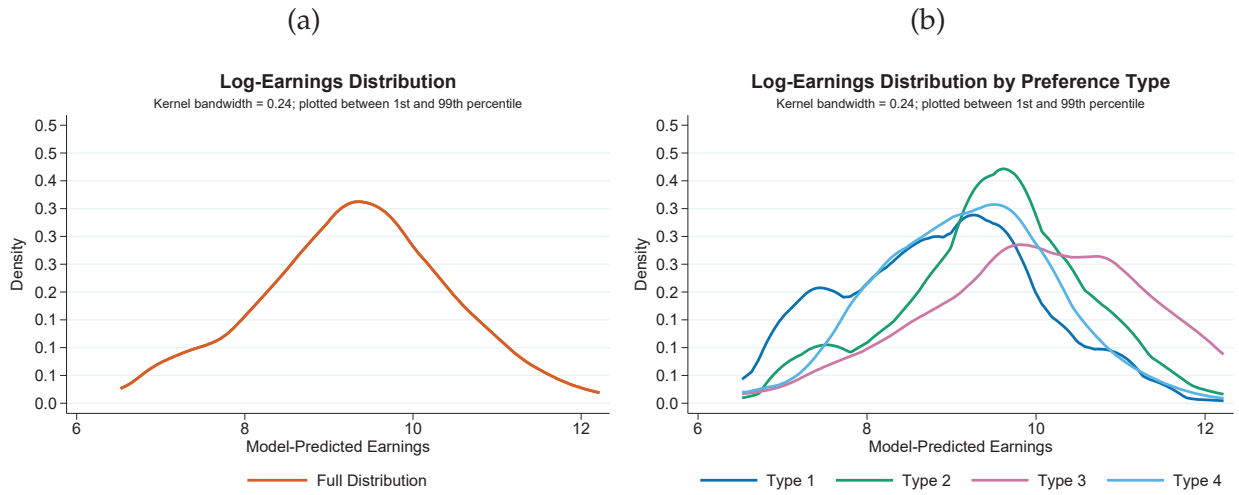
Table A1 characterizes the unobserved types. All types have somewhat equal share in the data, ranging from 19 percent for type 3 to 30 percent for type 4. We also report share of individuals with a posterior of greater than 0.9 conditional on being of that type, indicating how "confident" the model is of each individual's assignment. We see that the model is quite confident, but that there is some degree of uncertainty about an individual's categorization.

Type 3 is a higher socioeconomic status (SES) than the rest of the population, having higher *PAA* scores and with 39 percent combing from private high schools. Type 4 is a slightly lower SES, with 27 percent coming from a private high school and 43 percent being male. Type 2 is predominately male at over 94 percent. This group also has the lowest level of average test scores. Finally, Type 1 has average test scores and a much lower share coming from private high schools; 36 percent of this type is male.

These patterns are confirmed by looking at Figure 2.5, which shows the ex-post earnings distribution of unobserved types at their assigned options. These density estimates

use the same bandwidth and weight each individual's predicted Y_{ij}^q by the posterior of being from that type $\tilde{\Pi}_{q|i}$. Not only does Type 3 have a higher share of coming from a private high school, but their expected income distribution is shifted to the right as well. However, we see that Type 2, which is male dominated and has the lowest test scores, also has an income distribution that is shifted to the right, indicating that they end up with high returns at their assigned options despite low test scores.

Figure 2.5: Earnings Distribution



Notes: Plots kernel density estimates of expected log-earnings in (a) the unconditional distribution and (b) by each latent type (b). Density estimates are computed using data between the 1st and 99th percentile of the unconditional expected earnings distribution. To estimate kernel densities by latent type, we take each observed applicant's type-specific earnings Y_{ij}^q and weight the density estimate by the posterior probability the individual is of that type $\tilde{\Pi}_{q|i}$. All density estimates use an Epanechnikov kernel and a common bandwidth determined by a one-fifth of a standard deviation in the unconditional earnings distribution.

2.5.3 Variation in Preferences and Earnings: Do Students Value College-Major Earnings Gains?

We now investigate the distribution of mean preferences and earnings gains across programs, separating these out by latent preference types. We compute the mean utility for

program j for type q by averaging over the posteriors:

$$\delta_j^q = \frac{1}{\overline{\Pi}_q} \frac{1}{N} \sum_{i=1}^{N} U_{ij}^q \tilde{\Pi}_{q|i} \tag{2.5.1}$$

which we then demean within each application year. We only focus on intensive margin differences between programs in this analysis. The mean earnings gains α_j^q are computed similarly. We also compute a common component by weighting by the population type shares: $\bar{\delta}_j = \sum_{q=1}^{Q} \delta_j^q \overline{\Pi}_q$. For earnings, it follows that the average treatment effect for the population (i.e., the average effects on earnings for someone drawn randomly in the population) is $ATE_j = \bar{\alpha}_j = \sum_{q=1}^{Q} \alpha_j^q \overline{\Pi}_q$

The variation in mean earnings and preferences for each type and the overall population is summarized in Table 2.2. There is substantial dispersion in preferences and earnings gains across the population and within each type. For example, the variance in average program earnings effects within types ranges from 1.15 for Type 1 to 2.32 for Type 3.

To investigate whether these components are vertically or horizontally differentiated, we compute how correlated each type-specific mean preference is with the population-wide component that weights by population type shares. This is similar to performing an "ensemble" variance decomposition in two components.[29]

Table 2.2 shows that mean preferences are vertically differentiated, with the common component explaining 67% of the variance in utility for programs across types. In contrast, earnings exhibit a substantial amount of horizontal differentiation, with only 22% of earnings is explained by the common component. Indeed, we see that the variance in the common component of earnings effects is substantially smaller than each type individually, being about 12% of the magnitude of type 3. This indicates that students tend to "agree" on what is a desirable program, but have comparative advantage across

[29] Specifically, as $\delta_j^q = \bar{\delta}_j + \left(\delta_j^q - \bar{\delta}_j\right) = \bar{\delta}_j + \tilde{\delta}_j^q$, then $Var(\delta_j^q) = Cov(\delta_j^q, \bar{\delta}_j) + Cov(\delta_j^q, \tilde{\delta}_j^q)$.

programs in terms of earnings value added.

Do students value these heterogeneous earnings gains? While we cannot directly estimate the "structural" preference for earnings gains (captured in the parameter ϕ) without more assumptions or additional variation, we can relate estimated indirect utility with estimated returns at each program.

Our estimates indicate that there is a moderate preference for earnings value added and that this varies across student types. Across the population, mean utility and ATE have a correlation of 0.250 or a regression coefficient of 0.7. This varies across the student population, strengthening slightly when we look across student types. with Type 2 shows the strongest correlation of 0.45, implying a large regression coefficient of 0.8, helping to explain this types higher observed ex-post earnings. In contrast, Type 1 and Type 2 have the weakest association between preferences and earnings gains. This underscores the flexibility of our model – the implied preference for earnings is left unrestricted in our framework as well as the fact that students differ in the preferences for earnings across the distribution.

This finding indicates that there is some scope for students to respond to information on program value added. This stands in contrast to work on test score match effects, which has tended to find a limited association or perhaps reverse Roy selection on earnings gains (e.g. Walters [2018], Abdulkadiroglu et al. [2020]). Additionally, the fact that we find a strong relationship shows that our earnings estimates are presumably picking up "real" returns.

Table 2.2: Variation in Preferences and Earnings Effects

	Common	Type 1	Type 2	Type 3	Type 4
$\text{Var}(\bar{\delta}_j)$	2.335	1.965	4.000	5.070	3.138
Share Common	0.666				
$\text{Var}(\bar{\alpha}_j)$	0.286	1.151	1.242	2.319	0.571
Share Common	0.220				
$\text{Cor}(\bar{\delta}_j, \bar{\alpha}_j)$	0.250	0.339	0.452	0.414	0.217
Type Share		0.243	0.266	0.194	0.297

Notes: Decomposition of the variance in type-specific mean utility and earnings gains compared to unconditional mean utility and average treatment effects. δ_j^q and α_j^q are computed according to Equation 2.5.1. Variances and covariances are unweighted. The "common share" is computed according to Footnote 29.

Finally, we turn to a decomposition of the expected earnings distribution at the individual level in Table 2.3. We follow the decompositions discussed in Section 2.3. The total variance in this distribution is 1.31, where the full distribution is shown in Panel (a) of Figure 2.5. We see that ability bias explains 32.4 percent of the variance in earnings, match effects/sorting on gains explains 66.4 percent (with 67 percent of this share being observable), and average effects explain only 1.2 percent.

Looking at the full decomposition that breaks out the covariance terms, ATE explains a larger share at 20.2 percent of earnings. In other words, while ATE has relatively little association with earnings on average, it is still a quantitatively important component. What explains this difference is that ATE is negatively correlated with the unobservable match effect. This means that programs that have a high degree of sorting on gains for students who attend them (i.e., TOT) have low average gains for the typical student.

Table 2.3: Decomposition of Individual Expected Earnings

	Full Decomposition		Ensemble Decomposition	
	(1) Variance	(2) Share	(3) Covariance	(4) Share
Total Variance in Expected Earnings	1.310		1.310	
Variance of Ability	0.418	0.319	0.424	0.324
Variance of ATE	0.265	0.202	0.016	0.012
Variance of Observable Match	0.598	0.457	0.590	0.450
Variance of Unobservable Match	0.548	0.418	0.281	0.214
2Cov(Ability,ATE)	0.027	0.021		
2Cov(Ability,Observable Match)	-0.015	-0.011		
2Cov(Ability,Unobservable Match)	-0.001	-0.001		
2Cov(ATE,Observable Match)	0.003	0.002		
2Cov(ATE,Unobservable Match)	-0.528	-0.403		
2Cov(Observable Match,Unobservable Match)	-0.006	-0.004		

Notes: Variance decompositions of individual expected earnings at assigned options computed according to Equation 2.3.7, conditioning on the observed application lists. This decomposition excludes students assigned to the outside option. The first two columns report the variance components and their shares, accounting for covariance terms. The last two columns report an ensemble decomposition computed by looking at the share of the variance in the outcome explained by the covariance between predicted earnings and each component, ignoring cross-covariance terms. There are no error terms as these are computed from model-predicted objects.

2.6 Decomposition of Treatment Effects

In this section, we decompose treatment effects by considering assignments other than those occurring in the data. Table 2.4 shows the decomposition of the expected earnings distribution into student ability, program ATE, observable match effects, and unobservable match effects. These results differ from Table 2.3 in that we average across individuals within a program.

The program-level results echo the findings from Table 2.3. Looking at the ensemble decomposition, we see that ability bias explains 38.5 percent of the variation in program average earnings; match effects/sorting on gains accounts for 64.8 percent (of which 23.4 percent is observable); ATE accounts for -3.2 percent, being uncorrelated with average program earnings.

There are two key differences from the individual-level results. First, the relative magnitude of observable and unobservable match effects is different. This indicates that observable match effects are responsible for more within-program earnings dispersion

while unobservables are more important across programs.

Table 2.4: Decomposition of Program-Level Effects

	Full Decomposition		Ensemble Decomposition	
	(1) Variance	(2) Share	(3) Covariance	(4) Share
Total Variance in Program Expected Earnings	0.273		0.273	
Variance of Ability	0.099	0.363	0.105	0.385
Variance of ATE	0.286	1.047	-0.009	-0.032
Variance of Observable Match	0.043	0.158	0.042	0.152
Variance of Unobservable Match	0.422	1.544	0.136	0.496
2Cov(Ability,ATE)	0.004	0.015		
2Cov(Ability,Observable Match)	-0.003	-0.010		
2Cov(Ability,Unobservable Match)	0.010	0.037		
2Cov(ATE,Observable Match)	-0.006	-0.022		
2Cov(ATE,Unobservable Match)	-0.589	-2.152		
2Cov(Observable Match,Unobservable Match)	0.005	0.020		

Notes: Variance decompositions of program-level average expected earnings at assigned options, conditioning on observed application lists. The first two columns report the variance components and their shares, accounting for covariance terms. The last two columns report an ensemble decomposition computed by looking at the share of the variance in the outcome explained by the covariance between predicted earnings and each component, ignoring cross-covariance terms. There are no error terms as these are computed from model-predicted objects.

Second, in the full decomposition, ATE becomes more important at the program level, having a variance that is larger than the variance in average program earnings. Again, ATE is a quantitatively important object but has a small role in explaining observed earnings dispersion. There is a *negative correlation* between average unobservable match effects and program ATE. This shows that programs that tend to have high average earnings do actually have large causal effects, but only on a subset of the population.

What these decompositions mean for policy is at this point ambiguous. Sorting on gains and sorting into programs based on ability is particularly important for generating observed program effects. However, the wide variation in predicted ATEs also shows that there is a lot to potentially be gained by targeting programs with high ATEs. In the next section, use our choice model to predict how students will change their applications in response to different available options, allowing their earnings to change according to these model predictions.

2.7 Counterfactual Experiments

2.7.1 Alternate Admissions Policy

The first counterfactual policy we analyze keeps the number of available seats fixed to the baseline scenario in the data, but alters the process by which the seats are allocated. At baseline, seats are allocated according to program-specific index that is determined by ex-ante weights set on *PAA* score components. In all counterfactuals, we hold the random draw of unobserved types and preference shocks fixed.

First, we consider adjusting this index to a lottery mechanism that admits students using a randomly drawn index unique to each program. This approach does worse relative to the baseline on several dimensions. There are fewer seats that end up actually filled, falling 0.7pp to 98.9 percent. Average earnings in both the population and among assigned students also falls by 10.7 percent, from $26,887 at baseline to $24,002.

We also consider lotteries with a local regional bonus.[30] While this regional bonus slightly increases the share of students matched to a program relative to the baseline lottery, it does even worse on earnings. Average earnings are lowered by 15.3 percent relative to the baseline. As students display a strong sensitivity to distance, this is due to students selecting programs in their home region that have worse earnings outcomes.

[30]Specifically, we set the lottery offer index for each program to have the same range as the baseline admission index and consider an extra bonus of 50 points for every program within the student's home region.

Table 2.5: Effects of Alternate Admissions Policies

	Share Applicants Matched	Share Seats Filled	Average Earnings	Average Earnings (Assigned)
Baseline	79.8	99.6	26887	29832
Lottery	79.2	98.9	24002	26364
Scores	79.9	99.7	27343	30978
Earnings	79.5	99.3	62903	74197
Lottery w/ Region Bonus	79.3	99.0	22786	24839

Notes: Each row is a different counterfactual policy to allocate program seats, fixing program capacity across counterfactuals. "Baseline" is the baseline counterfactual that draws reports and earnings at estimated parameters and runs the mechanism to get new assignments. "Random" considers a scenario where admission to oversubscribed programs is broken with a random lottery. "Scores" are where admissions is done through a uniform index across programs based on the average of verbal and quantitative *PAA* scores. "Earnings" considers a scenario where the index to get into a program depends on a students earnings at that program. "Random w/ Bonus" considers a random admission index but with a bonus for schools in a student's own region. The random index is set to be equivalent to a 50 point bonus at the scale of the *PAA* scores.

The counterfactual results thus far indicate that the current admissions system does reasonably well. However, we do find that there are still potentially substantial earnings gains left on the table in the current admission scheme. We consider a counterfactual admissions policy that allocates seats based on a student's potential earnings at that program ("Earnings"), setting $Index_{ij} = \mathbb{E}[Y_{ij}|X_i, \Omega_i^F, R_i]$. In this case, student earnings rise substantially – more than doubling to $62,903. However, this counterfactual is of course infeasible and takes the extrapolations from the model very seriously. For example, we do not account for the fact that the return to an engineering degree will go down if we double the number engineers.

A more reasonable counterfactual is to allocate seats to each program *according to a uniform* admission index, rather than having program specific indices based on the weights the program sets ("Scores"). In this case, average earnings increase slightly by 1.7 percent to $27,343. This indicates that setting admission criteria based on *PAA* test scores increases earnings, but use of program-specific weights may actually slightly lower earnings.

2.7.2 Seat Reallocation Policy

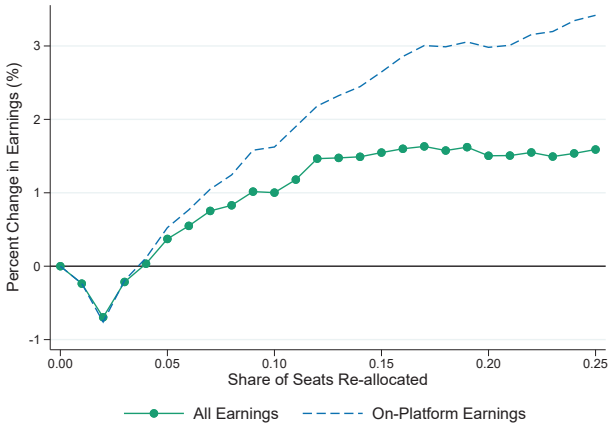
What would be the earnings effects of changing program capacity, holding the admission policy fixed? In this section, we choose a share of the total number of seats to reallocation. Given a desired share of seats to reallocate, we delete seats at the programs that do the "worst" on some index of earnings gains (ATE, average earnings, etc.). Then, we take the tonal number of deleted seats and double the capacities at each program that performs the best on the desired index until all seats are reallocated. This procedure is repeated within each admission cycle. Again, draws of the unobservables are held fixed across evaluations of the mechanism.

Figure 2.6 reallocates based on ATE – expanding programs where a student randomly drawn from the student population would be expected to perform best. This policy initially lowers earnings as programs with match effects are deleted but eventually raises student earnings. We see that after reallocating 20 percent of the seats, all student earnings are 1.8 percent higher with earnings for assigned students increasing by 3 percent. However, this reallocation policy ends up lowering student utility, as desirable programs are eliminated

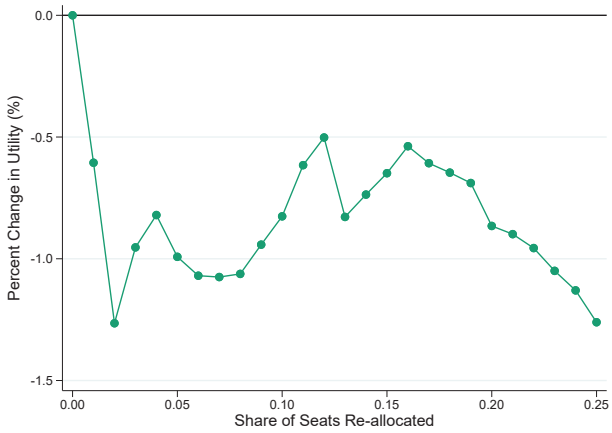
Figure 2.7 reallocates based on average earnings. This reallocation index will blend ability bias, ATE, and sorting gains, but is potentially easier for policymakers to implement. This policy raises earnings more than allocating based on ATE, with the relative difference peaking at around reallocating 10 percent of seats based. At that share, student earnings are 2.2 percent higher when reallocating based on average earnings but only 1 percent higher when reallocating based on ATE. Moreover, this approach increases overall student utility. This indicates that simply using average earnings to assign seats does reasonably well, which is good news for observational research designs. This is because these biased estimators reflect a substantial amount of sorting gains, and reallocating based on average earnings increases seats at program with high treatment on the treated.

Figure 2.6: Re-allocating Seats from Low to High-ATE Programs

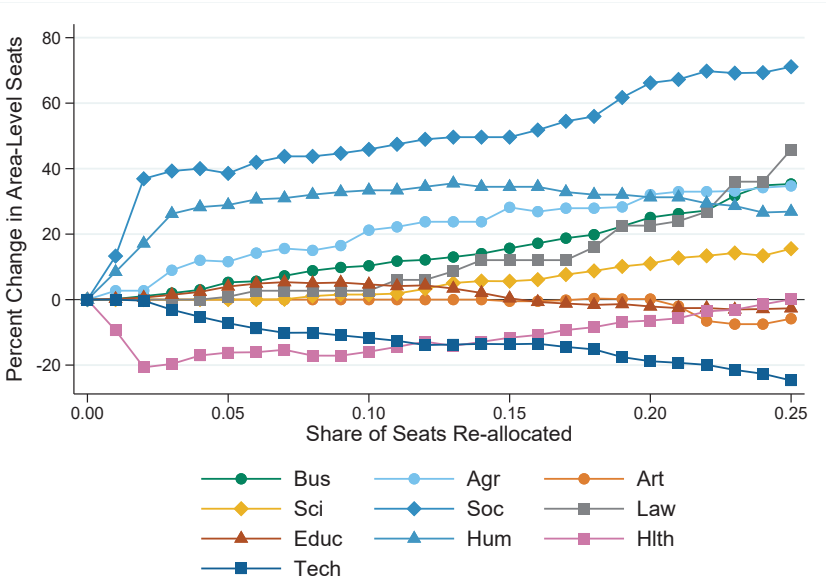
(a) Change in Predicted Earnings



(b) Change in Utility



(c) Percent Change in Seats by Broad Area of Study



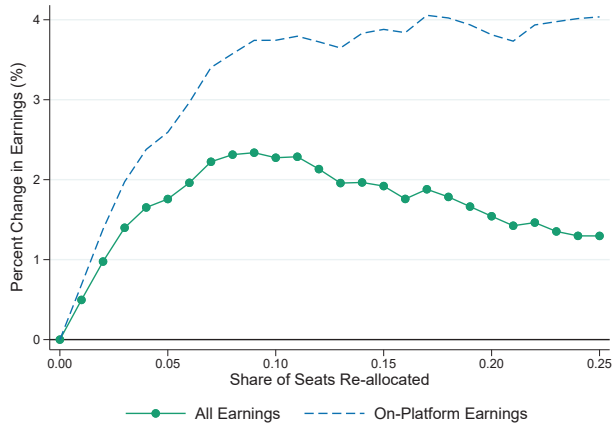
Notes: Each dot shows counterfactual scenarios where a different share of seats has been eliminated at the lowest ATE_j programs and capacities at the highest ATE_j programs are doubled until all seats have been reallocated. Figure (a) shows changes in average earnings for students assigned on-platform and for the entire population of students. Figure (b) shows changes in average utility, measured relative to the outside option. Figure (c) shows the percent change in available seats by broad field of study as a results of the reallocation policy.

However, as policy makers move deeper into the pool of potential enrollees, this metric for setting capacity becomes less useful. For example, at 20% of seats reallocated, the improvement in average earnings indexing by average program earnings is lower than reallocating based on ATE. This means that observational earnings are good for more marginal policy changes, but that ATE is a relevant parameter when policies drive deeper into the student population.

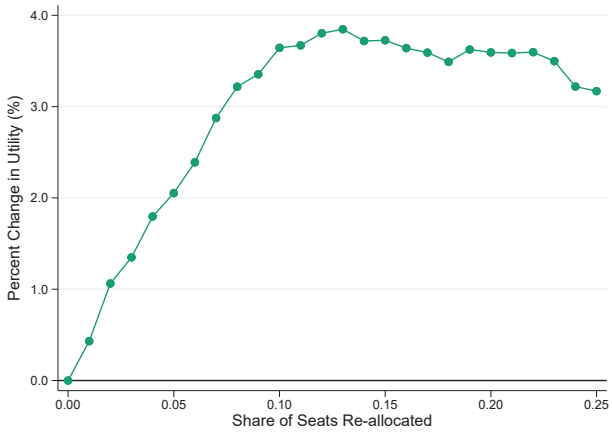
It is interesting to compare which fields of study are shifted under each reallocation policy, as they show substantial differences. The percent change in the number of seats by broad field of study is shown in Panel (c) of Figure 2.6 and Figure 2.7. ATE ends up increasing the number of seats in the Humanities, Social Sciences, and Agricultural programs while reducing enrollment in Technology and Health. In contrast, reallocating based on average earnings has a strong increase in enrollment in Health programs, with smaller increases in Technology and Business. This implies that the average student (in this context) benefits from studying these fields more than the fields we would typically associate with large returns. The average student would see lower earnings in Health and Technology *despite these being high return fields*. This is because these fields of study are associated with larger match effects – indicating that only certain students with sufficient preparation will perform well in them.

Figure 2.7: Re-allocating Seats from Low to High-Average Earning Programs

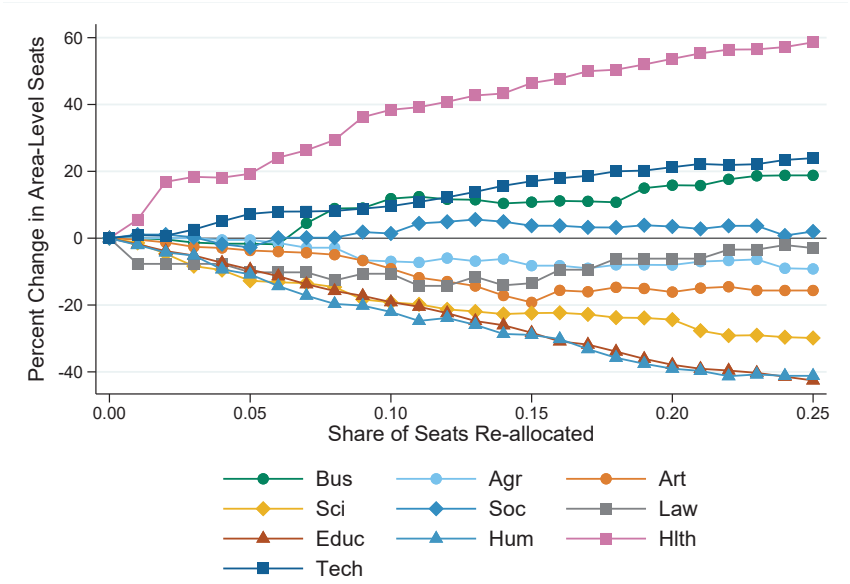
(a) Change in Predicted Earnings



(b) Change in Utility



(c) Percent Change in Seats by Broad Area of Study



Notes: Each dot shows counterfactual scenarios where a different share of seats has been eliminated at the lowest average (predicted) earning programs and capacities at the highest average (predicted) earning programs are doubled until all seats have been reallocated. Figure (a) shows changes in average earnings for students assigned on-platform and for the entire population of students. Figure (b) shows changes in average utility, measured relative to the outside option. Figure (c) shows the percent change in available seats by broad field of study as a results of the reallocation policy.

2.8 Conclusion

We develop a rich framework of student sorting into college majors on the basis of heterogeneous returns and preferences. We exploit the structure of Chilean of higher education, which allocates students according to a centralized mechanism. Our empirical approach provides a unified framework that ties together a variety of reduced form estimates. As such, we are able to extrapolate away from observed data and treatment effects in order to analyze how a alternative policies would impact students across the ability distribution, not just at the margin of admission.

Our results indicate a key role for sorting on observable and unobservable match effects in generating observed earnings dispersion across individuals and college majors. There is a significant amount of ability bias, but the majority of why some college majors earn more than others is due to causal returns that are specific to certain subpopulations of students. This result is substantially different than what is found in the K-12 literature in the United States that focuses on test score gains (Abdulkadiroglu et al. [2020]). In contrast, programs that have high average treatment effects are not ones we would typically associate with high average earnings, which includes programs in the social sciences and humanities. We find that this matters for counterfactual college admission and capacity policies. While these results may be specific to our context, it provides some caution against policies that uniformly push students into high earnings majors – one-size-fits-all recommendations may be costly. The results also show that there may need to be a significant amount of pre-college human capital investment needed for students to take advantage of high-return majors.

While our framework is well-suited to analyzing rich sources of sorting into college majors on the basis of heterogeneous returns and preferences, there are some policies we cannot analyze that are important for the design of post-secondary education. For example, our paper does not deal with college financing – tuition costs are absorbed into main effects in the utility equation. We also do not account for the supply side

of higher education. As a result, we cannot analyze policies that change the cost of college or account for college-side investments in program quality, college-side changes in admissions criteria (Kapor [2024]), or changes in tuition costs. Our model is also static and it may be that interactions between major choice and subsequent decisions to graduate (Heckman et al. [2018], Arcidiacono et al. [2024]) or interactions between college major and later careers (Eckardt [2023], Schulze [2024]) are important.

Bibliography

Atila Abdulkadiroglu, Parag A. Pathak, Jonathan Schellenberg, and Christopher R. Walters. Do parents value school effectiveness? *American Economic Review*, 110(5):1502–1539, 2020.

Atila Abdulkadiroglu, Joshua Angrist, Parag A. Pathak, Jonathan Schellenberg, and Christopher R. Walters. Breaking ties: Regression discontinuity design meets market design. *Econometrica*, 90(1):117–151, 2022.

Joseph G. Altonji, Erica Blom, and Costas Meghir. Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers. *Annual Review of Economics*, 4(1):185–223, 07 2012. URL `https://ideas.repec.org/a/anr/reveco/v4y2012p185-223.html`.

Peter Arcidiacono. Ability sorting and the eturns to college major. *Journal of Econometrics*, 121:343–375, 2004.

Peter Arcidiacono and John Bailey Jones. Finite mixture distributions, sequential likelihood, and the em algorithm. *Econometrica*, 71(3):933–946, 2003.

Peter Arcidiacono, Esteban Aucejo, Arnaud Maurel, and Tyler Ransom. College attrition and the dynamics of information revelation. *Working Paper, Duke University*, 2024.

Nano Barahona, Caué Dobbin, and Sebaastián Ortero. Affirmative action in centralized college systems. *Working Paper, Columbia University*, 2023.

Zachary Bleemer. Top percent policies and the return to postecondary selectivity. *Working Paper, Princeton University*, 2024.

David Card. The causal effect of education on earnings. *Handbook of Labor Economics*, 3: 1801–1863, 1999.

Stacy Berg Dale and Alan B. Krueger. Estimating the payoff to attending a more selective college: An application of selection on observables and unobservables. *Quarterly Journal of Economics*, 117(4):1491–1527, 2002.

Stacy Berg Dale and Alan B. Krueger. Estimating the effects of college characteristics over the career using administrative earnings data. *Journal of Human Resources*, 49(2): 323–358, 2014.

A.P. Dempster, N. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society*, 39:1–38, 1977.

Dita Eckardt. Training specificity and occupational mobility: Evidence from german apprenticeships. *Working Paper, University of Warwick*, 2023.

Gabrielle Fack, Julien Grenet, and Yinghua He. Beyond truth-telling: Preference estimation with centralized school choice and college admissions. *American Economic Review*,

109(4):1486–1529, 2019.

Zvi Griliches. Estimating the returns to schooling: Some econometric problems. *Ecnometrica*, 45(1):1–22, 1977.

Justine Hastings, Christopher Neilson, and Seth Zimmerman. Are some degrees worth more than others? evidence from college admission cutoffs in chile. *NBER Working Paper 19241*, 2013.

J Heckman and B. Singer. A method for minimzing the impact of distributiona assumptions in econometric models for duration data. *Econometrica*, 52(2):271–320, 1984.

J Heckman, Lance J. Lochner, and Petra E. Todd. Earnings functions, rates of return and treatment effects: The mincer equation and beyond. *Handbook of the Economics of Education*, 1, 2006.

James J. Heckman and Sergio Urzua. Comparing iv with structural models: What simple iv can and cannot identify. *Journal of Ecnometrics*, 156, 2010.

James J. Heckman, John Eric Humphries, and Gregory Veramendi. Returns to education: the causal effects of education on earnings, health, and smoking. *Journal of Political Economy*, 126(S1):S197–S246, 2018.

Mark Hoekstra. The effect of attending the flagship state university on earnings: A discontinuity-based approach. *Review of Economics and Statistics*, 91(4):717–724, 2009.

Chang-Tai Hsieh and Miguel Urquiola. The effects of generalized school choice on achievement and stratification: Evidence from chile's voucher program. *Journal of Public Economics*, 90, 2006.

Adam Kapor. Transparency and percent plans. *Working Paper, Princeton University*, 2024.

Adam Kapor, Mohit Karnani, and Christopher Neilson. Aftermarket frictions and the cost of off-platform options in centralized assignment mechanisms. *Journal of Political Economy*, 2024.

Lars Kirkebøen, Edwin Leuven, and Mogstad Magne. Field of study, earnings, and self selection. *The Quarterly Journal of Economics*, 131(3):1057–1112, 2016.

Michael Lovenheim and Jonathan Smith. Returns to different postsecondary investments: Institution type, academic programs, and credentials. *Handbook of the Economics of Education*, 6:187–318, 2023.

Jack Mountjoy. Community colleges and upward mobility. *American Economic Review*, 112(8):2580–2630, 2022.

Jack Mountjoy. Marginal returns to public universities. *NBER Working Paper no. 32296*, 2024.

Jack Mountjoy and Brent R. Hickman. The returns to colleges(s): Relative value-added and match effects in higher education. *NBER Working Paper no. 29276*, 2021.

Karl Schulze. Re-skilling, up-skilling, and the role of education in the adjustment to economic shocks. *Working Paper, Princeton University*, 2024.

Christopher R. Walters. The demand for effective charter schools. *Journal of Political Economy*, 126(6):2179–2223, 2018.

Christopher R. Walters and Patrick Kline. Evaluating public programs with close substitutes: The case of head start. *The Quarterly Journal of Economics*, 131:1795–1848, 2016.

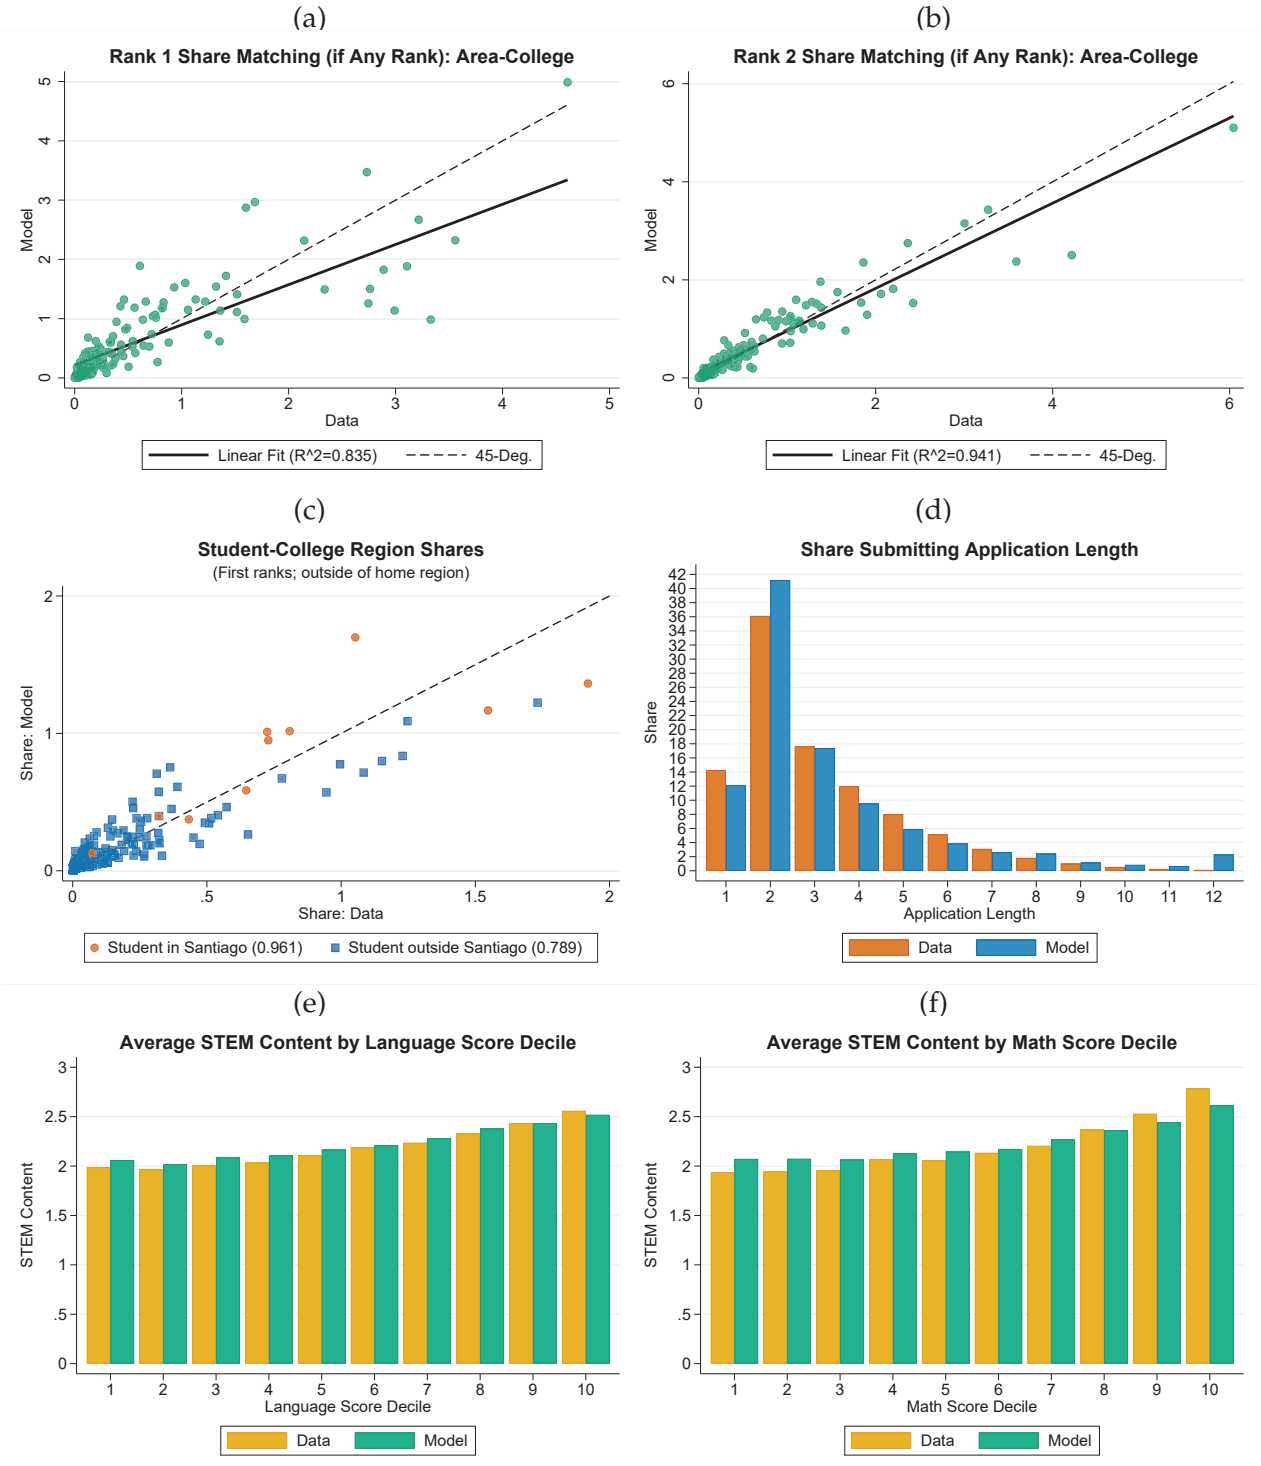
Robert J. Willis and Sherwin Rosen. Education and self-selection. *Journal of Political*

Economy, 87:S7–S36, 1979.

Seth Zimmerman. The returns to college admission for academically marginal students. *Journal of Labor Economics*, 32(4):711–754, 2014.

Appendix A: Additional Figures and Tables

Figure A1: Detailed Fit of the Choice Model



Notes: Application portfolios are restricted to credible applications on which the model is fit, as is described in the text. Panel (a) and (b) show the shares of each institution-by-broad area of study that appear in application rank 1 and 2, respectively. Panel (c) shows origin-by-destination region shares for application rank 1, separately by if the student's origin region is Santiago or not. Panel (d) shows the number of credible applications submitted. Panels (e) and (f) plot average STEM content (semesters of STEM) by deciles of the language and mathematics score on the PAA.

Table A1: Parameters: Prior Distribution and Ex-Post Characteristics of Types

	Type 1	Type 2	Type 3	Type 4
A. Type Distribution				
Overall Share	0.243	0.266	0.194	0.297
Posterior $> 0.9 \mid$ Type	0.430	0.414	0.353	0.328
B. Ex-Post Demographics				
E[Math Score $\mid$ Type]	2.647	2.473	3.362	3.060
E[Lang Score $\mid$ Type]	2.524	2.586	3.465	2.962
P[Male $\mid$ Type]	0.357	0.942	0.670	0.425
P[Private $\mid$ Type]	0.103	0.215	0.392	0.265
C. Parameters of Prior Distribution				
Constant		0.612	-0.110	-0.404
		(1.120)	(1.127)	(1.229)
Low Math Score		-0.177	-0.832	-0.569
		(0.855)	(1.537)	(0.794)
Low Read Score		0.114	-1.608	-0.446
		(0.758)	(1.615)	(0.733)
High Math Score		0.163	1.217	0.588
		(0.899)	(1.055)	(0.922)
High Read Score		-0.061	1.520	0.666
		(0.732)	(0.952)	(0.727)
Female		-3.486	-1.619	-0.483
		(0.679)	(0.938)	(0.817)
Private School		1.011	1.506	0.893
		(0.987)	(1.177)	(0.855)
Metro Region		0.521	-0.256	1.034
		(0.971)	(0.859)	(1.009)

Notes: Panel A shows the distribution of each type in the population. Panel B shows ex-post means of observable characteristics within each type using the formula: $\mathbb{E}[w_i|Typeq] = \frac{\sum_{i=1}^{N} w_i P_{q|i}}{\sum_{i=1}^{N} P_{q|i}}$, where w_i is the covariate of interest and $P_{q|i}$ is the posterior probability of i being type q at the estimated parameters. Panel C shows parameters that determine the prior distribution of unobserved types a s determined in Equation 2.3.4. Coefficients on Type 1 are normalized to zero to fix location. Standard errors are in parentheses.

Table A2: Parameters: Field of Study Reduced-Form Utility and Earnings Effects

	Preferences		Earnings	
	Coef.	*Std. Error*	*Coef.*	*Std. Error*
Accounting Auditor	3.193	0.969	1.915	3.427
Agroindustry Engineering	1.345	1.432	2.795	3.904
Agronomy	2.557	1.390	2.693	3.808
Anthropology and Archaeology	3.172	1.626	-1.464	3.805
Aquaculture and Fisheries Technician	2.011	1.581	0.688	3.858
Architecture	2.163	1.981	-1.458	4.694
Biochemistry	3.066	1.031	-1.086	2.855
Biology	4.002	1.021	-0.881	3.045
Chemistry and Pharmacy	3.828	1.404	0.829	3.688
Chemistry, Bachelor of Chemistry	2.402	0.922	-0.942	2.996
Civil Construction	1.669	0.996	-1.803	3.081
Civil Engineering	2.089	1.007	-3.000	3.005
Civil Engineering-Chemistry	0.904	0.996	-1.629	2.865
Civil Engineering-Computation and Information Systems	2.195	0.987	-1.662	2.991
Civil Engineering-Electrical	1.090	1.003	-2.163	3.043
Civil Engineering-Industrial	1.040	0.989	-1.972	3.097
Civil Engineering-Mechanical	0.998	1.025	-2.631	2.995
Civil Engineering-Metallurgical	-0.130	1.345	-6.663	4.783
Civil Engineering-Mining	0.483	0.990	-1.869	3.056
Civil Engineering-Public Works	0.973	1.076	-4.208	3.513
Commercial Engineering	3.228	1.032	1.735	3.559
Computer Technician	2.891	1.077	-1.656	3.206
Dentistry	3.917	1.329	-0.618	3.276
Design	2.567	1.910	-1.201	4.076
Electronic Civil Engineering	0.931	1.045	-2.338	3.104
Engineering Degree in Cartography	1.163	0.959	-3.351	2.341
Engineering Degree in Computing and Information Systems	2.229	0.947	-1.888	2.721
Engineering Degree in Construction	1.141	1.549	-7.527	6.336
Engineering Degree in Electricity	1.280	0.967	-2.364	2.781
Engineering Degree in Electronics	1.102	1.054	-4.471	2.975
Engineering Degree in Mathematics and Statistics	0.895	1.021	-2.297	3.236
Engineering Degree in Sound	1.379	1.012	-0.913	2.689
Engineering Degree in Wood Science	1.346	1.535	1.623	3.608
Engineering in Aquaculture and Fisheries	1.378	1.506	1.163	3.516
Engineering in Chemistry	0.480	1.159	-4.735	3.059
Engineering in Mining and Metallurgy	0.655	1.222	-4.644	3.527
Engineering, General	2.955	0.952	-0.696	3.118
Fine Arts	1.113	1.623	1.650	4.296
Food Engineering	0.571	1.002	-1.927	2.643
Forestry Engineering	2.503	1.470	1.894	3.531
Geography	2.086	1.355	-3.051	4.349
Geology	2.305	1.206	-1.327	3.388
Graphic Design	2.596	1.818	0.491	3.692
History	4.183	1.283	-0.328	4.607
Industrial Design	2.490	1.166	-2.120	2.329
Industrial Engineering	2.027	1.179	-2.056	2.464
Industrial Maintenance Technician	1.556	1.057	-3.058	3.203
Interpretation of Music, Singing or Dance	-1.237	1.567	0.467	3.947
Journalism	3.994	0.719	0.425	1.091
Kinesiology	4.498	1.432	0.487	2.853
Law	8.050	1.601	2.397	6.896
Library Science	2.212	1.264	-2.143	2.331
Literature	1.701	0.902	1.122	1.840
Mathematics and Statistics	2.960	0.924	-0.904	3.133
Mechanical Engineering	1.523	0.957	-2.185	2.433
Medical Technology	4.034	1.107	0.964	3.163
Medicine	3.114	1.550	-2.936	5.139
Nursing	3.752	1.273	0.962	3.234
Nutrition and Diet	3.662	1.434	0.903	3.217
Obstetrics and Pediatrics	3.626	1.323	1.186	3.477
Occupational Therapy	2.997	1.599	1.211	3.155
Other Civil Engineering	0.376	1.021	-1.090	2.980
Other Education Professionals	3.374	0.681	-0.062	2.994
Other Engineering	0.933	0.931	-2.053	2.388
Other Technology Professionals	2.573	0.977	-2.875	2.612
Other Technology Technicians	1.901	1.048	-1.990	3.238
Pedagogy in Arts and Music	1.555	0.569	0.897	2.572
Pedagogy in Basic Education	2.846	0.588	0.430	2.589
Pedagogy in Elementary Education	2.577	0.671	0.224	2.489
Pedagogy in History, Geography and Social Sciences	4.151	0.634	1.349	2.804
Pedagogy in Languages	1.535	0.657	1.209	2.771
Pedagogy in Mathematics and Computation	3.025	0.690	-0.428	3.355
Pedagogy in Philosophy and Religion	2.148	0.674	1.600	3.027
Pedagogy in Physical Education	2.018	0.594	0.256	2.838
Pedagogy in Sciences	3.006	0.663	0.093	3.204
Pedagogy in Spanish	2.705	0.607	1.412	2.865
Pedagogy in Special Education	2.597	0.628	-0.021	2.611
Philosophy	2.616	0.982	1.888	2.125
Psychology	2.802	1.439	-1.442	3.464
Public Administration and Political Science	3.073	1.531	-1.842	3.957
Social Work	3.541	1.443	-2.315	3.534
Sociology	3.304	1.496	-0.472	3.901
Speech Therapy	3.990	1.525	1.347	3.256
Technical Degree in Automization and Industrial Control	1.175	1.087	-2.844	3.457
Technical Degree in Automotive Mechanics	1.764	1.173	-2.978	4.189
Technical Degree in Chemistry (Analysis and Industrial)	3.330	1.166	-0.977	3.441
Technical Degree in Construction and Public Works	1.547	1.093	-3.578	3.196
Technical Degree in Electricity and Industrial Electricity	1.815	1.099	-4.136	3.228
Technical Degree in Electronics and Industrial Electronics	1.967	1.094	-2.513	3.745
Technical Degree in Food	2.022	1.106	-2.281	3.044
Technical Degree in Human Resources and Personnel Management	3.388	1.096	1.416	3.402
Technical Degree in Industrial Maintenance	3.546	1.244	-0.305	4.467
Technical Degree in Industrial Processes	4.301	1.148	-2.894	3.430
Technical Degree in Instrumentation, Automation, and Industrial Control	3.140	0.871	2.712	3.289
Technical Degree in Logistics	3.161	0.928	0.249	2.959
Technical Degree in Nursing	1.449	1.140	-3.189	3.895
Technical Degree in Photography	1.465	1.177	-2.337	3.775
Technical Degree in Publicity	1.676	0.903	-1.175	3.378
Technician in Forestry or Wood Industry	2.704	1.554	-0.393	3.777
Technician in Technical and Industrial Drawing	2.079	1.071	-2.130	3.123
Telecommunications Technician	2.375	1.229	-1.929	3.913
Tourism and Hospitality Administration	3.213	1.340	0.925	2.741
Translation and Interpretation	1.649	0.885	1.419	2.060
Veterinary Medicine	2.423	1.487	2.999	4.431

Notes: Parameters relating to college majors that govern the reduced form preferences in Equation 2.3.5 and earnings in Equation 2.3.2. The excluded category is "Humanities" whose effect is absorbed by the on-platform earnings parameters reported in Table A4.

Table A3: Parameters: Institution Reduced-Form Utility and Earnings Effects

	Preferences		Earnings	
	Coef.	*Std. Error*	*Coef.*	*Std. Error*
Pontificia Universidad Catolica De Chile	-0.503	0.222	0.023	0.056
Universidad De Concepcion	-0.780	0.271	-0.276	0.102
Pontificia Universidad Catolica De Valparaiso	-0.620	0.280	-0.193	0.104
Universidad Tecnica Federico Santa Maria	-1.642	0.271	-0.133	0.193
Universidad De Santiago De Chile	-1.019	0.178	0.058	0.146
Universidad Austral De Chile	-0.828	0.477	-0.447	0.224
Universidad Catolica Del Norte	-0.507	0.546	-0.483	0.322
Universidad De Valparaiso	-0.937	0.265	-0.263	0.153
Universidad Metropolitana De Ciencias De La Educacion	-0.954	0.307	-0.141	0.180
Universidad Tecnologica Metropolitana	-1.907	0.209	-0.023	0.276
Universidad De Tarapaca	-0.200	0.520	-0.658	0.340
Universidad Arturo Prat	-0.794	0.562	-0.758	0.395
Universidad De Antofagasta	-0.686	0.555	-0.553	0.267
Universidad De La Serena	-0.621	0.528	-0.392	0.208
Universidad De Playa Ancha	-0.878	0.395	-0.262	0.171
Universidad De Atacama	-1.301	0.561	-0.314	0.419
Instituto Profesional De Chillan	-1.251	0.291	-0.267	0.161
Universidad Del Bio-bio	-1.400	0.296	-0.327	0.208
Universidad De La Frontera	-0.750	0.496	-0.438	0.227
Instituto Profesional De Valdivia	-0.984	0.598	-0.667	0.349
Universidad De Los Lagos	-1.125	0.509	-0.506	0.291
Universidad De Magallanes	-1.588	0.570	-0.566	0.309
Universidad De Talca	-0.619	0.465	-0.353	0.206

Notes: Parameters relating to institutions that govern the reduced form preferences in Equation 2.3.5 and earnings in Equation 2.3.2. Institution effects are measured relative to "Universidad De Chile", whose effect is absorbed by the on-platform earnings parameters reported in Table A4.

Table A4: Parameters: Heterogeneity in Course Content and Broad Area of Study Effects, Outside Option, Mean On-platform Earnings, and Other Preference Parameters

	Constant		Type 2		Type 3		Type 4		Low Math Score		Low Read Score		High Math Score		High Read Score		Female		Private School		Metro Region	
	Preferences	Earnings	Preferences	Earnings	Preferences	Earnings	Preferences	Earnings	Preferences	Earnings	Preferences	Earnings	Preferences	Earnings	Preferences	Earnings	Preferences	Earnings	Preferences	Earnings	Preferences	Earnings
Large City Region	-0.500		0.798		1.744		0.771															
	(0.594)		(0.658)		(0.790)		(1.174)															
log(Distance) - Same Region	-0.082		0.047		0.047		0.033															
	(0.098)		(0.142)		(0.195)		(0.155)															
log(Distance) - Different Region	-0.258		-0.104		0.060		-0.587															
	(0.078)		(0.125)		(0.115)		(0.167)															
Outside Option	2.265	9.239	2.052	-0.085	2.935	0.566	0.546	-1.430	-0.193	-1.338	-0.153	-0.532	0.238	-1.733	0.295	-0.874	-0.291	-0.112	0.356	-0.020	0.572	2.204
	(0.876)	(0.352)	(0.759)	(0.367)	(0.970)	(0.720)	(1.007)	(0.398)	(0.288)	(0.186)	(0.274)	(0.171)	(0.262)	(0.257)	(0.328)	(0.337)	(0.378)	(0.254)	(0.277)	(0.171)	(0.556)	(0.488)
Mean On-Platform Earnings		9.530		3.736		0.795		1.555		-0.149		-0.139		0.417		0.129		-0.149		0.117		-0.791
		(1.584)		(3.084)		(1.961)		(2.040)		(0.098)		(0.129)		(0.126)		(0.122)		(0.126)		(0.112)		(0.436)
STEM Share			0.260	-1.504	0.458	-0.845	0.240	-0.748	0.019	-0.147	0.025	0.247	-0.007	0.210	0.017	-0.024	0.007	-1.702	0.006	-0.831		
			(0.265)	(1.602)	(0.408)	(1.404)	(0.658)	(1.546)	(0.324)	(2.232)	(0.304)	(2.386)	(0.310)	(2.942)	(0.310)	(2.938)	(0.228)	(1.991)	(0.140)	(2.279)		
HUMA Share			-1.030	-3.449	0.415	-1.702	-0.174	-1.391	0.029	-0.165	0.017	0.121	4.015e-04	-0.481	0.023	-0.153	0.008	-1.010	-0.004	-1.245		
			(0.644)	(1.785)	(0.328)	(1.784)	(0.614)	(1.476)	(0.336)	(2.452)	(0.403)	(2.121)	(0.288)	(2.756)	(0.418)	(2.915)	(0.259)	(1.899)	(0.219)	(2.382)		
Total Credits			0.276	-0.062	0.808	1.541	0.214	-0.384	-0.010	-0.230	-0.006	0.257	-0.007	-0.181	-0.003	0.081	-0.002	0.302	-0.003	0.048		
			(0.225)	(0.942)	(0.329)	(1.206)	(0.297)	(0.782)	(0.195)	(0.675)	(0.173)	(0.856)	(0.154)	(0.803)	(0.168)	(0.887)	(0.147)	(0.683)	(0.108)	(0.692)		
Selectivity			0.223	0.298	0.781	0.150	0.173	0.408	-0.008	-0.292	-0.024	-0.169	-0.030	-0.758	-0.006	0.170	-0.006	0.176	-0.012	-0.309		
			(0.298)	(0.232)	(0.405)	(0.241)	(0.313)	(0.453)	(0.156)	(0.699)	(0.138)	(0.820)	(0.137)	(0.906)	(0.138)	(0.699)	(0.104)	(0.677)	(0.102)	(0.852)		
Business Administration			0.861	-3.877	1.116	-0.181	0.125	-1.599	0.022	-1.260	0.009	-1.279	0.019	0.008	0.015	-2.135	0.006	-1.017	-0.023	0.318		
			(0.783)	(3.608)	(1.009)	(2.589)	(1.209)	(2.775)	(0.397)	(3.015)	(0.327)	(3.743)	(0.306)	(2.584)	(0.309)	(3.047)	(0.204)	(1.765)	(0.160)	(3.975)		
Agriculture			2.689	-7.223	1.463	-4.329	0.637	-4.028	0.081	-0.067	-0.005	-0.351	-0.003	0.843	0.008	-2.088	0.017	-1.325	-0.002	-1.164		
			(1.337)	(4.634)	(1.814)	(4.135)	(1.440)	(3.866)	(0.495)	(3.264)	(0.466)	(3.315)	(0.427)	(3.739)	(0.464)	(4.669)	(0.337)	(2.569)	(0.267)	(4.048)		
Art and Architecture			2.920	-2.382	0.490	-1.836	0.632	-1.884	0.037	0.583	-1.107e-04	0.578	0.001	1.201	-0.027	1.216	-0.002	-1.402	0.025	0.112		
			(1.949)	(3.557)	(2.075)	(3.013)	(1.785)	(4.256)	(0.796)	(4.460)	(0.640)	(6.894)	(0.706)	(5.211)	(0.650)	(3.598)	(0.500)	(3.601)	(0.338)	(3.566)		
Natural Science			0.072	-3.786	-0.566	-0.770	-0.130	-0.445	0.045	-0.133	0.010	-0.556	0.007	0.099	0.016	-0.095	-0.001	1.605	-0.013	0.663		
			(0.935)	(3.791)	(1.348)	(3.039)	(1.102)	(2.808)	(0.443)	(2.745)	(0.363)	(3.070)	(0.353)	(3.690)	(0.339)	(2.709)	(0.270)	(2.257)	(0.235)	(3.379)		
Social Science			0.042	-2.174	0.510	-0.696	-0.098	-0.096	0.073	1.825	0.041	1.099	-0.020	1.057	0.050	0.396	0.029	-0.616	-0.013	1.183		
			(2.292)	(3.432)	(1.188)	(2.002)	(1.310)	(2.408)	(0.555)	(4.146)	(0.484)	(3.487)	(0.473)	(4.371)	(0.484)	(3.191)	(0.279)	(3.183)	(0.209)	(3.211)		
Law			1.640	-4.039	-3.194	-2.063	-2.607	-1.792	0.097	-0.350	0.093	1.125	0.092	0.336	-0.002	0.394	0.033	-1.198	0.038	2.017		
			(2.782)	(5.362)	(2.347)	(3.748)	(2.205)	(5.057)	(1.294)	(6.848)	(0.893)	(6.753)	(1.030)	(5.172)	(0.893)	(7.359)	(0.671)	(5.451)	(0.520)	(6.395)		
Education			1.210	-2.374	-0.081	1.605	0.544	-0.110	0.010	-0.893	0.028	-0.040	-0.010	0.384	0.016	-0.490	0.006	-0.923	-0.002	0.465		
			(0.933)	(2.763)	(1.090)	(1.311)	(0.744)	(1.541)	(0.309)	(2.354)	(0.249)	(2.322)	(0.258)	(2.243)	(0.248)	(1.943)	(0.136)	(1.643)	(0.125)	(2.320)		
Health			-1.448	-5.100	0.163	-1.055	-0.458	-0.936	0.040	-1.714	0.049	-0.561	0.007	-9.539e-04	0.064	-0.893	0.012	-0.064	-0.051	1.161		
			(2.087)	(3.617)	(1.477)	(2.819)	(1.234)	(2.658)	(0.679)	(3.212)	(0.588)	(2.957)	(0.517)	(2.955)	(0.550)	(2.762)	(0.392)	(2.243)	(0.315)	(3.155)		
Technology			3.116	-2.695	0.976	-0.435	0.813	-0.359	0.030	0.442	-0.001	0.196	0.011	0.658	-0.003	-0.221	-0.011	0.813	0.029	0.449		
			(0.830)	(4.203)	(0.980)	(2.965)	(0.985)	(2.826)	(0.291)	(1.861)	(0.266)	(1.793)	(0.236)	(2.004)	(0.264)	(1.591)	(0.169)	(1.325)	(0.252)	(1.693)		

Notes: Additional parameters that govern the reduced form preferences in Equation 2.3.5 and earnings in Equation 2.3.2. Note that the main effects for course content and broad area of study parameters are absorbed into the major effects reported in Table A2. Reference category for broad area of study is "Humanities'. Standard errors are reported in paretheses.

Chapter 3

The Incidence of Mass Layoff Events: Evidence from Brazil

3.1 Introduction[1]

State and local governments often spend significant amounts of money to attract new local employers. For example, *The New York Times* estimates that state and local governments spend $80 billion dollars in tax incentives and direct subsidies to specific companies as of 2012 (Story [2012]). In times of crisis, governments often use bailouts to prevent firm and plant closures, as was the case of the automobile industry in the United States during the Great Recession when the federal government lent billions of dollars in an attempt to prevent key manufacturers from shutting down. Justification for such government intervention in the economy rets on the idea that there is a "multiplier effect" of local employment – each job directly created (or destroyed) by a policy or economic shock will result in additional, indirectly created (or destroyed) jobs. Empirical evidence often documents a "scarring" effect of local economic shocks, suggesting that a failure to prevent a local economic collapse may have permanent consequences.

[1]This chapter is co-authored with *Sharada Dharmasankar*.

The propagation of local economic shocks is an equilibrium outcome that can be generated by several economic forces. These forces may in turn depend on the institutional details of a specific labor market. Understanding different mechanisms that may generate such spillover effects is key for policy and the effectiveness of the aforementioned firm-specific incentives. For example, if flagship firms raise the productivity of nearby firms through agglomeration effects or production networks, then intervention that targets specific firms may help to stave off catastrophic "domino" effects. In contrast, if declines in local consumer demand drive total effects, then other policies that subsidize consumers rather than firms may be more preferred. If workers see earnings declines due to more limited outside options, then stronger institutional protections may help to insulate them from economic shocks. Since the seminal work of Blanchard and Katz [1992], many studies have assessed the impacts of local labor market shocks. However, work that is able to cleanly separate the direct and indirect effects of economic shocks and investigate mechanisms is much more limited. Moreover, there is little work in the context of developing countries, where the types of shocks experienced and labor market regulations in place may be substantially different.

This paper studies how economic shocks propagate through a local labor market by studying the effects of mass layoff events on all workers in the labor market. We extend the empirical approach of Gathmann et al. [2020] (henceforth, GHS) to the context of Brazil and examine heterogeneity in effects that may consistent with different mechanisms. By leveraging matched employer-employee data, we are able to identify sharp layoff events that affect a specific labor market and cleanly identify which workers are directly versus indirectly exposed to a layoff event. Our main empirical approach is to identify a comparison region for each region experiencing a large mass layoff event at a flagship employer in the formal sector. By matching to comparison regions with similar industrial composition and economic characteristics as layoff regions, we create a comparison group of regions that should be subject to similar macro-shocks but do not

experience a layoff event, thereby plausibly allowing the comparison group to serve as a counterfactual for if the layoff regions never experienced the mass layoff shock.

Our main finding is that mass layoff events permanently scar local labor markets with some evidence of spillover effects. Specifically, we estimate that a mass layoff results in a 10.1 percent reduction in local employment 4 years after a layoff event with 5.2 percent due to reductions in employment for workers not at the layoff firm. Importantly, we find that this spillover effect only occurs with a delay and that total employment losses on impact are about the same size as the size of the layoff – patterns consistent with the layoff event being a true shock that is idiosyncratic to the local economy. Our estimates imply a multiplier effect of 0.5: with every job lost in the layoff event causing an additional 0.5 extra job losses. These effects are larger than GHS, who estimate a slightly smaller multiplier effect of about 0.4.[2]

Following workers across regions shows a limited role for out-migration in generating spillover effects. While we find strong out-migration effects where directly displaced workers eventually move out of event regions, we find no effect on spillover workers. This implies that our headline employment effects on spillover workers are driven by moves into non-employment or the informal sector within a given labor market. Moreover, it suggests that the limited mobility responses to negative shocks that we see in more developed contexts (e.g. Bound and Holzer [2000]) may also be present in developing economies.

We find modest but sometimes statistically insignificant effects on earnings outcomes that are stronger on spillover workers. In particular, we find a 10.8 percent effect on total earnings in a local labor market (i.e., the wage bill) with 7.8 percent being driven by spillover workers, indicating a larger multiplier of 0.7. We estimate that workers from the layoff firm suffer total earnings declines of 13 percent four years after the layoff event

[2]Using German districts, GHS estimate a 3.7 percent decline in local employment following a mass layoff with 1.4 percent due to a spillover effect. German districts are slightly larger than Brazilian microregions – GHS report that a layoff region in their context has an average pre-layoff employment of 79,987 while we find 56,692 in pre-layoff employment in our Brazilian microregions.

and a 3 percent decline in log-wages conditional on working, as compared to spillover workers who see declines of 4 percent on total earnings and 1.4 percent on log-wages. This indicates that a third to a half of the full displacement effect is driven by common declines in earnings across a local labor market.

Finally, we use the rich worker heterogeneity and panel structure of the data to investigate the incidence of the the spillover effects of layoffs on individuals. While much more mixed, these worker-level results broadly show strong evidence of declines in consumer demand, modest evidence for agglomeration economies, and limited evidence for declines in worker bargaining power. In particular, the strongest spillover effects that we estimate are on earnings in the non-tradeable sector, which suffer impacts on log-earnings that account for the entire estimated spillover effect but do not see employment declines. In addition, we do find strong declines on total earnings for spillover workers in the same 2-digit industry as the layoff firm, which is driven by participation rather than impacts on wage rates, suggesting agglomeration and network effects. In contrast, we find that non-layoff workers with similar skills to displaced workers, as defined by occupation-education cells, see trends in all outcomes that are similar to workers in control regions. This points toward the incidence of the spillover effects of mass layoffs being by workers in the non-tradeable sector or the same narrow industry as the layoff firm, but having distinct sets of skills from workers separated in a layoff event.

Related Literature: At least as far back as Jacobson et al. [1993], a rich literature studies how worker displacement and mobility shape earnings and employment outcomes. Earlier studies using non-administrative sources of data [Ruhm, 1991] find that displaced workers may not necessarily experience negative long-term employment effects, but that there are shocks to earnings. More recent works examine sources of displacement effects using matched employer-employee data [Couch and Placzek, 2010, Eliason and Storrie, 2006, Goolsbee and Syverson, 2019, Flaaen et al., 2019, Davis and von Wachter, 2011]. Included among these, Lachowska et al. [2020], Jarosch [2023], Schmieder et al.

[2023] document different sources of displacement effects, including post-displacement employer characteristics, shocks from the displacement event itself, and human capital losses. Lachowska et al. [2020] finds that the persistence of displacement effects is mostly driven by declines in wage rates with a substantial role for lost employer-specific premiums.

Almost all of these papers examine labor market trajectories for long-tenure workers who are directly impacted by displacement (e.g., those who are laid off). Yet, as displaced workers are likely to interact with non-displaced workers who may be in closely-related sectors, non-displaced workers may also compete for jobs with displaced ones. To our knowledge, only two other papers, Gathmann et al. [2020] and Moghadam [2023] quantify the spillover effects of displacement events. Moghadam [2023] uses data from California to see how layoff events affect geographically close establishments, finding employment drops in nearby establishments. Gathmann et al. [2020] use a matching design where regions hit by mass layoff events are compared to observably similar regions in Germany that did not experience a layoff event, finding that firms in the same sector and region as the a layoff plant experienced persistent negative spillover effects. They document heterogeneity in the effects by worker age, finding that younger workers are able to move from labor markets shocked by displacement effects, thus shielding them from negative employment or earnings effects. We contribute to the literature by extending this analysis: aside from revisiting the research question in a new setting (Brazil), we assess heterogeneous impacts of mass layoffs across different worker types and decompose sources of spillovers across the following three areas: worsening conditions in local labor market conditions, competition between displaced workers and non-laid-off workers, and changes in worker outside options for employment.

As the layoff events we consider occur at the granular establishment level, our paper also connects to literature studying the effects of policies and shocks on local labor markets. Common examples include spatially-targeted policies that encourage economic

[Busso et al., 2013, Arefeva et al., 2024] or physical growth [Fu and Gregory, 2019], local minimum wages [Dube and Lindner, 2021, Jardim et al., 2022], or shocks like plant openings/closures [Greenstone et al., 2010, Moretti, 2004]. Such events raise a common concern of spillover or indirect effects that need to be measured: as local (labor) markets share boundaries that may be closely linked, shocks directly impacting one sector or area may impact workers in another. Findings on the magnitude and existence of spillover effects are mixed; some studies acknowledge huge spatial or local spillovers, while others note that they are negligible. Gathmann et al. [2020] note two key sources of spillovers in local labor markets: an agglomeration effect caused by density of similar sectors in a particular geography [Greenstone et al., 2010, Kline and Moretti, 2014], and a multiplier effect that captures worsening economic conditions and a decline in local area income; this typically manifests through changes in demand for products in non-layoff sectors [Moretti, 2010, Moretti and Thulin, 2013, Faggio and Overman, 2014]. Another potential spillover source is the change in worker outside options [Beaudry et al., 2012, Caldwell and Harmon, 2019]. We incorporate as a channel in our decomposition of spillover effects by examining how workers in similar occupations across industries are affected by layoffs.

Finally, we contribute to a rich literature studying Brazilian labor markets and how workers and firms adapt and respond to shocks. These include policies in trade and their heterogeneous impacts [Sharma, 2023, Dix-Carneiro and Kovak, 2017], changes in minimum wages [Engbom and Moser, 2022, Derenoncourt et al., 2021], and hiring quotas [Szerman, 2023]. Many of these papers document the size and existence of spillover effects, but focus less on *why* effects are the way they are. Crucially, the prior literature highlights a key feature of the Brazilian labor market that makes it different from North American or European markets: there is a substantial informal sector that up to a third of workers participate in, and also who interact with formal sector workers [Ulyssea, 2018, Meghir et al., 2015]. Thus, informality is bundled into the potential outside options

individuals face when competing against a glut of laid-off workers in the same labor market. Ultimately, our contribution is the study of several local shocks in Brazil (mass layoffs), and specifically decomposing sources behind their indirect effects on non-laid off workers into several channels.

Organization: The structure of the paper is as follows. Section 3.2 discusses our source of matched employer-employee data and the context of the Brazilian labor market; Section 3.3 discusses how we identify mass layoff events and provides some descriptive evidence; Section 3.4 exhibits our approach to research design; Section 3.5 and 3.6 shows local labor market and worker-level results. Section 3.7 concludes.

3.2 Data and Context

3.2.1 RAIS Data

Our main data source is administrative data from *Relação Anual de Informações Sociais* (RAIS) for the years 1987 to 2017. The RAIS data contains the universe of employer-employee relationships in the formal sector between 1987 and 2017, reporting the firm and establishment of employment for each worker in each year. This mandatory survey is conducted by the Brazilian Ministry of Labor and Employment to administer social programs. Fines are levied for reporting inaccurate information, resulting in high-quality data. We focus on the years from 1994 onwards, when we see weekly hours and consistent industry codes. We combine aggregate establishments of the same firm within the same municipality to the firm-by-municipality level. We do this as (i) some establishment show a very small number of employees, (ii) to ensure that our later layoff definitions are not driven by moves across establishments within the same firm and municipality, and (iii) to not miss mass layoffs that may occur at multiple sites within the same town.

The RAIS data report average monthly earnings for formal workers in Brazil.[4] This earnings measure is gross and includes regular salary payments in addition to tips and bonuses.[5] Firms also report the start and end months of any employment relationship. We use this to compute total annual earnings and hourly wages using weekly contracted hours.[6]

We follow Szerman [2023] in classifying establishments by industry.[7] We often split out firms into the non-tradeable sector, which makes goods that are produced and consumed locally, versus the tradeable sector, which makes goods in industries that are subject to more external shocks. We define the tradeable sector as those firms in the Agricultural, Fishing, Extractive, or Manufacturing industries.

3.2.2 The Brazilian Economy

We describe a few economic and social trends in Brazil in this section. Brazil underwent several political and economic changes over the span of our analysis time horizon. This includes the resumption of democratic elections, trade liberalization[8], a financial crisis, and the tail end of a long stretch of hyperinflation [Alvarez et al., 2018]. Most relevant for our context are the changes that could have affected worker employment opportunities and social policy, which shape worker outside options; several relevant labor market

[4]The RAIS data report this measure as multiples of the monthly minimum wage. We convert this to total earnings in terms of constant December 2018 Brazilian reals (BRL) using data on nominal minimum wage levels and price indices from Engbom and Moser [2022], sourced from the Brazilian Institute of Applied Economic Research (Instituto de Pesquisa Economica Aplicada [2022]).

[5]The measure does exclude capital income and in-kind transfers from the employer, as well as earnings from informal employment.

[6]Some employment relationships report zero or very low monthly earnings.

[7]Using the first two digits of the National Classification of Economic Activities (CNAE) codes provided in the RAIS data, we categorize employers into one of the following sectors: Agriculture, Cattle, and Forestry; Fishing; Extractive Industries; Transformation Industries; Production and Distribution of Electricity, Gas, and Water; Construction; Wholesale and Retail Trade and Repair of Motor Vehicles and Household Products; Accommodation and Food Services; Transportation, Storage, and Communication; Financial Activities; Real Estate Activities; Public Administration, Defense, and Social Security; Education; Health and Social Services; Other Services Activities; Domestic Services; International Organizations and Other Extraterritorial Institutions.

[8]Beginning in the late 1980s and gaining momentum in the 1990s, initially high import tariffs were selectively eliminated, and other non-tariff trade barriers were also reduced [Alvarez et al., 2018].

reforms have been implemented during this time period. These labor market regulations are detailed in Section 3.2.3.

At the turn of the millennium, Brazil experiences the expansion of education and income assistance social programs, endeavors that could affect workforce composition and participation. One key initiative, *Bolsa Família*, is a public conditional cash transfer program given to poor families in Brazil if they ensured that their children were sent to school or vaccinated [Barros et al., 2010]. The program covered almost 25% of Brazil's total population in 2006 [Barros et al., 2010, Alvarez et al., 2018]. The early 2000s also saw the expansion of access to education in Brazil; grades completed increased in stark comparison to previous decades and education inequality declined [Barros et al., 2010]. Education spending also increase at the same time. This setting made it so a higher share of the Brazilian workforce had at least a high school degree [Alvarez et al., 2018].

3.2.3 Employment Regulations in Brazil

Brazil had a highly regulated labor market since at least the 1960s, though many of these regulations began loosening in the late 1980s to early 1990s. One key legislative example is the Wage Adjustment Act: in addition to determining minimum wages,[9] the government negotiated mandatory annual wage increases for workers. In 1995 (shortly after the Real was made the currency, also marking the end of a long period of hyperinflation), the Wage Adjustment Act was abandoned. As a consequence, the upward adjustment of wages was negotiated between employers and workers (rather than the government or labor unions). In 1997, the Brazilian government introduced temporary employment contracts, which lowered payroll taxes paid by employers and reduced dismissal costs [Amadeo et al., 2000, Alvarez et al., 2018]. In our context, these law changes justify the possibility of mass layoffs in Brazil, as well as motives laid off workers have to return to

[9] Alvarez et al. [2018] note that the Worker's Party implemented a series of accelerated minimum wage hikes beginning in 2003, with annual increases ranging from 10% to 20%.

work.

As noted in Britto et al. [2022], workers in Brazil are employed "at will": firms are free to dismiss workers without just cause, though workers are also allowed to leave voluntarily. However, there is a cost to dismissing workers without cause levied on firms: workers who are fired are entitled to a savings account which their employer pays into each month, and on top of that, are given a severance payment. The sum of the payments made when workers are laid off without cause amounts to 1.34 times the monthly wages of the worker per year of tenure.[10]

In addition to the severance pay given to workers dismissed without cause, displaced workers are also given unemployment insurance (this program has been in place since 1986). Unemployment insurance lasts between 3 to 5 months (the extension to 5 months was implemented in 1996). Eligible workers are formal sector employees dismissed without cause, and they receive a 100% replacement rate at the most generous end of the policy [Britto et al., 2022, Amadeo et al., 2000]. Upon benefit expiration, the last remaining option of income support from the government is the *Bolsa Família* cash transfer, which we describe above in Section 3.2.2.

3.3 Mass Layoffs in Brazil

3.3.1 Definition of Flagship Mass Layoff Events

We focus on flagship mass layoff events using a definition that closely follows Gathmann et al. [2020]. We depart from them by making additional restrictions to address pre-layoff increases in employment, which we think are necessary in order to account for the fact that many of the layoffs we identify are related to agricultural or resource extraction

[10]The Brazilian workforce is characterized by high levels of labor informality (between 40% an 60% by some measures). Employers of informal workers are not obliged to pay taxes or comply with the same labor market regulations that cover formal sector workers [Derenoncourt et al., 2021]. We note that our data does not cover informal workers, so we can't speak to how mass layoffs directly impact them or transitions between the formal and informal sectors.

industries. In contrast to the literature that focuses on displaced workers alone (e.g. Lachowska et al. [2020]), we focus on large employment losses at flagship employers in order to define clear treatment and control groups.[11] We find that our results are robust to including smaller mass layoff events or using alternate layoff definitions.

We define a flagship mass layoff event as an establishment-level reduction in employment in the tradeable sector of at least 500 workers that is sustained for two years at establishments that have at least 500 employees in the four years leading up the employment loss. To allow for 4 years of leads and lags around the event, we focus on events occurring between the years 1998 to 2013. To avoid the possibility that employment losses come off the heels of a large employment expansion, we exclude layoff events where the year prior to the mass layoff sees an increase in employment of more than half of the target year's employment losses.[12] As is standard in the mass layoff literature, we exclude events that we identify as mergers, spin-offs, or changes in firm identifiers.[13] We also make several restrictions at the region-level to ensure that the labor market is in a steady state and that layoffs are idiosyncratic, which follows from Gathmann et al. [2020].[14] Ultimately, we are left with 62 flagship mass layoff events with an average pre-layoff employment of 3,604 and average employment loss of 49.8 percent or 1,802 workers.[15]

[11] For example, using the RAIS data, Britto et al. [2022] defines a mass layoff event as an employment reduction of 33% at firms with 15 or more employees.

[12] In the event of multiple such events at the same firm, we focus on the first such event that our data documents.

[13] We define a candidate layoff event as a spinoff if, among those making a transition to a new employer following a layoff event: (i) more than 30 percent go the same firm or (ii) more than 70 percent go to the same 3 firms. Note that this restriction is made at the firm/enterprise-level, rather than the establishment-level to account for workers being distributed across a multi-unit firm following a layoff.

[14] Specifically, we exclude layoff events that are preceded by any smaller layoff that is more than half of the reference layoff event layoff size in the 2 years prior or larger than the layoff event years 3 and 4 prior. We also exclude layoff events that are immediately followed by a single firm expansion that is larger then the initial layoff amount. This final restriction should capture spinoff and merger events that are missed in the earlier steps.

[15] We combine mass layoff occurring in the same region but at different establishments. In this case, we assign the layoff event to the firm that sheds that largest number of workers.

3.3.2 Characteristics of Mass Layoff Events

We now outline the basic characteristics of the layoff events in order to validate our layoff definition and get a sense of how mass layoffs differ in Brazil relative to other contexts.

Figure 3.1 plots industry composition of mass layoff events across years. All layoffs occur in the manufacturing or resource extraction industries (see Table A1 for a granular view of the 5-digit industries that drive flagship mass layoff events).[16] Most of the mass layoff events we identify occur after the mid-2000s: 31 occur in 2008 or after. We see that many of the layoff events are related either to agriculture or the intermediate manufacture of primary materials and agricultural goods. At least 20 of these layoff events occur in the cultivation of sugarcane or at sugar mills.[17] However, there are also additional layoffs occurring in heavy industry, such as steel production, aircraft assembly, electronics, and other household goods. Below, we will parse our results into manufacturing and non-manufacturing firms.

Figure 3.2 depicts how mass layoffs are spatially distributed in Brazil across its 27 federal states, separating events by broad industry category. Panel 3.2a plots events for manufacturing establishments. Flagship mass layoffs in manufacturing are relatively less spatially concentrated the non-manufacturing layoffs, which, at their most severe, occur in contiguous states. The areas with the most manufacturing mass layoffs (2+ layoffs) tend to be more populated (the upper half of Brazilian states in terms of size) and contribute to substantive shares of Brazil's GDP (about 2% or higher). Manufacturing mass layoffs also tend to occur in areas where the economy is dominated by the industrial sector or the agricultural sector. The areas where non-manufacturing layoffs occur appear similar to those where the manufacturing layoffs are concentrated, with the exception

[16] We note that Brazil experienced substantial trade liberalization in the 1980s and 1990s with the reduction or elimination of tariffs or import quotas, which may play some role in the layoffs experienced in the manufacturing sector.

[17] The distribution of layoffs in sugarcane cultivation is relatively even across years, ranging from 1998 (the earliest) to 2013 (the latest). In 2013 alone, there are 3 such mass layoff events. 6 occur in 2008 or later and the remaining 8 occurring before 2008.

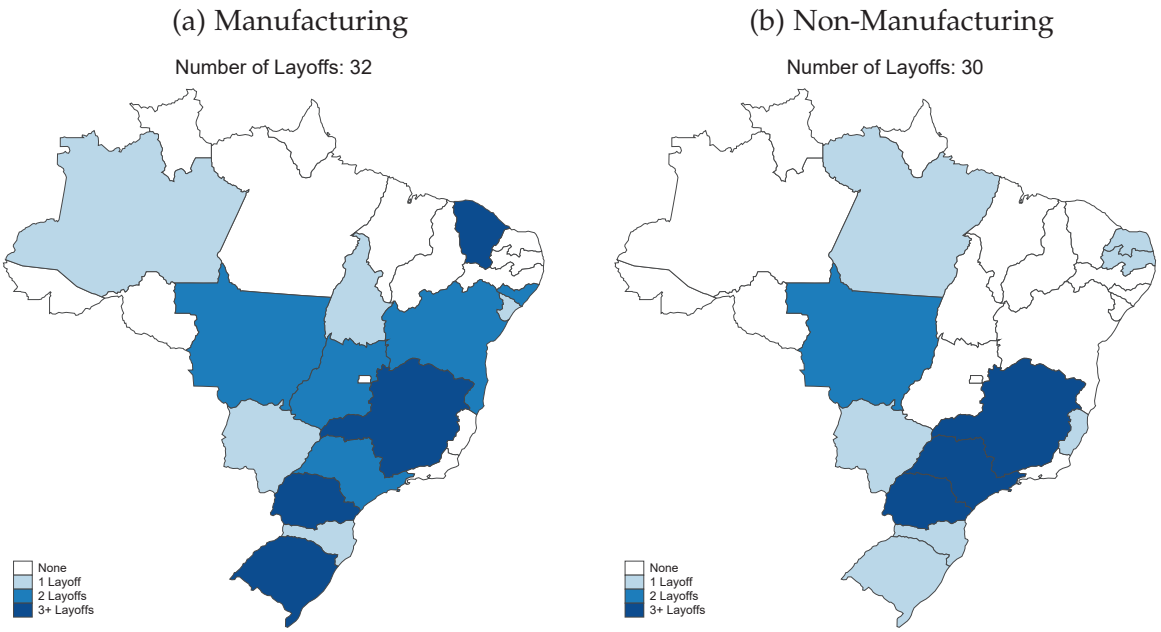
Figure 3.1: Distribution of Baseline Flagship Mass Layoff Events by Year and Broad Industry



Notes: This figure shows the count of the mass layoff events occurring in each year using the definition outlined in the main text.

that they tend to have large service sectors or large industrial sectors. There is some spatial correlation in where the non-manufacturing layoffs occur (see Panel 3.2b). The states of Minas Gerais and Paraná both have 3+ mass layoffs in each, manufacturing and non-manufacturing; São Paulo and Mato Grosso rank highly in both kinds of layoffs as well.

Figure 3.2: Locations of Mass Layoff Events by Broad Industry

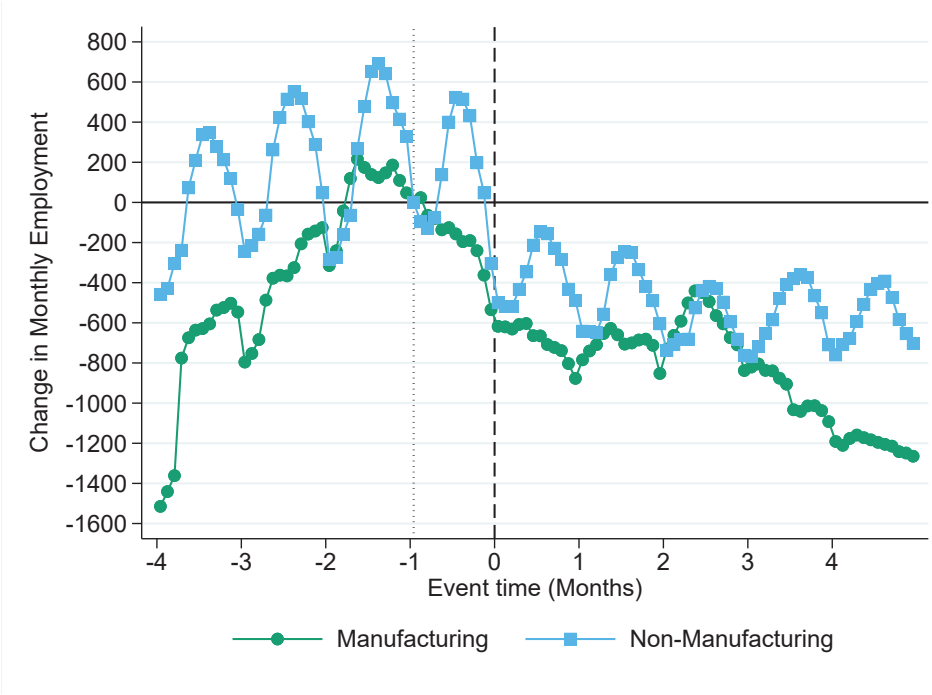


Notes: This figure shows the spatial distribution of flagship layoff events using the baseline definition, aggregating from microregions to the the level of federal states and breaking out by time period. See Figure A1 shows the spatial distribution of layoff events by year.

We now focus on establishment-level trends in employment surrounding mass layoff events. Figure 3.3 shows changes in monthly employment levels relative to December of the year one year prior to the layoff event. Both, manufacturing and non-manufacturing firms grow in the years leading up to the layoff event. In our research design, we account for this by matching layoff regions to control regions on total employment in years prior to the layoff (see Section 3.4.1 for details). This growth is more obvious for manufacturing firms, as the cyclical nature of monthly employment for non-manufacturing firms obscures the growth somewhat. Upon onset of the layoff, non-manufacturing firms trend downward in employment in the three months immediately after before picking up due

to seasonality. On average, though, there are between 400 and 600 fewer employees after the mass layoff event relative to December in the year before layoff. The patterns are starker for manufacturing firms. Immediately after the layoff, there is a steady decline in employment. After the layoff year start, manufacturing employment is permanently depressed with an average employment loss of 600 workers at the end of the layoff year that tapers to 1200 workers four years after the layoff event.

Figure 3.3: Monthly Employment at the Layoff Firm



Notes: Change in the number of workers employed at the flagship layoff firm in each month relative to December of the year one year prior to the layoff, averaged over the 62 layoff firms.

Taken together, flagship mass layoffs in Brazil tend to occur in a specific subset of industries closely tied to agriculture.[18] We document manufacturing layoffs in 14 of the federal states and in general, across wider geographies. Non-manufacturing layoffs

[18]This marks a key contribution and differentiating factor of our work. Papers studying displaced workers or mass layoffs in other more developed settings document that such displacement events are concentrated in different industries. For example, Gathmann et al. [2020] documents substantial heterogeneity in layoff industries but finds that they are relatively more concentrated in capital goods manufacturing, and Jacobson et al. [1993] notes that displaced workers are distributed across durable and non-durable goods manufacturing, with the former dominating.

occur in 11 states which are relatively contiguous.

The patterns documented in this section also help validate our definition of mass layoffs. Though restrictive, our categorization of flagship mass layoffs does indeed capture a sharp, one-time, substantial, and permanent shock to a firm's employment (evidenced especially by the depressed state of employment levels at the layoff firm after the event). In Section 3.4 we discuss the assumptions that allow us to use this shock to identify main and spillover effects of mass layoffs and check that our findings are robust to alternate layoff definitions in Table A3.

3.3.3 Characteristics of Workers at the Layoff Firm

We now describe the workers who are separated compared to retained in mass layoff events. This will further illustrate the nature of mass layoffs in Brazil as well as lend credence to our preferred layoff definition.

Table 3.1 summarizes traits of workers who are exposed to a layoff. The table examines workers who are directly laid off and those who remain employed by the flagship firm and further stratifies the sample by long-tenure workers. Note that the statistics are representative of a 5% random sample of individuals in the RAIS data.

We begin with Panel A of Table 3.1, which summarizes worker characteristics one year prior to the layoff event. In terms of demographics, there is little difference between the laid off and retained workers; they are largely the same in terms of race, gender, and age.[19] They are also similar in average weekly hours worked; this figure ranges from 41.6 hours per week to 44.1 hours per week.[20] We note also that long-tenure workers, both in total and as part of the layoff sample, are a minority share of the layoff firm's employment pool. They make up slightly under 10% of laid off workers and just about

[19] We do note that the long-tenure workers are older by three to four years in comparison to the average worker at the layoff; this is in line with our definition of long-tenure workers, who have to be employed at the layoff firm for three continuous years prior to the event.

[20] In 1988, the Brazilian constitution was amended to decrease the maximum weekly work hours allowable from 48 to 44, with an increase in overtime pay [Amadeo et al., 2000].

16% of total employment at the layoff firm. So, the patterns we see for all workers at the firm are relatively more driven by characteristic of short-tenure workers than long-tenure workers.[21]

Worker characteristics across the groups diverge when we consider education, annual wages, and occupation. Separated workers already begin to experience a drop in annual earnings at $t = -1$, making $R\$$ 9,937 on average compared to their retained counterparts, who make nearly triple that much at $R\$$ 25,343. Similar patterns hold when we look at the long-tenure workers at the layoff firm, where those who will ultimately be separated from the firm make substantially less than the long-tenure workers who will be retained. The pay differential is smaller between separated and retained long-tenure workers relative to the average worker. Laid off workers also work multiple jobs in the year prior to layoff (44.9% of all separated workers and 23.3% of long-tenure separated workers, compared to 9.0% of all retained workers and 2.2% of the long-tenure retained workers), perhaps to smooth the over the decreased earnings that they receive. This also points to alternative jobs as being a source of implicit insurance in our context.

Some characteristics are unique to the average laid-off worker (driven mostly by the shorter-tenure workers). For example, this group of individuals is the least educated (18.9% with a high school degree or more, compared to 25% to 33% of other worker groups). They are relatively low-paid compared to other workers in the layoff firm, including other separated workers, making below the median in terms of hourly wages. Finally, while there is heterogeneity in the occupations worked by employees at the layoff firm, the separated workers are more likely to be agricultural workers, particularly among the shorter-tenure workers.

[21] The RAIS data is unique in that it also provides a reason for job separations. As documented in Table A2, there is heterogeneity in the causes of job separations. Comparing the average displaced worker to long-tenure workers, similar shares of them are fired without any reason (5.2% versus 7.2%), with cause, or for some other reason; they both also have low but positive probabilities of being reassigned (this is specific to public sector workers). They differ in the probability of being fired indirectly for another reason - from the data documentation, indirect terminations are the result of workers initiating their separation from the firm.

Panel B of Table 3.1 summarizes worker characteristics in the year the flagship mass layoff event occurs. Separated workers experience large drops in annual earnings regardless of their tenure. Average retained worker earnings move very slightly. Some other notable data facts include occupation switches and mobility patterns among separated workers. These workers appear to respond immediately to layoff events: 58% of separated workers remain in the same occupation conditional on working, while the rest switch jobs. Only 51% of these workers remain in the same region upon separation, suggesting that migration is a possible response to being laid-off. Separations appear to occur early in each year, as only 12.5% of laid-off workers have any work at the layoff firm during the year of the event (corroborated by the patterns in Figure 3.3).

Finally, Panel C of Table 3.1 summarizes characteristics of workers at the layoff firm four years after the event. Interestingly, separated workers experience improvements in earnings (potentially due to gaining tenure at their new jobs and workforce re-entry) relative to retained workers in the year of the layoff.[22] More separated workers appear to have left the layoff firm's region over time, as well as their occupation at separation (48% of them remain in the same occupation as when they were laid off four years after the layoff, compared to 58% in the year of the layoff). Among separated workers who are back in the labor force, 14.1% of them return to the layoff firm (11.4% for long-tenure worker specifically). In general, earnings and employment circumstances improve slightly for separated workers relative to the year of the layoff, though recovery is not total, as they do not catch up to their pre-layoff earnings or employment rates. In an interesting contrast, retained worker circumstances change quite substantially. As noted above, they experience earnings declines relative to the year of the layoff year, and about one quarter of them do not work at all four years after the layoff. The same share moves to a new region and about 35.7% of them move occupations. Interestingly, there is also a slight increase in the share of retained workers who work multiple jobs, conditional on

[22]Retained workers witness declines in earnings (these declines are larger in magnitude for long-tenure workers).

working at all.

Table 3.1: Characteristics of Workers Exposed to Layoff

	All Workers in Layoff Firm		*Long Tenure Workers in Layoff Firm*	
	Separated in t=0	Retained in t=0	Separated in t=0	Retained in t=0
A. Baseline Characteristics in t=-1				
Age	32.0	33.0	36.1	37.2
Female (%)	27.8	28.8	29.4	24.0
Race: White (%)	49.5	46.1	46.6	43.8
Race: Indigenous (%)	0.1	0.2	0.0	0.1
Education: High School+ (%)	18.9	33.3	25.8	25.2
Education: College+ (%)	2.2	6.3	2.5	2.0
Annual Earnings	9937.6	25343.9	12869.8	17699.3
Weekly Hours	41.6	43.4	42.3	44.1
Hourly Wage	8.3	13.0	8.7	8.6
Within-Firm Hourly Wage Rank (%)	44.9	54.8	54.2	64.3
Working Multiple Jobs (%)	37.4	9.0	23.3	2.2
Occupation: Operatives and Production Workers (%)	11.0	10.6	10.5	9.7
Occupation: Drivers and Transport Operatives (%)	5.9	9.5	5.6	6.5
Occupation: Craftsmen, Technicians, and Machinists (%)	15.1	26.0	29.1	28.7
Occupation: Agricultural Workers (%)	62.1	40.5	44.2	43.7
Occupation: Clerks, Sales, and Services (%)	3.8	7.7	6.5	8.6
Occupation: Professional and Managers (%)	2.1	5.7	4.2	2.9
B. Characteristics in t=0				
Annual Earnings	6808.6	25788.4	7424.8	17133.5
Any Work in Year (%)	56.5	100.0	57.9	100.0
...At Layoff Firm (%)	12.5	100.0	4.4	100.0
...In Same Occupation (%)	58.0	94.8	52.6	96.1
...In Same Region (%)	51.0	95.8	61.1	97.6
...In Same City (%)	37.6	94.3	48.9	97.0
...In Multiple Jobs (%)	24.8	11.4	16.5	7.7
...Weekly Hours	42.9	43.0	43.1	43.5
...Hourly Wage	8.5	13.9	8.8	9.3
C. Characteristics in t=4				
Annual Earnings	9211.1	24536.2	10756.9	14775.2
Any Work in Year (%)	58.1	73.9	63.4	77.8
...At Layoff Firm (%)	14.1	52.5	11.4	59.1
...In Same Occupation (%)	46.0	64.3	46.4	69.0
...In Same Region (%)	46.0	74.3	65.5	82.3
...In Same City (%)	31.6	65.5	51.0	77.2
...In Multiple Jobs (%)	20.5	11.3	16.0	7.7
...Weekly Hours	42.6	43.2	43.0	43.5
...Hourly Wage	9.9	17.4	9.7	10.0
Number of Workers	5,666	5,438	554	1,291

Notes: This table reports characteristics of workers who are employed at the layoff firm in the year prior to the layoff. Statistics reported in italics are conditional on working, whereas annual earnings includes zeros. By construction, all workers are employed in panel A. Weekly hours are computed for each each month a worker works in the year, assuming they work 4 weeks-per-month. "High School+" and "College+" are not mutually exclusive. Long tenure workers are defined as workers who have been continuously employed at the same firm for 3 years in the year prior to the event. A worker is classified as separated in $t = 0$ if they are not employed at the layoff firm at any point in the year. Results are based on a 5 percent sample of the RAIS data.

Taken together, our summary statistics tell a story about the kinds of workers who are exposed to mass layoffs as well as the different ways that they respond following layoff events. There are heterogeneous responses to mass layoffs that we document by worker tenure; we also note that layoffs appear more targeted towards low-education workers.

Importantly, the effects of layoffs linger over time: we see their effects on geography and earnings even four years after the event. Our findings hint at causal impacts, but without finding appropriate comparison groups for layoff workers and layoff regions, we cannot exactly quantify such effects. In the next sections, we turn to describing our empirical strategy and describing threats to identifying both the regional and and worker-specific effects of mass layoff exposure. We take care to also discuss identification of the spillover effects: as the mass layoffs are shocks that hit entire labor markets, their effects are not necessarily isolated to displaced workers.

3.4 Empirical Strategy

The crux of our analysis involves identifying the causal impacts of sharp, unforeseen layoff events on local labor market conditions, as well as to explain the magnitudes and directions of our estimated impacts. For example, we want to measure how much non-laid off worker earnings and employment prospects may be impacted by a localized mass layoff. To estimate such effects, we would ideally randomly assign plant closures or flagship mass layoffs across different regions and economies, then measure how local workers were impacted. However, in the absence of a such an experiment, we turn to a research design leveraging event timing and a closely related comparison group prior to the layoff.

In this section, we motivate the matched event study that we employ to identify the main and spillover effects of mass layoffs and how constructing an appropriate comparison group helps us overcome identification challenges. We then describe our baseline event study specification, which captures event-time effects of mass layoffs on both directly separated workers as well as the spillover effects on non-laid off workers. Finally, to better understand mechanisms explaining our results, we turn to a worker-level analysis (also detailed below).

3.4.1 Matching Approach

One of the key challenges to identifying causal impacts of mass layoffs is the potential selection of firms and geographies into layoff events. Specifically, we are concerned layoffs are contemporaneous with other market-wide unobserved shocks. We investigate this in Table 3.2, which provides summary statistics of layoff firms and regions, as well as all regions at large in Brazil, and potential comparison regions.

Table 3.2 also shows that layoff regions regions are not directly comparable with the typical Brazilian region. In the third column of Table 3.2, "Layoff Region", we summarize characteristics of all *regions* which experience a mass layoff in the first and second years prior to the layoff. For example, layoff regions employ an average of 59,142 workers one year before the layoff occurs (Panel A) and two years prior to the layoff, 29.7% of firms were manufacturing firms. These characteristics differ somewhat from the average region. In the second column of Table 3.2, "All", we average characteristics of all non-mass-layoff regions over time, then average across the years prior to the layoff: for interpretation, we can think of these statistics as representing typical regions in Brazil. There are some striking differences between layoff regions and all other regions in Brazil. The average region had much larger employment levels in both years prior to the layoff, a more educated workforce (65% of workers in the typical region two years prior to the layoff had a high school degree or higher, compared to 44.6% of workers in the layoff region), and substantially lower shares of firms in agricultural or manufacturing sectors - in stark comparison to the layoff firm and layoff region compositions. As a result of this imbalance, our research design needs to account for these differences in size, employment, and industry structure.

For every layoff region in our dataset, we match a control region in the same year that has similar characteristics two to four years before the mass layoff event. As control regions are perfectly balanced with treated regions, this approach ensures that our event study results do not suffer from concerns about unbalanced, heterogeneous treatments

outlined in the "new differences-in-differences" literature.[23] This means that for a layoff event that occurs in 2010, we look at the characteristics of all potential control regions in 2010 and match using characteristics in 2006 to 2008. We adopt Mahalanobis distance matching for this rather than propensity score matching due to the large number of covariates and small number of potentially heterogeneous events.

We make several blocking restrictions to ensure appropriate comparisons. To avoid the prospect of additional spillovers on the control region, we restrict to matches outside the same federal state as the reference layoff event. We also exclude control microregions that experience a mass layoff event within 4 years of the reference mass layoff event.[24] Finally, to ensure that control regions are on the same trend, we only consider control regions whose employment levels in years four to two years prior to the layoff event are within 10 percent of the target layoff region. Note that we do not impose any restrictions on period $t = -1$.

Given these blocking restrictions, we use Mahalanobis matching that leverages the following variables: (i) average annual earnings, (ii) the standard deviation of annual earnings, (iii) average hourly earnings conditional on working, (iv) total employment, (v) total hours, (vi) the number of firms, (vii) the share of workers with a college education, (viii) the share of workers who are female, (ix) the share of workers under the age of 25, (x) the share of workers in an agricultural occupation, (xi) the share of employment in the largest three firms, and (xii) the share of employment in agriculture/resource extraction, manufacturing, retail trade/services, and the public sector. Average annual earnings (i) and total employment (iv) use information from years three and four prior to the layoff, with the rest of the variables focusing only on two years prior to the layoff.

[23] Concerns in that literature arise from heterogeneous treatment effects in the presence of heterogeneous event times. In this context, using later events as controls for earlier events can result in incorrect estimates. Our matching approach identifies a non-layoff region for every, using observables in the control region up until the reference layoff year. This means that our treatment and control regions are perfectly balanced, and our we identify an average treatment effect on the treated that correctly averages over treated regions given our identification assumptions.

[24] This means that two "ever mass layoff" event regions can be matched together, as long as there is a sufficient gap between the two events.

We orthogonalize these matching variables by the inverse of the variance matrix, which accounts for the fact that correlated variables may over-weight redundant information. Then, we match each treated region to the closest based on the Euclidean distance computed using these orthogonalized variables.

Table 3.2: Characteristics of Layoff Regions and Matched Comparison Regions

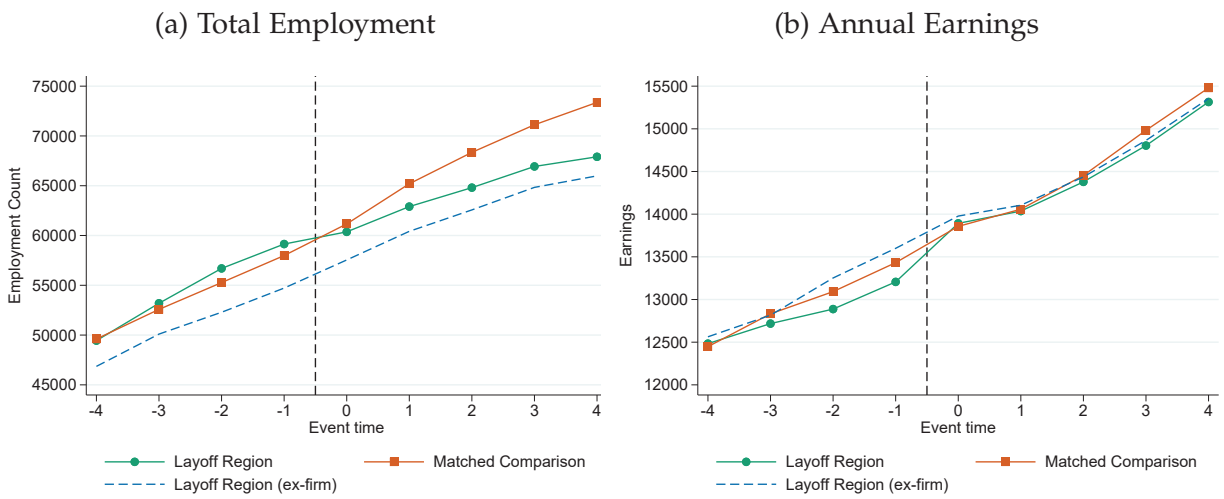
	Layoff Firm	All	Layoff Region	Comparison Region	p-value of t-test
A. Layoff Variables					
Employment in $\tau = -1$	3,604	90,528	59,142	57,987	
Employment Change in $\tau = 0$	-1,300	4,647	1,224	3,176	
Percent Change	-49.8	5.1	0.0	5.3	
Layoff is closure (%)	4.8				
B. Characteristics in $\tau = -2$					
Total Employment	3,765	85,803	56,692	55,271	0.91
Employment Growth ($\tau = -4$ to $\tau = -2$) (%)	18.0	5.4	7.2	6.1	0.20
Earnings	11179	21110	12888	13092	0.70
SD	12616	25574	14402	13924	0.47
Earnings Growth ($\tau = -4$ to $\tau = -2$) (%)	1.6	1.0	2.0	2.6	0.39
Number of Firms	1	3,935	2,848	3,235	0.49
Share in largest 3 firms (%)		25.6	23.6	16.9	0.00***
Female (%)	26.1	38.8	34.9	35.9	0.39
Age < 25 (%)	29.7	22.2	25.4	24.4	0.17
Age 25-49 (%)	62.0	66.7	64.5	65.4	0.07*
Age ≥ 50 (%)	8.3	11.1	10.2	10.2	0.89
Education: High School+ (%)	20.6	65.0	44.6	49.7	0.10
Education: College+ (%)	2.5	14.9	8.4	9.6	0.10
State: Rio de Janeiro or Sao Paulo (%)	21.0	40.0	21.0	9.7	0.08*
C. Industry Composition in $\tau = -2$					
Agriculture and Resources (%)	48.4	6.1	21.2	16.3	0.02**
Manufacturing (%)	51.6	18.6	29.7	23.9	0.02**
Utilities (%)	-	0.5	0.4	0.5	0.26
Construction (%)	-	5.9	2.6	3.0	0.39
Wholesale and Retail Trade (%)	-	18.6	15.8	20.1	0.00***
Transportation (%)	-	5.1	2.7	3.3	0.07*
Hotel and Restaurants (%)	-	3.6	1.8	2.7	0.00***
Health and Social Services (%)	-	3.2	2.0	2.3	0.24
Other Services (%)	-	0.0	0.0	0.0	0.17
Finance, Insurance, and Real Estate (%)	-	14.6	4.3	4.9	0.30
Education and Public Administration (%)	-	20.1	17.1	20.1	0.07*
Number of Layoff Events	62				

Notes: "Layoff firm" shows only summary statistics that can be computed across all layoff firms. "Layoff Region" are the summary statistics for regions experiencing a mass layoff and "Comparison Region" is the matched control. For every layoff region, we link it to the average summary statistics across all Brazilian microregions in the same year as the layoff. These summary statistics are displayed in column "All", which effectively re-weights the aggregate summary statistics to the distribution of the layoff years. The p-value reflects differences in characteristics between the layoff region and the matched comparison region.

Looking at the columns "Comparison Region" and "p-value of t-test" in Table 3.2, we see that our matching procedure does a reasonable job at selecting comparison microregions. In particular, the comparison region matches closely to the layoff areas in employment levels, earnings, and demographics. We verify closeness in employment

and earnings for years before $\tau = -2$ in Figure 3.4. Panel (a) illustrates employment levels for the layoff region, the layoff region excluding the layoff firm, and the matched comparison region. We note that the flagship layoff firm makes up a substantial share of regional employment even after the layoff event; employment losses at these firms are a huge fraction of local labor markets. Layoff firm workers are also paid slightly less than the other workers in their local labor market. Prior to the layoff, the matched comparison tracks very closely with the layoff region in employment levels and annual earnings.[25] After the layoff, earnings between layoff firm and other workers in the region converge to a level beneath that of the comparison group, suggesting that the layoff has spillover effects on other regional workers. The same holds true for employment.

Figure 3.4: Plots of Region-Level Employment and Earnings around Event Time



Notes: Figure (a) shows average total employment across layoff regions (green circles), matched comparisons (orange squares), and layoff regions removing the contribution of those employed at the layoff firm (dashed line). Figure (b) shows the same for mean earnings averaged across regions. The x-axis is years relative to the timing of the layoff event, which occurs at 0.

The matching procedure overall does a good job in finding comparison regions but misses somewhat on a few variables: notably, the share of employment in the largest three firms as well as some industry composition measures. While important, this is less of a concern for measuring the effects of mass layoffs, as it does not impact our results.

[25]We verify parallel trends in outcomes as well in our empirical results.

3.4.2 Empirical Specification: Local Labor Market-Level Analysis

Using our pairs of treatment and matched comparison regions, we compare local labor market outcomes in areas that suffered flagship mass layoffs to similar-looking microregions that are not impacted by mass layoffs. Our preferred specification "stacks" a balanced panel of outcomes for the layoff microregion with data from the same calendar years in the matched comparison region.[26] Letting j denote microregion, t denote event time, and c denote a set of matched comparisons, we run the following regression:

$$Y_{jtc} = \alpha_{ct} + \alpha_{jc} + \sum_{\tau \neq -1} D[t = t_c^0 + \tau]\delta_\tau \times LayoffRegion(c, j) + u_{jtc} \tag{3.4.1}$$

where t_c^0 is the calendar year of the mass layoff event for the layoff region (or placebo layoff event for the comparison region) and τ is event time relative to this event. We let τ range from $\tau = -4$ to $\tau = 4$ and include comparison-time fixed effects, α_{ct}, which account for aggregate shocks in years experienced relative to the mass layoff, and comparison-region fixed effects α_{jc} which control for time-invariant region factors within each comparison-treated unit set. By including α_{ct} and α_{jc}, this specification imposes the parallel trends assumption at *each* comparison level. Unless otherwise reported, we weight our estimates using total employment in the region two years prior to the layoff event.[27] Standard errors are clustered at the level of the matched comparison.

The coefficient of interest is δ_τ, which measures changes in the outcome Y_{jtc} in the layoff region that is not accounted for by trends in the comparison region. Under the assumption of parallel trends – that post-layoff trends in the comparison region trace out the counterfactual in the layoff region – this parameter estimates the average treatment effect on the treated (ATT). We ultimately obtain two sets of estimates for δ_t, which allow us to analyze both the main and spillover effects of mass layoffs at the aggregate,

[26] To be concrete, we have 62 layoff regions and 62 comparison regions for a total of 124 region-level panels of 9 observations each.

[27] Other outcomes and unweighted results are reported in Table 3.3. We prefer the weighted results since this more closely mirrors running an analysis at the worker level.

region level. To analyze the spillover effects, we take the series of the relevant regional outcomes and remove the contribution of workers employed at the layoff firm in either of the two years prior to the layoff from the entire series for the layoff microregion. Thus, we capture the average impacts of layoffs on non-separated workers who could still potentially interact with laid-off workers in the same local labor market. When we report δ_t for aggregate results, we make it clear whether we are looking at total effects or spillover effects.

The key identification assumption is that the comparison region forms a valid counterfactual for the treated region in the years after the layoff. Specifically, the comparison region must be a valid control for *both* the entire treated region (which averages the displaced, spillover, and other workers) as well as the regional spillover workers.[28] This is because we don't change comparison groups for different outcome constructions. Our matching strategy also aids in this, as the comparison group is similar to the treated regions ex-ante, and should difference out the time effects driven by common trends in both the treatment and control regions. We verify our identification assumption by checking for parallel trends in both, total and spillover effects.

In identifying spillover effects of a mass layoff event, a key source of bias is that a mass layoff event may be caused by an unobserved shock that affects both the layoff firm and other firms in the same local labor market. If an unobserved shock both directly causes a mass layoff event as well as broader employment declines, this will result in an overestimation of the spillover effect if the same shock is not picked up in the control region.[29]

We do not believe correlated shocks are a large source of bias for two reasons. First,

[28]The worker-level results presented below will directly address this concern in our results by comparing workers defined by pre-layoff covariates cells.

[29]For example, Uber's headquarters are located in San Francisco, a hub for high-tech companies. It is possible that overall declines in the tech industry both cause mass layoffs at Uber as well as broader declines in the tech industry, that affect California. In this example, our strategy seeks to compare employment declines in San Francisco to another control region that also has high employment in the tech sector and is thus subject to the same aggregate shocks.

control regions are matched based on the baseline industry composition, in addition to other characteristics. To the extent that the unobserved, correlated shock is driven by aggregate declines within a specific industry, this should be captured by the trajectory of the control region, which is subject to similar macro-shocks. Unlike Gathmann et al. [2020], we do not find that control regions suffer large employment declines at the same time. Second, in practice, we find that the spillover effect only shows up with a lag relative to the layoff event. This pattern is inconsistent with our results being driven by correlated shocks to the treated regions' local labor markets. This means that such an unobserved shock would have to cause the mass layoff in the treated region in a way that is not related to industry composition and also not show up as affecting the non-layoff firms until a year later. We find it difficult to rationalize this pattern with a violation of the exclusion restriction.[30]

3.4.3 Understanding Mechanisms: Worker-Level Analysis

We turn to an analysis of worker-level outcomes in order to quantify heterogeneity in the incidence of spillover effects. This analysis will also have the benefit of allowing us to look at the role that internal migration plays in the analysis specifically through estimating the impact of being a treated or a spillover worker on the likelihood an individual migrates or switches occupations altogether. Our empirical approach will use the trends in outcomes for observably similar workers in each comparison region to impute counterfactual outcomes for workers in the layoff region. The implicit model justifying this research design is that observably similar workers in the layoff and comparison re-

[30] Another source of bias is that the mass layoff event has spillover effects across micro-regions. leading to a violation of the stable unit treatment values assumption (SUTVA) implicit in the research design. Such a violation would lead to underestimation of the spillover effect due to a misattribution of differential trends in the control region. First, when selecting a matching control region, we select microregions that are outside of the same state as the control region. It our hope that this limits the potential for additional spillovers on the control region. Second, in practice, we find that the employment trends in the control regions are on average smooth around the layoff event, which we would not expect in the case of additional spillovers on the control region.

gions would have had similar trends in equilibrium skill prices in the absence of the mass layoff. For computational convenience, these results will only use a 5% sample of the data.

A key focus of the worker-level analysis is to understand how mass layoffs impact different types of workers. As a result our key parameters of interest will be various worker characteristics interacted with event-time coefficients. At the highest-level, we will break out effects into workers employed at the layoff firm in any of the 4 years prior to the layoff event ("layoff workers") from other workers not directly employed by the layoff firm ("spillover workers"). In our preferred specification, we will control for trends common to observably similar treated and control units using pre-layoff worker characteristics. This will relax the parallel trends assumption to hold only within these worker cells.

As before, we "stack" balanced panels of microdata on workers in both the layoff region and the comparison microregions using data from 4 years prior to the event and 4 years after the event. Our sample is defined as any worker who appears in the formal sector of a region two years before a layoff event (or placebo layoff event in the case of comparisons regions) and is employed in $\tau \in \{-4, -3, -2\}$. Letting i denote worker, j denote microregion, t denote calendar time, and c to denote a matched comparison set, we run the following regression:

$$Y_{ijct} = \alpha_i + \alpha_{ct} + \boldsymbol{X}_i\boldsymbol{\theta}_{ct} + \sum_{\tau \neq -1} D[t = t_c^0 + \tau] \left(\delta_\tau^{Layoff} D_{ic}^{Layoff} + \delta_\tau^{Spillover} D_{ic}^{Spillover} \right) + u_{ijct} \tag{3.4.2}$$

As before, t_c^0 is the calendar year of the (placebo) mass layoff event for the layoff (comparison) region and τ is event time relative to this event. D_{ic}^{Layoff} is an indicator variable for being a layoff worker in the layoff microregion and $D_{ic}^{Spillover}$ is an indicator for being a spillover worker in the layoff region, both being employed in the layoff microregion in at least $\tau = -2$. Importantly, we include an individual fixed effect α_i to control for

time-invariant person-level characteristics that are separate from the layoff effects.[31]

The key parameters of interest are δ_τ^{Layoff} and $\delta_\tau^{Spillover}$, which measure how changes relative to the individual's baseline in $\tau = -1$ differs in the layoff region as compared to the control within observably similar worker cells defined by $\boldsymbol{X_i}$. Key to our approach will be to include these pre-layoff covariates and allow for the trends of these covariates to vary with the comparison set and year through $\boldsymbol{\theta_{ct}}$. These covariates include both pre-layoff earnings and other economic outcomes, characteristics of the pre-layoff employer, and fine demographic cells.[32]

3.5 Local Labor Market-Level Analysis

3.5.1 Baseline Results

We present results for both, the total and spillover effect estimates of δ_t for a battery of outcomes in Figures 3.5 and 3.6. The employment outcomes in Figure 3.5 are the log of the total regional employment, log of the total regional hours worked, and log total regional employment. The earnings outcomes in Figure 3.6 are the log of the total wage bill, average annual earnings, and hourly earnings in Brazilian reals (R$). The complete set of parameter estimates with comparison-group clustered standard errors are reported in Table 3.3. Our preferred specifications weight each observation by microregion employment at $\tau = -2$, though we present unweighted estimates for completeness. All estimated effects are presented relative to $\tau = -1$.

Figure 3.5 plots the effects of flagship mass layoff events on employment outcomes in

[31] There are a few individuals who can end up in multiple comparison events c. We denote the individual fixed effect as α_i for clarity.

[32] Specifically, these variables consist of: (i) quartiles of pre-layoff earnings interacted with pre-layoff earnings growth between $\tau = -3$ and $\tau = -2$; (ii) quartiles of total annual hours; (iii) 5-year age bins interacted with gender and indicators for educational attainment (<HS,HS+,College+); (iv) indicators for firm tenure in $\tau = -2$ of $(0, 0-3, 3)$ years; (v) 1-digit industry interacted with indicators for employer size of $(\leq 10, 11-100, 101+)$; (vi) 1-digit occupation interacted with education, where terciles/quartiles of hours and earnings variables are computed using the distribution in $\tau = -2$ in the layoff region and defined within the comparison set.

local labor markets as well as the standard errors. The total effects of flagship mass layoffs on regional employment variables are presented in teal, and the separated spillover effects are plotted in orange. The difference between the estimated total and spillover effects captures the direct impacts of layoffs on separated workers across the treated microregions.

The first striking thing about Figure 3.5 is the lack of substantial pretrends across both total and spillover effect estimates for the employment outcomes. One notable exception is the log of employment for the region in total, where the layoff event's impact at $\tau = -4$ is $-.025$ log points and is statistically significant at the 5 percent level. In examining the spillover effects, there are also no economically or statistically significant pretrends. This demonstration of parallel trends helps validate our identification strategy. In particular, it suggests that our selected comparison regions do a good job in terms of tracking treated region employment outcomes for *both*, the region as a whole and for non-laid off workers.

Across all employment outcomes, we find that flagship mass layoffs have strong, sharp effects on local labor market health. Both displaced workers and non-displaced workers are affected by layoffs. We note that layoff workers are immediately affected in same year as the layoff event ($\tau = 0$), but that spillover workers are affected in the year after the layoff event ($\tau = 1$). This additionally lends credibility to our identification strategy: if the changes in outcomes were driven by some localized shock other than the layoff, we would expect that the total effects and spillover effects move at the same time. This is not the cause, suggesting that it is the flagship mass layoff that drives changes in employment outcomes.

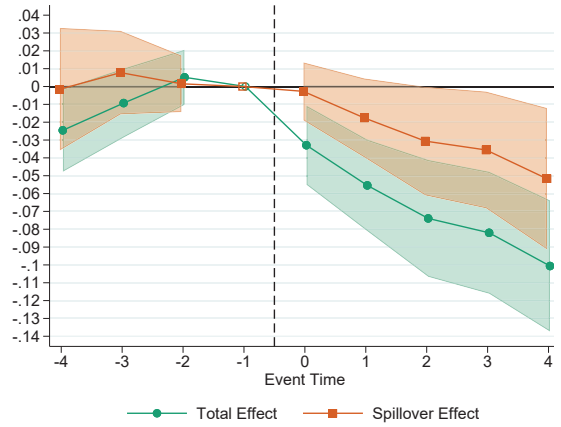
We focus now on Panel (c) of Figure 3.5, the total employment effects (corresponding to the "log(Employment)", "log(Hours)", and "Employment" columns in Table 3.3). At the onset of the mass layoff, total employment in the layoff region goes down 1,952 workers on average, and it continues persistently declining, with no plateau – four years after

the layoff event, employment levels are down 6,618 workers relative to if no layoff occurred. The spillover group, which consists of all workers in the microregion not directly separated during the layoff event, also experiences consistently declining employment levels after the mass layoff. In the first year after the layoff, employment for this group drops by 350 workers, though this swiftly gains momentum (and also grows in statistical significance), growing to 4,111 workers in this group no longer working four years after the layoff. The spillover effects initially account for about 43% of the total mass layoff effect ($\frac{1,510}{3,444} = .43$ at $\tau = 1$), though this share grows to 62% four years afterwards. One potential reason that the spillover effects explain an increasing share of the total impact of mass layoffs could be that separated workers leave the region. Indeed, Table 3.1 shows that 37 percent of separated workers leave the city where the layoff occurred within a year.

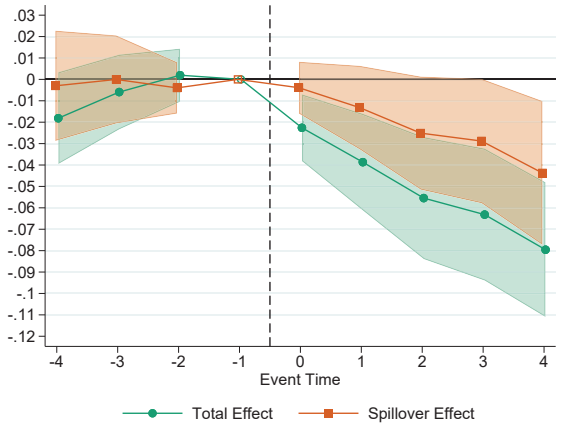
These estimates imply a multiplier effect of 0.5: with every job lost in the layoff event causing an additional 0.5 extra job losses. These effects are larger than GHS, who estimate a slightly smaller multiplier effect of about 0.4.

Figure 3.5: Effects of Mass Layoffs on Employment in the Local Labor Market

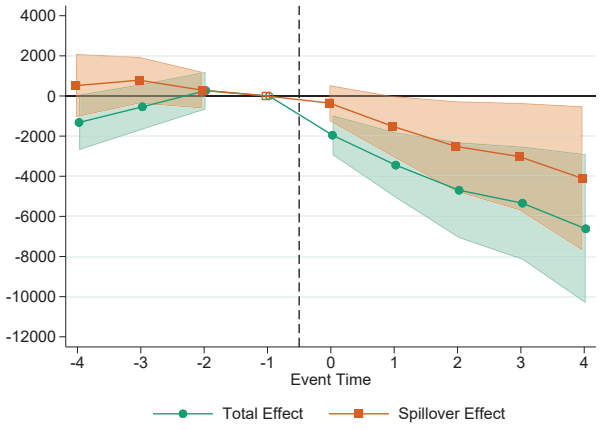
(a) log(Total Employment) (Weighted)



(b) log(Total Hours) (Weighted)



(c) Total Employment (Unweighted)

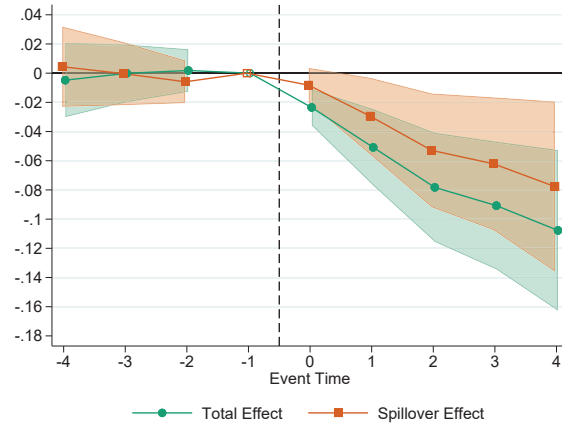


Notes: Plots coefficients of δ_t from the specification outlined in Equation 3.4.1. Point estimates for other outcomes are reported in Table 3.3. Panel (a) plots log-employment weighted by pre-layoff employment. Panel (c) plots unweighted regressions on total employment. The "Total Effect" coefficients are estimated using data on all workers in the layoff and comparison regions, whereas the "Spillover Effect" excludes workers from the layoff region who worked at the layoff region in $t = -1$ or $t = -2$. Standard errors are clustered at the matched comparison level.

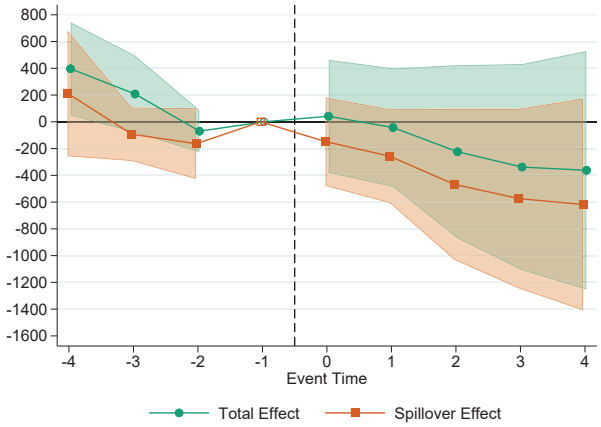
We now turn to examining how mass layoffs shape earnings outcomes across affected regions. Figure 3.6 reports the total log-wage bill for workers in the layoff region, average annual earnings, and the average hourly wages for laid-off workers and spillover workers in the microregions. Strikingly, average hourly earnings, as well as annual earnings, are not substantially impacted for either the region as a whole or the spillover group separately. Interestingly, for worker earnings, it appears that the spillover group actually responds immediately in $\tau = 0$, while the total affect seems to experience a lag. Also interesting are that the point estimates on worker earnings are lower for the spillover group than for the total group (implying the spillover workers have larger earnings declines following a flagship mass layoff than displaced workers). However, the effects across the groups are statistically indistinguishable and confidence intervals engulf both sets of δ_τ estimates for earnings.

Figure 3.6: Effects of Mass Layoffs on Earnings in the Local Labor Market

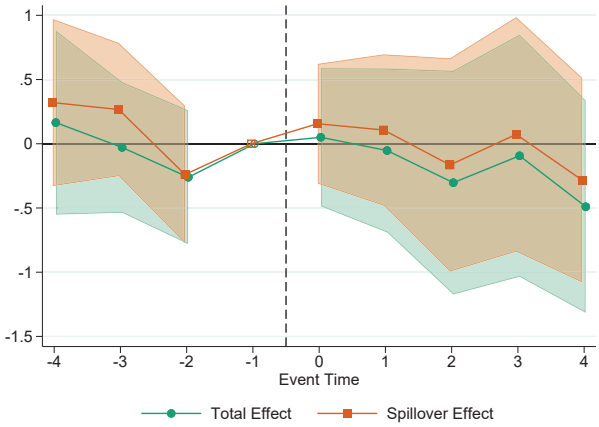
(a) Log-Wage Bill



(b) Average Annual Earnings



(c) Average Hourly Earnings



Notes: Plots coefficients of δ_t from the specification outlined in Equation 3.4.1. Point estimates for other outcomes are reported in Table 3.3. All estimates are weighted by pre-layoff employment. The "Total Effect" coefficients are estimated using data on all workers in the layoff and comparison regions, whereas the "Spillover Effect" excludes workers from the layoff region who worked at the layoff region in $t = -1$ or $t = -2$. Standard errors are clustered at the matched comparison levels.

The starkest results come when we examine the log-wage bill. Here, the effect timing is more in line with what we saw from the employment estimates in Figure 3.5: the total effects occur immediately, while the spillover effects lag a year later. The spillover effects after the layoff event closely track the main effects once they become more noticeable at $\tau = 1$. These wage bill estimates imply a larger multiplier effect of 0.7, which is larger than the employment effects.

Table 3.3: Local Labor Market Event Study Results

	log(Employment)		log(Hours)		log(Wage Bill)		Employment		Average Earnings (if work)		Average Hourly Wage (if work)	
	Total	Spillover	Total	Spillover	Total	Spillover	Total	Spillover	Total	Spillover	Total	Spillover
A. Weighted												
Event ($\tau = -4$)	-0.025**	-0.002	-0.018	-0.003	-0.005	0.004	-	-	396.547**	214.271	0.165	0.321
	(0.012)	(0.018)	(0.011)	(0.013)	(0.013)	(0.014)			(178.033)	(240.806)	(0.365)	(0.332)
Event ($\tau = -3$)	-0.009	0.008	-0.006	-0.000	-0.000	-0.000	-	-	207.799	-92.659	-0.028	0.267
	(0.010)	(0.012)	(0.009)	(0.010)	(0.010)	(0.011)			(147.138)	(102.759)	(0.260)	(0.265)
Event ($\tau = -2$)	0.005	0.002	0.002	-0.004	0.002	-0.006	-	-	-69.794	-163.623	-0.261	-0.240
	(0.008)	(0.008)	(0.006)	(0.006)	(0.007)	(0.007)			(81.087)	(135.472)	(0.266)	(0.273)
Event ($\tau = -1$)	-	-	-	-	-	-	-	-	-	-	-	-
Event ($\tau = 0$)	-0.033***	-0.003	-0.023***	-0.004	-0.024***	-0.008	-	-	41.669	-147.430	0.050	0.156
	(0.011)	(0.008)	(0.008)	(0.006)	(0.006)	(0.006)			(215.466)	(169.292)	(0.274)	(0.238)
Event ($\tau = 1$)	-0.055***	-0.018	-0.039***	-0.013	-0.051***	-0.029**	-	-	-42.988	-257.138	-0.053	0.107
	(0.013)	(0.011)	(0.012)	(0.010)	(0.013)	(0.013)			(225.891)	(179.186)	(0.325)	(0.301)
Event ($\tau = 2$)	-0.074***	-0.031*	-0.055***	-0.025*	-0.078***	-0.053***	-	-	-223.647	-467.392	-0.303	-0.165
	(0.017)	(0.016)	(0.015)	(0.013)	(0.019)	(0.020)			(329.988)	(288.131)	(0.444)	(0.423)
Event ($\tau = 3$)	-0.082***	-0.036**	-0.063***	-0.029*	-0.091***	-0.062***	-	-	-338.255	-573.780*	-0.093	0.073
	(0.017)	(0.017)	(0.016)	(0.015)	(0.022)	(0.023)			(392.556)	(342.363)	(0.480)	(0.466)
Event ($\tau = 4$)	-0.101***	-0.052**	-0.080***	-0.044**	-0.108***	-0.078**	-	-	-362.404	-617.785	-0.490	-0.285
	(0.019)	(0.020)	(0.016)	(0.017)	(0.028)	(0.030)			(454.525)	(404.144)	(0.422)	(0.406)
B. Unweighted												
Event ($\tau = -4$)	-0.012	0.033**	-0.007	0.020	0.003	0.024	-1320.210*	515.968	261.099	-52.472	0.009	0.579
	(0.014)	(0.015)	(0.013)	(0.014)	(0.017)	(0.019)	(706.636)	(806.511)	(195.995)	(215.825)	(0.559)	(0.563)
Event ($\tau = -3$)	-0.003	0.027**	0.003	0.015	0.003	0.011	-540.629	786.161	111.127	-183.936*	0.082	0.643
	(0.012)	(0.013)	(0.011)	(0.012)	(0.013)	(0.013)	(581.647)	(595.523)	(101.823)	(104.731)	(0.406)	(0.400)
Event ($\tau = -2$)	0.008	0.000	0.008	-0.001	0.008	-0.001	266.387	286.581	21.037	-8.288	-0.105	-0.034
	(0.010)	(0.010)	(0.008)	(0.008)	(0.010)	(0.010)	(479.632)	(467.531)	(81.267)	(83.384)	(0.244)	(0.273)
Event ($\tau = -1$)	-	-	-	-	-	-	-	-	-	-	-	-
Event ($\tau = 0$)	-0.052***	-0.005	-0.033***	-0.004	-0.032***	-0.009	-1952.355***	-350.935	261.170***	-47.081	-0.098	0.169
	(0.009)	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(507.126)	(458.073)	(80.233)	(73.349)	(0.262)	(0.243)
Event ($\tau = 1$)	-0.073***	-0.014	-0.048***	-0.008	-0.056***	-0.022*	-3444.919***	-1510.581*	203.768**	-121.361	-0.263	0.081
	(0.015)	(0.014)	(0.015)	(0.013)	(0.014)	(0.013)	(832.983)	(770.468)	(87.185)	(83.224)	(0.326)	(0.291)
Event ($\tau = 2$)	-0.093***	-0.026	-0.068***	-0.019	-0.080***	-0.039**	-4696.532***	-2513.903**	150.559	-184.902*	-0.646	-0.392
	(0.019)	(0.017)	(0.018)	(0.016)	(0.019)	(0.018)	(1217.117)	(1148.194)	(107.181)	(107.259)	(0.459)	(0.448)
Event ($\tau = 3$)	-0.100***	-0.027	-0.077***	-0.022	-0.093***	-0.046**	-5342.097***	-3020.839**	46.284	-286.561**	-0.233	0.053
	(0.022)	(0.019)	(0.020)	(0.019)	(0.021)	(0.019)	(1437.492)	(1363.962)	(132.072)	(128.683)	(0.421)	(0.404)
Event ($\tau = 4$)	-0.104***	-0.027	-0.081***	-0.023	-0.095***	-0.045**	-6618.016***	-4111.984**	61.003	-292.974**	-0.875**	-0.559
	(0.022)	(0.021)	(0.020)	(0.019)	(0.022)	(0.021)	(1894.722)	(1834.842)	(148.931)	(145.972)	(0.431)	(0.415)
N (Events)	62	62	62	62	62	62	62	62	62	62	62	62
R-Squared	0.997	0.997	0.997	0.997	0.998	0.998	0.998	0.998	0.986	0.985	0.897	0.890
Adj. R-Squared	0.992	0.993	0.993	0.993	0.993	0.994	0.996	0.996	0.964	0.961	0.731	0.713

Notes: Estimates of coefficients of δ_τ from the specification outlined in Equation 3.4.1. Panel A weights by microregion employment two years prior to the mass layoff event. The "Total" coefficients are estimated using data on all workers in the layoff and comparison regions, whereas the "Spillover" excludes workers from the layoff region who worked at the layoff region in $\tau = -1$ or $\tau = -2$. Standard errors are clustered at the matched comparison level.

3.5.2 Heterogeneity

We now turn to characterizing the heterogeneous regional impacts of mass layoffs based on layoff characteristics. Results are presented in Table 3.4, where each column runs the previous analysis on a subset of layoffs based on the severity of the layoff and the layoff firm's characteristics.

The first two columns under the heading "All" present the baseline results for mass layoffs on regional log(Employment). These results correspond to the weighted estimates displayed in the first two columns of Table 3.3. Events are also split by severity and whether or not they occurred in manufacturing, resulted in a closure, or took place in a large microregion.

As in the baseline results, we note that there are minimal pretrends in the regional heterogeneity analyses with only $\tau = -4$ demonstrating any statistical significance (this only shows up for severe layoffs or layoffs occurring in large microregions).

Table 3.4: Heterogeneity in Employment Effects of Mass Layoffs at the Local Labor Market Level

	All		Severe		Manufacturing		Non-Manufacturing		Closure		Large Region	
	Total	Spillover	Total	Spillover	Total	Spillover	Total	Spillover	Total	Spillover	Total	Spillover
Event Region ($\tau = -4$)	-0.025**	-0.002	-0.027*	0.020	-0.027	-0.013	-0.022	0.012	0.024	0.032	-0.034**	-0.021
	(0.012)	(0.018)	(0.014)	(0.016)	(0.019)	(0.028)	(0.014)	(0.016)	(0.059)	(0.054)	(0.014)	(0.019)
Event Region ($\tau = -3$)	-0.009	0.008	-0.009	0.020	-0.017	-0.007	0.000	0.025*	0.029	0.035	-0.015	-0.003
	(0.010)	(0.012)	(0.014)	(0.016)	(0.014)	(0.017)	(0.013)	(0.014)	(0.025)	(0.022)	(0.012)	(0.013)
Event Region ($\tau = -2$)	0.005	0.002	0.008	-0.001	0.006	0.001	0.004	0.002	0.024	0.024	0.000	0.000
	(0.008)	(0.008)	(0.011)	(0.011)	(0.013)	(0.013)	(0.009)	(0.009)	(0.009)	(0.011)	(0.010)	(0.010)
Event Region ($\tau = -1$)	-	-	-	-	-	-	-	-	-	-	-	-
Event Region ($\tau = 0$)	-0.033***	-0.003	-0.066***	-0.017	-0.013	0.012	-0.055***	-0.019*	-0.018	-0.010	-0.023*	-0.003
	(0.011)	(0.008)	(0.011)	(0.010)	(0.014)	(0.010)	(0.010)	(0.010)	(0.063)	(0.058)	(0.013)	(0.010)
Event Region ($\tau = 1$)	-0.055***	-0.018	-0.094***	-0.031	-0.043***	-0.009	-0.069***	-0.027	-0.048	-0.039	-0.044***	-0.019
	(0.013)	(0.011)	(0.023)	(0.021)	(0.016)	(0.012)	(0.021)	(0.021)	(0.103)	(0.099)	(0.014)	(0.013)
Event Region ($\tau = 2$)	-0.074***	-0.031*	-0.111***	-0.038	-0.060***	-0.022	-0.089***	-0.041	-0.074	-0.063	-0.061***	-0.034*
	(0.017)	(0.016)	(0.032)	(0.030)	(0.018)	(0.016)	(0.029)	(0.029)	(0.097)	(0.093)	(0.019)	(0.019)
Event Region ($\tau = 3$)	-0.082***	-0.036**	-0.125***	-0.047	-0.067***	-0.025	-0.099***	-0.047	-0.066	-0.056	-0.073***	-0.044**
	(0.017)	(0.017)	(0.031)	(0.030)	(0.020)	(0.019)	(0.030)	(0.030)	(0.100)	(0.096)	(0.020)	(0.020)
Event Region ($\tau = 4$)	-0.101***	-0.052**	-0.135***	-0.052*	-0.092***	-0.047*	-0.111***	-0.057*	-0.095	-0.084	-0.099***	-0.068***
	(0.019)	(0.020)	(0.030)	(0.030)	(0.023)	(0.027)	(0.032)	(0.032)	(0.116)	(0.112)	(0.023)	(0.023)
N (Events)	62	62	42	42	32	32	30	30	3	3	28	28
R-Squared	0.999	0.999	0.996	0.997	0.999	0.999	0.996	0.996	0.996	0.996	0.998	0.998
Adj. R-Squared	0.996	0.996	0.991	0.991	0.997	0.998	0.990	0.990	0.984	0.984	0.996	0.996

Notes: Estimates of coefficients of δ_τ from the specification outlined in Equation 3.4.1, broken out by groups of layoffs. Estimates in this table are effects on log-Employment and weighted by microregion employment two years prior to the mass layoff event. The "Total" coefficients are estimated using data on all workers in the layoff and comparison regions, whereas the "Spillover" excludes workers from the layoff region who worked at the layoff region in $\tau = -1$ or $\tau = -2$. "Severe" layoffs are defined as layoffs where displaced workers account for more than 2 percent of pre-layoff microregion employment. "Large Region" is defined as a layoff occurring in a microregion with more than 50,000 workers in the RAIS data. Standard errors are clustered at the matched comparison level.

Turning our focus to "Severe" layoffs, we first note that these account for two-thirds of all the layoff events defined in our sample. The total effect is immediate and apparent: in the year a layoff takes place, regional employment drops by approximately 6.6%; this grows with each passing years before ultimately getting to a 13.5% decline due to the mass layoff four years out. Here, spillover effects contribute to between one-fifth and one-third of the effects that we estimate. Comparing the severe spillover effects to the aggregate spillover effects, it appears that severe effects rather closely track to baseline results.

Studying the main and spillover effects of mass layoffs on manufacturing versus non-manufacturing firms allows us to test whether layoffs in the agricultural and resource extraction industries deliver different results. Perhaps surprisingly, we get somewhat results regardless of whether the layoff firm is a manufacturing firm or a non-manufacturing firm. For both kinds of layoffs, we find that there are strong, persistent declines in log(Employment) for microregions as a whole. Total effects are larger in magnitude for non-manufacturing sectors than for manufacturing. We do not document significant spillover effects in either of these sectors specifically until $\tau = 4$.

Finally, firm closures do not appear to be driving our aggregate results in any major way (and indeed, make up less than 5% of layoff events defined and identified in our sample) but large regions do. Layoff events in larger microregions track the total effects of all layoff events closely, but spillover effects in these areas are slightly more severe than the aggregate results

3.6 Worker-Level Analysis

We now turn to results from the worker-level specification outlined in Equation 3.4.2. The worker-level analysis complements the aggregate analysis by allowing us to investigate the incidence of the layoff shock. We can also look at person-level movement across

firms and geography, whereas the previous analysis treated the region as the unit of analysis. Examining heterogeneity in which workers are affected by the layoff shock will help us get a sense of the mechanisms by which local labor markets adjust to shocks. For example, if workers in industries that are closely related to the layoff industry are most affected, then this may indicate agglomeration economies. In contrast, workers from the non-tradeable sector are most affected, then this is evidence for declines in the demand for locally produced goods. Finally, if effects are concentrated on workers who are tightly connected across skill groups, then this may reflect declines in worker's outside options.

Table 3.5 shows results for total annual earnings (including zeros), total annual hours (including zeros), and log-wages for those working. All regressions control for individual fixed effects which account for compositional differences from individuals dropping out the sample. Column (1) shows the results splitting the event coefficient into directly affected layoff workers and the indirectly affected spillover workers. Figure 3.7 plots the effects for a variety of outcomes. Column (2) of Table 3.5 investigates heterogeneity in the spillover effect by running a regression with joint interactions between the spillover event coefficient and five sources of heterogeneity, which can be potentially correlated with each other. Table A4 shows this running each heterogeneity regression individually, which ignores potential correlation among the interactions.

In general, these results are much noisier and there are sometimes pretrends depending on the outcome, even with the rich battery of controls that we include. This is especially true for the workers who are directly affected by the layoff. However, to analyze the sources of spillover effects, we just need for the rest of the local labor market to be on a similar trend to the control group, which we often find is the case across outcomes.

Column (1) of Table 3.5 and Figure 3.7 shows that we tend to find statistically insignificant effects on most outcomes. The most notable exception is on wages, where

Panel (d) of Figure 3.7 shows declines of 1.4 percent four years after the mass layoff event for workers not involved in the layoff itself, relative to the individual's baseline in $\tau = -1$. This is mirrored in the overall earnings effects, which show a large R$522 (347) decline, although this is statistically insignificant. The other notable finding is that there does not appear to be a strong migration response, as shown in Panel (e) of Figure 3.7. When comparing to the overall displaced worker effect of R$1472, the spillover worker effect implies that 35 percent of the overall displaced worker effect is due to impacts across the local labor market. However, this does not account for reduced in-migration into a region, which likely rationalizes the results.

Column (2) examines heterogeneity in the incidence of the spillover effect. This shows that the overall earnings declines four years after the layoff are driven by workers (i) in the same city as the layoff firm (-575 with SE=363), (ii) in the industry as the layoff firm (-641 with SE=378), (iii) or in the non-tradeable industry (-720 with SE=395). Interestingly, there are little impacts on workers with similar skills – we find that a 10 percent increase in exposure to an a worker's skill and occupation group only causes a R$22 (40) loss in earnings. In general, we see small effects on hours across all sources of heterogeneity, although there are insignificant effects on both the directly affected layoff workers and the spillover workers.

Looking to log-wages, which more accurately capture the price of skill but may be more subject to selection issues, we see that impacts are primarily driven by workers in the non-tradeable sector. This implies a reduction of the price of skill for workers at these firms. This would be consistent with declines in local demand for locally produced consumer goods driving effects. When comparing to the overall displaced worker effect of 2.9 percent, the spillover worker effect implies that 48 percent of the overall displaced worker effect on log-wages is due to impacts across the local labor market. Taken together, these worker-level results broadly show strong evidence of declines in consumer demand, modest evidence for agglomeration economies, and little evidence for declines

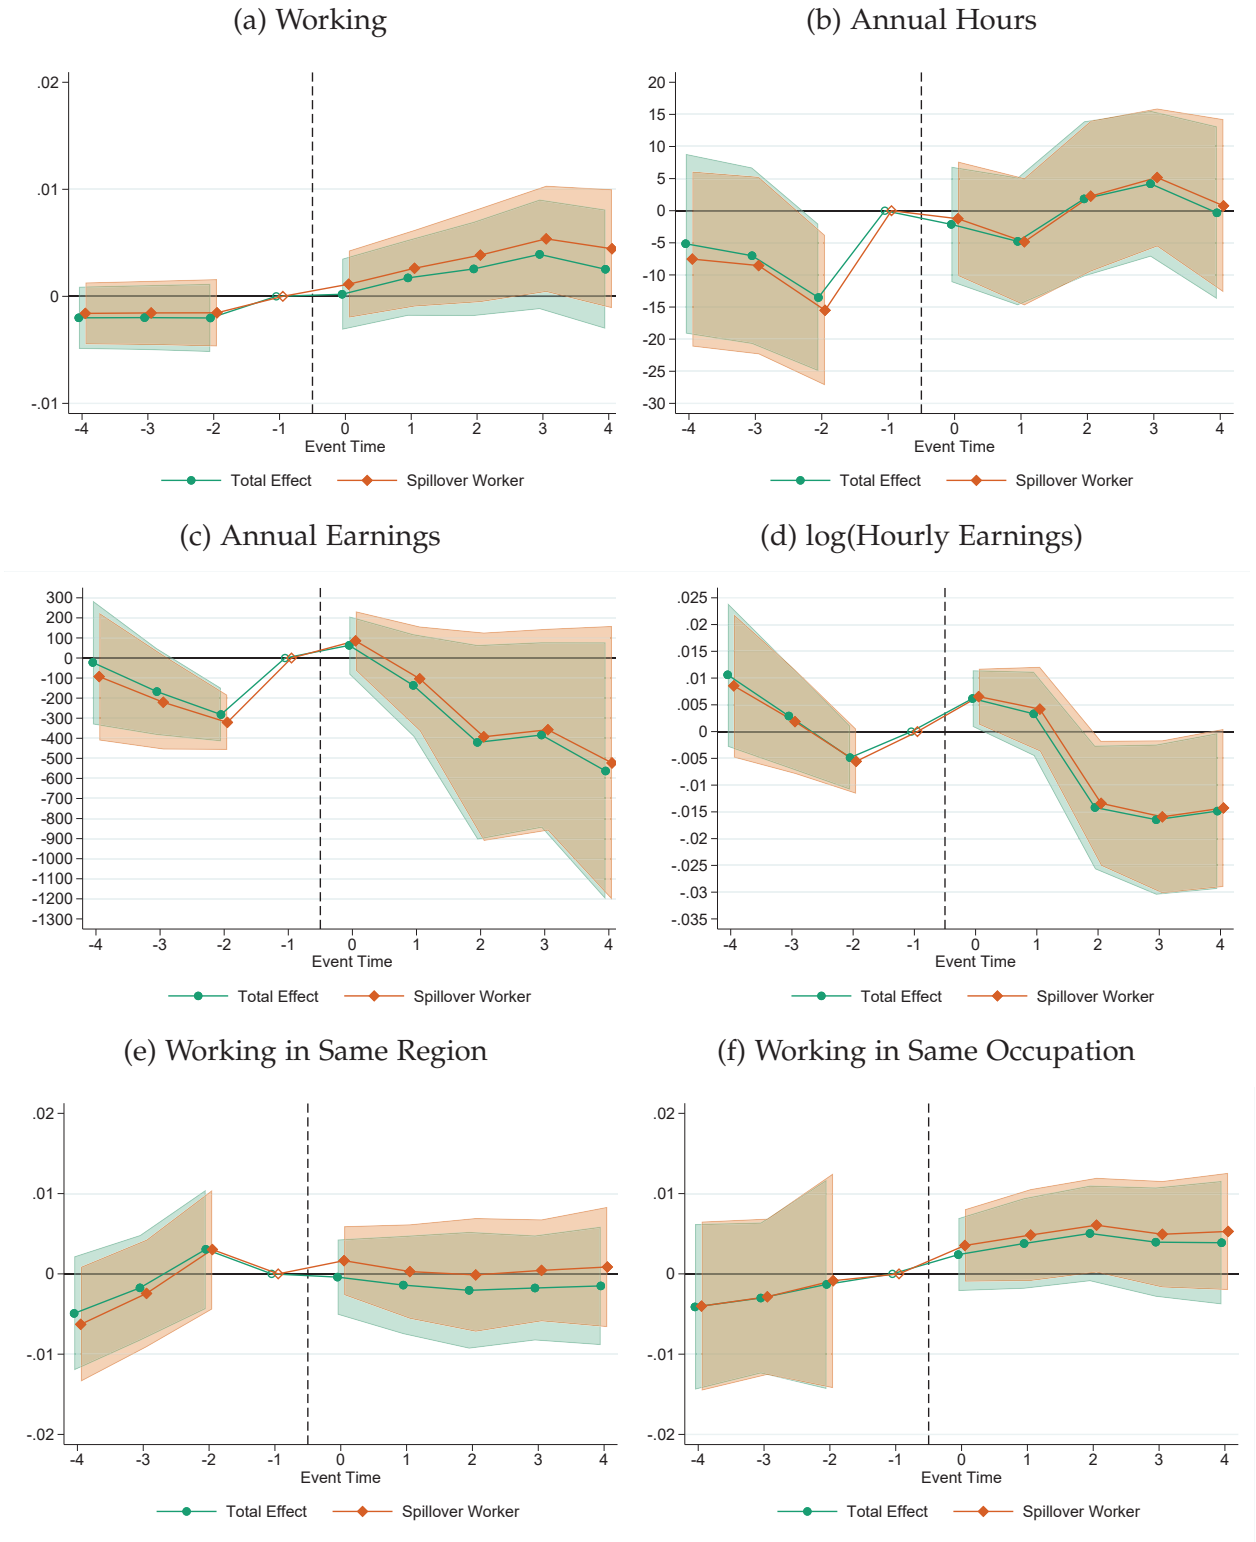
in worker bargaining power.

Table 3.5: Worker-Level Estimates and Incidence of Spillover Effects

	(1)		(2)						
	Layoff	Spillover	Layoff	Spillover	× Layoff City	× Layoff Industry	× Non-Tradeables	× Skill Exposure	× Firm Exposure
A. Annual Earnings									
Event ($\tau = -4$)	1579.843**	-92.572	1432.628*	97.914	108.837	-829.919**	-304.054	4.864	-1.500
	(690.876)	(162.343)	(725.949)	(323.091)	(370.887)	(351.862)	(219.214)	(5.461)	(22.228)
Event ($\tau = -3$)	1039.199	-220.540*	937.795	177.630	-103.632	-810.700	-513.402**	3.315	-2.165
	(748.792)	(120.143)	(756.333)	(188.063)	(160.320)	(500.884)	(193.279)	(5.867)	(30.352)
Event ($\tau = -2$)	599.187	-320.786***	553.886	-276.793	-271.938**	-184.609	49.298	1.753	-2.202
	(401.416)	(70.066)	(397.081)	(196.833)	(116.298)	(311.718)	(195.785)	(2.947)	(18.020)
Event ($\tau = -1$)	-	-	-	-	-	-	-	-	-
Event ($\tau = 0$)	-445.922	85.650	-505.487	228.963	-35.953	-476.299**	-195.188	3.200	1.108
	(284.624)	(75.364)	(303.677)	(171.166)	(138.487)	(228.991)	(212.746)	(2.966)	(18.594)
Event ($\tau = 1$)	-880.509***	-103.296	-1003.331***	432.458*	-179.152	-725.794**	-632.168**	-1.043	5.350
	(295.447)	(132.839)	(322.875)	(250.468)	(184.405)	(300.458)	(281.116)	(3.545)	(21.309)
Event ($\tau = 2$)	-1051.686***	-392.429	-1217.075***	204.564	-216.910	-888.487***	-688.107**	-1.355	17.082
	(377.364)	(264.553)	(443.089)	(261.061)	(361.787)	(303.109)	(340.226)	(3.357)	(27.028)
Event ($\tau = 3$)	-981.996**	-357.268	-1057.108**	266.936	-479.705	-665.114**	-650.106*	-1.484	15.059
	(448.859)	(256.256)	(522.177)	(322.768)	(365.273)	(267.184)	(339.817)	(3.850)	(24.435)
Event ($\tau = 4$)	-1472.021***	-522.971	-1520.573***	175.545	-575.738	-641.783*	-720.379*	-2.206	20.292
	(430.896)	(347.702)	(517.748)	(322.997)	(363.290)	(378.511)	(395.152)	(4.005)	(27.718)
N (Events)	62	62	62	62	62	62	62	62	62
N (Workers)	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620
B. Annual Hours									
Event ($\tau = -4$)	49.193	-7.526	48.600	-12.366	12.047*	-0.851	0.699	0.261	-4.173*
	(32.395)	(6.944)	(34.462)	(13.699)	(6.887)	(24.675)	(17.257)	(0.231)	(2.207)
Event ($\tau = -3$)	27.827	-8.543	24.135	-2.868	10.090	-6.427	-10.621	0.059	-3.591*
	(26.403)	(7.045)	(30.997)	(8.989)	(6.430)	(24.579)	(11.913)	(0.262)	(1.883)
Event ($\tau = -2$)	31.329*	-15.507**	33.714	-24.013**	6.995	11.355	7.900	0.070	-1.016
	(15.928)	(5.969)	(20.933)	(11.368)	(6.852)	(22.814)	(11.260)	(0.194)	(1.624)
Event ($\tau = -1$)	-	-	-	-	-	-	-	-	-
Event ($\tau = 0$)	-21.903	-1.247	-19.211	9.960	-0.546	-8.215	-17.042*	0.074	-0.866
	(21.094)	(4.544)	(19.983)	(10.397)	(5.828)	(14.424)	(10.075)	(0.190)	(1.255)
Event ($\tau = 1$)	-3.301	-4.841	-9.144	7.243	-2.551	-23.101	-13.593	-0.063	0.939
	(24.834)	(5.054)	(24.343)	(11.096)	(7.678)	(17.028)	(11.893)	(0.178)	(1.244)
Event ($\tau = 2$)	-7.858	2.265	-18.924	15.149	-1.717	-24.562	-12.084	-0.168	0.083
	(26.834)	(5.987)	(27.321)	(11.136)	(7.795)	(18.709)	(11.120)	(0.174)	(1.355)
Event ($\tau = 3$)	-17.910	5.170	-27.151	17.005	-7.254	-22.820	-9.299	-0.159	1.094
	(26.838)	(5.481)	(27.017)	(11.234)	(9.245)	(18.285)	(11.777)	(0.184)	(1.239)
Event ($\tau = 4$)	-25.118	0.760	-32.118	11.611	-5.410	-34.604	-11.147	-0.028	3.456***
	(24.998)	(6.875)	(26.248)	(13.421)	(8.007)	(23.877)	(12.925)	(0.188)	(1.260)
N (Events)	62	62	62	62	62	62	62	62	62
N (Workers)	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620
C. Log-Hourly Wage									
Event ($\tau = -4$)	0.058***	0.009	0.053***	0.015	0.001	-0.017	-0.009	0.000	-0.000
	(0.013)	(0.007)	(0.018)	(0.010)	(0.009)	(0.013)	(0.008)	(0.000)	(0.001)
Event ($\tau = -3$)	0.027**	0.002	0.023	0.012*	-0.002	-0.013	-0.012*	-0.000	-0.001
	(0.013)	(0.005)	(0.017)	(0.006)	(0.006)	(0.013)	(0.006)	(0.000)	(0.001)
Event ($\tau = -2$)	0.011	-0.006*	0.012	0.000	-0.008**	0.003	-0.005	-0.000	-0.001**
	(0.009)	(0.003)	(0.010)	(0.006)	(0.004)	(0.008)	(0.005)	(0.000)	(0.000)
Event ($\tau = -1$)	-	-	-	-	-	-	-	-	-
Event ($\tau = 0$)	-0.003	0.007**	-0.005	0.005	-0.010***	-0.012	0.006	0.000**	-0.001
	(0.009)	(0.003)	(0.011)	(0.004)	(0.004)	(0.008)	(0.006)	(0.000)	(0.001)
Event ($\tau = 1$)	-0.018*	0.004	-0.017	0.008	-0.001	-0.008	-0.006	0.000**	-0.001
	(0.010)	(0.004)	(0.011)	(0.006)	(0.008)	(0.009)	(0.006)	(0.000)	(0.001)
Event ($\tau = 2$)	-0.033***	-0.013**	-0.031**	-0.002	0.003	-0.005	-0.017**	0.000	-0.001*
	(0.011)	(0.006)	(0.012)	(0.006)	(0.009)	(0.012)	(0.007)	(0.000)	(0.001)
Event ($\tau = 3$)	-0.030***	-0.016**	-0.024**	-0.003	0.002	0.002	-0.021***	0.000	-0.002**
	(0.010)	(0.007)	(0.011)	(0.007)	(0.009)	(0.012)	(0.008)	(0.000)	(0.001)
Event ($\tau = 4$)	-0.029**	-0.014*	-0.024*	-0.001	-0.006	0.009	-0.017*	-0.000	-0.001*
	(0.011)	(0.008)	(0.013)	(0.008)	(0.011)	(0.013)	(0.009)	(0.000)	(0.001)
N (Events)	62	62	62	62	62	62	62	62	62
N (Workers)	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948

Notes: Estimates of coefficients from the specification outlined in Equation 3.4.2. Sample consists of all workers in a layoff or comparison microregion who are employed in the formal sector in $\tau = -2$. (1) and (2) are separate regressions with each panel showing a different outcome and columns showing different interactions with the event coefficients. (1) shows the main layoff and spillover worker estimates and (2) considers exhaustive interactions with the spillover worker. See Table A4 for these interaction effects separately included. Layoff workers are defined as workers in the layoff region who are employed at the mass layoff firm in any of the 4 years prior to the layoff; spillover workers are defined as the complement for workers in the layoff microregion sample. Standard errors are clustered at the matched comparison levels.

Figure 3.7: Effects of Mass Layoffs on Worker Outcomes



Notes: Plots estimates of coefficients from the specification outlined in Equation 3.4.2. Sample consists of all workers in a layoff or comparison microregion who are employed in the formal sector in $\tau = -2$. The orange line plots estimates of $\delta_{\tau}^{Spillover}$, the gray line plots estimates of δ_{τ}^{Layoff}, and the teal line plots estimates that combine the two coeffiicents. See column (1) of Table A4 for reported point estimates. Layoff workers are defined as workers in the layoff region who are employed at the mass layoff firm in any of the 4 years prior to the layoff; spillover workers are defined as the complement for workers in the layoff microregion sample. Standard errors are clustered at the matched comparison level.

3.7 Conclusion

This paper studies how economic shocks propagate through a local labor market by studying the spillover effects of mass layoff events. We study this in the context of a developing country that has strong labor market institutions, a large informal sector, and is subject to different sorts of shocks than more developed countries. Our empirical approach extends Gathmann et al. [2020], identifying large flagship layoff events and matching to a comparison region on a similar economic trajectory and with a similar demographic makeup and industrial composition. These comparison regions should be subject to similar aggregate shocks but do not experience a large plant closure or layoff. We build on Gathmann et al. [2020] by showing heterogeneity in the spillover effect using a worker-level analysis.

Our main finding is that mass layoff events permanently scar local labor markets with evidence of spillover effects. Specifically, we estimate that a mass layoff results in a 10.1 percent reduction in local employment 4 years after a layoff event with 5.2 percent due to reductions in employment for workers not at the layoff firm. This indicates a multiplier effect of about 0.5, with larger effects when looking at the total wage bill. We find some effects on earnings, especially for workers in the non-tradeable sector where log-wages decline in addition to participation. While much more mixed, these worker-level results broadly show strong evidence of declines in consumer demand, modest evidence for agglomeration economies, and limited evidence for declines in worker bargaining power.

Future work will seek to: (a) complement the migration results by using population-level measure of regional employment , (b) use a more continuous measure of industry exposure to identify agglomeration effects, (c) more precisely investigate the role of labor market institutions by stratifying regions based on measures of pre-layoff informality. In addition, we think that the interpretation of the estimates would benefit greatly from a theoretical framework that accounts for direct and indirect effects on workers and

expected heterogeneity from multiple channels.

Bibliography

Jorge Alvarez, Felipe Benguria, Niklas Engbom, and Christian Moser. Firms and defline in earning inequality in brazil. *American Economic Journal: Macroeconomics*, 10(1):149–189, 2018.

Edward J. Amadeo, Indermit S. Gill, and Marcelo C. Neri. Brazil: The pressure points in labor regulation. 2000.

Alina Arefeva, Morris A. Davis, Andra C. Ghent, and Minseon Park. The effect of capital gains taxes on business creation and employment: The case of opportunity zones. *Working Paper*, 2024.

Ricardo Barros, Mirela de Carvalho, Samuel Franco, and Rosane Mendonça. Markets, the state, and the dynamics of inequality: The case of brazil. *UNDP Discussion Paper*, 2010.

Paul Beaudry, David A. Green, and Benjamin Sand. Does industrial composition matter for wages? a test of search and bargaining theory. *Econometrica*, 80(3):1063–1104, 2012.

Oliver Jean Blanchard and Lawrence F. Katz. Regional evolutions. *Brookings Papers on Economic Activity*, 1992(1):1–75, 1992.

John Bound and Harry J. Holzer. Demand shifts, population adjustments, and labor market outcomes during the 1980s. *Journal of Labor Economics*, 18(1):20–54, 2000.

Diogo G.C. Britto, Paolo Pinottti, and Breno Sampaio. The effect of job loss and unemployment insurance on crime in brazil. *Econometrica*, 90(4):1393–1423, 2022.

Matias Busso, Jesse Gregory, and Patrick Kline. Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review*, 103(2):897–947, 2013.

Sydnee Caldwell and Nikolaj Harmon. Outside options, bargaining, and wages: Evidence from coworker networks. *Working Paper*, 2019.

Brantly Callaway and Pedro H.C. Sant'Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, pages 200–230, 2021.

Kenneth A. Couch and Dana W. Placzek. Earnings losses of displaced workers revisited. *American Economic Review*, 100(1):572–589, 2010.

Steven J. Davis and Till M. von Wachter. Recessions and the cost of job loss. *National Bureau of Economic Research Working Paper no. 17638*, 2011.

Ellora Derenoncourt, François Gérard, Lorenzo Lagos, and Claire Montialoux. Racial inequality, minimum wage spillovers, and the informal sector. *Working Paper*, 2021.

Rafael Dix-Carneiro and Brian K. Kovak. Trade liberalization and regional dynamics. *American Economic Review*, 107(10):2908–2946, 2017.

Arindrajit Dube and Attila Lindner. City limits: What do local-area minimum wages

do? *Journal of Economic Perspectives*, 31(1):27–50, 2021.

Marcus Eliason and Donald Storrie. Lasting or latent scares? swedish evidence on the long-term effects of job displacement. *Journal of Labor Economics*, 24(4):831–856, 2006.

Niklas Engbom and Christian Moser. Earnings inequality and the minimum wage: Evidence from brazil. *American Economic Review*, 112(12):3803–3847, 2022.

Giulia Faggio and Henry Overman. The effect of public sector employment on local labour markets. *Journal of Urban Economics*, 79:91–107, 2014.

Aaron Flaaen, Matthew D. Shapiro, and Isaac Sorkin. Reconsidering the consequences of worker displacements: Firm versus worker perspective. *American Economic Journal: Macroeconomics*, 11(2):193–227, 2019.

Chao Fu and Jesse Gregory. Estimation of an equilibrium model with externalities: Post-disaster neighborhood rebuilding. *Econometrica*, 87(2):387–421, 2019.

Christina Gathmann, Ines Helm, and Uta Schoenberg. Spillover effects of mass layoffs. *Journal of the European Economic Association*, 18(1):427–468, 2020.

Austan Goolsbee and Chad Syverson. Monpsony power in higher education: A tale of two tracks. *National Bureau of Economic Research working paper no. 26070*, 2019.

Michael Greenstone, Richard Hornbeck, and Enrico Moretti. Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. *Journal of Political Economy*, 118(3):536–598, 2010.

Instituto de Pesquisa Economica Aplicada. Salário mínimo vigente (current minimum wage), 1985–2018. http://www.ipeadata.gov.br/ExibeSerie.aspx?serid=1739471028, 2022. Accessed: July 22, 2021.

Louis Jacobson, Robert J. LaLonde, and Daniel Sullivan. Earnings losses of displaced workers. *American Economic Review*, 83(4):685–709, 1993.

Ekaterina Jardim, Mark C. Long, Robert Plotnick, Emma van Inwegen, Jacob Vigdor, and Hilary Wething. Minimum-wage increases and low-wage employment: Evidence from seattle. *American Economic Journal: Economic Policy*, 14(2):263–314, 2022.

Gregor Jarosch. Searching for job security and the consequences of job loss. *Econometrica*, 1(3):903–942, 2023.

Patrick Kline and Enrico Moretti. Local economic development, agglomeration economies, and the big push: 100 years of evidence from the tennessee valleyy authority. *The Quarterly Journal of Economics*, 129(1):275–331, 2014.

Marta Lachowska, Alexandre Mas, and Stephen A. Woodbury. Sources of displaced workers' long-term earnings losses. *American Economic Review*, 110(10):3231–3266, 2020.

Costas Meghir, Renata Narita, and Jean-Marc Robin. Wages and informality in developing countries. *American Economic Review*, 105(4):1509–1546, 2015.

Roozbeh Moghadam. Spillover effects of mass layoffs on neighboring firms. *Working Paper*, 2023.

Enrico Moretti. Workers' education, spillovers, and productivity: Evidence from plant-level production functions. *American Economic Review*, 94(3):656–690, 2004.

Enrico Moretti. Local multipliers. *American Economic Review: Papers & Proceedings*, 100: 1–7, 2010.

Enrico Moretti and Per Thulin. Local multipliers and human capital in the united states and sweden. *Industrial and Corporate Change*, 22(1):339–362, 2013.

Christopher J. Ruhm. Are workers permanently scarred by job displacements? *American Economic Review*, 81(1):319–324, 1991.

Johannes F. Schmieder, Till von Wachter, and Jörg Heining. The costs of job displacement over the business cycle and its sources: Evidence from germany. *American Economic Review*, 113(5):1208–1254, 2023.

Garima Sharma. Monopsony and gender. *Working Paper*, 2023.

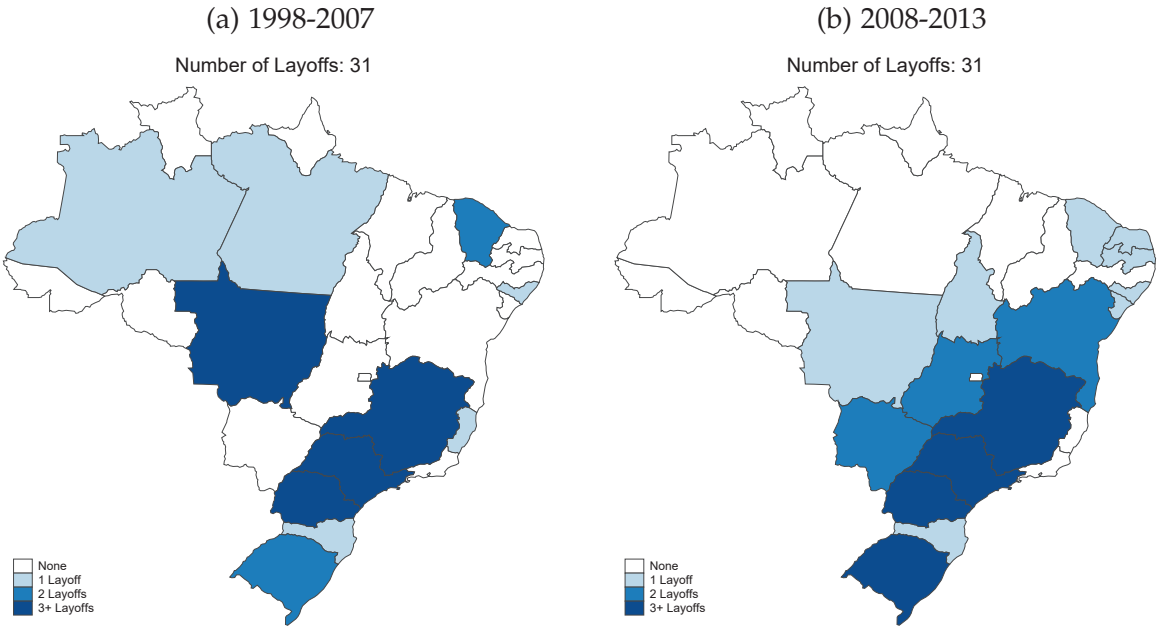
Louise Story. As companies seek tax deals, governments pay a high price. *The New York Times*, 2012. URL https://archive.nytimes.com/www.nytimes.com/interactive/2012/12/01/us/government-incentives.html.

Christiane Szerman. The employee costs of corporate debarment in public procurement. *American Economic Journal: Applied Economics*, 15(1):411–441, 2023.

Gabriel Ulyssea. Firms, inormality, and development: Theory and evidence from brazil. *American Economic Review*, 108(8):2015–2047, 2018.

Appendix A: Additional Figures and Tables

Figure A1: Locations of Mass Layoff Events by Time Period



Notes: This figure shows the spatial distribution of flagship layoff events using the baseline definition, aggregating from microregions to the the level of federal states and breaking out by time period. See Figure 3.2 for the spatial distribution by the industry of the layoff event.

Table A1: Flagship Layoff Industries

Industry	Broad Industry	Number of Layoffs
Cultivation of Sugarcane	Agriculture and Resources	14
Manufacture of Leather Shoes	Manufacturing	6
Sugar Mills	Manufacturing	6
Cultivation of Coffee	Agriculture and Resources	5
Alcohol Production	Manufacturing	3
Cultivation of Citrus Fruits	Agriculture and Resources	3
Cultivation of Other Permanent Crop Products	Agriculture and Resources	3
Slaughter of Cattle, Preparation of Meat Products	Manufacturing	3
Cultivation of Other Temporary Crop Products	Agriculture and Resources	2
Manufacture of Laminated Wood and Plywood Sheets, Pressed or Agglomerated Wood	Manufacturing	2
Processing, Milling, and Preparation of Other Plant-Origin Products	Manufacturing	2
Tobacco Product Manufacturing	Manufacturing	2
Construction and Assembly of Aircraft	Manufacturing	1
Cultivation of Soybeans	Agriculture and Resources	1
Cultivation of Vegetables, Legumes, and Other Horticultural Products	Agriculture and Resources	1
Livestock Related Services Activities, Except Veterinary Activities	Agriculture and Resources	1
Manufacture of Basic Electronic Material	Manufacturing	1
Manufacture of Electrical Material for Vehicles Except Batteries	Manufacturing	1
Manufacture of Furniture Predominantly Made of Wood	Manufacturing	1
Manufacture of Plastic Shoes	Manufacturing	1
Manufacture of Stoves, Refrigerators, and Washing and Drying Machines for Household Use	Manufacturing	1
Production of Long Steel Products	Manufacturing	1
Slaughter of Poultry and Other Small Animals and Preparation of Meat Products	Manufacturing	1

Notes: This table report the number of mass layoff events in each 5-Digit industry as defined by the first edition of the National Classification of Economic Activities (CNAE 1.0) from the Instituto Brasileiro de Geografia e Estatística (IBGE). Broad industry relates these definitions to the broad groupings reported in other results reported below.

Table A2: Stated Reason for Job Separation

	All Displaced Workers	Long Tenure Displaced Workers
No Stated Reason (%)	5.2	7.2
Fired: with Cause (%)	2.7	4.0
Fired: Other Reason (%)	70.9	76.0
Indirect: with Cause (%)	0.1	0.0
Indirect: Other Reason (%)	18.4	7.8
Indirect: Reassignment (%)	2.5	4.9
Death or Retirement (%)	0.2	0.2
Number of Workers	5,666	554

Notes: This table reports the reason for job separation for workers who are employed at the layoff firm in the year prior to the event but have no employment at the firm in the subsequent (i.e., layoff) year. Long tenure workers are defined as workers who have been continuously employed at the same firm for 3 years in the year prior to the event. See Table 3.1 for other characteristics of these layoff workers. Results are based on a 5 percent sample of the RAIS data.

Table A3: Robustness of Local Labor Market Analysis

	Baseline		Alt. Spec. 1		Alt. Defn. 1		Alt. Defn. 2		Alt. Defn. 3		Alt. Matching	
	Total	Spillover	Total	Spillover	Total	Spillover	Total	Spillover	Total	Spillover	Total	Spillover
A. Weighted												
Event Region ($\tau = -4$)	-0.025**	-0.002	-0.014	0.010	-0.035**	0.003	-0.018	0.005	-0.012	0.004	-0.018	0.004
	(0.012)	(0.018)	(0.012)	(0.016)	(0.014)	(0.014)	(0.011)	(0.016)	(0.011)	(0.011)	(0.012)	(0.018)
Event Region ($\tau = -3$)	-0.009	0.008	-0.000	0.017	-0.019	0.006	-0.007	0.010	-0.009	0.003	-0.016	0.001
	(0.010)	(0.012)	(0.009)	(0.011)	(0.012)	(0.013)	(0.010)	(0.012)	(0.008)	(0.009)	(0.011)	(0.012)
Event Region ($\tau = -2$)	0.005	0.002	0.010*	0.005	-0.002	-0.006	0.003	-0.001	-0.007	-0.009	0.005	0.002
	(0.008)	(0.008)	(0.006)	(0.006)	(0.009)	(0.008)	(0.008)	(0.008)	(0.007)	(0.007)	(0.008)	(0.008)
Event Region ($\tau = -1$)	-	-	-	-	-	-	-	-	-	-	-	-
Event Region ($\tau = 0$)	-0.033***	-0.003	-0.035***	-0.003	-0.049***	-0.016	-0.043***	-0.013	-0.037***	-0.017**	-0.039***	-0.009
	(0.011)	(0.008)	(0.010)	(0.008)	(0.010)	(0.009)	(0.012)	(0.009)	(0.007)	(0.007)	(0.012)	(0.009)
Event Region ($\tau = 1$)	-0.055***	-0.018	-0.052***	-0.013	-0.066***	-0.021	-0.070***	-0.033***	-0.041***	-0.016	-0.055***	-0.018
	(0.013)	(0.011)	(0.011)	(0.009)	(0.017)	(0.016)	(0.013)	(0.011)	(0.011)	(0.011)	(0.013)	(0.012)
Event Region ($\tau = 2$)	-0.074***	-0.031*	-0.065***	-0.022	-0.081***	-0.030	-0.089***	-0.047***	-0.045***	-0.018	-0.064***	-0.021
	(0.017)	(0.016)	(0.014)	(0.013)	(0.023)	(0.021)	(0.016)	(0.015)	(0.012)	(0.012)	(0.018)	(0.017)
Event Region ($\tau = 3$)	-0.082***	-0.036**	-0.075***	-0.029*	-0.086***	-0.031	-0.100***	-0.055***	-0.043***	-0.013	-0.070***	-0.023
	(0.017)	(0.017)	(0.015)	(0.015)	(0.025)	(0.023)	(0.017)	(0.016)	(0.013)	(0.013)	(0.020)	(0.019)
Event Region ($\tau = 4$)	-0.101***	-0.052**	-0.089***	-0.041**	-0.094***	-0.035	-0.116***	-0.068***	-0.054***	-0.023	-0.073***	-0.024
	(0.019)	(0.020)	(0.018)	(0.019)	(0.027)	(0.025)	(0.017)	(0.018)	(0.015)	(0.015)	(0.022)	(0.024)
B. Unweighted												
Event Region ($\tau = -4$)	-0.012	0.033**	-0.007	0.038**	-0.022	0.032*	0.003	0.049***	0.003	0.026**	-0.003	0.042***
	(0.014)	(0.015)	(0.015)	(0.015)	(0.016)	(0.017)	(0.012)	(0.014)	(0.011)	(0.011)	(0.012)	(0.013)
Event Region ($\tau = -3$)	-0.003	0.027**	0.003	0.032***	-0.015	0.020	0.007	0.033***	-0.001	0.015	-0.015	0.015
	(0.012)	(0.013)	(0.011)	(0.011)	(0.014)	(0.014)	(0.011)	(0.011)	(0.009)	(0.009)	(0.013)	(0.013)
Event Region ($\tau = -2$)	0.008	0.000	0.011	0.003	0.003	-0.005	0.009	0.001	0.000	-0.003	0.002	-0.005
	(0.010)	(0.010)	(0.007)	(0.007)	(0.010)	(0.010)	(0.009)	(0.009)	(0.008)	(0.008)	(0.009)	(0.009)
Event Region ($\tau = -1$)	-	-	-	-	-	-	-	-	-	-	-	-
Event Region ($\tau = 0$)	-0.052***	-0.005	-0.055***	-0.007	-0.060***	-0.011	-0.060***	-0.010	-0.037***	-0.006	-0.059***	-0.013
	(0.009)	(0.008)	(0.008)	(0.007)	(0.011)	(0.010)	(0.008)	(0.008)	(0.007)	(0.007)	(0.010)	(0.010)
Event Region ($\tau = 1$)	-0.073***	-0.014	-0.070***	-0.011	-0.077***	-0.014	-0.084***	-0.021	-0.043***	-0.006	-0.071***	-0.013
	(0.015)	(0.014)	(0.013)	(0.011)	(0.015)	(0.014)	(0.012)	(0.013)	(0.011)	(0.011)	(0.016)	(0.015)
Event Region ($\tau = 2$)	-0.093***	-0.026	-0.083***	-0.016	-0.100***	-0.029	-0.103***	-0.033*	-0.046***	-0.005	-0.079***	-0.012
	(0.019)	(0.017)	(0.017)	(0.015)	(0.020)	(0.018)	(0.016)	(0.017)	(0.012)	(0.012)	(0.021)	(0.020)
Event Region ($\tau = 3$)	-0.100***	-0.027	-0.095***	-0.023	-0.109***	-0.032	-0.116***	-0.040*	-0.035**	0.010	-0.084***	-0.012
	(0.022)	(0.019)	(0.020)	(0.017)	(0.023)	(0.021)	(0.019)	(0.021)	(0.015)	(0.015)	(0.024)	(0.022)
Event Region ($\tau = 4$)	-0.104***	-0.027	-0.102***	-0.026	-0.117***	-0.036	-0.124***	-0.045*	-0.042**	0.004	-0.079***	-0.001
	(0.022)	(0.021)	(0.023)	(0.021)	(0.024)	(0.022)	(0.020)	(0.023)	(0.018)	(0.018)	(0.024)	(0.023)
N (Events)	62	62	62	62	68	68	91	91	81	81	62	62
R-Squared	0.997	0.997	0.994	0.994	0.995	0.996	0.997	0.997	0.998	0.998	0.997	0.997
Adj. R-Squared	0.992	0.993	0.993	0.993	0.988	0.989	0.993	0.993	0.994	0.994	0.992	0.993

Notes: Estimates of coefficients of δ_τ from the specification outlined in Equation 3.4.1. Panel A weights by microregion employment two years prior to the mass layoff event. The "Total" coefficients are estimated using data on all workers in the layoff and comparison regions, whereas the "Spillover" excludes workers from the layoff region who worked at the layoff region in $\tau = -1$ or $\tau = -2$. "Alt. Spec. 1" uses the baseline layoff definition and matching approach, but controls for a year effect across all comparisons instead of time-by-comparison set interactions. "Alt. Defn. 1" alters the baseline layoff definition by dropping the restriction on pre-layoff employment growth at the layoff firm and making the restriction concerning post-layoff firm expansions to two years after the layoff (which is the definition Gathmann et al. [2020] use in Germany). "Alt. Defn. 2" alters the baseline layoff definition by using all layoffs of at least 200 displaced workers at firms whose pre-layoff employment makes up at least 0.1% of the municipality. "Alt. Defn. 3" alters the baseline layoff definition to use all layoffs of at least 200 displaced workers at firms whose pre-layoff employment makes up at least 0.1% of the municipality, adds the pre-layoff employment growth restriction at the layoff firm to apply to years 3 and 4 prior to a layoff (employment growth must be less than 50% of the layoff size in each year), and drops the restriction concerning large single firm expansions post-layoff. "Alt. Matching" uses the baseline matching definition but matches on a propensity score that uses the same matching variables with linear effects, a year trend, own-quadratic effects, interactions with calendar time, total employment, and average earnings; interactions between total employment and averages earnings in $\tau = -4$ and $\tau = -2$. Standard errors are clustered at the matched comparison level in all regressions.

Table A4: Worker-Level Estimates and Incidence of Spillover Effects: Separate Interactions

	(1)	(2)		(3)			(4)		(5)	
	Spillover	Spillover	× Layoff City	Spillover	× Layoff Industry	× Non-Tradeables	Spillover	× Skill Exposure	Spillover	× Firm Exposure
A. Annual Earnings										
Event ($\tau = -4$)	-92.572	-92.572	108.837	97.914	-829.919**	-304.054	97.914	4.864	97.914	-1.500
	(162.343)	(162.343)	(370.887)	(323.091)	(351.862)	(219.214)	(323.091)	(5.461)	(323.091)	(22.228)
Event ($\tau = -3$)	-220.540*	-220.540*	-103.632	177.630	-810.700	-513.402**	177.630	3.315	177.630	-2.165
	(120.143)	(120.143)	(160.320)	(188.063)	(500.884)	(193.279)	(188.063)	(5.867)	(188.063)	(30.352)
Event ($\tau = -2$)	-320.786***	-320.786***	-271.938**	-276.793	-184.609	49.298	-276.793	1.753	-276.793	-2.202
	(70.066)	(70.066)	(116.298)	(196.833)	(311.718)	(195.785)	(196.833)	(2.947)	(196.833)	(18.020)
Event ($\tau = -1$)	-	-	-	-	-	-	-	-	-	-
Event ($\tau = 0$)	85.650	85.650	-35.953	228.963	-476.299**	-195.188	228.963	3.200	228.963	1.108
	(75.364)	(75.364)	(138.487)	(171.166)	(228.991)	(212.746)	(171.166)	(2.966)	(171.166)	(18.594)
Event ($\tau = 1$)	-103.296	-103.296	-179.152	432.458*	-725.794**	-632.168**	432.458*	-1.043	432.458*	5.350
	(132.839)	(132.839)	(184.405)	(250.468)	(300.458)	(281.116)	(250.468)	(3.545)	(250.468)	(21.309)
Event ($\tau = 2$)	-392.429	-392.429	-216.910	204.564	-888.487***	-688.107**	204.564	-1.355	204.564	17.082
	(264.553)	(264.553)	(361.787)	(261.061)	(303.109)	(340.226)	(261.061)	(3.357)	(261.061)	(27.028)
Event ($\tau = 3$)	-357.268	-357.268	-479.705	266.936	-665.114**	-650.106*	266.936	-1.484	266.936	15.059
	(256.256)	(256.256)	(365.273)	(322.768)	(267.184)	(339.817)	(322.768)	(3.850)	(322.768)	(24.435)
Event ($\tau = 4$)	-522.971	-522.971	-575.738	175.545	-641.783*	-720.379*	175.545	-2.206	175.545	20.292
	(347.702)	(347.702)	(363.290)	(322.997)	(378.511)	(395.152)	(322.997)	(4.005)	(322.997)	(27.718)
N (Events)	62	62	62	62	62	62	62	62	62	62
N (Workers)	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620
B. Annual Hours										
Event ($\tau = -4$)	-7.526	-7.526	12.047*	-12.366	-0.851	0.699	-12.366	0.261	-12.366	-4.173*
	(6.944)	(6.944)	(6.887)	(13.699)	(24.675)	(17.257)	(13.699)	(0.231)	(13.699)	(2.207)
Event ($\tau = -3$)	-8.543	-8.543	10.090	-2.868	-6.427	-10.621	-2.868	0.059	-2.868	-3.591*
	(7.045)	(7.045)	(6.430)	(8.989)	(24.579)	(11.913)	(8.989)	(0.262)	(8.989)	(1.883)
Event ($\tau = -2$)	-15.507**	-15.507**	6.995	-24.013**	11.355	7.900	-24.013**	0.070	-24.013**	-1.016
	(5.969)	(5.969)	(6.852)	(11.368)	(22.814)	(11.260)	(11.368)	(0.194)	(11.368)	(1.624)
Event ($\tau = -1$)	-	-	-	-	-	-	-	-	-	-
Event ($\tau = 0$)	-1.247	-1.247	-0.546	9.960	-8.215	-17.042*	9.960	0.074	9.960	-0.866
	(4.544)	(4.544)	(5.828)	(10.397)	(14.424)	(10.075)	(10.397)	(0.190)	(10.397)	(1.255)
Event ($\tau = 1$)	-4.841	-4.841	-2.551	7.243	-23.101	-13.593	7.243	-0.063	7.243	0.939
	(5.054)	(5.054)	(7.678)	(11.096)	(17.028)	(11.893)	(11.096)	(0.178)	(11.096)	(1.244)
Event ($\tau = 2$)	2.265	2.265	-1.717	15.149	-24.562	-12.084	15.149	-0.168	15.149	0.083
	(5.987)	(5.987)	(7.795)	(11.136)	(18.709)	(11.120)	(11.136)	(0.174)	(11.136)	(1.355)
Event ($\tau = 3$)	5.170	5.170	-7.254	17.005	-22.820	-9.299	17.005	-0.159	17.005	1.094
	(5.481)	(5.481)	(9.245)	(11.234)	(18.285)	(11.777)	(11.234)	(0.184)	(11.234)	(1.239)
Event ($\tau = 4$)	0.760	0.760	-5.410	11.611	-34.604	-11.147	11.611	-0.028	11.611	3.456***
	(6.875)	(6.875)	(8.007)	(13.421)	(23.877)	(12.925)	(13.421)	(0.188)	(13.421)	(1.260)
N (Events)	62	62	62	62	62	62	62	62	62	62
N (Workers)	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620	1,838,620
C. Log-Hourly Wage										
Event ($\tau = -4$)	0.009	0.009	0.001	0.015	-0.017	-0.009	0.015	0.000	0.015	-0.000
	(0.007)	(0.007)	(0.009)	(0.010)	(0.013)	(0.008)	(0.010)	(0.000)	(0.010)	(0.001)
Event ($\tau = -3$)	0.002	0.002	-0.002	0.012*	-0.013	-0.012*	0.012*	-0.000	0.012*	-0.001
	(0.005)	(0.005)	(0.006)	(0.006)	(0.013)	(0.006)	(0.006)	(0.000)	(0.006)	(0.001)
Event ($\tau = -2$)	-0.006*	-0.006*	-0.008**	0.000	0.003	-0.005	0.000	-0.000	0.000	-0.001**
	(0.003)	(0.003)	(0.004)	(0.006)	(0.008)	(0.005)	(0.006)	(0.000)	(0.006)	(0.000)
Event ($\tau = -1$)	-	-	-	-	-	-	-	-	-	-
Event ($\tau = 0$)	0.007**	0.007**	-0.010***	0.005	-0.012	0.006	0.005	0.000**	0.005	-0.001
	(0.003)	(0.003)	(0.004)	(0.004)	(0.008)	(0.006)	(0.004)	(0.000)	(0.004)	(0.001)
Event ($\tau = 1$)	0.004	0.004	-0.001	0.008	-0.008	-0.006	0.008	0.000**	0.008	-0.001
	(0.004)	(0.004)	(0.008)	(0.006)	(0.009)	(0.006)	(0.006)	(0.000)	(0.006)	(0.001)
Event ($\tau = 2$)	-0.013**	-0.013**	0.003	-0.002	-0.005	-0.017**	-0.002	0.000	-0.002	-0.001*
	(0.006)	(0.006)	(0.009)	(0.006)	(0.012)	(0.007)	(0.006)	(0.000)	(0.006)	(0.001)
Event ($\tau = 3$)	-0.016**	-0.016**	0.002	-0.003	0.002	-0.021***	-0.003	0.000	-0.003	-0.002**
	(0.007)	(0.007)	(0.009)	(0.007)	(0.012)	(0.008)	(0.007)	(0.000)	(0.007)	(0.001)
Event ($\tau = 4$)	-0.014*	-0.014*	-0.006	-0.001	0.009	-0.017*	-0.001	-0.000	-0.001	-0.001*
	(0.008)	(0.008)	(0.011)	(0.008)	(0.013)	(0.009)	(0.008)	(0.000)	(0.008)	(0.001)
N (Events)	62	62	62	62	62	62	62	62	62	62
N (Workers)	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948	1,611,948

Notes: Estimates of coefficients from the specification outlined in Equation 3.4.2. Sample consists of all workers in a layoff or comparison microregion who are employed in the formal sector in $\tau = -2$. (1), (2), (3), (4), and (5) are separate regressions with each panel showing a different outcome and columns showing different interactions with the event coefficients. Unlike Table 3.5, this shows influence of each interaction separately rather than jointly. Point estimates for layoff workers are excluded from this table which are reported in Table 3.5. Layoff workers are defined as workers in the layoff region who are employed at the mass layoff firm in any of the 4 years prior to the layoff; spillover workers are defined as the complement for workers in the layoff microregion sample. Standard errors are clustered at the matched comparison levels.