

UCLA

UCLA Electronic Theses and Dissertations

Title

Essays in Urban and Public Economics

Permalink

<https://escholarship.org/uc/item/47d7s3d7>

Author

Sauers, Zachary

Publication Date

2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Essays in
Urban and Public Economics

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

by

Zachary Louis St George Sauer

2023

© Copyright by
Zachary Louis St George Sauers
2023

ABSTRACT OF THE DISSERTATION

Essays in
Urban and Public Economics

by

Zachary Louis St George Sauer
Doctor of Philosophy in Management
University of California, Los Angeles, 2023
Professor Paola Giuliano, Chair

In Chapter 1 of this dissertation, I study the effect of student debt on the post-schooling migration decisions of high-skill workers in the United States. Over the past 40 years, the U.S. has experienced significant skill-based geographic sorting, with high-skill workers increasingly concentrating in large cities. During the same period, the growth of student debt has far exceeded the rate of inflation. In this chapter, I document a link between these two facts. I first estimate the causal effect of student debt on post-schooling location choices by exploiting an expansion of federal student loan limits. Using a difference-in-differences framework, I find that \$10,000 of additional debt increases the probability that individuals locate in large metropolitan counties by 6.5 percentage points. By incorporating student debt into a standard spatial equilibrium model, I find that the rise in student debt from 1980 to 2019 can account for 5-19 percent of the increase in skill-based sorting over this period. Counterfactual simulation of three policy proposals – debt forgiveness, tuition-free college, and income-driven repayment – show that only income-driven repayment can eliminate distortions to location choices while improving welfare.

The remaining chapters focus on public-sector employee pension systems in the United States and the over \$3 trillion in debt associated with them. In Chapter 2, I consider the political economy problem of setting pension benefit levels, where politicians balance the demands of general voters and public-sector unions. More specifically, I empirically show that expanded collective bargaining rights for public-sector employees significantly increased pension plan generosity in the 20th century, and is associated with higher levels of unfunded liabilities in the 21st century. I also provide descriptive evidence that increased plan generosity resulted in higher levels of unfunded liabilities because local governments shirked their expected contributions to pension funds and made overly optimistic assumptions on investment returns, both of which were made possible by systematic information asymmetries around public pension plans.

In Chapter 3, I examine the implications of public-sector pension debt for the local economy. I exploit plausibly exogenous shocks to the reported levels of unfunded pension liabilities in a difference-in-differences framework to investigate the speed and extent to which debt shocks are capitalized into house prices. I find that increases in public debt depress local house prices relatively quickly (within 9 months). Additionally, this effect is driven by responses in the price of single-family homes, owners of which may be more likely to rely on public goods that are subject to cuts following spikes in reported pension under-funding.

The dissertation of Zachary Louis St George Sauers is approved.

Melanie Sharon Wasserman

Romain T Wacziarg

Christian Dippel

Paola Giuliano, Committee Chair

University of California, Los Angeles

2023

To my parents, Louis and Sandra, and my sister, Valerie.

TABLE OF CONTENTS

1 Student Debt and High-Skill

Worker Location Choice	1
1.1 Introduction	1
1.2 Background on Student Borrowing	9
1.2.1 Federal Revisions of Stafford Loan Limits	11
1.3 Identifying the Effect of Student Debt on Location Choice	12
1.3.1 Data	12
1.3.2 Empirical Strategy	15
1.3.3 Identifying Constrained Borrowers	18
1.3.4 Results	20
1.4 Mechanisms	24
1.4.1 Data	26
1.4.2 Urban Preference Mechanism	27
1.4.3 Association Driven by Individuals Facing High Urban Wage Premiums	29
1.5 Spatial Equilibrium Model	31
1.5.1 Labor Supply	32
1.5.2 Labor Demand	34
1.5.3 Housing Market	35
1.5.4 Government	36
1.5.5 Education and Debt Choice	37
1.5.6 Equilibrium	38

1.6	Calibration	39
1.6.1	Parameters in Labor Supply	39
1.6.2	Parameters in Labor Demand	41
1.6.3	Parameter in Housing Market	41
1.6.4	Parameters in Education/Debt Choice	41
1.7	Counterfactual Analysis	42
1.7.1	Counterfactual 1: Debt Reversed to 1980 Level	44
1.7.2	Counterfactual 2: Debt Forgiveness	45
1.7.3	Counterfactual 3: Income-Driven Repayment	46
1.7.4	Counterfactual 4: Tuition-Free Postsecondary Education	47
1.8	Conclusion	48
Appendices		80
1.A	Early Adulthood Migration Drives Skill-Based Sorting	80
1.B	Data: U.S. Census	86
2 Does Increased Union Power Cause Pension Under-Funding in the Public Sector?		111
2.1	Introduction	111
2.2	Background on Public-Sector Union Bargaining	116
2.2.1	Pension Plan Finances 101: Actuarial Accounting	117
2.2.2	How Funding Gaps Emerge	118
2.2.3	The Role of Politics	120
2.3	Data on Public Sector Collective Bargaining Laws	122
2.3.1	Variation in E-PSCBLD	124

2.4	Data on Public-Sector Pension Plans	126
2.4.1	Public Plans Data (PPD): 2001 – 2017	126
2.4.2	ASPP-PERS Data: 1957 – 2017	127
2.5	Empirical Strategy	130
2.6	Results	131
2.7	Conclusion	134
Appendices		144
2.A	How Valuable are Pension Promises When They are Under-Funded? .	144
2.B	Additional Background on the Public Sector Collective Bargaining Law Data Set	145
2.C	Extending the PSCBLD	146
3 Shocking Debt Illusion: Actuarial Pension Plan Adjustments and House Price Capitalization		154
3.1	Introduction	154
3.2	Background on Public-Sector Union Bargaining	157
3.2.1	Pension Plan Finances 101: Actuarial Accounting	157
3.2.2	Adjustments to the AAR	158
3.3	Data	159
3.4	Empirical Strategy	163
3.5	Results	164
3.6	Conclusion	166

LIST OF FIGURES

1.1	Skill Differential in Migration Choices Over the Life Cycle	67
1.2	Inflation-Adjusted Cost of Post-Secondary Education	68
1.3	Time to Degree Completion	68
1.4	Share of Loan Recipients on Income-Independent Repayment Plans . . .	69
1.5	Defining Constrained Borrowers	70
1.6	First-Year Borrowing	71
1.7	Effect of Increased Loan Access on First-Year Borrowing	72
1.8	Effect of Increased Loan Access on Cumulative Borrowing Through Year 4	73
1.9	Effect of Increased Loan Access on Year 6 County Population Density (logs)	74
1.10	Effect of Increased Loan Access on Probability of Year 6 County Being 80 th Percentile by Population Density	75
1.11	Effect of Increased Loan Access on Probability of Year 6 County Being Top Metro RUCC	76
1.12	Skill-Based Geographic Sorting: Change in High Skill Share Since 1980 .	77
1.13	Role of Student Debt in Post-1980 Geographic Sorting: Partial Equilibrium	78
1.14	Role of Student Debt in Post-1980 Geographic Sorting: Full Equilibrium	78
1.15	Income-Driven Repayment Counterfactual Map: Partial Equilibrium . . .	79
1.16	Income-Driven Repayment Counterfactual Map: Full Equilibrium	79
1.17	Association Between Being High-Skill and Probability of Changing Locations	92
1.18	Association Between Being High-Skill and Probability of Living in Top Metropolitan County (RUCC-Defined)	93

1.19	Association Between Being High-Skill and Percentile of County Population Density	94
1.20	Association Between Being High-Skill and Probability of Living in 80 th Percentile Population Density County	95
1.21	Association Between Being High-Skill and Probability of Living in 95 th Percentile Population Density County	96
1.22	Association Between Being High-Skill and Probability of Moving to Top Metropolitan County (RUCC-Defined)	97
1.23	Association Between Being High-Skill and Probability of Moving to Higher Population Density County	98
1.24	Association Between Being High-Skill and Probability of Moving to 80 th Percentile Population Density County	99
1.25	Association Between Being High-Skill and Probability of Moving to 95 th Percentile Population Density County	100
1.26	Association Between Student Debt (\$1,000s, log) and Probability of Post-College County Location Being in Top X Percentile Pop. Density	101
2.1	Historical Variation of Collective Bargaining Laws by Region	140
2.2	Historical Variation of Collective Bargaining Laws by Occupational Group	141
2.3	Trends in Pension Plan Membership	142
2.4	Trends in Pension Plan Generosity	143
3.1	Geographic Distribution of AAR Adjustments	170
3.2	How AAR Adjustments Affect Funding Ratio	171

LIST OF TABLES

1.1	Federal Borrowing Limits by Academic Year and Level (Class)	51
1.2	Federal Borrowing Limits by Entry Cohort	51
1.3	Student Characteristics by Treatment Status and Cohort	52
1.4	Student Characteristics Comparison: Difference-in-Differences Specification	53
1.4	Student Characteristics Comparison: Difference-in-Differences Specifica- tion (<i>continued</i>)	54
1.5	Difference-in-Differences Main Results	55
1.5	Difference-in-Differences Main Results (<i>continued</i>)	56
1.6	Difference-in-Differences: Great Recession Robustness	57
1.6	Difference-in-Differences: Great Recession Robustness (<i>continued</i>)	58
1.7	Persistence of Location Choices	59
1.8	Student Debt and Post-College Location Choice	60
1.8	Student Debt and Post-College Location Choice (<i>continued</i>)	61
1.9	Student Debt and Post-College Location Choice: Heterogeneity by Urban Wage Premium	62
1.9	Student Debt and Post-College Location Choice: Heterogeneity by Urban Wage Premium (<i>continued</i>)	63
1.10	Model Results: Skill-Based Geographic Sorting	64
1.11	Model Results: Consumer Welfare	65
1.12	Model Results: Aggregate Output	66
1.13	Importance of Student Debt in for Career Outcomes	88
1.14	Educational Migration Differentials: Overall Mobility	89

1.15	Educational Migration Differentials: Live in ‘Urban’	90
1.16	Educational Migration Differentials: Move to ‘Urban’	91
2.1	Funding Status of Pension Plans: 2001–2017	135
2.2	Collective Bargaining Rights and Actuarial Funding Ratio	136
2.3	Collective Bargaining Rights and Increased Plan Generosity	137
2.4	Collective Bargaining Rights and Government Contributions	138
2.5	Collective Bargaining Rights and Government Optimism	139
2.6	Measuring Public-Sector Collective Bargaining Rights	146
3.1	Funding Status of Pension Plans: 2001–2017	167
3.2	AAR Shocks and House Price Capitalization: DID Baseline	168
3.3	AAR Shocks and House Price Capitalization: Heterogeneity by Housing Type	169

ACKNOWLEDGMENTS

I am incredibly grateful to my dissertation committee members, Christian Dippel, Romain Wacziarg, Melanie Wasserman, and especially Paola Giuliano, my committee chair. Paola Giuliano's insightful guidance and advice were essential for this work. Your endless positivity and curiosity were contagious, and the enthusiasm I felt after our meetings kept me going through the tough parts. I am very thankful to Melanie Wasserman for her brilliance and patience in always guiding me in the right direction. Christian Dippel provided the foundation for how I conduct research and I am grateful for his generosity in sharing it. He is also the co-author and was instrumental for the work in Chapters 2 and 3. I am incredibly fortunate to have benefitted from the economic wisdom of Romain Wacziarg. I also benefitted immensely from discussions with other faculty at Anderson, especially Gregor Schubert, without whom I do not think the structural section in my first chapter would have been possible. I also thank Jasper Hoek for encouraging me to run with my ideas in the early days, and for being a valuable mentor in the later ones.

I am also thankful to the other Ph.D. students in Anderson, the UCLA Economics department, and elsewhere that I have had the pleasure of interacting with over the last six years. A special thanks to Edward Kim, Matias Giacobasso, Sebastian Ottinger, Thomas Groesbeck, Caitlin Hegarty, James O'Neill, Logan Spencer, Anh Nguyen, Misha Galashin, Elisabetta Campagna, Aleksandr Gevorkian, Yangkeun Yun, and Jesper Böjeryd. Ariadna Jou and Tomás Guanzioli, our surf and donuts club was the highlight of many weeks. I am not sure I would have made it through without the friendship of Sumit Shinde. I also thank him for explaining concepts from first-year when I discovered that unfortunately yes, it was something that I would end up using in research.

Thank you to the Broad Street Boys, Dr. Samuel Amsterdam, Dr. Joseph

Kotinsly, and Dr. William Sgrignoli who, against all odds and most opinions, ended up being immensely positive influences in my life, and showed me it was cool to be smart. I would also like to thank all the people who balanced out the unbearable parts with the best parts: Caitlin Quattrocchi, Pablo Aranzabal, Christopher Mandhe, Micky McBride, Trace Levinson, George Jiranek, and Elizabeth Shaw Doppelt. Finally, thank you to the wonderful people who supported and encouraged me throughout my life, especially Joseph Quattrocchi, A. Susan Goldy, and Jane Yeatter. To the people I inevitably forgot: I'm sorry, but thank you!

Additionally, this work would not have been possible without the countless people who helped make the data available to me. I would like to thank Evan White, Eduard Tomany, Gary Gremaux, April Chang, and the rest of the California Policy Lab staff for their immense help and patience in my work with the UCCCP dataset. Taara Cason and Bethany Consic from the Institute of Education Studies and Harry Morall, Freddy Lopez, Kevin Kurata, and Karla Zapeda from UCLA were all incredibly helpful in my work with the Education Longitudinal Study of 2002. I am also grateful for the help of Linda Eggenberger and other members of the Institute of Social Research at the University of Michigan in my work with the Panel Study of Income Dynamics. Joshua Mimura provided excellent research assistance and was crucial for the data collection in Chapter 2.

Lastly, I gratefully acknowledge receiving financial support from the UCLA Graduate Division, the UCLA Global Economics and Management area, the UCLA Ziman Center for Real Estate, and the California Policy Lab during my time completing this research.

VITA

- 2015 Bachelor of Science in Economics, University of Pittsburgh, Pittsburgh, Pennsylvania.
- 2015–2017 Research Assistant, Board of Governors of the Federal Reserve System, Washington, D.C.
- 2017–2021 Academic Fellowship, Graduate Division, UCLA, Los Angeles, California.
- 2018–2023 Teaching Assistant, Global Economics and Management Department, UCLA Anderson School of Management, Los Angeles, California.
- 2019–2021 Research Assistant, Prof. Paola Giuliano, Global Economics and Management Department, UCLA Anderson School of Management.
- 2020 Candidate of Philosophy in Management, UCLA Anderson School of Management.
- 2020 Ziman Center for Real Estate Fellow, UCLA Anderson School of Management.
- 2021 Graduate Fellow, California Policy Lab, Berkeley, California.
- 2022–2023 Dissertation Year Fellowship, Graduate Division, UCLA.

CHAPTER 1

Student Debt and High-Skill

Worker Location Choice

1.1 Introduction

Over the past 40 years, there has been a significant amount of skill-based geographic sorting in the United States, with high-skill workers increasingly concentrating in large, dense locations (Diamond 2016; Moretti 2013; Berry and Glaeser 2005). As a result, by 2020, the average high-skill worker share in top labor markets was 57 percent while it was just 37 percent in rural areas – double the difference in 1980. This sorting has been of interest to researchers and policy-makers alike due to the implications at the national and local level for productivity, inequality, housing markets, and voting patterns (Diamond 2016; Glaeser 2008; Scala and Johnson 2017). While skill-based spatial sorting is now a well-known fact, I document the thus-far overlooked fact that migration decisions in early adulthood drive this divergence. As shown in Panel (a) of Figure 1.1, high-skill workers are significantly more likely to move to higher-density locations in early adulthood relative to low-skill workers. This results in differential location patterns that stabilize by age 30, as shown in Panel (b). Together, these patterns suggest that factors weighing heavily on individuals in early adulthood could contribute to aggregate sorting patterns.

In this paper, I investigate one such factor and show that it distorts location choices of high-skill workers: student debt. The cost of postsecondary education

in the U.S. has increased fourfold since 1980 and students have increasingly turned to loans as a way to finance their education. I show that increased student debt pushes high-skill workers to large, densely-populated areas where it is widely known that wages and costs tend to be higher (Glaeser 2008). The proposed mechanism by which student debt affects post-college location choice is that it makes borrowers more elastic to nominal wages than local prices. This stems from the structure of student debt repayment plans which, for the majority of borrowers, are fixed and independent of local prices and income.

To build intuition for this mechanism, consider individuals with and without student debt (D) that choose location based only on nominal wages (W) and local prices (P), i.e., they maximize $\ln\left(\frac{W-D}{P}\right)$ – a simplified version of the standard Rosen (1979)-Roback (1982) framework. For individuals with no debt, location choice is equally elastic to nominal wages and local prices. When individuals carry student debt, the elasticity with respect to wages ($\frac{1}{1-\frac{D}{W}}$) is strictly greater than the elasticity with respect to prices (-1), and it increases with the level of debt. This is because debt is repaid from the dollar gap between nominal wages and housing costs, which is typically higher in urban areas even if the real standard of living is lower. Individuals with student debt then maximize utility by choosing to locate in areas with relatively higher nominal wages.

In the first part of the paper, I estimate the causal effect of student debt on post-college location choice. Since eligibility for federal student loans is nearly universal, uptake could be endogenous to personal, social, and economic characteristics. As a result, isolating the causal effect of student debt is challenging. To solve this problem, I exploit an increase of federal student loan limits in 2008 and 2009 that created cohort variation in the maximum amount individuals could borrow. I employ a difference-in-differences approach that compares post-schooling location choices of students who enrolled in college before and after the loan limit increases across two

groups: (1) students unconstrained by limits, and (2) students likely to be constrained by limits. The key identifying assumption is that, in the absence of loan limit increases, differences in post-college location choices of constrained and unconstrained students would be similar across cohorts.¹

I implement this strategy using newly available data in the University of California Consumer Credit Panel (UCCCP). The UCCCP covers the complete population of individuals with credit reports who lived in California between 2004 and 2019 – containing all credit reports for these individuals throughout this period, even when they reside outside of California. Precise location is observed throughout while detailed account information allows me to observe educational borrowing from 2002 to 2019. I first employ the difference-in-differences analysis to confirm that increased access to educational loans resulted in additional borrowing. I find that increased loan limits resulted in an average of \$2,600 in additional loans for constrained students relative to unconstrained students through year 4 of borrowing – the last year of postsecondary education for the majority of students.²

I then turn to the location choices of student borrowers in post-college years. I find that \$10,000 in additional debt increases the population density of a borrower’s initial post-college county by 8.5 percent. Focusing on migration to top urban areas, I find that \$10,000 in additional borrowing increases the probability of locating in counties classified as being in the 95th percentile by population density by 6.2 percentage points or 11.4 percent from the mean. Indicative of lasting effects on migration choices, I show that this effect is persistent through year 9 from first education loan, the last year with available data. These results are robust to alternative classifications of counties provided by the Department of Agriculture’s Rural-Urban Continuum Codes.

1. This empirical approach is similar to Black et al. (2020), who examine financial well-being outcomes.

2. This result is similar to Black et al. (2020), who find roughly \$1,800 of additional borrowing in a national sample.

In the second part of the paper, I provide descriptive evidence consistent with high-skill workers differentially valuing the higher nominal wages of cities. For this analysis, I use data from the Education Longitudinal Study of 2002 (ELS), a longitudinal survey of 2002 high school seniors that, unlike the UCCCP, provides a rich set of controls to perform heterogeneity analysis. This analysis stems from a testable implication of the proposed mechanism: the association between student debt and urban post-college location choice should be higher for individuals who face steeper urban wage premiums. Using American Community Survey microdata containing nominal wages, location, and degree fields for high-skill workers, I compute urban wage premiums by degree field and map them onto the universe of academic majors in the ELS. As predicted by the proposed mechanism, the association between student debt and urban post-college location choice is driven by individuals who major in fields with above-median urban wage premiums even after accounting for spatial variation in the distribution of industries. For individuals who major in degree fields with below-median urban wage premiums, there is no significant relationship between student debt and urban post-college location choice.

One alternative explanation for the identified effect of student debt is that preference for urban locations may be correlated with willingness to accept student debt. This would require individuals make college, student debt, and post-college location choices at least partially as a joint decision. I present two descriptive facts from the ELS that suggest this mechanism does not play a large role. First, I show that the association between student debt and urban post-college location choice is robust to controlling for high school and college county urbanicity as proxies for urban preference. Second, as a large portion of college graduates live and work near their postsecondary institution (Conzelmann et al. 2022), this implies that individuals making a joint decision would likely have the same college and post-college location. I show that the documented association between student debt and urban post-college

location choices is robust to analysis of only the subpopulation of individuals who left their college county – a group which may be making post-college location choices more independently of college and debt decisions.

How important is student debt for aggregate sorting patterns? What effect may policy proposals such as student loan forgiveness, a free college option, or widespread adoption of income-driven repayment plans have? To study these counterfactual questions, in the third part of the paper I develop and calibrate a spatial equilibrium model that enables me to quantify the role of student debt in skill-based sorting and simulate outcomes under various policy proposals. Workers in the model choose skill and debt levels endogenously and have heterogeneous preferences over locations. Locations differ by skill-specific productivity and amenity levels. Housing markets differ across locations due to heterogeneity in their elasticity of housing supply. To conduct my analysis with complete coverage of the US, I estimate new 1990 Commuting Zone-level elasticities following the methods of Saiz (2010) and Howard, Liebersohn, et al. (2018).

After calibrating the model using parameters from the literature, I present simulated outcomes under various counterfactual scenarios. I present results in full equilibrium, where prices and wages respond to population changes, and partial equilibrium, where prices and wages remain fixed and changes in outcomes reflect only a labor supply response via migration. By remaining agnostic about wage and price responses, the partial equilibrium results offer short-term predictions while full equilibrium results, which account for spillover effects on low-skill workers through price and wage adjustments, may more accurately depict long-term outcomes. In both, location- and skill-specific amenities remain fixed.

I begin the counterfactual analysis by estimating that the growth of student debt accounts for 3.5-4.5 percent of the increase in skill-based sorting from 1980 to 2019 (11.5-18.7 percent in partial equilibrium). I then turn to assessing the impact of vari-

ous policy proposals. Debt forgiveness is estimated to decrease the level of skill-based sorting by 2.4-3.4 percent (8.1-14.0 in partial equilibrium). This is accompanied by a small decline in welfare with high-skill workers gaining from the reduction in debt and the elimination of the location distortions while low-skill workers are worse off due to the introduction of a uniform tax to fund the policy. Widespread adoption of income-driven repayment plans, which do not distort location choices, results in similar reductions to skill-based geographic sorting; however, all workers now experience an increase in welfare as location choices adjust to fully reflect preferences without imposing a tax on low-skill workers. The introduction of tuition-free post-secondary education has a negligible net effect on skill-based sorting as two opposing forces interact: the rising national share of high-skill workers are still drawn to cities by higher productivity, but their location choices are no longer distorted by debt repayment. Welfare increases for high-skill workers as they no longer bear the full cost of education while it declines for low-skill workers as they assume some of the cost. For debt forgiveness and income-driven repayment, aggregate output declines by a fraction of a percent as high-skill workers express preferences to locate in lower productivity regions. As tuition-free college increases the share of high-skill workers, it increases aggregate production by 1.5-2 percent.

These findings have important implications for policy-makers at both the national and local level. In 2022, the Biden Administration announced the nation's largest student loan forgiveness program, covering up to \$20,000, and is pushing for increased adoption of income-driven repayment plans. Implementing a free college option is also part of some political platforms. As these policies are novel in the United States, the findings in this paper provide valuable insights about implications for internal migration and labor markets. Additionally, these findings provide a new avenue to address the so-called 'brain drain' from rural areas to urban ones – the phenomenon where top-performing students leave rural areas rather than enhancing the local stock of

college-educated adults. While leaders have frequently employed place-based policies to address this problem with mixed success (Neumark and Simpson 2015) and unclear implications for aggregate welfare (Kline and Moretti 2014; Glaeser and Gottlieb 2008), reducing outflows of highly productive workers by switching to income-driven repayment plans for student debt offers a promising new strategy to boost economic activity in rural areas via the supply-side.

Related Literature. My work builds on a number of papers exploring the causes of skill-based geographic sorting. There is robust evidence that the primary driver is changes in local productivity (Berry and Glaeser 2005; Moretti 2013) resulting in part from skill-biased agglomeration economies (Baum-Snow, Freedman, and Pavan 2018; Giannone 2017). Amenity differences have also been documented as an important contributing factor (Shapiro 2006; Albouy et al. 2016) that amplifies productivity-driven sorting (Diamond 2016). I add to this literature by identifying student debt as new driver of skill-based spatial sorting that, unlike the previously-documented factors, is the result of a policy decision governing the structure of debt repayment rather than underlying economic forces.

This paper also adds to a growing literature examining the implications of student debt for post-graduation outcomes. While student debt is considered a good investment for most people (Barrow and Malamud 2015; Oreopoulos and Petronijevic 2013) even in the face of growing debt (Avery and Turner 2012), conclusions about its effects on post-schooling life are less clear. There is increasing evidence that student borrowers do not behave as a standard life-cycle model would suggest, i.e., student debt should have a minimal effect on post-college decisions because the ratio of debt to the present discounted value of lifetime earnings is small (less than 1 percent). Instead, student debt has been found to affect many key decisions in adulthood: household formation via homeownership (Mezza et al. 2020; Bleemer et al. 2021) and co-residence with parents (Dettling and Hsu 2018); family formation via

marriage and fertility (Gicheva 2016); job choice (Di Maggio, Kalda, and Yao 2020; Rothstein and Rouse 2011); and income (Gervais and Ziebarth 2019; Luo and Mongey 2019; Bettinger et al. 2019).^{3,4} While nearly all find significant effects of student debt on the respective outcome, the magnitude and direction varies by sample and empirical strategy. I contribute to this literature by identifying a missing dimension, across which many of these outcomes vary significantly, that could account for the mixed results: geography.

Methodologically, the empirical work in this paper builds on those in the student debt literature that use experimental or quasi-experimental variation. The most closely related is Black et al. (2020), who use the same policy change to reflect on educational attainment and post-college financial well-being. Other empirical work exploits variation from grant aid, tuition, and bankruptcy regulations. Additionally, this paper augments a small but growing number of studies in the student debt literature employing large consumer credit panels (Chakrabarti, Gorton, and Lovenheim 2020; Black et al. 2020). The structural component of this paper builds on the original Rosen (1979)-Roback (1982) framework as well as more recent spatial equilibrium models like Hsieh and Moretti (2019) and Diamond (2016). The proposed mechanism is similar to that in Albert and Monras (2017), who find that immigrants concentrate in expensive U.S. cities because remittances to origin countries reduce sensitivity to local price levels. In my setting, debt repayments at a national price, i.e. fixed across locations, make student borrowers relatively less sensitive to local price levels.

In the following section, I discuss relevant institutional details on student borrow-

3. See as well: household formation (Black et al. 2020; Akers and Chingos 2014; Houle and Berger 2015; Gicheva and Thompson 2015; Chakrabarti, Gorton, and Lovenheim 2020; Bleemer et al. 2014), family formation (Shao 2014), job choice (Krishnan and Wang 2019), income (Di Maggio, Kalda, and Yao 2020; Chapman 2015; Minicozzi 2005; Weidner 2016).

4. Contradictions of the life-cycle model are attributed to debt aversion (Burdman 2005; Callender and Jackson 2005; Field 2009) or credit constraints after college (Rothstein and Rouse 2011; Gicheva and Thompson 2015).

ing and the policy change that generates the identifying variation. Section 1.3 outlines the empirical strategy for identifying the effect of student debt on post-college location choice and presents the results. In Section 1.4, I provide supporting evidence for the proposed mechanism. Section 1.5 embeds the proposed mechanism in a spatial equilibrium model, which is calibrated in Section 1.6. Section 1.7 simulates outcomes under various policy proposals. Section 1.8 concludes.

1.2 Background on Student Borrowing

The cost of postsecondary education in the United States has increased rapidly over the past 40 years. As shown in Figure 1.2, inflation-adjusted tuition and fees are now four times the 1980 level. Students have absorbed these rising prices by taking out additional debt - with loans per student increasing nearly fivefold in real terms over the same period. In 2019, the typical bachelor's degree-holder leaves school with roughly \$30,000 in debt. Federal lending dominates the landscape of student borrowing for post-secondary education, comprising 88-93% of all educational loans over the past decade (Baum et al. 2019).⁵ Under the umbrella of federal lending, Stafford Loans account for roughly two thirds of borrowing (Baum et al. 2019). Stafford Loans have historically been provided by one of two federal lending programs: the Federal Family Education Loan (FFEL) Program, authorized as part of the Higher Education Act of 1965 (HEA), and the Federal Direct Loan Program, created as a 1992 amendment to the HEA. Though the source of the funds has varied over time, the function of Stafford Loans has remained consistent from the student's perspective.⁶

5. This has been remarkably stable over time. In 1998-99, the first period with nonfederal borrowing data, federal loans accounted for 91% of all education borrowing (Baum et al. 2019).

6. The FFEL was eliminated in 2010 as part of the Health Care and Education Reconciliation Act of 2020. After this date, nearly all federal lending is through the Direct Loan Program. Stafford Loans under the Direct Loan Program were issued directly from the Department of Education, while loans under the FFEL Program were issued by private sector institutions and guaranteed by the federal government. There is no practical difference from the student's perspective.

Undergraduate students have essentially uniform access and terms for Stafford Loans. To qualify for Stafford Loans or any other type of federal student aid, students must complete the Free Application for Federal Student Aid (FAFSA), an annual form which collects demographic and financial information. This information, which includes assets and income, pertain to the student and their household for dependent students. The Department of Education uses the information provided in the FAFSA to determine a student's eligibility across the two types of Stafford Loans: subsidized and unsubsidized. Subsidized Stafford Loans are need-based and do not accrue interest while the student is enrolled. Unsubsidized Stafford Loans are not based on financial need and do accrue interest while the student is pursuing their degree. Although the package of offered Stafford Loans varies by student, this has limited implications for cumulative borrowing over the duration of an individual's enrollment.⁷ The more meaningful constraint and determinant of cumulative borrowing is the federal student loan limit, which governs the maximum allowable loans over both types of Stafford Loans for a given academic level (i.e., freshmen, sophomore, or upper level).

The majority of college graduates are awarded their degree in 4 years or less and repay educational debt using plans with fixed payments across space. As shown in Figure 1.3, over 60% of students who attained a Bachelor's degree in 2017 did so in 4 years or less. This increases to 85% by year 5. Repayment typically begins after graduation and a grace period of 6 months. Repayment periods are typically up to 10 years for single loans and between 10 and 30 years for consolidation loans. Income-driven repayment plans were first introduced in 2009; however, as of 2020, roughly 70% of borrowers are still on traditional repayment plans with amounts set

7. The value of the in-school interest subsidy varies by entry year and duration of schooling. Subsidized Stafford Loans had a slightly lower interest rates from 2008 to 2013. Black et al. (2020) estimate that this subsidy ranges from \$34 to \$82 for a \$1,000 loan when repayment starts one year after origination.

independent of income or location (Figure 1.4). Durante et al. (2017) find an average monthly payment of \$393 among those actively making payments in 2016.

1.2.1 Federal Revisions of Stafford Loan Limits

The borrowing limits for Stafford Loans can only be adjusted via federal legislation – something that has only been happened twice in the 21st Century. The first was with the Higher Education Reconciliation Act of 2005, which went into effect in the 2007-08 academic year. The second adjustment occurred as part of the Ensuring Continued Access to Student Loans Act of 2008, which took effect starting in the 2008-09 academic year. Both changes increased the loan limits. The experienced borrowing limits are reported by academic year and level in Table 1.1. The first wave of increases only impacted individuals in their freshmen and sophomore years, while the second adjustment increased limits for all academic levels. Although aggregate cumulative limits were adjusted as well (last Column in Table 1.1), these limits would never constrain a borrower who attains their degree in 4 years further than academic year-x-level limits.⁸

Table 1.2 shows how these changes impacted students by entry cohort. While the first wave of loan limit increases took effect in the 2007-08 academic year, individuals in earlier cohorts can still be affected if they are in school when adjustments took effect. This combination of staggered introduction over time and uneven increases across academic levels generates the identifying variation used in the following section. For example, students who entered in the 2005-06 academic year may have experienced a \$2,000 increase in their year 4 borrowing limit if they were still enrolled through year 4. The increased loan availability phases in over the 2005-06

8. Considering cohorts 2002-3 through 2012-3 (the sample used for identification in Section 1.3) and individuals who enroll in postsecondary education for up to 6 years, aggregate limits would only constrain borrowers further than academic year-x-level limits for individuals in the 2002-03 cohort who are enrolled and borrow the maximum amount each year through year 6.

through 2008-09 cohorts and peaks at an additional \$9,875 in borrowing ability (last column of Table 1.2).

1.3 Identifying the Effect of Student Debt on Location Choice

In this section, I estimate the effect of student debt on post-schooling location choice by exploiting variation in student borrowing driven by a policy change that increased the maximum amount students are able to borrow for postsecondary education from Federal sources. I estimate this effect for the full population of California student borrowers in the recently developed University of California Consumer Credit Panel (UCCCP) assembled by the California Policy Lab.

1.3.1 Data

The UCCCP is an individual-level longitudinal dataset following roughly 60 million consumers with credit reports on a quarterly basis since 2004. While not the first credit panel used in the literature (Black et al. 2020; Chakrabarti, Gorton, and Lovenheim 2020), it is significantly larger than the Federal Reserve Bank of New York’s Consumer Credit Panel (approx. 13 million) and the Consumer Financial Protection Bureau’s Consumer Credit Panel (approx. 5 million). The underlying credit histories are sourced from Experian, one of the three nationwide credit bureaus. The UCCCP is composed of two samples: a nationally representative 2 percent sample of U.S. adult consumer with credit records, and a full ‘sample’ consisting of 100 percent of Californians with credit histories. The California sample includes all consumers with credit reports who lived in California between 2004 and 2019.⁹ This includes those who originated in the state, moved to California for college, or resided there after

9. Though the UCCCP continually adds archives up to the present, I limit the sample to 2019 to avoid confounding factors arising from the COVID-19 pandemic.

schooling. For individuals that meet this inclusion criteria, the UCCCP includes all available reports regardless of location in any given quarter. Although both samples provide the same coverage of variables, this analysis uses the California sample due to its size.

The UCCCP includes information on tradeline-level account information, credit scores, location, and demographics of consumers. Tradelines include student loans, auto loans, credit cards, mortgages, and other forms of credit. Data on tradelines include account opening date, account type, account condition (open, closed, in deferment, in repayment, etc.), principal amount (for loans), borrowing limits (for credit cards), and latest balance amount among others. Geographic information consists of 5-digit ZIP codes sourced from tradeline mailing addresses, which I then map to county-level characteristics for analysis.¹⁰ Demographic information includes gender, month and year of birth, and education codes.¹¹

Since the UCCCP does not include enrollment information, I use information on student loans contained in credit histories to infer entry cohorts and build a dataset of borrower-x-year since entry observations. To do so, I assume that the first academic year an individual is observed opening a student loan is the first year that they enter school, i.e. their entry cohort.^{12,13} Since each quarterly archive of the UCCCP

10. ZIP codes mapped to counties using a crosswalk from the U.S. Department of Housing and Urban Development.

11. Demographic information is often limited in credit reports because federal law prohibits discrimination in credit transactions on the basis of race, ethnicity, religion, sex, age, marital status, or receipt of public assistance. Education codes are modeled/estimated by the credit bureau using sample surveys.

12. Academic year defined as July through June and denoted as the calendar year it ends. For example, the 2003-2004 academic year, denoted just by 2004, includes loans opened from July 1st, 2003 through June 30th, 2004.

13. Most individuals that ever borrow to finance their college education do so in their first year. Black et al. (2020) estimate that 73 percent of all dependent undergraduates in the 2016 National Postsecondary Student Aid Study who ever took out student loans and graduated in 2016 borrowed in their first year.

provides a snapshot of an individual's full credit report at a given point in time, I am able to construct first-year borrowing information for cohorts back to 2002 using 2004 archives.¹⁴ Loan transfers and consolidations in addition to changing tradeline identifiers make linking and tracking the exact evolution of borrowing beyond first year unreliable, particularly for cohorts before 2004. To solve this problem, I measure a borrower's cumulative borrowing at any given point as the sum of all active loans. While eliminating the need to trace all transfers and consolidations, this limits observations to first-year borrowing and cumulative borrowing through years 4, 5, and 6 from first education loan for all individuals in cohorts 2002 to 2013.

I restrict the population in two ways to target borrowers financing first-time undergraduate education. First, I only include borrowers who open their first education loan between the ages of 16 and 20. Second, I exclude all borrowers whose first-year loans exceed the Federal student loan borrowing limits for first year undergraduate students in a given academic year. The intention of both restrictions is to reduce the inclusion of individuals who first borrow for a graduate degree (and thus face a different labor market), individuals who first borrow in upper academic levels of undergraduate education, and those that are independent students.

To create a balanced panel, I only include individuals from the resulting sample that are observed at a minimum through year 6 from first loan and for up to 9 years. As shown in Figure 1.3, a vast majority of borrowers who attain an undergraduate degree do so in 4 or 5 years. I ensure all individuals are observed through year 6 because I will consider this the start of when location information in credit archives reflect their post-college location. The resulting dataset contains roughly 940,000 student borrowers who entered in 2002 through 2013 cohorts and are observed for up to 9 years after entry.

14. As I am using the 2004 archive, this approach conveniently avoids the problem that credit histories prior to 2004 often suffer from incomplete reporting of student loans.

1.3.2 Empirical Strategy

The empirical strategy for identifying the causal effect of student debt on post-college location choices compares outcomes for students likely to be constrained and unlikely to be constrained (unconstrained) by original borrowing limits in years before and after the federal loan limit increases. The identifying assumption is that, in the absence of loan limit increases, differences in post-college location choices of constrained and unconstrained students would be similar across cohorts. I begin the empirical analysis with an event study framework which allows me to analyze dynamics over the cohorts, such as differences in baseline characteristics, cumulative borrowing ('first stage'), and location choices ('reduced form'). I then move on to the difference-in-differences analysis to get the main estimate of interest: the effect of additional borrowing on post-college location choices. I am also able to examine persistence through year 9 from initial borrowing in this specification.

The event study framework is given by:

$$\begin{aligned}
 Y_{isc} = & \alpha + \beta_1 \text{Constrained}_i + \sum_{c \neq 2005} \beta_2^c [\mathbf{1}[\text{Cohort}_i = c] \times \text{Constrained}_i] \\
 & + \mathbf{X}_i' \boldsymbol{\beta}_x + \delta_c + \delta_s + \epsilon_{isc},
 \end{aligned} \tag{1.1}$$

where Y_{isc} is the outcome of interest (baseline characteristic, cumulative borrowing through a certain year, or post-college location choice) for individual i who attended college in state s and first borrowed as part of cohort c . When considering post-college location choices, urbanicity of a location is defined in three ways: (1) a continuous measure of a county's population density (in logs); (2) an indicator function for if the county is in the 95th percentile by population density; and (3) an indicator function for if the county is classified in the top category (metropolitan area with a population of at least 1 million) using the Rural-Urban Continuum Codes (RUCC). Whether individuals are likely or unlikely to be constrained by original borrowing limits is

captured by the binary variable Constrained_i . The coefficients of interest, β_2^c , capture the interaction between the indicator for being constrained and an indicator being in cohort c . The limited demographic variables available in credit reports are included in X_i . These include sex, age when individual was first issued an education loan, and characteristics of the borrower’s credit history prior to entry.¹⁵ I also include cohort (δ_c) and college state fixed effects (δ_s) to capture common trends and unobserved characteristics that may vary by state of college attendance.¹⁶ Errors are clustered by college state.¹⁷

In the event study specification and the difference-in-differences framework to follow, the final cohort designated as unaffected by the loan limit increases contains those who entered in 2005. This reflects the variation in possible cumulative borrowing through year 4 as shown in Table 1.2. This choice, consistent with Black et al. (2020), was based on evidence that the majority of students who attain a Bachelor degree do so in 4 years or less, as illustrated in Figure 1.3. To the extent that the share of individuals staying more than four years experience additional loan limits prior to 2005, the event study estimates would show pre-trends and the difference-in-differences estimates would underestimate the full effect of additional borrowing.¹⁸

The main advantage of the event study framework is to provide insight into the

15. This includes an indicator for the existence of a prior credit report, an indicator for a credit score, the credit score, and indicators for the most common accounts in early adulthood, auto loans and credit cards.

16. College states other than California exist in the data as the UCCCP tracks Californians with a credit report prior even when they leave the state. The UCCCP sample also includes individuals who moved to California after college (with historical credit reports from time outside the state). College state FEs are included to capture selection bias involved with the two, but results are also robust to restricting the sample to only individuals who are in California during postsecondary education.

17. College state is captured by location in year 3 from first educational loan. This closely follows the framework of Black et al. (2020), although they focus on post-college financial outcomes.

18. The UCCCP does not include information on enrollment or degree-level. As a result, I am not able to distinguish increases in post-year 4 borrowing as additional borrowing for an undergraduate degree, new borrowing for a graduate or professional degree, or restructuring of existing loans.

identifying assumption of parallel trends in the difference-in-differences specification. The parallel trends assumption in this application is that in the absence of loan limit increases, differences in post-college location choices of constrained and unconstrained students would be similar across cohorts. Although this assumption is untestable, it suggests that the difference in outcomes between groups should be similar across the untreated cohorts, i.e., pre-2005 cohorts. As shown in Figure 1.7-1.8 and Figure 1.9-1.11 for borrowing and location outcomes, respectively, there is little evidence of pre-trends. The only variable displaying differences between constrained and unconstrained borrowers that vary significantly from the 2005 cohort level are for borrowing outcomes (‘first stage’), all of which are small in magnitude and do not suggest a clear pattern. There are no pre-trends in location outcomes. While not conclusive, this provides evidence that is consistent with the parallel trends assumption being valid in this setting.

The main estimates for the effect of student debt on post-college location outcomes comes from the following difference-in-differences specification aggregating cohorts based on treatment status:

$$\begin{aligned}
 Y_{isc} = & \alpha + \beta_1 \text{Constrained}_i + \beta_2 [\mathbf{1}[\text{Cohort}_i > 2005] \times \text{Constrained}_i] \\
 & + \mathbf{X}_i' \boldsymbol{\beta}_x + \delta_c + \delta_s + \epsilon_{isc}
 \end{aligned}
 \tag{1.2}$$

where the coefficient of interest, β_2 , captures the effect of being constrained across all cohorts. The rest of the specification, including covariates, fixed effects, and error clustering is the same as in the event study framework. To avoid distortion from year-specific fluctuations that affect constrained and unconstrained borrower outcomes equally, I continue to include flexible cohort year fixed effects rather than a simple indicator for the post-policy change period. This specification provides an estimate that is a weighted average of the effects for each treated cohort.

1.3.3 Identifying Constrained Borrowers

The empirical strategy identifying the effect of student debt on post-schooling location choice relies on the ability to classify students as likely or unlikely to be or have been constrained by *pre-policy-change* borrowing limits. I do so using observed loans in a student's first year of borrowing.¹⁹ For students who entered in years unaffected by loan limit increases for first-year borrowing (pre-2008), I classify students who borrow exactly at the loan limit of \$2,625 to be likely constrained. For students who experienced increased borrowing limits in their first year, they are classified as likely to have been constrained if they borrow at or above the original limit of \$2,625. In all years, students who borrow below the original first-year limit are classified as unlikely to be constrained. This classification system is depicted in Figure 1.5.

This strategy assumes that students who borrow exactly at the loan limit in years before the increases would have borrowed more if given the possibility. Although it is impossible to verify this assumption using the data in the UCCCP, the distribution of first year borrowing across cohorts indicates that it is likely to be true for the majority of borrowers at the limit. As shown in Figure 1.6, the borrowing distribution for individuals who first took out education loans in 2002-2007 academic years has a large mass exactly at the limit that year (\$2,625). For individuals who first borrowed in the 2008 academic year, when the loan limit increased to \$3,500, the largest mass in the distribution of first year borrowing shifts to the new limit. This shift to the new limit is also observed for borrowers in 2009-2013 entry cohorts for whom the first year loan limit was \$5,500. This bunching and the quick shifts of the distribution to new limits suggest students are constrained at limits and would borrow more when

19. This is the same classification strategy used in Black et al. (2020). They estimate that 73 percent of all dependent undergraduate students who ever borrowed and graduated in 2016 borrowed in their first year.

given the possibility.²⁰

Characteristics of constrained and unconstrained borrowers in the UCCCP are reported in Table 1.3. Federal laws prohibit creditors from discriminating against applicants on the basis of many personal characteristics, including sex, race, color, religion, and marital status. As a result, credit histories, and thus the UCCCP, contain little demographic information; however, I am able to observe age at first education loan and sex along with variables to characterize borrower's credit profiles. As shown in Row 1 of Table 1.3, the sample contains borrowers that were a little older than 18 at the time of their first education loan. Consistent with national statistics showing women make up a majority of recent postsecondary degree recipients, women make up slightly over half of the sample. As for credit characteristics of the UCCCP sample, about 30-40 percent had a credit report prior to opening an education loan, with this share increasing in later years. The share of individuals with a credit score, which is only calculated after 3 to 6 months of credit activity, also increases over the sample period from about 10 percent to 30 percent. For borrowers with a credit score, the average is in the low 600s, which falls in the 'fair' category. About 20 percent of individuals have a credit card account and very few have an auto loan.

The difference-in-differences framework allows me to compare the baseline characteristics of borrowers presented in Table 1.3 across treatment groups and the sample period. To do so, I modify Equation 1.2 only by omitting any baseline characteristics. The results of this analysis are reported in Table 1.4. As shown, individuals who are constrained in the post-period tend to be slightly older (5 days) and a slightly higher percentage are women (1.4 percentage points). There is no significant variation in the presence of a credit report, though constrained borrowers in the post-period are

20. Another possibility is that students may be likely to accept financial aid as it is packaged by schools and nearly all four-year institutions include the maximum available Stafford Loans in aid packages (Marx and Turner 2019). As argued by Black et al. (2020), this channel still induces additional borrowing and the empirical strategy still produces a causal estimate of student debt on student outcomes.

less likely to have a credit score and have a slightly lower credit score when present – both indicative that constrained borrowers in the post period may have less financial experience than in the pre-period; however, a larger share have credit cards. Although differences are small relative to variable means, I control for all of these characteristics in both the event study and difference-in-differences framework to reduce any bias introduced by changes in the sample population.

The estimates provided by the event study and difference-in-differences specifications reflect the effect of increased access to educational loans, but there are limits to the interpretation. Due to the setup and strategy for identifying the constrained status of borrowers, estimates reflect effects for students *already enrolled and borrowing*. One concern is that increased borrowing limits enabled students to enroll in postsecondary education that were previously too credit constrained to enroll; however, Marx and Turner (2019) find only minor effects of increased borrowing ability on enrollment. Another concern is the implications for school choice. While UC-CCP data does not allow me to identify enrollment institutions, Black et al. (2020) find no evidence that increased access to borrowing led to more transfers from community college to four-year institutions in a sample of Texas students. Lastly, it is possible that students who are defined as unconstrained by first-year borrowing become constrained in subsequent years. To the extent that this occurs, results will underestimate the true effect.

1.3.4 Results

1.3.4.1 The Effect of Increased Loan Limits on Borrowing

Increased access to educational loans resulted in higher debt for students who were students likely to have been constrained by original loan limits. The results of the event study specification in Equation 1.1 when examining first-year borrowing are

reported in Figure 1.7. As expected since first-year loan limits remain at the original level, the effect of increased borrowing limits on first-year borrowing is flat through the 2007 cohort. The estimated effect on first-year borrowing then increases sharply to around \$2,000 by the 2010 cohort. This is similar to the magnitude increase of the expansion in first-year loan limits shown in Table 1.2, suggesting constrained students take near-full advantage of additional credit in the first year.

Increased loan access also resulted in additional cumulative borrowing through year 4 from entry – the last year of postsecondary enrollment for the majority of students. Estimated coefficients from the event study specification are illustrated in Figure 1.8. Cumulative borrowing of constrained relative to unconstrained students increases starting with the 2008 cohort and levels out around \$4,000 by the 2010 cohort. Columns 1 and 2 in Table 1.5 present the corresponding difference-in-differences results when examining first-year and cumulative borrowing through year 4. The effect of increased loan access on first year and cumulative year 4 borrowing are \$1,214 and \$2,600, respectively.²¹ This is roughly a quarter of the increase one might expect if all constrained students fully take advantage of higher loan limits. This could reflect misclassification of some individual’s constrained status or changes in this status over the enrollment period. These explanations would both bias results down when examining borrowing and location choices, but do not pose a threat to identification.

1.3.4.2 The Effect of Increased Loan Limits on Post-College Migration

The additional debt for constrained borrowers caused them to choose initial post-college destinations with higher population densities. Figure 1.9 shows the event study coefficients when considering the effect of increased loan limits on county population density in year 6 from entry. As shown, there is a significant increase in year 6

21. Black et al. (2020) estimate that an increase of \$1,800 through year 4 in a national credit panel.

county population density for all post-period cohorts except 2007 and 2008. Column 3 of Table 1.5 shows the aggregate estimate from the difference-in-differences specification. Additional loan access increased the year 6 county population density by 2.2 percentage points. Scaling by the additional debt estimated in the previous section, this suggests that \$1,000 in additional student loans increased the population density of a borrower's year 6 county by 5.7 percent.

The additional debt for constrained borrowers also resulted in increased probability of locating in *top* urban areas after school. Figure 1.10 shows the event study estimates when considering the probability of locating in counties classified as being in the 95th percentile by population density. Figure 1.11 shows estimates from the same regression when considering the probability of locating in counties classified as a top metropolitan area with a population of more than 1 million, as specified by the Rural-Urban Classification Codes. A similar pattern to that when examining population density emerges with higher precision when considering these outcomes. Difference-in-differences estimation (Column 4 of Table 1.5) finds that a \$1,000 increase in debt caused constrained borrowers to be 4.2 percentage points more likely to locate in a county in the 95th percentile by population density – a roughly 8 percent increase from the mean. Similarly, a \$1,000 increase in student debt increased the probability of locating in a RUCC-defined top metropolitan area by 4.4 percentage points – a 6 percent increase from the mean. Columns 6 through 8 show that these results remain consistent when limiting the sample to individuals who reported a college address in California (year 3 from entry).

The Great Recession, officially starting in December of 2007 and ending in June of 2009, likely affected students differently depending on their academic level. Cohorts 2004 through 2009 all experienced the crisis at different points during their postsecondary education (assuming a 4-year degree). Event study results present some unexpected cohort heterogeneity that may reflect this fact. First, estimated

effects on location outcomes for 2007 and 2008 cohorts are consistently lower despite additional borrowing for constrained individuals in the 2008 cohort. One possible explanation is that individuals entering in these years, when the effects of the recession were most acute, may have adjusted their career goals. Liu, Sun, and Winters (2019) find that individuals who began postsecondary education in recession years were less likely to major in business and finance, both of which are fields that have relatively high urban wage premiums (as explored further in Section 1.4). If this phenomenon is widespread in 2007 and 2008 cohorts, it is possible that urban areas, which would normally attract student debtholders because of higher nominal wages, did not offer higher wages for these individuals and this is why the effect is dampened for these cohorts. Unfortunately, I cannot control for or explore this potential channel as college major is not observed in the UCCCP.

Another puzzling feature of the event study estimates is that constrained borrowers in the 2006 cohort showed little increase in student debt through year 4, but subsequently exhibited increased urban post-college location choice. One possible explanation is a change in enrollment behavior. Long (2014) find a decline in full-time enrollment accompanied by an increase in part-time enrollment during the Great Recession.²² As the 2006 cohort was only exposed to increased borrowing limits in year 4, it is possible that many had to shift to part-time enrollment, drop out, or extend the time to graduation. In the case of the latter, they may have benefited from additional borrowing capacity when they resumed full-time enrollment. Unfortunately, I cannot observe enrollment intensity or academic level so it is impossible for me to unpack this channel in this setting.²³

22. Van Horn et al. (2012) find that most individuals who were no longer enrolled as full-time students report the inability to afford the cost of college as the main reason.

23. The main issue is that additional borrowing in year 5, 6, etc. could be used to finish an undergraduate degree or for a graduate/profession degree. To some extent, Column 8 of Table 1.4 suggests little effect of additional borrowing on education attainment. Nevertheless, all specification include controls for education level reported in credit histories.

While data limitations prevent me from digging into these dynamics further in the current setup, subpopulation analysis allows me to minimize any confounding factors introduced by the economic downturn. Specifically, I drop all cohorts affected by the Great Recession (2004-2009) from the difference-in-differences estimation. Results of this restricted analysis are presented in Table 1.6. As shown, the results are very similar albeit larger given the omission of cohorts with only partial increases in loan access.

In Table 1.7, I explore the persistence of the effect of additional borrowing on the location choice of constrained borrowers. For all three outcomes measures characterizing destination counties, the effects are persistent through year 9 after entry (the last year with available data). While I cannot determine graduation dates at the individual level, this would typically be 5 years after exiting school for most individuals. This provides some initial evidence for how student debt may impact location choices in the long run as well.

1.4 Mechanisms

While results in the previous section show that additional student debt causes individuals to locate in urban areas at higher rates, data limitations prevent me from examining potential mechanisms within that empirical framework. To that end, there are two primary potential explanations for the identified effect of student debt on post-college location choice. The first is the proposed mechanism that student debt makes individuals more sensitive to nominal wages than local price levels. The second is that individuals who have a preference for urban areas are willing to borrow more for postsecondary education. This could be the case if, for example, individuals who aspire to live and work in New York City after school would also like to attend college in a large city, where tuition and related costs tend to be higher and may require additional borrowing. A necessary component of this alternative mechanism

is that individuals are making debt, college, and post-college location choices at least partially as a joint decision. In this section, I first show that the association between student debt and urban post-college location choice is robust to controlling for the second potential mechanism – suggesting it plays at most a minor role. Second, I present descriptive evidence supporting the proposed mechanism that student debt makes borrowers more sensitive to nominal wage differences in location choices.

To better understand the proposed mechanism and produce a testable implication of it, consider individuals with and without student debt (D) that choose a location j from choice set J based only on nominal wages (W_j) and local prices (P_j):²⁴

$$\max_{j \in J} \ln \left(\frac{W_j - D}{P_j} \right) \quad (1.3)$$

$$\epsilon_w = \frac{1}{1 - \frac{D}{W_j}} \geq 1, \quad \epsilon_p = -1$$

When individuals consider location options, one step is the consideration of the partial derivative with respect to wages and prices. For individuals with no student debt, their location choices are equally elastic to nominal wages and local price levels. In this case, the standard result holds: individuals locate in places with the highest real wages. When individuals carry student debt, the elasticity with respect to wages is strictly greater than the elasticity with respect to prices. As a result, individuals with student debt choose to locate in higher nominal wage locations. Additionally, this framework shows that, among individuals with the same level of student debt, urban locations should only be relatively more attractive if they provide significantly higher nominal wages for that individual. This is the testable implication I will examine further.

24. This framework is a simplified version of the standard Rosen (1979) and Roback (1982) frameworks.

1.4.1 Data

To explore these potential mechanisms, I exploit rich data from the Education Longitudinal Study of 2002 (ELS). The main advantage of using the ELS instead of the UCCCP is that it provides a rich set of individual-level characteristics to perform heterogeneity analysis on the association between student debt and urban post-college location choice. It contains a nationally representative sample of over 16,000 10th graders in 2002 and 12th graders in 2004. The ELS includes a survey of students, their parents, and school officials with the stated goal of understanding student trajectories from high school through postsecondary education and into the workforce. Individuals are surveyed 4 times: 2 times during high school years (2002 and 2004); 1 time two years after prospective high school graduation during typical college-going years (2006); and 1 time 8 years after prospective high school graduation (2012), which would typically be 4 years after postsecondary education is completed. At each point the survey documented their location at the county level. The ELS also collects information on total amount borrowed for postsecondary education. Demographic information includes parental education/occupation, ability (via high school grade point average), college major, sex, and race as additional factors that may influence migration choices.

To reduce unobserved heterogeneity, I limit the ELS sample in two ways. First, I examine only outcomes of those with a postsecondary degree – omitting both individuals who end the observation period with higher or lower levels of educational attainment. As labor market opportunities vary significantly by education level, this sample restriction aims to eliminate this confounding factor. Second, I limit the sample to individuals who borrowed a positive amount to fund postsecondary education, i.e. the intensive margin. This was done to reduce unobserved heterogeneity as there is likely a larger amount of unobserved factors that enter into the borrowing decision on the extensive margin than on the intensive margin. Finally, both restrictions offer

the added benefit of constructing a sample that is similar to the UCCCP sample used in Section 1.3. The resulting sample includes roughly 2,200 individuals.

1.4.2 Urban Preference Mechanism

I begin by showing a robust association between student debt and urban post-college location choice in the ELS sample. I do so by estimating the following cross-sectional regression:

$$\begin{aligned} \text{Post-College Cty. Urban}_i = & \alpha + \beta_1 \text{Debt}_i + \mathbf{X}'_i \boldsymbol{\beta}_x + \mathbf{1}[\text{Origin Cty. Urban}_i] \\ & + \mathbf{1}[\text{College Cty. Urban}_i] + \gamma_{\text{Degree Field}(i)} + \epsilon_i \end{aligned} \quad (1.4)$$

where the outcome of interest is the urban classification of the post-college location of an individual. The urbanicity of a location is defined in three ways: (1) a continuous measure of a county’s population density (in logs); (2) an indicator function for if the county is in the 80th percentile by population density; and (3) an indicator function for if the county is classified in the top category (metropolitan area with a population of at least 1 million) using the Rural-Urban Continuum Codes (RUCC). The coefficient of interest, β_1 , captures the association between urban post-college location choice and student debt (in logs).

To partially account for urban preferences, I control for the urban classifications of high school and college locations based on population density or RUCC. To account for geospatial variation in industry concentration and the resulting variation in available employment opportunities for individuals with different academic majors, I include fixed effects for postsecondary degree field as reported in the ELS (γ). I also account for a rich set of individual characteristics ($X'_i \beta_x$), including father’s education, father’s income, race, sex, and ability (via high school GPA). The regression is weighted based on the ELS sample design with errors clustered at the high school level.

To provide a baseline, I first estimate the coefficients from Equation 1.4 when not including the proxies for urban preference. These results are presented in Columns 1-3 of Table 1.8. As Column 1 shows, student debt has a positive and significant relationship with post-college county population density. A 10 percent increase in student debt is associated with a 1.4 percent increase in post-college county population density. The positive relationship between student debt and urban post-college location choice is also present when we consider locating in top urban areas. As shown in Column 2, a 10 percent increase in student debt is associated with a 0.3 percent increase in the probability of locating in a county in the 80th percentile of population density. Results in Column 3, where urbanicity is defined as the top category in the RUCC, are similar in magnitude but estimated with lower precision.

Controlling for urban preference reduces the magnitude of these estimates, but only accounts for roughly one third of the association between student debt and urban post-college location. Estimated coefficients from Equation 1.4 when including the proxies for urban preference are presented in Columns 4-6 of Table 1.8. As shown, a 10 percent increase in student debt is still associated with a 1.1 percent increase in post-college county population density. The association is also still present when we consider locating in top urban areas. As shown in Column 5, a 10 percent increase in student debt is associated with a 0.2 percent increase in the probability of locating in a county in the 80th percentile of population density.²⁵ Results in Column 6, where urbanicity is defined as the top category in the RUCC, are qualitatively similar but now lack precision.

These results provide two indications that the urban preference mechanism is unlikely to be the primary driver of the results in Section 1.3. First, as expected, both origin and college county urban indicators have a strong association with post-college urban location; however, a robust relationship between student debt and post-college

25. As the 80th percentile threshold is arbitrary, I adjust the cutoff in Figure 1.26.

location choice is still found despite including these proxies for urban preference. Second, as noted above, the urban preference mechanism requires that individuals make college, student debt, and post-college location choices at least partially as a joint decision. Since Conzelmann et al. (2022) have shown that a large portion of college graduates live and work near their postsecondary institution, this implies that individuals making a joint decision that includes post-college location would likely have the same college and post-college location. This suggests one possible way to address this concern would be to examine individuals who *did not* locate in the same place for postsecondary education and post-college life. These individuals may be making their post-college location choices more independently of college and debt decisions. Columns 7 through 9 in Table 1.8 conduct the same analysis on the subpopulation of individuals who left their college county. As shown, the results are qualitatively similar and slightly larger although not significantly so.²⁶

1.4.3 Association Driven by Individuals Facing High Urban Wage Premiums

The proposed mechanism for why additional student debt would cause individuals to locate in urban areas at higher rates is that student debt makes individuals more sensitive to nominal wage differences across locations than differences in local price levels. As noted above, an implication of this mechanism is that the association between student debt and urban post-college location choice should be higher for individuals that face steeper urban wage premiums. To test this, I estimate the wage premium associated with increases in population density for each degree field reported

26. The same result holds when considering the subpopulation of individuals that left their college commuting zone, which perhaps is a more accurate measure of the entire local labor market where the college is located.

in the ACS 2010 sample.²⁷ By mapping ELS-reported postsecondary academic major to ACS degree field, I can then determine the urban wage premium that each student faces upon graduation.

Empirically, I expand on Equation 1.4 by interacting student debt with an indicator function for whether an individual entered a field with an above- or below-median urban wage premium:

$$\begin{aligned}
 \text{Post-College Cty. Urban}_i = & \alpha + \beta_1^0 \text{Debt}_i \times \mathbf{1}[\text{Below Median Urban Wage Premium}_i] \\
 & + \beta_1^1 \text{Debt}_i \times \mathbf{1}[\text{Above Median Urban Wage Premium}_i] \\
 & + \mathbf{X}_i' \boldsymbol{\beta}_x + \mathbf{1}[\text{Origin Cty. Urban}_i] \\
 & + \mathbf{1}[\text{College Cty. Urban}_i] + \gamma_{\text{Degree Field}(i)} + \epsilon_i \quad ,
 \end{aligned} \tag{1.5}$$

where β_1^0 and β_1^1 capture the association between student debt and urban post-college location choice for individuals facing below- and above-median urban wage premiums by field, respectively. The rest of the specification remains the same as in Equation 1.4 with the exception of the degree field fixed effects which now capture ACS rather than ELS postsecondary degree fields. Though the divisions change slightly, these fixed effects still capture spatial variation in employment opportunities by degree field.

The results of this analysis are reported in Table 1.9. Columns 1, 4, and 7 repeat the results in Table 1.8, while columns 2, 5, and 8 conduct this analysis on the sample that have majors mapped to ACS degree fields. Results are nearly identical, suggesting there is no issue of sample selection in matching rates by ELS major. Columns 3, 6, and 9 report the results from the estimation of Equation 1.5 on the

27. I derive wage premium as the coefficient on PUMA population density interacted with postsecondary degree field in a regression of nominal wages on sex, age, race fixed effects, and degree field fixed effects. Sample includes employed prime working-age (22-54) individuals making a positive income. ACS degree field only recorded for Bachelor's degree holders.

three measures of post-college location urbanicity. As shown, the positive relationship between student debt and urban post-college location is driven by individuals facing above-median urban wage premiums. For individuals who enter fields with below-median urban wage premiums, additional student debt is not associated with any increased likelihood of locating in urban areas. This is consistent with the proposed mechanism that student debt makes individuals more sensitive to spatial variation in nominal wages.

1.5 Spatial Equilibrium Model

In this section, I introduce a spatial equilibrium model that delivers the empirical regularities presented so far and additionally affords a framework to simulate counterfactual outcomes under various policy interventions. The setup is similar to the canonical Rosen (1979) and Roback (1982) framework, but I allow for heterogeneity in workers' productivity and location preferences as well as cities' productivity and housing supply. To focus on the role of student debt, I minimize departures from 'standard' spatial equilibrium models in the literature, and follow recent versions from Diamond (2016), Hsieh and Moretti (2019) and Albert and Monras (2017).

The model has J regions that differ by low- and high-skill productivity, amenity level, and housing supply. Firms in all locations produce an identical good that is freely traded between regions at no cost. Individuals consume this tradable good as well as a non-tradable local good, which for simplicity I will call housing. The housing sector generates the congestion force in the model, as housing prices and population have a positive relationship.

In the rest of the section, I imbed the basic mechanism outlined in Equation 1.3 in a comprehensive discrete location choice model with two periods. In period 1, individuals choose a level of education/skill and take out student debt if they choose to

become high-skill workers. In period 2, individuals choose a location, work, consume local and tradable goods, and pay off student debt. In period 1, the model only has one component: a combined education and debt choice by individuals. In period 2, the model has three components: labor supply (location choice), labor demand, and the housing market. A banking sector spans both periods to facilitate borrowing for education. To allow for policy counterfactuals involving debt relief and tuition-free postsecondary education, a ‘federal’ government institutes a tax to cover the associated cost. Because decisions in period 1 are made recursively based on utility in period 2, I present period 2 setup first followed by period 1.

1.5.1 Labor Supply (Period 2)

Individuals enter period 2 as low or high-skill workers ($g \in \{l, h\}$). If they are part of the high-skill group, they also carry a positive student debt (D_g), which is constant for all individuals in the group. Student debt is zero for all low-skill individuals. Workers choose location j , consumption of a tradable good C_T (numeraire good), and consumption of a non-tradable good, C_{NT} (at price p_j). The utility of individual i from group g locating in j is given by:

$$\ln U_{igj} = (1 - \beta) \ln C_T + \beta \ln C_{NT} + \ln A_{gj} + \ln \epsilon_{ij} , \quad (1.6)$$

where $(1 - \beta)$ denotes the expenditure share devoted to tradable goods. A_{gj} is the utility derived from local amenities in location j , and is group-specific. Finally, individuals have an idiosyncratic location preference ϵ_{ij} that has a Fréchet distribution with inverse shape parameter $\alpha \geq 0$, which governs the variance of the idiosyncratic taste shocks. C_{NT} represents the consumption of housing and other non-tradable goods which need to be consumed in location j . For simplicity, I will call this housing.

Individuals maximize utility subject to the following budget constraint:

$$C_T + p_j C_{NT} + D_g \leq W_{gj}, \quad (1.7)$$

where wages, W_{gj} , vary by group (skill level) and location. The demand for each good is given by:

$$C_T = (1 - \beta)(W_{gj} - D_g) \quad , \quad C_{NT} = \beta \left(\frac{W_{gj} - D_g}{p_j} \right) \quad (1.8)$$

Plugging the optimal demand functions into the utility function, the indirect utility of living in location j is:

$$\ln V_{igj} = \ln V_{gj} + \ln \epsilon_{ij} = \kappa + \ln(W_{gj} - D_g) - \beta \ln p_j + \ln A_{gj} + \ln \epsilon_{ij}, \quad (1.9)$$

where $\kappa = \beta \ln \left[\beta (1 - \beta)^{\frac{1-\beta}{\beta}} \right]$. Indirect utility can be broken into a group-specific valuation of location j , V_{gj} , and the individual idiosyncratic preference for locating in j , ϵ_{ij} .

Given the distribution of the idiosyncratic taste parameter, the share of workers in group g locating in j is equal to:

$$\pi_{gj} = \frac{(V_{gj})^\alpha}{\sum_k (V_{gk})^\alpha} = \left(\frac{V_{gj}}{V_{gJ}} \right)^\alpha \quad (1.10)$$

where $V_{gJ} = \left[\sum_{j' \in J} (V_{gj'})^\alpha \right]^{\frac{1}{\alpha}}$,

where V_{gJ} represents the expected value, or welfare, of being in this economy for a worker in group g . The shape parameter α on the idiosyncratic preferences governs the elasticity of migration with respect to changes in indirect utility of locations.²⁸

28. Previous work, including Diamond (2016) and Bound and Holzer (2000), has found differ-

Assuming each worker inelastically supplies one unit of labor, the overall supply of low- and high-skill labor in location j is given by:

$$L_j = \pi_{lj}N_{lJ} \tag{1.11}$$

$$H_j = \pi_{hj}N_{hJ}, \tag{1.12}$$

where N_{lJ} and N_{hJ} are total low- and high-skill workers in the economy, respectively.

I, the econometrician, observe wages (W_{gj}), low- and high-skill population (L_j and H_j), price levels (p_j), and high-skill debt levels (D_g). Exogenous amenities (A_{gj}) and workers' idiosyncratic taste for each location (ϵ_{ij}) are unobserved. Parameters to be calibrated are the worker expenditure share devoted to non-tradable goods, β , and the shape parameter for the idiosyncratic location preference, α .

1.5.2 Labor Demand (Period 2)

Each location j has a single firm that produces the tradable goods with a production function that combines low- and high-skill labor as the only inputs. The output in location j is given by:

$$Y_j = (\theta_{lj}L_j^\rho + \theta_{hj}H_j^\rho)^{\frac{1}{\rho}} \tag{1.13}$$

$$\theta_{gj} = \exp(\epsilon_{gj})$$

where firms combine low- and high-skill labor as imperfect substitutes in production with a constant elasticity of substitution, $\frac{1}{1-\rho}$.²⁹ Skill-specific productivity, θ_{lj} and θ_{hj} , differ across locations and are determined exogenously. Labor markets are assumed to

ences between high- and low-skill workers in migration elasticity with respect to real wages. In the counterfactual analysis, I simulate outcomes with both uniform and skill-specific shape parameters, α_g .

29. This production function is prominent in the literature examining wage inequality and its relation to supply of high- and low-skill labor (Katz and Murphy 1992; Katz and Autor 1999; Acemoglu 2002; Diamond 2016).

be perfectly competitive such that wages equal the marginal product of labor. Profit maximization leads to the following demand for low- and high-skill labor in location j :

$$W_{lj} = \theta_{lj} L_j^{\rho-1} (\theta_{lj} L_j^\rho + \theta_{hj} H_j^\rho)^{\frac{1-\rho}{\rho}} \quad (1.14)$$

$$W_{hj} = \theta_{hj} H_j^{\rho-1} (\theta_{lj} L_j^\rho + \theta_{hj} H_j^\rho)^{\frac{1-\rho}{\rho}} \quad (1.15)$$

I observe wages (W_{gj}), low, and high-skill population (L_j and H_j). Local exogenous productivity, ϵ_{lj} and ϵ_{hj} , are unobserved. The parameter governing the elasticity of substitution between high- and low-skill labor, ρ , needs to be calibrated.

1.5.3 Housing Market (Period 2)

In each location, the supply of housing is produced using land for homes (T_j), which is a fixed factor, and the tradable good (Y_j^T) according to the following production function.³⁰

$$Y_j^{NT} = \zeta_j^{-\zeta_j} (Y_j^T)^{\zeta_j} (T_j)^{1-\zeta_j} \quad , \quad (1.16)$$

where $1 - \zeta_j$ is the weight of land in the production of housing. I assume land is owned by absentee landlords.³¹ This results in the following housing supply equation:

$$Y_j^{NT} = p_j^{\gamma_j} T_j, \quad (1.17)$$

where $\gamma_j = \frac{\zeta_j}{1-\zeta_j}$ is the housing supply elasticity. Note that γ_j differs across locations, which could capture a combination of factors such as limits on the amount of land

30. Housing production function also used in Albert and Monras (2017). Resulting housing supply identical to reduced-form version used in Hsieh and Moretti (2019).

31. A common assumption in the literature. See Albert and Monras (2017), Diamond (2016), Hsieh and Moretti (2019), and Eeckhout, Pinheiro, and Schmidheiny (2014).

or land use regulations (Saiz 2010). Cities with limited land or cumbersome land use regulations have a lower γ_j (lower elasticity), while cities with few limits on available land or land use have a higher γ_j (higher elasticity).

Total demand for housing is given by the sum of the local demands of individuals in each location. Local housing prices are implicitly defined by market clearing in each location:

$$p_j^{\gamma_j} T_j = L_j \beta \frac{W_{lj}}{p_j} + H_j \beta \frac{W_{hj} - D_h}{p_j} \quad (1.18)$$

This equation captures one difference between my model and standard spatial equilibrium models (such as Hsieh and Moretti (2019)): the demand for housing in each location depends on the size *and* skill composition of the population.³²

As the econometrician, I observe price levels (p_j), low- and high-skill population (L_j and H_j), and high-skill debt levels (D_h). Location-specific housing supply elasticities (γ_j) need to be estimated and land available for homes (T_j) is unobserved.

1.5.4 Government (Period 2)

In the baseline model, there is no government; however, certain policy counterfactuals involving student debt forgiveness and the elimination of tuition require a government to facilitate them. In each case, there is an aggregate cost, T , that will be funded by uniform taxation of nominal wages at rate τ .³³ The government budget constraint is given by:

$$T = \sum_j (L_j W_{lj} + H_j W_{hj}) * \tau \quad (1.19)$$

It is important to note that, unlike debt repayment, the tax is proportional to wages

32. This is similar to the model in Albert and Monras (2017), where demand for housing depends on location-specific size and *immigrant* composition of the population.

33. In the case of debt forgiveness, $T = D_h * N_{h,J}$. When considering a free college option, $T = D_h * N_{h,J}^{CF}$ for the counterfactual high-skill population.

and, therefore, does not distort location choices. This can be seen by considering the group-specific indirect utility of living in location j (Equation 1.9) under a policy eliminating debt and instituting a tax:

$$\ln V_{gj} = \kappa + \ln(W_{gj} * (1 - \tau)) - \beta \ln p_j + \ln A_{gj} \quad (1.20)$$

As shown, τ is separable from wages and thus cancels out in Equation 1.10 determining worker shares in each location.

1.5.5 Education and Debt Choice (Period 1)

In period 1, individuals choose their skill level for period 2 employment. There is no cost to becoming a low-skill worker. To become a high-skill worker, individuals must pay a monetary cost, D , that is paid by borrowing in period 1 and repayment in period 2. There is also an idiosyncratic utility cost, z_i , which follows a Frechét distribution with inverse shape parameter $\mu \geq 0$. This utility cost captures two traits: (1) an individual's ability to become a high-skill worker (e.g., aptitude); and (2) an individual's ability (including willingness) to take on student debt. The joint education-debt decision is made based on the period 2 expected value of being in this economy as a worker of each type ($g \in \{l, h\}$):

$$U_i^g = V_{gJ} + z_i * \mathbf{1}[g = h], \quad (1.21)$$

where V_{gJ} is the same as in Equation 1.10 and represents expected value of being in the period 2 economy as a worker in group g . An individual chooses to become a high-skill worker if $U_i^h > U_i^l$.

Given the distribution of the idiosyncratic preferences, the share of individuals

that become high-skill workers can be represented as:

$$s_h = \frac{(e^{U_i^h})^\mu}{(e^{U_i^l} + e^{U_i^h})^\mu} \quad (1.22)$$

As shown, μ , governs the elasticity of worker type choice with respect to changes in expected utility of being a high-skill worker. This includes changes induced by adjustments to student debt – an important detail for the counterfactual exercises.

As the econometrician, I do not observe U_i^H , U_i^L , μ , or z_i ; however, as discussed further in the calibration and counterfactual analysis sections, the relevant dimension is $\frac{\partial s_h}{\partial D}$.

Banking Sector (Periods 1 and 2). For simplicity, the banking sector offers credit to all individuals at 0 interest rate.

1.5.6 Equilibrium

Definition I. *The spatial equilibrium is defined as follows:*

1. *Workers decide where to live and how much to consume of each good.*
2. *Firms decide how many workers to hire to maximize profits.*
3. *Landlords decide how much housing to supply.*
4. *Tradable goods, labor, and housing markets clear.*
5. *Government budget constraint satisfied.*

1.6 Calibration

The location choice set is defined as the 722 1990-defined Commuting Zones (CZ) in the contiguous United States.³⁴ Related prior literature has used metropolitan statistical areas, but this does not allow for full coverage of the U.S. and would omit migration across the dimension of interest: the urban-rural divide. There are two main sources of data for the structural analysis: the 2015-2019 American Community Survey (ACS)³⁵ and the National Center for Education Statistics (NCES). ACS provides the baseline levels of wages and population for low- and high-skill workers. Local prices are derived from housing costs in ACS, following the approach of Moretti (2013) for 1990 Commuting Zones. Baseline levels of student debt among high-skill workers was estimated from the NCES data on cumulative borrowing for bachelor's degree-holders.³⁶

1.6.1 Parameters in Labor Supply

There are two parameters that need to be calibrated in labor supply: worker expenditure share devoted to non-tradable goods, β , and the shape parameter governing idiosyncratic location preference and migration elasticity, α . The expenditure share devoted to non-tradable goods, β , will be taken from the literature and set to 0.6.³⁷ Migration elasticity, α , is taken from the literature as well, but there is far less consensus. Most estimates center around 4, but some are as low as 0.4 and others as high

34. Alaska and Hawaii excluded due to the unique nature of their labor markets and the discontinuous costs associated with moving to either location.

35. Downloaded from IPUMS (Ruggles et al. 2022). See Section 1.B for more information.

36. Payments based on typical repayment period (10 years) and interest rate (4.99%) for federal student loans.

37. Hsieh and Moretti (2019) find β of 0.6, which is used by Monras and Albert (2022). Moretti (2013) finds local good expenditure of 0.59. Diamond (2016) uses 0.62, supported by analysis of the Consumer Expenditure Survey.

as 12.8.³⁸ The breadth of these estimates likely reflect the paper-specific variation and estimation strategy in each. Higher values may also reflect elasticity over a longer time horizon to some extent. Additionally, the literature has found that migration elasticity may vary by skill group, with low-skill workers being less mobile than high-skill workers.³⁹ As this parameter is a crucial component to generating counterfactual population distributions, the simulations below will report results under 2 values of migration elasticity. First, I consider homogeneous migration elasticity across skill groups and set it around the median in the literature, 4. I then consider heterogeneity in this parameter by setting low- and high-skill migration elasticity to the values estimated in Hornbeck and Moretti (2018).

Exogenous amenities (A_{gj}) are estimated from Equation 1.10. Specifically, by setting one location (k) as the reference location, we can consider the share of individuals in skill group g locating in j relative to k :

$$\ln \left(\frac{\pi_{gj}}{\pi_{gk}} \right) = \alpha \ln \left(\frac{\frac{W_{gj} - D_g}{P_j^\beta}}{\frac{W_{gk} - D_g}{P_k^\beta}} \right) + \alpha \ln \left(\frac{A_{gj}}{A_{gk}} \right) \quad (1.23)$$

Population shares, wages, and debt are all observed, allowing me to derive amenities as the residual. This allows for a perfect fit of the location choice data. In counterfactual analysis, amenities will remain fixed.

38. Hsieh and Moretti (2019) and Hornbeck and Moretti (2018) estimate 3.3. Suárez, Carlos, and Zidar (2016) estimate 0.75-4.2; Caliendo, Dvorkin, and Parro (2019) estimate 2; Albert and Monras (2017) estimate 12.8; and Monras (2018) estimate 0.4.

39. Hornbeck and Moretti (2018) estimates $\alpha_l = 2.6$ and $\alpha_h = 6.7$; Diamond (2016) estimates $\alpha_l = 2.1$ and $\alpha_h = 4$, but that low-skill migration elasticity is only significant at the 10-percent level. Bound and Holzer (2000) find that college workers' migration is elastic to local demand, but that low-skill workers are inelastic.

1.6.2 Parameters in Labor Demand

The parameter governing the elasticity of labor substitution in production, ρ , will be taken from the literature. Diamond (2016) and Card (2009) provide estimates using the same production framework, finding estimates of $\rho = 0.4$ and $\rho \in (0.3, 0.6)$.⁴⁰ For the counterfactual exercises, I will use ρ equal to 0.4, which implies an elasticity of labor substitution of 1.7. With ρ , I estimate skill-specific productivities for each location j using labor demand (Equation 1.14) and baseline data on wages and population. Exogenous productivity will remain fixed in counterfactual simulations.

1.6.3 Parameter in Housing Market

There are two elements of the housing market that need to be calibrated or estimated: elasticity of housing supply in each location (γ_j) and the land available for housing (T_j). Location-specific elasticities of housing supply, γ_j , are calculated by commuting zone similar to Saiz (2010) and Howard, Liebersohn, et al. (2018). The land available for housing, T_j , is calculated using the housing market clearing condition (Equation 1.18) and observed low- and high-skill population, wages, debt, and price levels.

1.6.4 Parameters in Education/Debt Choice

For the counterfactual analysis, the parameter of interest is the elasticity of high-skill worker population share with respect to required student debt ($\frac{\partial s_h}{\partial D}$). Although estimates of this elasticity exist in the literature, the counterfactual analysis below only requires an estimate for the change in high-skill worker share associated with one discrete change in debt: a free college option. As most of the literature estimates this elasticity from marginal changes in the cost of attendance, a complete elimination

40. Katz and Autor (1999) also provides similar estimates using earlier data.

of cost is likely to have an effect that is different from estimates extrapolated from this literature. For this reason, I calibrate the change in high-skill worker share associated with the introduction of a free college option from the literature examining this exact scenario. Kane (2003) exploits discontinuities in the California Grant program eligibility, which is a program that provides enough aid to cover the cost of attending University of California or California State University institutions. They estimate that a free college option increases enrollment by 3-4 percentage points. In a similar analysis examining expansion of the Georgia HOPE Scholarship Program, which waives tuition and fees for eligible students, Dynarski (2008) finds that a free college option increases college completion by 3 percentage points.

1.7 Counterfactual Analysis

With the model fully calibrated, I now use it to estimate counterfactual outcomes under four scenarios. First, I simulate outcomes if the average debt required for post-secondary education had stayed constant at the 1980 level. The purpose of doing so is to estimate the contribution of the growth in student debt to post-1980 skill-based geographic sorting. In the remaining counterfactual simulations, I estimate the effect of 3 policy proposals that address student debt on sorting, consumer welfare, and aggregate output: (1) debt forgiveness for *existing* borrowers⁴¹; (2) income-driven repayment plans for *existing* borrowers; and (3) the institution of tuition-free postsecondary education (accessible to all individuals). When considering debt forgiveness or the elimination of tuition, I pay for the policy with a uniform tax on all individuals.

To analyze effects on skill-based geographic sorting for each counterfactual, I use the model to simulate the population distribution then discuss the effect on the *level*

41. The debt forgiveness program proposed by the Biden administration forgives up to \$20,000 for student debt-holders, which is paid for by taxpayers. In my counterfactual analysis, I consider complete debt relief.

of skill-based geographic sorting. I do so by calculating the correlation between the counterfactual high-skill worker shares and the population density in each commuting zone.⁴² This approach, consistent with the literature’s method of documenting skill-based sorting, provides a concise and transparent comparison to observed sorting. The observed levels of spatial sorting are presented in Figure 1.12. In the case of setting debt to the 1980 level, I also consider the effect on the level *relative to 1980*.

The model also enables me to reflect on aggregate outcomes of output and consumer welfare when considering the three policy proposals. The change in aggregate output can be calculated by summing across location-specific production (Equation 1.13) under baseline and counterfactual population estimates. The change in aggregate consumer welfare is measured using the expected value of being in the baseline and counterfactual economy for a worker in each group (V_{gJ} in Equation 1.10):

$$\widehat{W} = \frac{N_l}{N} \widehat{V}_{lJ} + \frac{N_h}{N} \widehat{V}_{hJ} \quad (1.24)$$

The results of the counterfactual analysis, reported in Table 1.10-1.12, are broadly split into two panels for partial and full equilibrium simulations separately as both offer unique advantages and together provide a more complete picture. Partial equilibrium results allow for only a labor supply response via migration, but keep wages and local prices fixed. By remaining agnostic about wage and price responses, both of which have been scarcely estimated in the literature for rural areas, the partial equilibrium results offer the most straightforward analysis. Additionally, partial equilibrium results provide an understanding of short-term outcomes. Full equilibrium results, in contrast, allow for wage and local price responses. In doing so, they account for spillover effects on low-skill workers and perhaps more accurately depict long-term outcomes. In both, location- and skill-specific amenities remain fixed. Individuals

42. Results consistent across alternative measures of urbanicity, including local wages and prices.

can adjust their joint education-debt decisions in both considerations of tuition-free postsecondary education. Within each, I simulate the model under various values of the most important parameter: migration elasticity to real wages, as discussed in Section 1.6.

Counterfactual simulations that hold wages and prices fixed (partial equilibrium) may overestimate the effect of resorting for two possible reasons. First, to the extent that prices in initially low-population areas respond to a rising population, this would discourage large growth in rural areas depending on location-specific housing supply elasticity. As a result, partial equilibrium solutions would overestimate the counterfactual resorting. Second, the shape of skill-specific demand curves. In the case of traditional downward-sloping demand for skill workers, wages for high-skill workers would increase in areas where the high-skill population declines – in turn, reducing the outflow. This would mean that partial equilibrium counterfactuals are an overestimate of resorting.⁴³

1.7.1 Counterfactual 1: Debt Reversed to 1980 Level

I begin the counterfactual analysis by examining the role of rising student debt in post-1980 skill-based geographic sorting. The partial equilibrium results for this analysis are reported in Panel A of Table 1.10, Columns 1 and 2. Column 1 shows the percentage reduction in post-1980 sorting due to a decrease in debt to the 1980 level. In other words, these estimates reflect the share of post-1980 sorting that can be attributed to the growth of student debt over this period. The estimates range from 11.5 percent to 18.7 percent in the partial equilibrium solutions. Since wages

43. Alternatively, demand for high-skill workers could be upward-sloping depending on the strength of agglomeration economies in production. For simplicity, the model presented here does not incorporate agglomeration economies, but Diamond (2016) offers a framework to do so. If high-skill worker demand exhibits strong enough agglomeration economies, wages for high-skill workers would decline in urban areas as high-skill workers relocate. As a result, partial equilibrium results would *underestimate* the reduction in sorting.

and prices remain fixed, there are no spillovers to low-skill workers and variation in partial equilibrium solutions reflect only adjustments to migration elasticity for high-skill workers: higher elasticity equates to a larger exodus from urban areas when debt is decreased. Column 2 reports the estimated impact on the *level* of 2019 geographic sorting. The implications for geographic sorting are directionally the same but smaller in magnitude when considering full equilibrium solutions (see Panel B of Table 1.10, Column 1 and 2). The growth of student debt explains 3.5-4.5% of the post-1980 skill-based geographic sorting in this case – a smaller adjustment that reflects the dampening effects of rising rural prices and urban high-skill wages over the long-term. These results are also depicted in Figure 1.13 and 1.14 under the heterogeneous migration elasticities provided by Hornbeck and Moretti (2018). As shown, the observed growth of student debt resulted in the largest increases in high-skill worker shares in dense urban commuting zones, while rural areas experience a decrease in their high-skill workers shares relative to the counterfactual where student debt remains at the 1980 level in real terms.

1.7.2 Counterfactual 2: Debt Forgiveness

Next, I examine the model-predicted outcomes for the first of three counterfactual policies: debt forgiveness for *existing borrowers*. To facilitate debt forgiveness, the government levies a uniform tax on all individuals that is a fixed percentage of wages. As the tax is a constant share of income, it does not distort the spatial allocation of low or high-skill workers. Predictions for the effect on the *level* of skill-based geographic sorting in 2019 are reported in Column 3 of Table 1.10. In partial equilibrium, debt forgiveness significantly reduces sorting by 8.1-14.0 percent depending on the assumed migration elasticities. In full equilibrium simulations, this percentage drops to a still significant 2.4-3.4 reduction. Debt forgiveness and the resulting migration decisions increase welfare for high-skill workers, but comes at the cost of low-skill workers who

now pay a tax without the benefit of debt relief (Column 1 in Table 1.11). There is also a reduction in aggregate output as high-skill workers relocate to lower-productivity areas, as shown in Table 1.12, but these results should be interpreted with caution as production functions outside of urban areas have been poorly estimated in the literature.

1.7.3 Counterfactual 3: Income-Driven Repayment

The second counterfactual policy proposal I consider is switching all high-skill workers to income-driven repayment plans. To do so without reducing the aggregate debt level, I calibrate the repayment rate such that aggregate debt repayments remain fixed at the baseline level.⁴⁴ The resulting effect on skill-based geographic sorting is reported in Column 4 of Table 1.10. In partial equilibrium, the results are the same as in the counterfactual of debt forgiveness – as expected given that both entail a shift to debt repayment methods that do not distort location choices. In the full equilibrium solutions, the picture is a bit different: the reduction in spatial sorting is lower than under debt forgiveness. When considering debt forgiveness, the costs of education are paid uniformly by all; however, with income-driven repayment, the costs remain with high-skill workers. This does not distort high-skill worker location choices, but it does affect prices in the full equilibrium solution. It lowers prices in areas with larger high-skill worker shares (see Equation 1.18), which in turn reduces the exodus of high-skill workers. This dampening effect is stronger when high-skill workers are more mobile than low-skill workers. The maps provided in Figure 1.15 and 1.16 provide insight on how resorting affects commuting zone high-skill worker shares across the country. As shown, the largest declines are in commuting zones

44. In real life, income-driven repayment plans are capped at the standard repayment amount, often extend repayment periods, and involve some measure of debt relief for remaining balances at the end of the repayment period. To fully capture these details, heterogeneity in debt as well as a multi-period model is needed.

containing large metropolitan areas like New York City, Los Angeles, and Austin, Texas. The largest increases are seen in the Midwest.

The effects on welfare are also different under this counterfactual policy as shown in Column 2 of Table 1.11. The most notable difference is that they are an order of magnitude smaller than when considering debt forgiveness – a fact that reflects the absence of a shift of the education costs from high to low-skill workers. The remaining changes in welfare are driven by migration choices that reflect preferences. In both partial and general equilibrium, welfare for high-skill workers increases as the distortion of debt repayments is removed. In partial equilibrium, there is no effect on low-skill workers as prices and wages remain the same. In full equilibrium, low-skill workers actually see an increase in welfare as they are less ‘crowded out’ from urban areas. There is also a reduction in aggregate output as high-skill workers relocate to lower-productivity areas, as shown in Table 1.12.

1.7.4 Counterfactual 4: Tuition-Free Postsecondary Education

In the final counterfactual simulation, I consider the policy proposal of eliminating tuition for postsecondary education. The cost will instead be borne by all via a uniform tax rate. Individuals will be able to adjust their education decisions, i.e. additional individuals can choose to become high-skill workers that would not otherwise have done so when debt is required. The resulting effect on skill-based geographic sorting is reported in Column 5 of Table 1.10. In partial equilibrium, the reduction in sorting ranges from 5.7-11.6% while there is virtually no change in full equilibrium. This reflects how two opposing forces interact: a rising national share of high-skill workers are still drawn to cities by higher productivity, but their location choices are no longer distorted by debt repayment. Welfare increases for high-skill workers as they no longer bear the full cost of education while it decreases for low-skill workers as they assume some of the cost. Output increases as the aggregate supply of high-skill

labor, which is more productive, increases.

1.8 Conclusion

In this paper, I have shown that student debt distorts the location choices of high-skill workers, increasing the probability that they locate in dense, urban areas. I estimated this effect using a policy change that increased federal student loan limits, a recently developed large consumer credit panel, and a difference-in-differences framework. I then presented descriptive evidence supporting the proposed mechanism that student debt makes individuals more sensitive to nominal wages in location decisions. More specifically, I showed that the association between student debt and urban post-college residency is driven by those who attain degrees with higher urban wage premiums, i.e., those for whom urban areas offer relatively large nominal wage gains. I also provided evidence that correlation between urban preference and student debt is unlikely to be the primary driver of the identified effect.

To quantify the impact on aggregate spatial sorting, I incorporated student debt into a standard spatial equilibrium model and ran counterfactual simulations to show that the growth of student debt since 1980 accounted for 3.5-4.5 percent of skill-based geographic sorting over the same period. Estimates based on partial equilibrium analysis, which remains agnostic about price and wage responses, suggest this share could be 11.5-18.7 percent. The model also enabled me to reflect on the implications of various policy proposals under consideration by policy-makers, including debt forgiveness, income-driven repayment plans, and a free college option. While all policies eliminated the distortion of high-skill location preferences caused by traditional student debt repayment, only income-driven repayment did so while strictly increasing welfare.

The findings in this paper represent the first examination of the impact of student

debt on individual location choices and the aggregate spatial distribution of high-skill workers. Regarding the magnitude of aggregate counterfactual estimates, I would like to highlight three caveats to the model predictions. First, the model assumes homogeneous high-skill workers in both productivity and debt. While it is clear that student debt is not evenly distributed (Looney and Yannelis 2018), there is limited knowledge about the full distribution and potential correlation with worker productivity. If individuals who borrow the most are differentially productive in urban areas and earn higher wages as result, the model would overestimate resorting.

Second, the model assumes low- and high-skill productivity are independent of the skill-mix in a location despite robust evidence in the literature to the contrary (Diamond 2016; Moretti 2011). While the presence of agglomeration economies among high-skill workers in cities would mean the model underestimates out-migration of high-skill workers from urban areas in counterfactual simulations, little is known about how productivity responds to a changing population outside of major metropolitan areas. Due to the uncertainty around production in rural areas, partial equilibrium results may present the most straightforward predictions as they do not impose estimates of city labor demand functions on rural areas. Finally, the model maintains a fixed migration elasticity within skill groups despite knowledge of varying mobility over demographic characteristics like age. While I did not incorporate endogenous migration elasticity into the model for simplicity, recent evidence suggests that migration choices of debt-holders are not entirely path-dependent and borrowers do indeed respond to debt relief (Di Maggio, Kalda, and Yao 2020).

The focus of this paper is on empirically documenting the effect of student debt on location choices of high-skill workers and providing a qualitative understanding of how it impacts aggregate sorting and welfare. In future research, I plan to extend this work to better understand how student debt interacts with optimal sorting patterns by developing a structural framework that more accurately captures production re-

sponses in urban and rural environments. Additionally, future work should examine the implications of the COVID-19 pandemic, which resulted in a decoupling of wages and location, for student debt-driven spatial sorting.

Table 1.1: Federal Borrowing Limits by Academic Year and Level (Class)

Academic Year	Freshmen	Sophomore	Upper Level	Aggregate Limit
2006-07 and earlier	\$2,625	\$3,500	\$5,500	\$23,000
2007-08	\$3,500	\$4,500	\$5,500	\$23,000
2008-09 and later	\$5,500	\$6,500	\$7,500	\$31,000

NOTE: Federal Stafford Loan limits for dependent undergraduate students. Limits apply to the sum of both subsidized and unsubsidized Stafford Loans. Independent students can borrow an additional \$4,000 in Freshmen and Sophomore years, and an additional \$5,000 in upper levels. Aggregate limits for independent students were double the dependent student limit through 2007-08 academic year, and has been \$57,500 since.

Table 1.2: Federal Borrowing Limits by Entry Cohort

Entry Cohort	Individual Level Limits				4-Year Sum	
	Freshmen	Sophomore	Junior	Senior+	Limit	Relative 2004-05
2004-05 and earlier	\$2,625	\$3,500	\$5,500	\$5,500	\$17,125	-
2005-06	\$2,625	\$3,500	\$5,500	\$7,500	\$19,125	\$2,000
2006-07	\$2,625	\$4,500	\$7,500	\$7,500	\$22,125	\$5,000
2007-08	\$3,500	\$6,500	\$7,500	\$7,500	\$25,000	\$7,875
2008-09 and later	\$5,500	\$6,500	\$7,500	\$7,500	\$27,000	\$9,875

NOTE: Federal Stafford Loan limits for dependent undergraduate students. Limits apply to the sum of both subsidized and unsubsidized Stafford Loans. Independent students can borrow an additional \$4,000 in Freshmen and Sophomore years, and an additional \$5,000 in upper levels.

Table 1.3: Student Characteristics by Treatment Status and Cohort

Variable	Unconstrained		Constrained	
	2002-2005	2006-2013	2002-2005	2006-2013
Age at Entry	18.6	18.6	18.3	18.4
(s.d.)	(0.9)	(0.9)	(0.7)	(0.7)
$\mathbb{1}[\text{Female}]$	0.57	0.57	0.53	0.55
<u>Before First Education Loan:</u>				
$\mathbb{1}[\text{Credit Report}]$	0.42	0.48	0.32	0.36
$\mathbb{1}[\text{Credit Score}]$	0.12	0.35	0.11	0.25
Credit Score	624	628	654	651
(s.d.)	(78)	(85)	(66)	(82)
$\mathbb{1}[\text{Credit Card}]$	0.27	0.24	0.20	0.18
$\mathbb{1}[\text{Auto Loan}]$	0.002	0.002	0.002	0.001
Number of Students	49,242	93,067	136,383	661,102

NOTE: Sample includes individuals in UCCCP California sample that first took out education loans between July 1, 2001 and June 30, 2013 (i.e., 2002-2013 cohorts) and were less than 21 years old at that time. Sample was also restricted to those who had borrowed a cumulative amount in their first year that was at or below the Federal Stafford Loan limit for first year borrowers in their respective cohort (see Table 1.2). Finally, individuals had to maintain a credit report through year 6 from entry.

Table 1.4: Student Characteristics Comparison: Difference-in-Differences Specification

Variable	Before First Education Loan (Entry)			
	(1) Age at Entry	(2) 1[Female]	(3) 1[Credit Report]	(4) 1[Credit Score]
1[Constrained] × 1[Cohort>2005]	0.015* (0.007)	0.014*** (0.004)	-0.003 (0.005)	-0.050*** (0.004)
Dep. Variable Mean	18.385	0.549	0.367	0.233
Number of Students	939,794	939,794	939,794	939,794

*** p< 0.01, ** p< 0.05, * p< 0.1.

NOTE: Sample includes individuals in UCCCP California sample that first took out education loans between July 1, 2001 and June 30, 2013 (i.e., 2002-2013 cohorts), were less than 21 years old at that time, borrowed a cumulative amount in their first year that was at or below the Federal Stafford Loan limit for first year borrowers in their respective cohort (see Table 1.2), and maintained a credit report through year 6 from entry. Coefficients from DID specification in Equation 1.2 on respective outcome variable with only year 3 state FEs as controls. Standard errors (clustered by year 3 state) are reported in parentheses.

Table 1.4: Student Characteristics Comparison: Difference-in-Differences Specification (*continued*)

Variable	Before First Education Loan (Entry)			Last Obs. [†]
	(5) Credit Score	(6) 1[Credit Card]	(7) 1[Auto Loan]	(8) Education
1[Constrained] × 1[Cohort>2005]	-5.0** (2.0)	0.015*** (0.002)	-2.9E-5 1.2E-4	0.003 (0.004)
Dep. Variable Mean	647.0	0.195	1.5E-3	2.055
Number of Students	939,794	939,794	939,794	939,794

*** p< 0.01, ** p< 0.05, * p< 0.1.

[†] All individuals are observed through at least year 6 and 76% are observed through year 9 from first education loan.

NOTE: Sample includes individuals in UCCCP California sample that first took out education loans between July 1, 2001 and June 30, 2013 (i.e., 2002-2013 cohorts), were less than 21 years old at that time, borrowed a cumulative amount in their first year that was at or below the Federal Stafford Loan limit for first year borrowers in their respective cohort (see Table 1.2), and maintained a credit report through year 6 from entry. Coefficients from DID specification in Equation 1.2 on respective outcome variable with only year 3 state FEs as controls. Standard errors (clustered by year 3 state) are reported in parentheses.

Table 1.5: Difference-in-Differences Main Results

Variable	First Stage: Borrowing	
	(1)	(2)
	Year 1	Year 4 (Cumulative)
$\mathbf{1}[\text{Constrained}] \times$ $\mathbf{1}[\text{Cohort} > 2005]$	1,214.32*** (10.46)	2,600.25*** (149.47)
$\mathbf{1}[\text{Constrained}]$	1,000.40*** (22.70)	2,829.30*** (239.77)
Dep. Variable Mean	3,548.18	13,940.23
Dep. Variable S.D.	(1,392.74)	(10,530.49)
Number of Students	939,794	939,794
Cohort FEs	✓	✓
Demographic Controls	✓	✓
Credit History Controls	✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Sample includes individuals in UCCCP California sample that first took out education loans between July 1, 2001 and June 30, 2013 (i.e., 2002-2013 cohorts), were less than 21 years old at that time, borrowed a cumulative amount in their first year that was at or below the Federal Stafford Loan limit for first year borrowers in their respective cohort (see Table 1.2), and maintained a credit report through year 6 from entry. Coefficients from DID specification in Equation 1.2 on respective outcome variable. Standard errors (clustered by year 3 state) are reported in parentheses.

Table 1.5: Difference-in-Differences Main Results (*continued*)

Second Stage: Year 6 Location						
Variable	Full Sample			Year 3 California Sample		
	(3) Pop. Density (logs)	(4) 1[Top 5% Pop. Density]	(5) 1[Top RUCC Metro]	(6) Pop. Density (logs)	(7) 1[Top 5% Pop. Density]	(8) 1[Top RUCC Metro]
1[Constrained] × 1[Cohort>2005]	0.022*** (0.008)	0.016*** (0.002)	0.017*** (0.003)	0.026** (0.011)	0.017*** (0.004)	0.020*** (0.003)
1[Constrained]	0.090*** (0.008)	0.016*** (0.002)	0.018*** (0.002)	0.099*** (0.010)	0.017*** (0.003)	0.018*** (0.003)
Dep. Variable Mean	6.762	0.545	0.748	6.793	0.562	0.785
Dep. Variable S.D.	1.493	-	-	1.424	-	-
Number of Students	939,794	939,794	939,794	624,847	624,847	624,847
Cohort FEs	✓	✓	✓	✓	✓	✓
Demographic	✓	✓	✓	✓	✓	✓
Credit History	✓	✓	✓	✓	✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Sample includes individuals in UCCCP California sample that first took out education loans between July 1, 2001 and June 30, 2013 (i.e., 2002-2013 cohorts), were less than 21 years old at that time, borrowed a cumulative amount in their first year that was at or below the Federal Stafford Loan limit for first year borrowers in their respective cohort (see Table 1.2), and maintained a credit report through year 6 from entry. Coefficients from DID specification in Equation 1.2 on respective outcome variable. Standard errors (clustered by year 3 state) are reported in parentheses.

Table 1.6: Difference-in-Differences: Great Recession Robustness

Variable	First Stage: Borrowing	
	(1) Year 1	(2) Year 4 (Cumulative)
$\mathbb{1}[\text{Constrained}] \times \mathbb{1}[\text{Cohort} > 2005]$	1,989.48*** (45.74)	4,070.99*** (163.00)
$\mathbb{1}[\text{Constrained}]$	955.97*** (25.82)	2,952.70*** (246.23)
Dep. Variable Mean	4,091.53	14,445.09
Dep. Variable S.D.	1,417.86	10,229.86
Number of Students	500,020	500,020
Cohort FEs	✓	✓
Demographic	✓	✓
Credit History	✓	✓

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

NOTE: Sample includes individuals in UCCCP California sample that first took out education loans between July 1, 2001 and June 30, 2003 (i.e., 2002-2003 cohorts) or between July 1, 2009 and June 30, 2013 (i.e., 2009-2013 cohorts), were less than 21 years old at that time, borrowed a cumulative amount in their first year that was at or below the Federal Stafford Loan limit for first year borrowers in their respective cohort (see Table 1.2), and maintained a credit report through year 6 from entry. Coefficients from DID specification in Equation 1.2 on respective outcome variable. Standard errors (clustered by year 3 state) are reported in parentheses.

Table 1.6: Difference-in-Differences: Great Recession Robustness (*continued*)

Variable	Second Stage: Year 6 Location					
	Full Sample			Year 3 California Sample		
	(3) Pop. Density (logs)	(4) 1[80 th Percentile Pop. Density]	(5) 1[Top RUCC Metro]	(6) Pop. Density (logs)	(7) 1[80 th Percentile Pop. Density]	(8) 1[Top RUCC Metro]
1[Constrained] × 1[Cohort>2005]	0.038*** (0.009)	0.024*** (0.003)	0.027*** (0.005)	0.037** (0.018)	0.023*** (0.006)	0.029*** (0.005)
1[Constrained]	0.088*** (0.010)	0.013*** (0.003)	0.011** (0.005)	0.093*** (0.016)	0.013** (0.006)	0.010** (0.005)
Dep. Variable Mean	6.776	0.550	0.758	6.797	0.564	0.787
Dep. Variable S.D.	1.478	-	-	1.420	-	-
Number of Students	500,020	500,020	500,020	352,638	352,638	352,638
Cohort FEs	✓	✓	✓	✓	✓	✓
Demographic	✓	✓	✓	✓	✓	✓
Credit History	✓	✓	✓	✓	✓	✓

*** p< 0.01, ** p< 0.05, * p< 0.1.

NOTE: Sample includes individuals in UCCCP California sample that first took out education loans between July 1, 2001 and June 30, 2003 (i.e., 2002-2003 cohorts) or between July 1, 2009 and June 30, 2013 (i.e., 2009-2013 cohorts), were less than 21 years old at that time, borrowed a cumulative amount in their first year that was at or below the Federal Stafford Loan limit for first year borrowers in their respective cohort (see Table 1.2), and maintained a credit report through year 6 from entry. Coefficients from DID specification in Equation 1.2 on respective outcome variable. Standard errors (clustered by year 3 state) are reported in parentheses.

Table 1.7: Persistence of Location Choices

Variable	(1) Pop. Density (logs)	(2) 1[80 th Percentile Pop. Density]	(3) 1[Top RUCC Metro]
Year 6 (Full Sample, N=939,794)	0.022*** (0.008)	0.016*** (0.002)	0.017*** (0.003)
Year 6 (N=715,067)	0.018* (0.010)	0.015*** (0.004)	0.014*** (0.004)
Year 7 (N=715,067)	0.014 (0.011)	0.015*** (0.003)	0.015*** (0.005)
Year 8 (N=715,067)	0.021* (0.011)	0.017*** (0.003)	0.017*** (0.005)
Year 9 (N=715,067)	0.021** (0.010)	0.018*** (0.002)	0.019*** (0.004)
Cohort FEs	✓	✓	✓
Demographic Controls	✓	✓	✓
Credit History Controls	✓	✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Sample includes individuals in UCCCP California sample that first took out education loans between July 1, 2001 and June 30, 2013 (i.e., 2002-2013 cohorts), were less than 21 years old at that time, borrowed a cumulative amount in their first year that was at or below the Federal Stafford Loan limit for first-year borrowers in their respective cohort (see Table 1.2), and maintained a credit report through year 6 from entry. In all but the first row, the sample is limited to those observed through year 10. Coefficients from DID specification in Equation 1.2 on the respective outcome variable. Standard errors (clustered by year 3 state) are reported in parentheses.

Table 1.8: Student Debt and Post-College Location Choice

	Post-College County		
	Full Sample		
	(1) Pop. Density (logs)	(2) 1[80 th Percentile Pop. Density]	(3) 1[Top RUCC Metro]
Student Debt (\$thds, log)	0.1407*** [0.0478]	0.0312*** [0.0109]	0.0270* [0.0144]
Origin County 1[80 th Percentile Pop. Density]			
1[Top RUCC Metro]			
College County 1[80 th Percentile Pop. Density]			
1[Top RUCC Metro]			
Observations	2,210	2,210	2,210
R-squared	0.1767	0.1466	0.1232
Dep. Mean	6.6460	0.8269	0.5954
Student Loan Mean	3.1660	3.1660	3.1660
Student Loan S.D.	0.9567	0.9567	0.9567
Demographic Controls	✓	✓	✓
ELS Major FEs	✓	✓	✓

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

NOTE: Dependent variable pertains to post-college county and classifies them based on varying definitions of urban. Columns (1), (4), and (7) use a continuous measure of population density (in logs); Columns (2), (5), and (8) use an indicator for whether the county is in the top 20 percentile of counties by population density; and Columns (3), (6), and (9) use an indicator for whether the county is classified as a top metropolitan area with a population of 1 million+ using the Rural-Urban Continuum Codes. Origin and college county location controls are similar indicator functions based on the stated urban criteria. Demographic controls include sex, race, father's education, father's occupation, and high school GPA. ELS major fixed effects are reported academic major in postsecondary education. Columns (1)-(6) use the full ELS sample of individuals who attained a postsecondary degree. Columns (7)-(9) restrict this sample to individuals who left their college county after graduation. Number of observations rounded to nearest 10 per IES restricted-use guidelines. Robust standard errors clustered by sample design (survey base-year high school) in parenthesis.

SOURCE: U.S. Department of Education, National Center for Education Statistics, Education Longitudinal Study of 2002 (ELS:2002), "Base Year", 2002, "First Follow-up", 2004, "High School Transcripts", 2005, "Second Follow-up", 2006, and "Third Follow-up", 2012.

Table 1.8: Student Debt and Post-College Location Choice (*continued*)

	Post-College County					
	Full Sample			“Leavers” (College Cty. \neq Post-College Cty.)		
	(4) Pop. Density (logs)	(5) 1[80 th Percentile Pop. Density]	(6) 1[Top RUCC Metro]	(7) Pop. Density (logs)	(8) 1[80 th Percentile Pop. Density]	(9) 1[Top RUCC Metro]
Student Debt (\$thds, log)	0.1054** [0.0415]	0.0207** [0.0092]	0.0072 [0.0100]	0.1359** [0.0634]	0.0337*** [0.0126]	0.0185 [0.0148]
Origin County						
1[80 th Percentile Pop. Density]	1.2865*** [0.1040]	0.3750*** [0.0295]		1.5182*** [0.1295]	0.4556*** [0.0347]	
1[Top RUCC Metro]			0.4601*** [0.0280]			0.5399*** [0.0310]
College County						
1[80 th Percentile Pop. Density]	0.8474*** [0.1066]	0.2542*** [0.0316]		0.4461*** [0.1295]	0.0810** [0.0321]	
1[Top RUCC Metro]			0.3421*** [0.0259]			0.1306*** [0.0282]
Observations	2,210	2,210	2,210	1,470	1,470	1,470
R-squared	0.3548	0.4201	0.5378	0.3827	0.4069	0.4837
Dep. Mean	6.6460	0.8269	0.5954	6.5087	0.7952	0.5959
Student Loan Mean	3.1660	3.1660	3.1660	3.1902	3.1902	3.1902
Student Loan S.D.	0.9567	0.9567	0.9567	0.9280	0.9280	0.9280
Demographic Controls	✓	✓	✓	✓	✓	✓
ELS Major FEs	✓	✓	✓	✓	✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: See table notes on the previous page. Number of observations rounded to nearest 10 per IES restricted-use guidelines. Robust standard errors clustered by sample design (survey base-year high school) in parenthesis.

SOURCE: U.S. Department of Education, National Center for Education Statistics, Education Longitudinal Study of 2002 (ELS:2002), “Base Year”, 2002, “First Follow-up”, 2004, “High School Transcripts”, 2005, “Second Follow-up”, 2006, and “Third Follow-up”, 2012.

Table 1.9: Student Debt and Post-College Location Choice: Heterogeneity by Urban Wage Premium

	Post-College County		
	Pop. Density (logs)		
	(1) Full Sample	(2) ACS-Matched Sample	(3) ACS-Matched Sample
Student Debt (\$thds, log)	0.1054** [0.0415]	0.1062** [0.0533]	
× 1[Below Median Urban Wage Prem.]			0.0292 [0.0697]
× 1[Above Median Urban Wage Prem.]			0.1434** [0.0717]
Origin County			
1[80 th Percentile Pop. Density]	1.2865*** [0.1040]	1.2715*** [0.1052]	1.2732*** [0.1054]
1[Top RUCC Metro]			
College County			
1[80 th Percentile Pop. Density]	0.8474*** [0.1066]	0.7887*** [0.1302]	0.7989*** [0.1286]
1[Top RUCC Metro]			
Observations	2,210	1,700	1,700
R-squared	0.3548	0.2942	0.2951
Dependent Variable Mean	6.646	6.7556	6.7556
Student Loan Mean	3.1660	3.2344	3.2344
Student Loan S.D.	0.9567	0.9270	0.9270
Demographic Controls	✓	✓	✓
ELS Major FEs	✓		
ACS Degree Field FEs		✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Dependent variable pertains to post-college county and classifies them based on varying definitions of urban. Columns (1), (4), and (7) use a continuous measure of population density (in logs); Columns (2), (5), and (8) use an indicator for whether the county is in the top 20 percentile of counties by population density; and Columns (3), (6), and (9) use an indicator for whether the county is classified as a top metropolitan area with a population of 1 million+ using the Rural-Urban Continuum Codes. Origin and college county location controls are similar indicator functions based on the stated urban criteria. Demographic controls include sex, race, father’s education, father’s occupation, and high school GPA. ELS major fixed effects are reported academic major in postsecondary education. Columns (1)-(6) use the full ELS sample of individuals who attained a postsecondary degree. Columns (7)-(9) restrict this sample to individuals who left their college county after graduation. Number of observations rounded to nearest 10 per IES restricted-use guidelines. Robust standard errors clustered by sample design (survey base-year high school) in parentheses.

SOURCE: U.S. Department of Education, National Center for Education Statistics, Education Longitudinal Study of 2002 (ELS:2002), “Base Year”, 2002, “First Follow-up”, 2004, “High School Transcripts”, 2005, “Second Follow-up”, 2006, and “Third Follow-up”, 2012.

Table 1.9: Student Debt and Post-College Location Choice: Heterogeneity by Urban Wage Premium (*continued*)

	Post-College County					
	1[80 th Percentile Pop. Density]			1[Top RUCC Metro]		
	(4) Full Sample	(5) ACS-Matched Sample	(6) ACS-Matched Sample	(7) Full Sample	(8) ACS-Matched Sample	(9) ACS-Matched Sample
Student Debt (\$thds, log)	0.0229** [0.0092]		0.0072	0.0040 [0.0100]		
× 1[Below Median Urban Wage Premium]		[0.0108]	0.0099 [0.0194]		[0.0119]	-0.0207 [0.0216]
× 1[Above Median Urban Wage Premium]			0.0292** [0.0130]			0.0159 [0.0144]
Origin County						
1[80 th Percentile Pop. Density]	1.2865*** [0.0295]	0.3482*** [0.0326]	0.3485*** [0.0326]			
1[Top RUCC Metro]				0.4601*** [0.0280]	0.4652*** [0.0337]	0.4649*** [0.0337]
College County						
1[80 th Percentile Pop. Density]	0.8474*** [0.0316]	0.2342*** [0.0377]	0.2359*** [0.0377]			
1[Top RUCC Metro]				0.3421*** [0.0259]	0.3165*** [0.0306]	0.3179*** [0.0307]
Observations	2,210	1,700	1,700	2,210	1,700	1,700
R-squared	0.4201	0.3576	0.3580	0.5378	0.4938	0.4948
Dependent Variable Mean	0.8269	0.8444	0.8444	0.5954	0.6235	0.6235
Student Loan Mean	3.1660	3.2344	3.2344	3.1660	3.2344	3.2344
Student Loan S.D.	0.9567	0.9270	0.9270	0.9567	0.9270	0.9270
Demographic Controls	✓	✓	✓	✓	✓	✓
ELS Major FEs	✓			✓		
ACS Degree Field FEs		✓	✓		✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: See table notes on the previous page. Number of observations rounded to nearest 10 per IES restricted-use guidelines. Robust standard errors clustered by sample design (survey base-year high school) in parenthesis.

SOURCE: U.S. Department of Education, National Center for Education Statistics, Education Longitudinal Study of 2002 (ELS:2002), “Base Year”, 2002, “First Follow-up”, 2004, “High School Transcripts”, 2005, “Second Follow-up”, 2006, and “Third Follow-up”, 2012.

Table 1.10: Model Results: Skill-Based Geographic Sorting

Model	Counterfactual				
	1980 Debt Level		Debt Forgiveness	Income-Driven Repayment	Free College Option
	(1)	(2)	(3)	(4)	(5)
	Δ Sorting (Rel. to 1980)	Δ Sorting (2019 Level)	Δ Sorting (2019 Level)	Δ Sorting (2019 Level)	Δ Sorting (2019 Level)
Panel A: Partial Equilibrium					
Homogeneous Migration Elasticity	-11.5%***	-6.7%***	-8.1%***	-8.1%***	-5.7%***
Heterogeneous Migration Elasticity [†]	-18.7%***	-11.2%***	-14.0%***	-14.0%***	-11.6%***
Panel B: Full Equilibrium					
Homogeneous Migration Elasticity	-3.5%***	-2.0%***	-2.4%***	-2.4%***	0.0%
Heterogeneous Migration Elasticity [†]	-4.5%***	-2.8%***	-3.4%***	-2.2%***	0.2%***

*** p < 0.01, ** p < 0.05, * p < 0.1.

[†] Elasticity estimates from Hornbeck and Moretti 2018.

NOTE: Reported values reflect the change in the coefficient on commuting zone population density high-skill worker share in a regression weighted by commuting zone population when comparing counterfactual to observed data.

Table 1.11: Model Results: Consumer Welfare

Model	Counterfactual								
	(1) Debt Forgiveness			(2) Income-Driven Repayment			(3) Tuition-Free College		
	Δ Aggregate	Δ HS	Δ LS	Δ Aggregate	Δ HS	Δ LS	Δ Aggregate	Δ HS ^{††}	Δ LS ^{††}
Panel A: Partial Equilibrium									
Homogeneous Migration Elasticity	-0.64%	2.06%	-2.40%	0.05%	0.13%	0.00%	-0.67%	1.93%	-2.53%
Heterogeneous Migration Elasticity [†]	-0.64%	2.07%	-2.41%	0.05%	0.12%	0.00%	-0.72%	1.93%	-2.54%
Panel B: Full Equilibrium									
Homogeneous Migration Elasticity	-1.08%	1.52%	-2.78%	0.12%	0.18%	0.09%	-0.66%	1.93%	-2.53%
Heterogeneous Migration Elasticity [†]	-1.10%	1.52%	-2.79%	0.11%	0.17%	0.08%	-0.71%	1.94%	-2.53%

[†] Elasticity estimates from Hornbeck and Moretti 2018.

^{††} For tuition-free college counterfactual, reported value reflects individual worker changes, i.e. not accounting for changing skill-mix of the general population.

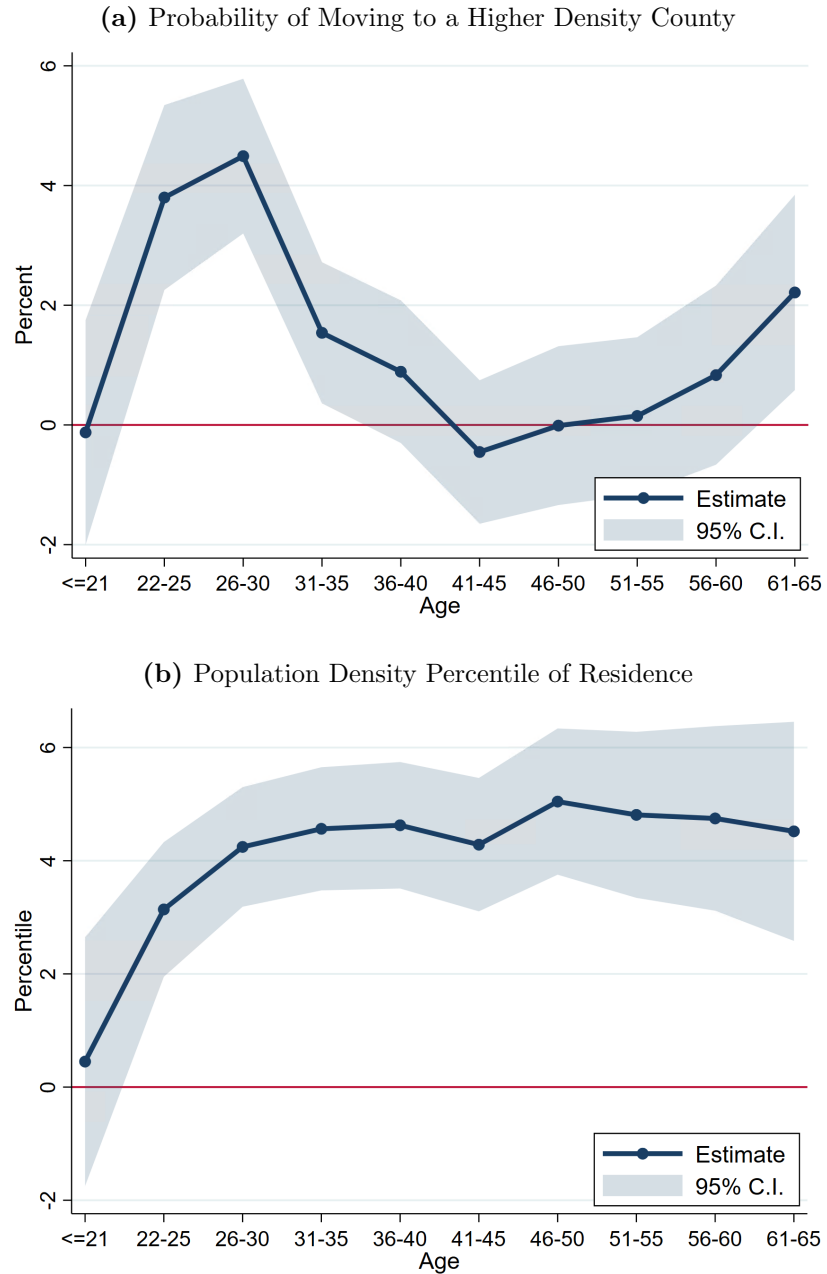
NOTE: Reported values reflect the change in aggregate consumer welfare, the change in welfare for high-skill workers, and the change in welfare for low-skill workers.

Table 1.12: Model Results: Aggregate Output

Model	Counterfactual		
	(1)	(2)	(3)
	Debt Forgiveness	Income-Driven Repayment	Tuition-Free College
	Δ Output (2019 Level)	Δ Output (2019 Level)	Δ Output (2019 Level)
Panel A: Partial Equilibrium			
Homogeneous Migration Elasticity	-0.30%	-0.30%	1.72%
Heterogeneous Migration Elasticity [†]	-0.52%	-0.52%	1.48%
Panel B: Full Equilibrium			
Homogeneous Migration Elasticity	-0.09%	-0.10%	1.99%
Heterogeneous Migration Elasticity [†]	-0.10%	-0.11%	1.97%


[†] Elasticity estimates from Hornbeck and Moretti 2018.

Figure 1.1: Skill Differential in Migration Choices Over the Life Cycle




NOTE: Panel (a) shows the coefficients from a linear probability model of an indicator for moving to a county with a higher population density relative to previous county of residence on the interaction between being a high-skill worker and age categories. Panel (b) shows the coefficients from the same regression when using the percentile of residence county population density as a dependent variable. Both estimated using a PSID sample of individuals born between 1940 and 1979 and observed biennially 1968 to 2019. Excludes waves when individuals enrolled in postsecondary institutions. Regressions include wave, cohort, and demographic controls. See 1.A for a complete discussion. SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).

Figure 1.2: Inflation-Adjusted Cost of Post-Secondary Education




SOURCE: Loan data from College Board; includes undergraduate average federal and nonfederal loans per full-time equivalency student. College tuition and fees index from U.S. Bureau of Labor Statistics.

Figure 1.3: Time to Degree Completion



NOTE: Time to Bachelor degree completion among individuals who attain a degree.
 SOURCE: U.S. Department of Education, National Center for Education Statistics (2019). Baccalaureate and Beyond (B&B:16/17): A First Look at the Employment and Educational Experiences of College Graduates, 1 Year Later (NCES 2019-106), Table 2.


Figure 1.4: Share of Loan Recipients on Income-Independent Repayment Plans



NOTE: Income-independent plans are all plans not classified as income-driven repayment plans (income-contingent, income-sensitive, income-based, Pay As You Earn (PAYE), and Revised Pay As You Earn (RePAYE)). All data excludes borrowers in default, in-school, or in grace periods before repayment begins. Recipient counts are based at loan level. As a result, recipients may be counted multiple times across varying loan statuses.


SOURCE: Data on Federal Student Loan Portfolio provided by Federal Student Aid; original source is National Student Loan Data System.

Figure 1.5: Defining Constrained Borrowers



NOTE: Students are classified as likely to be constrained if: (1) they borrow exactly at the original first-year borrowing limit (\$2,625) in years up to the 2007 academic year; or, (2) if they borrow between the original limit (\$2,625) and the new limit in their given entry year. The bar for constrained borrowers in years prior to 2008 appears larger than the single point (\$2,625) in this graph for illustrative purposes.


Figure 1.6: First-Year Borrowing



NOTE: This distribution reflects the cumulative sum of loans in the first year of borrowing for each individual.

SOURCE: University of California Consumer Credit Panel, California Sample.


Figure 1.7: Effect of Increased Loan Access on First-Year Borrowing



NOTE: Coefficients from the estimation of Equation 1.1 with the dependent variable being cumulative education loans in year 1 of borrowing. Confidence intervals reflect standard errors clustered by year 3 state.

SOURCE: University of California Consumer Credit Panel, California Sample.


Figure 1.8: Effect of Increased Loan Access on Cumulative Borrowing Through Year 4



NOTE: Coefficients from the estimation of Equation 1.1 with the dependent variable being cumulative education loans through year 4 of borrowing. Confidence intervals reflect standard errors clustered by year 3 state.

SOURCE: University of California Consumer Credit Panel, California Sample.


Figure 1.9: Effect of Increased Loan Access on Year 6 County Population Density (logs)



NOTE: Coefficients from the estimation of Equation 1.1 with the dependent variable being (log) population density for the county of residence in year 6 from entry. Confidence intervals reflect standard errors clustered by year 3 state.


SOURCE: University of California Consumer Credit Panel, California Sample.

Figure 1.10: Effect of Increased Loan Access on Probability of Year 6 County Being 80th Percentile by Population Density



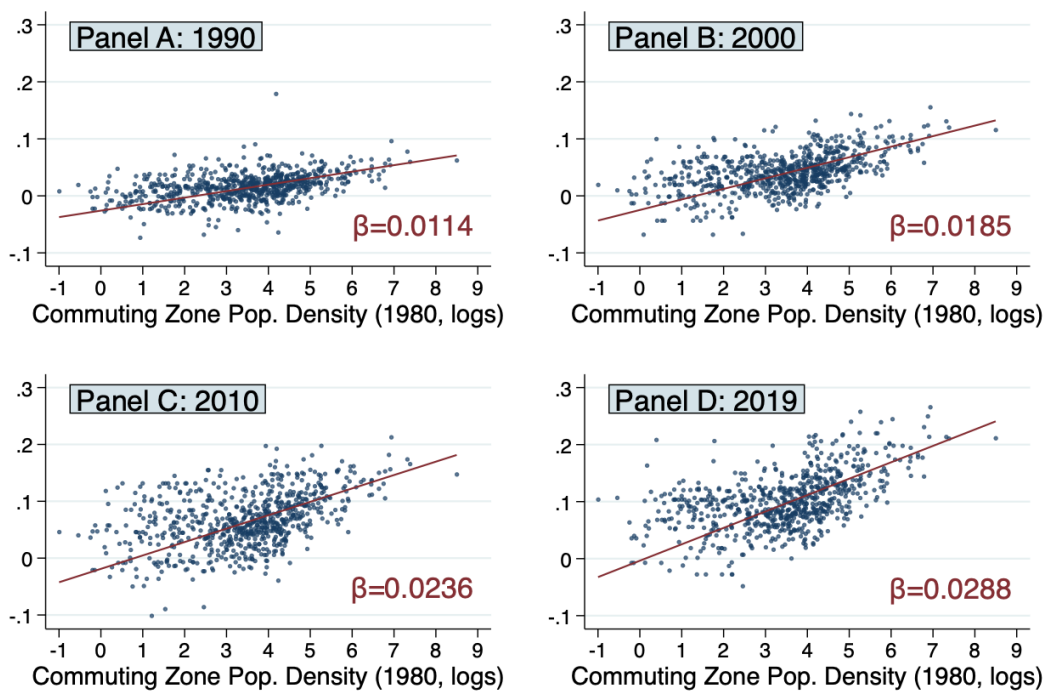
NOTE: Coefficients from the estimation of Equation 1.1 with the dependent variable being an indicator function for county of residence in year 6 from entry being in the 80th percentile by county population density. Confidence intervals reflect standard errors clustered by year 3 state.
 SOURCE: University of California Consumer Credit Panel, California Sample.

Figure 1.11: Effect of Increased Loan Access on Probability of Year 6 County Being Top Metro RUCC



NOTE: Coefficients from the estimation of Equation 1.1 with the dependent variable being an indicator function for county of residence in year 6 from entry being in the top metropolitan category (1m+) of the RUCC. Confidence intervals reflect standard errors clustered by year 3 state.
 SOURCE: University of California Consumer Credit Panel, California Sample.

Figure 1.12: Skill-Based Geographic Sorting:
Change in High Skill Share Since 1980



NOTE: X-axis population reflects total population; skill shares are among employed, working-age (25-54). Points are 1990-defined commuting zones. Best fit line weighted by commuting zone total population in 1980.

SOURCE: U.S. Census 1980 5% state, 1990 5% state, 2000 5%, 2010 ACS, and 2019 5-year ACS.

Figure 1.13: Role of Student Debt in Post-1980 Geographic Sorting:
Partial Equilibrium

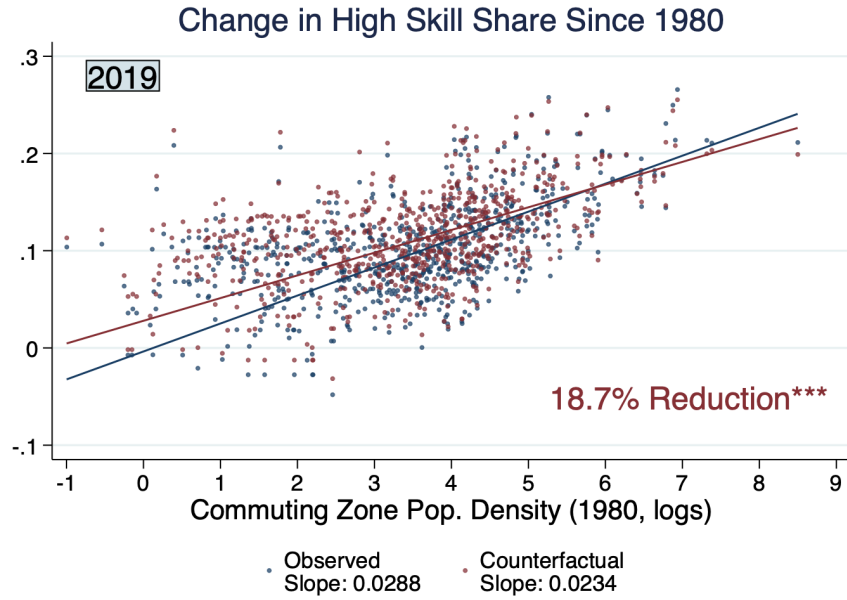


Figure 1.14: Role of Student Debt in Post-1980 Geographic Sorting:
Full Equilibrium




Figure 1.15: Income-Driven Repayment Counterfactual Map:
Partial Equilibrium

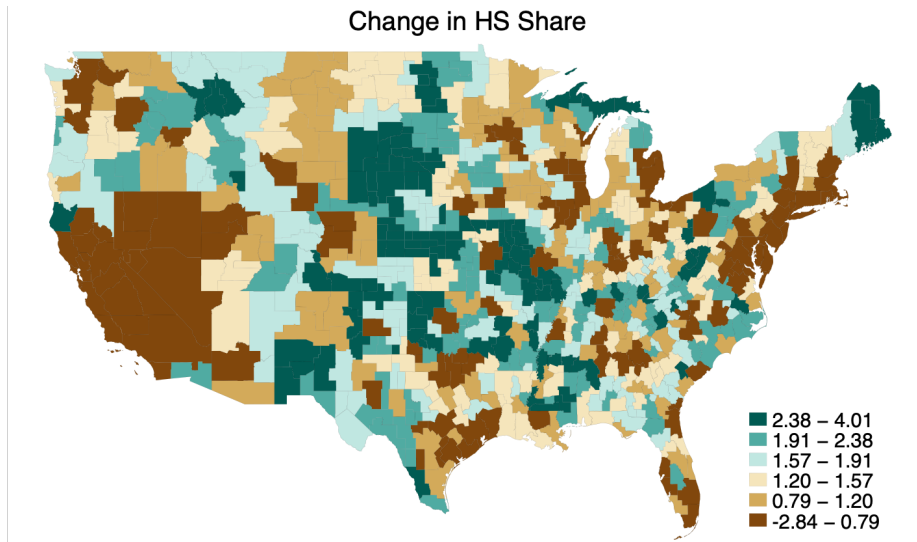
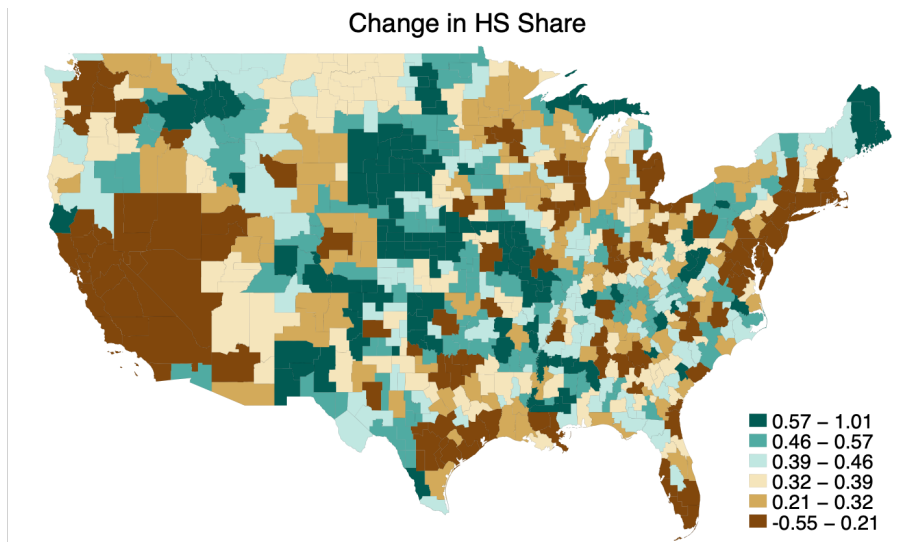


Figure 1.16: Income-Driven Repayment Counterfactual Map:
Full Equilibrium



APPENDICES

1.A Early Adulthood Migration Drives Skill-Based Sorting

In this section, I show that differences between low and high skill workers in mobility and spatial distribution, particularly across the urban-rural divide, are driven by migration choices in early adulthood. This finding, which is absent of any quantitative connection to student debt, suggests that factors weighing heavily on individuals in early adulthood could impact key initial migration choices that are persistent into later adulthood. While a standard life-cycle model would suggest student debt, which is a small share of lifetime income, should have little impact on post-college choices like migration and career, debt repayments as a share of income are largest in early adulthood, and Rothstein and Rouse 2011 find that student debt causes individuals to choose initially higher-salary jobs.⁴⁵ This departure from the expected behavior of life-cycle agents is consistent with evidence from the Education Longitudinal Study of 2002 showing that additional student debt is positively correlated with recent graduates taking a job outside of their field of study, having to work more than one job at the same time, and having to work more hours than desired (see Table 1.13). If student debt impacts early adulthood migration choices as well, which I will identify in Section 1.3, the findings in this section suggest it could have a persistent effect on the spatial distribution of workers.

I proceed in this section by tracking the location choices of all adults in the Panel Study of Income Dynamics that were born between 1940 and 1979, observing their

45. They attribute this departure from standard life-cycle model predictions to credit constraints after schooling, but cannot rule out debt aversion.

location choices biennially from 1968 to 2019.^{46,47} I classify workers by high (college-educated) and low (less than college degree) skill using their educational attainment through age 40. First, I show that high skill individuals (college graduates) are generally more geographically mobile than non-college educated individuals over their lifetime, a result consistent with previous literature⁴⁸, but that this overall mobility differential masks significant heterogeneity over the life cycle. High skill workers are more mobile in early adulthood, but less mobile in later periods – suggesting that early migration choices are particularly persistent for high skill workers. I then examine educational mobility differentials over the life cycle in relation to skill-based geographic sorting across various definitions of the urban-rural divide. When considering migration from one PSID wave to the next, high skill individuals are 2-3 percent more likely than low skill individuals to move from a ‘rural’ county to an ‘urban’ one for waves in which the individual is younger than 30 (mean sample probability of 2.2 percent).

1.A.0.1 Data

The PSID is a longitudinal survey with a nationally representative sample that collects location information at a high frequency and allows me to understand migration across the urban-rural divide over the life cycle. The survey has been conducted since 1968, making it the longest-running longitudinal survey in the world. The sample follows original households and their descendants with occasional ‘refresher’ samples

46. Sample was limited by date of birth for three reasons: (1) so that each cohort is observed for a considerable period of their adult life (20 years); (2) at least once before the age of 30; and (3) so that they are individuals in their prime working years during the time period of interest (1980-2019) for skill-based geographic sorting. Results are robust to the addition of all cohorts in the PSID.

47. The PSID conducted surveys annually from 1968 through 1997, then biennially. To provide consistency across the sample, the 1968 through 1997 subsample was trimmed so that location is observed at most biennially starting from an individual’s first year in the sample.

48. See Ladinsky 1967; Greenwood 1975; Wozniak 2010; Malamud and Wozniak 2012.

to ensure the PSID remains representative of the U.S. population. Households were surveyed annually through 1997 and biennially thereafter. To maintain consistency across the sample, the 1968 through 1997 data was trimmed so that individuals are observed at most biennially starting from an individual's first year in the sample. I limit the sample population to the reference person and spouses of households, i.e. the adult population.⁴⁹ Furthermore, I restrict the sample to adults age 18 to 65 that were born between 1940 and 1979 for three reasons: (1) each cohort is observed for a considerable period (20 years) of their adult life; (2) at least once before the age of 30 ('early adulthood'); and (3) they are individuals in their prime working years during the time period of interest (1980-2019) for skill-based geographic sorting. The resulting sample includes approximately 7,000 individuals, observed on average for 25 years (12.5 waves).

In addition to location, the PSID collects information on a variety of other topics including education, employment, income, and wealth. Location is observed at the county level for each wave. Individuals are also asked to report the "county where they grew up", hereinafter referred to as the origin location – a factor that has been shown to be important in migration decisions (Diamond 2016). The PSID also includes a wealth of other information on an individual's personal characteristics, including sex, race, education attainment, and marital status. Although I only highlight a few here that are likely to impact mobility and migration behavior, a variety of other demographic and economic variables are included in different robustness specifications below.

49. Individuals that are descendants of PSID families establish independent family units when they live in a separate household and are above the age of 18 (underage individuals can be independent households after 1993, but I continue prior definition for consistency). Individuals in an institution, including a college dormitory, are not labelled independent households.

1.A.0.2 Overall Mobility

I use the following framework to estimate educational differentials in general mobility:

$$\mathbb{1}[\text{Moved}_{it}] = \alpha + \beta_1 \mathbb{1}[\text{Skill}_i = \text{High}] + \sum_c \delta_c \mathbb{1}[\text{Age}_{it} \in c] + \mathbf{X}_i' \boldsymbol{\beta}_x + \gamma_t + \gamma_c + \mu_{it} + \epsilon_{it} , \quad (1.25)$$

where $\mathbb{1}[\text{Moved}_{it}]$ is an indicator for whether an individual changed county of residence from wave $t - 1$ to wave t . The coefficient of interest, β_1 , captures whether an individual is a high skill (college educated) worker. Fixed effects control for age of individual in wave t by age categories (δ_c), wave (γ_t), and cohort (γ_c). As high and low skill workers likely establish independent households in the PSID at different ages⁵⁰, I also include a fixed effect for the number of waves that an individual has been an independent household (μ_{it}). This would capture common trends such as if there is frequent movement while initially setting up an independent household. A set of demographic characteristics are captured by the vector \mathbf{X}_i . These include sex, race, and marital status. Coefficients are estimated using a linear probability model with standard errors clustered by individual.

To understand how educational mobility differentials differ across the life cycle, I augment this framework by interacting the indicator for being a high skill worker with age categories:

$$\mathbb{1}[\text{Moved}_{it}] = \alpha + \sum_c \beta_1^c (\mathbb{1}[\text{Skill}_i = \text{High}] \times \mathbb{1}[\text{Age}_{it} \in c]) + \sum_c \delta_c \mathbb{1}[\text{Age}_{it} \in c] + \mathbf{X}_i' \boldsymbol{\beta}_x + \gamma_t + \gamma_c + \mu_{it} + \epsilon_{it} , \quad (1.26)$$

where all other aspects of the estimating equation remain the same.

Results from the estimation of Equations 1.25 and 1.26 using a linear probability

50. Individuals enrolled in post-secondary institutions are not considered independent households, and are thus not included in this sample until after graduation.

model are reported in Table 1.14. As shown in Column 1, high skill individuals are more geographically mobile than low skill workers overall. In any given wave, the probability that a high skill individual changed counties is 1.6 percentage points higher than for a low skill individual – a meaningful difference relative to a mean location change probability of 28.2 percent. This overall effect hides significant heterogeneity over the life cycle, as shown in Column 2. High skill workers are equally mobile for ages less than 22, then are significantly more mobile than low skill workers from 22 to 35. From age 36 to 56, geographic mobility of high skill workers is less than low skill workers, though this effect is insignificant outside of the 46-50 age range. These results can also be seen in Figure 1.17, Panel (a). As shown in Panel (b), these results are not driven by any single cohort.

1.A.0.3 Skill-Based Geographic Sorting

I next turn to analyzing mobility across the urban-rural divide. I use the same framework as above, but classify locations by four measures of ‘urban’: (1) indicator for top metropolitan county as defined by the Rural-Urban Continuum Codes⁵¹; (2) percentile of county population density (decade-specific); (3) indicator for county in 80th percentile of population density (decade-specific); and (4) indicator for county in 95th percentile of population density (decade-specific). As origin location has been shown to play a significant role in location preferences⁵², I also augment the previous specifications by controlling for origin location urbanicity using the same measure as the dependent variable. The educational migration differential along the urban-rural

51. A roughly decade-specific measure. RUCC were initially developed in 1974 and have been updated in 1983, 1993, 2003, and 2013. Codes used in this analysis are RUCC codes closest to PSID wave. Top metropolitan counties in the RUCC have a population of over 1 million.

52. See, for example, Molloy, Smith, and Wozniak 2011.

divide is estimated using the following framework:

$$\begin{aligned} \mathbb{1}[\mathbf{Live\ in\ 'Urban'}_{it}] &= \alpha + \beta_1 \mathbb{1}[\text{Skill}_i = \text{High}] + \\ &\quad \sum_c \delta_c \mathbb{1}[\text{Age}_{it} \in c] + \mathbb{1}[\text{Origin 'Urban'}_i] + \mathbf{X}'_i \boldsymbol{\beta}_x + \\ &\quad \gamma_t + \gamma_c + \mu_{it} + \epsilon_{it} \ , \end{aligned} \quad (1.27)$$

where all other variables are as defined following Equation 1.25. Coefficients are estimated using a linear probability model with standard errors clustered by individual.

To understand how migration differentials across the urban-rural divide vary by age, I interact the indicator for individual skill level with age:

$$\begin{aligned} \mathbb{1}[\mathbf{Live\ in\ 'Urban'}_{it}] &= \alpha + \sum_c \beta_1^c (\mathbb{1}[\text{Skill}_i = \text{High}] \times \mathbb{1}[\text{Age}_{it} \in c]) + \\ &\quad \sum_c \delta_c \mathbb{1}[\text{Age}_{it} \in c] + \mathbb{1}[\text{Origin 'Urban'}_i] + \mathbf{X}'_i \boldsymbol{\beta}_x + \\ &\quad \gamma_t + \gamma_c + \mu_{it} + \epsilon_{it} \ , \end{aligned} \quad (1.28)$$

where all other aspects of the estimating equation remain the same. Taking advantage of the biennial frequency in the PSID data, I also conduct this analysis for moving to urban areas. For a movement indicator to be 1, an individual has to living in a 'rural' county in wave $t - 1$ and then in an 'urban' county in wave t based on a given definition of urban and rural.

Results of this analysis on living in and moving to urban areas are reported in Tables 1.15 and 1.16, respectively. As shown in Columns 1 through 4 of Table 1.15, high skill individuals live in urban counties at higher rates than low skill individuals when using all 4 definitions of urban. This is consistent with the skill-based geographic sorting detailed in Figure 1.12. In Columns 5-8 of Table 1.15 and Figures 1.18-1.21, I decompose this sorting by age category. For all definitions of urban, high skill workers are not living in urban counties at higher rates than low skill workers when

they are less than 22. Between ages 22 and 35, this changes significantly as high skill workers increasingly choose to live in urban counties. Following age 35, the educational differential with respect to living in urban areas levels out and remains constant for the rest of the prime working-age years. Panel (b) of each figure shows that these estimates are not driven by any specific cohort.

The results in Table 1.16 and Figures 1.22-1.25 more directly examine this sorting by observing movement to urban counties. As shown in Columns 1-4, high skill individuals are significantly more likely move to urban counties than low skill individuals. In Columns 5-8 and the accompanying figures, it is clear that migration in early adulthood (ages 22-35) drives this effect while there is minimal difference in migration across the urban-rural divide between skill groups after 35. Interestingly, the coefficients for living in and migrating to counties in the 80th percentile by population density are smaller when compared to living in and migrating to counties in the 95th percentile. This suggests that skill-based geographic sorting, and particularly that which is driven by early adulthood migration, is most prominent along the division between counties that are a top urban county and those that are not.

1.B Data: U.S. Census

To characterize skill-based sorting and run counterfactual simulations, I utilize data from the following U.S. Census samples: 1980 5 percent, 1990 5 percent, 2000 5 percent, 2010 ACS, and 2019 5-year ACS.⁵³ In each year, I characterize a location's labor force skill mix by examining a sub-population restricted to: prime working-age (25-54), employed, earning a positive wage, and living in the contiguous U.S. Using this sample, I create measures of high and low skill populations (and high skill share as a function of the two) and population density in each location and year. For

53. All download from IPUMS Ruggles et al. 2022.

regressions weighted by total population, this includes all individuals in a location and was downloaded from IPUMS NHGIS (Manson et al. 2022). Additionally, I use data on wages and rents to estimate various measures of a location's 'urbanicity' as described in the next section.

Table 1.13: Importance of Student Debt in for Career Outcomes

	(1) 1[Took Job Outside of Field of Study]	(2) 1[Have to Work >1 Job At a Time]	(3) 1[Have to Work More Hours Than Desired]
Student Debt (\$thds, log)	0.086*** [0.014]	0.085*** [0.013]	0.122*** [0.014]
Observations	2,080	2,080	2,080
R-squared	0.141	0.125	0.144
Dep. Mean	0.363	0.255	0.346
Student Loan Mean	3.171	3.171	3.171
Student Loan S.D.	0.953	0.953	0.953
Demographic Controls	✓	✓	✓
ELS Major FEs	✓	✓	✓

*** p< 0.01, ** p< 0.05, * p< 0.1

NOTE: Estimates from a linear probability model. Dependent variable is a survey question asking directly about the influence of student debt. Demographic controls include sex, race, father's education, father's occupation, and high school GPA. ELS major fixed effects are reported academic major in postsecondary education. Number of observations rounded to the nearest 10 per IES restricted-use guidelines. Robust standard errors clustered by sample design (survey base-year high school) in parentheses.

SOURCE: U.S. Department of Education, National Center for Education Statistics, Education Longitudinal Study of 2002 (ELS:2002), "Base Year", 2002, "First Follow-up", 2004, "High School Transcripts", 2005, "Second Follow-up", 2006, and "Third Follow-up", 2012.

Table 1.14: Educational Migration Differentials: Overall Mobility

	(1)	(2)
	I[Move]	I[Move]
I[High-Skill]	0.016** (0.005)	
× I[Age: <22]		0.005 (0.012)
× I[Age: 22-25]		0.051*** (0.011)
× I[Age: 26-30]		0.096*** (0.011)
× I[Age: 31-35]		0.033*** (0.011)
× I[Age: 36-40]		-0.007 (0.010)
× I[Age: 41-45]		-0.013 (0.010)
× I[Age: 46-50]		-0.035*** (0.011)
× I[Age: 51-55]		-0.013 (0.011)
× I[Age: 56-60]		0.002 (0.012)
× I[Age: 61-65]		0.022* (0.013)
Observations	86,447	86,447
Unique Individuals	6,983	6,983
R-squared	0.113	0.135
Dep. Variable Mean	0.282	0.282
Demographic Controls	✓	✓
Age Category FEs	✓	✓
Wave FEs	✓	✓
Cohort FEs	✓	✓
PSID Age FEs	✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Dependent variable is an indicator function for whether an individual changed counties from wave $t - 1$ to t . Based on the estimation of Equations 1.25 and 1.26. PSID Age is the number of waves an individual has been in the PSID sample. Robust standard errors clustered by individual in parentheses. Demographic controls include sex, race, and marital status.

SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).

Table 1.15: Educational Migration Differentials: Live in ‘Urban’

	(1) Top RUCC	(2) Pop. Density Percentile	(3) 80 th Percentile Pop. Density	(4) 95 th Percentile Pop. Density	(5) Top RUCC	(6) Pop. Density Percentile	(7) 80 th Percentile Pop. Density	(8) 95 th Percentile Pop. Density
I[High-Skill]	0.088*** (0.012)	4.376*** (0.508)	0.104*** (0.011)	0.127*** (0.014)				
× I[Age: <22]					-0.001 (0.025)	0.451 (1.121)	-0.013 (0.025)	0.001 (0.028)
× I[Age: 22-25]					0.041*** (0.014)	3.139*** (0.607)	0.071*** (0.013)	0.100*** (0.017)
× I[Age: 26-30]					0.095*** (0.013)	4.244*** (0.539)	0.094*** (0.012)	0.129*** (0.015)
× I[Age: 31-35]					0.095*** (0.013)	4.564*** (0.555)	0.110*** (0.012)	0.128*** (0.015)
× I[Age: 36-40]					0.096*** (0.014)	4.627*** (0.570)	0.118*** (0.013)	0.136*** (0.016)
× I[Age: 41-45]					0.098*** (0.015)	4.284*** (0.601)	0.113*** (0.014)	0.137*** (0.017)
× I[Age: 46-50]					0.095*** (0.017)	5.046*** (0.659)	0.122*** (0.015)	0.133*** (0.019)
× I[Age: 51-55]					0.087*** (0.019)	4.811*** (0.748)	0.112*** (0.017)	0.143*** (0.020)
× I[Age: 56-60]					0.091*** (0.021)	4.747*** (0.832)	0.105*** (0.020)	0.136*** (0.023)
× I[Age: 61-65]					0.092*** (0.025)	4.519*** (0.987)	0.101*** (0.023)	0.115*** (0.026)
Observations	86,447	86,447	86,447	86,447	86,447	86,447	86,447	86,447
Unique Individuals	6,983	6,983	6,983	6,983	6,983	6,983	6,983	6,983
R-squared	0.294	0.354	0.307	0.062	0.294	0.354	0.307	0.062
Dep. Variable Mean	0.449	83.297	0.700	0.380	0.449	83.297	0.700	0.380
Demographic Controls	✓	✓	✓	✓	✓	✓	✓	✓
Age Category FEs	✓	✓	✓	✓	✓	✓	✓	✓
Wave FEs	✓	✓	✓	✓	✓	✓	✓	✓
Cohort FEs	✓	✓	✓	✓	✓	✓	✓	✓
PSID Age FEs	✓	✓	✓	✓	✓	✓	✓	✓
Origin Location Control	✓	✓	✓	✓	✓	✓	✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Dependent variable in Columns 1, 3, 4, 5, 7, and 8 is an indicator function for whether an individual lived in an ‘urban’ county in wave t based on the following definitions of urban: metropolitan area of 1 million+ (Top RUCC), in the 80th percentile of counties by population density (wave-decade specific), and in the 95th percentile of counties by population density (wave-decade specific). Columns 2 and 6 dependent variable is county percentile in population density (wave-specific). Robust standard errors clustered by individual in parenthesis. SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).


Table 1.16: Educational Migration Differentials: Move to ‘Urban’

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Top RUCC	Pop. Density Percentile	80 th Percentile Pop. Density	95 th Percentile Pop. Density	Top RUCC	Pop. Density Percentile	80 th Percentile Pop. Density	95 th Percentile Pop. Density
I[High-Skill]	0.010*** (0.001)	0.014*** (0.003)	-0.002 (0.002)	0.019*** (0.005)				
× I[Age: <22]					0.001 (0.005)	-0.001 (0.010)	0.002 (0.006)	-0.005 (0.008)
× I[Age: 22-25]					0.025*** (0.005)	0.038*** (0.008)	0.015*** (0.006)	0.019*** (0.007)
× I[Age: 26-30]					0.028*** (0.004)	0.045*** (0.007)	0.007 (0.005)	0.023*** (0.007)
× I[Age: 31-35]					0.009*** (0.003)	0.015** (0.006)	-0.004 (0.004)	0.027*** (0.007)
× I[Age: 36-40]					0.004 (0.003)	0.009 (0.006)	-0.009** (0.004)	0.020*** (0.007)
× I[Age: 41-45]					0.004 (0.003)	-0.005 (0.006)	-0.009* (0.005)	0.011 (0.007)
× I[Age: 46-50]					0.004 (0.003)	-0.000 (0.007)	-0.004 (0.005)	0.018** (0.008)
× I[Age: 51-55]					-0.003 (0.004)	0.001 (0.007)	-0.009* (0.005)	0.013 (0.008)
× I[Age: 56-60]					0.011*** (0.004)	0.008 (0.008)	-0.003 (0.005)	0.023** (0.009)
× I[Age: 61-65]					0.000 (0.004)	0.022*** (0.008)	0.003 (0.006)	0.024** (0.009)
Observations	86,447	86,447	86,447	86,447	86,447	86,447	86,447	86,447
Unique Individuals	6,983	6,983	6,983	6,983	6,983	6,983	6,983	6,983
R-squared	0.017	0.101	0.014	0.013	0.018	0.102	0.015	0.014
Dep. Variable Mean	0.022	0.090	0.031	0.053	0.022	0.090	0.031	0.053
Demographic Controls	✓	✓	✓	✓	✓	✓	✓	✓
Age Category FEs	✓	✓	✓	✓	✓	✓	✓	✓
Wave FEs	✓	✓	✓	✓	✓	✓	✓	✓
Cohort FEs	✓	✓	✓	✓	✓	✓	✓	✓
PSID Age FEs	✓	✓	✓	✓	✓	✓	✓	✓
Origin Location Control	✓	✓	✓	✓	✓	✓	✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Dependent variable in Columns 1, 3, 4, 5, 7, and 8 is an indicator function for whether an individual moved to an ‘urban’ county in wave t based on the following definitions of urban: metropolitan area of 1 million+ (Top RUCC), in the 80th percentile of counties by population density (wave-decade specific), and in the 95th percentile of counties by population density (wave-decade specific). Columns 2 and 6 dependent variable is county percentile in population density (wave-specific). Robust standard errors clustered by individual in parenthesis. SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).

Figure 1.17: Association Between Being High-Skill and Probability of Changing Locations




NOTE: Figure (a) shows the association between being high-skill and the probability of changing county of residence in a given PSID wave relative to the previous wave. Figure (b) shows this correlation for each cohort separately. Graphical representation of the results in Table 1.14 Column 2.


SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).

Figure 1.18: Association Between Being High-Skill and Probability of Living in Top Metropolitan County (RUCC-Defined)

(a) Aggregate Estimate




(b) By Cohort



NOTE: Figure (a) shows the association between being high-skill and the probability of living in a top metropolitan county in a given PSID wave, as defined by the Rural-Urban Continuum Codes (decade-specific). Figure (b) shows this correlation for each cohort separately. Graphical representation of the results in Table 1.15 Column 5.

SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).

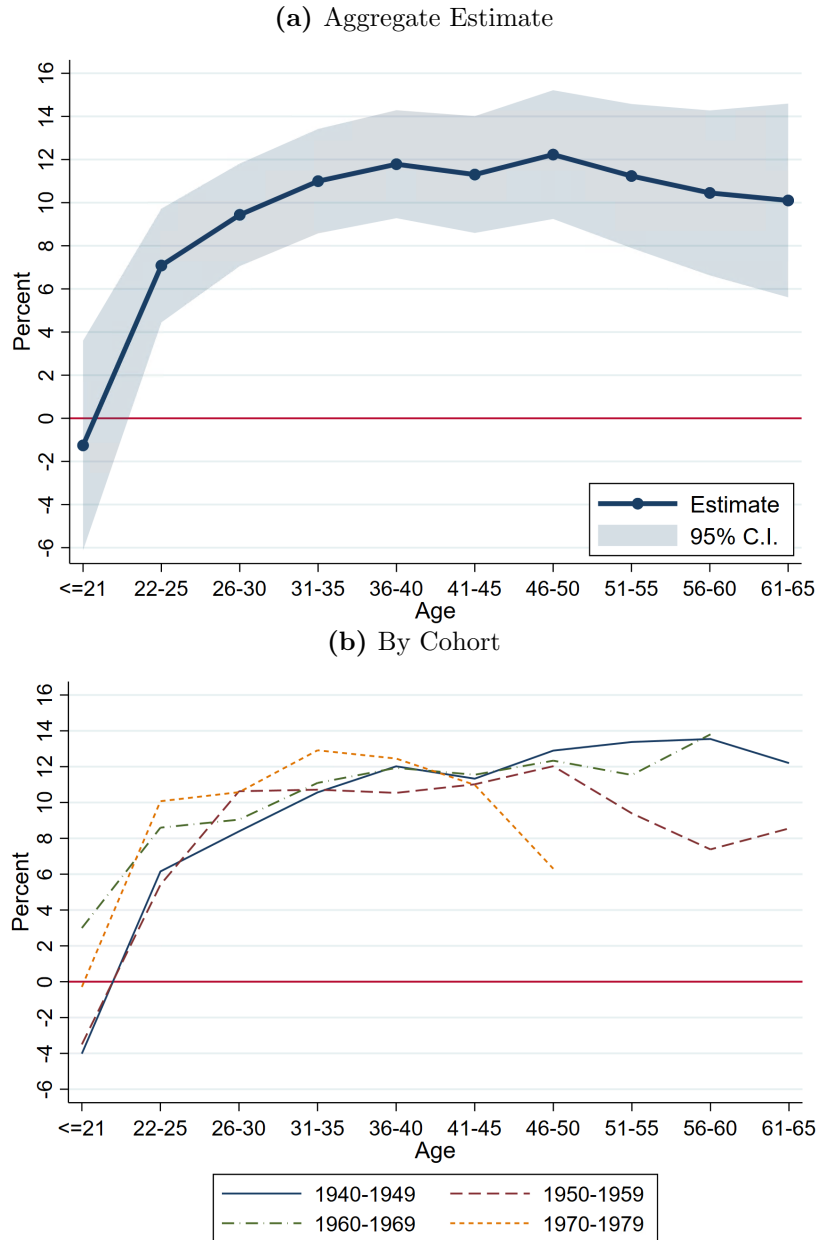
Figure 1.19: Association Between Being High-Skill and Percentile of County Population Density



NOTE: Figure (a) shows the association between being high-skill and the percentile of county population density (decade-specific) in a given PSID wave. Figure (b) shows this correlation for each cohort separately. Graphical representation of the results in Table 1.15 Column 6.

SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).


Figure 1.20: Association Between Being High-Skill and Probability of Living in 80th Percentile Population Density County



NOTE: Figure (a) shows the association between being high-skill and the probability of living in a county with a population density in the 80th percentile (decade-specific) in a given PSID wave. Figure (b) shows this correlation for each cohort separately. Graphical representation of results in Table 1.15 Column 7.

SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).

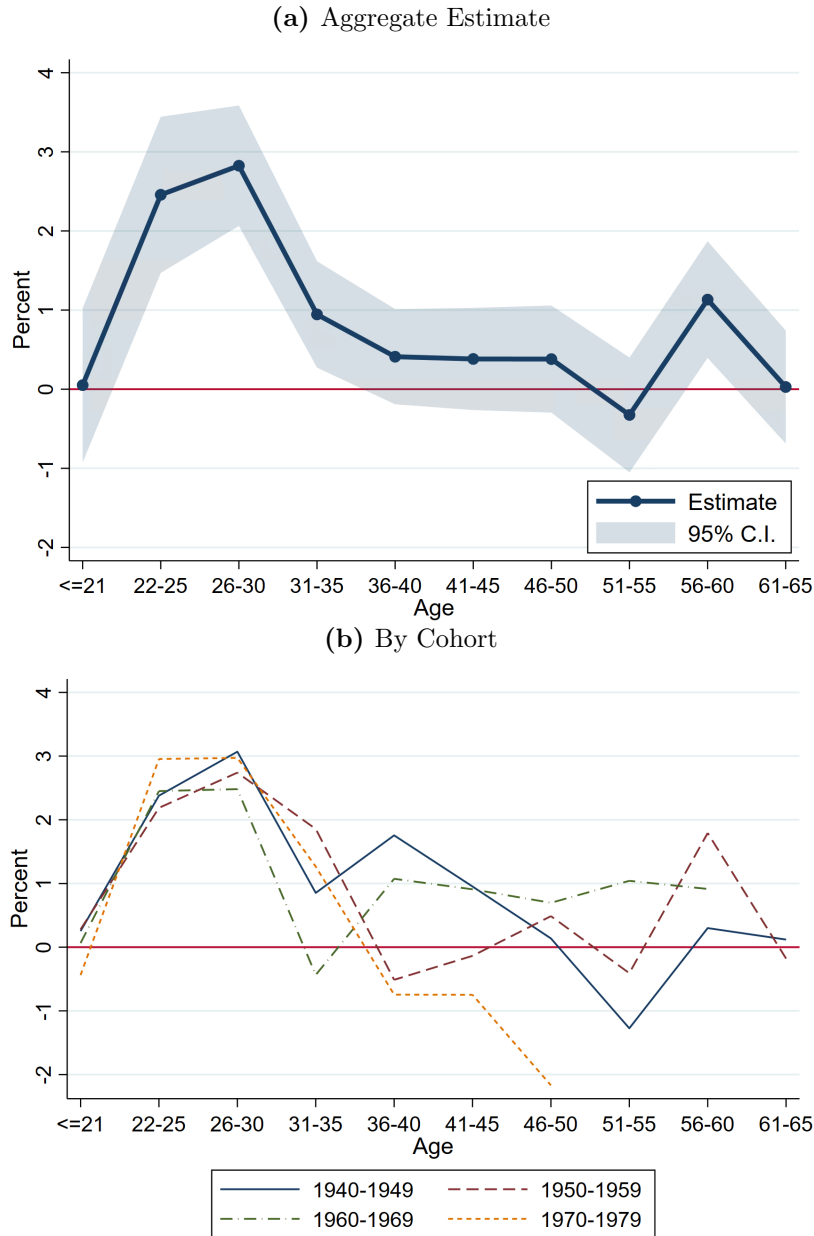
Figure 1.21: Association Between Being High-Skill and Probability of Living in 95th Percentile Population Density County



NOTE: Figure (a) shows the association between being high-skill and the probability of living in a county with a population density in the 95th percentile (decade-specific) in a given PSID wave. Figure (b) shows this correlation for each cohort separately. Graphical representation of results in Table 1.15 Column 8.

SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).


Figure 1.22: Association Between Being High-Skill and Probability of Moving to Top Metropolitan County (RUCC-Defined)



NOTE: Figure (a) shows the association between being high-skill and the probability of moving to a top metropolitan county in a given PSID wave, as defined by the Rural-Urban Continuum Codes (decade-specific). Figure (b) shows this correlation for each cohort separately. Graphical representation of results in Table 1.16 Column 5.

SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).

Figure 1.23: Association Between Being High-Skill and Probability of Moving to Higher Population Density County




NOTE: Figure (a) shows the association between being high-skill and the probability of moving to a county with a higher percentile of county population density (decade-specific) in a given PSID wave, relative to the previous county of residence. Figure (b) shows this correlation for each cohort separately. Graphical representation of results in Table 1.16 Column 6.

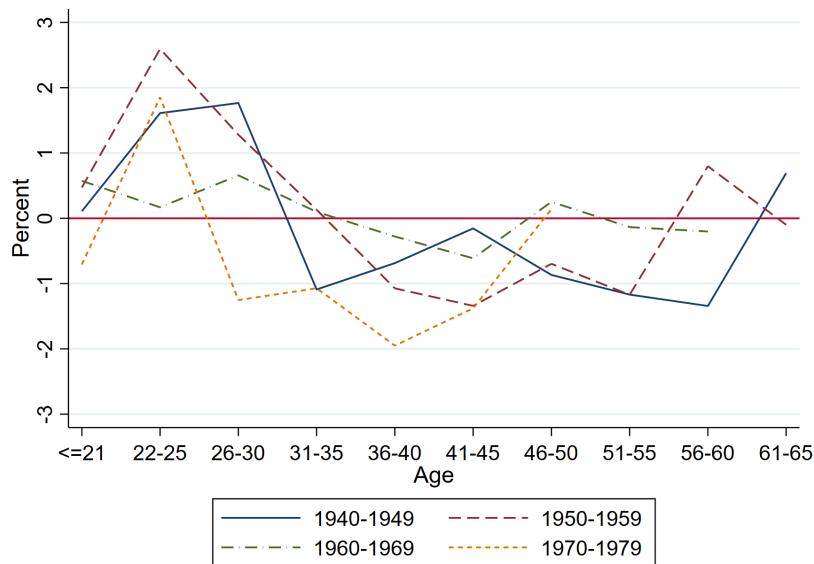
SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).

Figure 1.24: Association Between Being High-Skill and Probability of Moving to 80th Percentile Population Density County

(a) Aggregate Estimate




(b) By Cohort



NOTE: Figure (a) shows the association between being high-skill and the probability of moving to a county with a population density in the 80th percentile (decade-specific) in a given PSID wave. Figure (b) shows this correlation for each cohort separately. Graphical representation of results in Table 1.16 Column 7.

SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).


Figure 1.25: Association Between Being High-Skill and Probability of Moving to 95th Percentile Population Density County



NOTE: Figure (a) shows the association between being high-skill and the probability of moving to a county with a population density in the 95th percentile (decade-specific) in a given PSID wave. Figure (b) shows this correlation for each cohort separately. Graphical representation of results in Table 1.16 Column 8.

SOURCE: Panel Study of Income Dynamics, restricted use data. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, MI (2022).

Figure 1.26: Association Between Student Debt (\$1,000s, log) and Probability of Post-College County Location Being in Top X Percentile Pop. Density



NOTE: From estimation of Equation 1.4 including origin and college county location controls (defined using same population density threshold as dependent variable), demographic controls (sex, race, father’s education, father’s occupation, and high school GPA), and ELS major fixed effects. Confidence intervals reflect robust standard errors clustered by sample design (survey base-year high school). As shown, the coefficient is near zero and insignificant when considering 40th, 50th, 90th, and 95th percentile thresholds. It is positive, significant, and nearly identical when considering location choice across the 60th, 70th, and 80th percentiles. This is further suggestive evidence that student debt is pushing individuals who would locate in counties in the bottom 60th percentile of population density (perhaps one definition of “rural”) to locate in counties in the top 20th percentile.

SOURCE: U.S. Department of Education, National Center for Education Statistics, Education Longitudinal Study of 2002 (ELS:2002), “Base Year”, 2002, “First Follow-up”, 2004, “High School Transcripts”, 2005, “Second Follow-up”, 2006, and “Third Follow-up”, 2012.

References

- Acemoglu, Daron. 2002. "Technical Change, Inequality, and the Labor Market." Journal of Economic Literature 40 (1): 7–72.
- Akers, Beth, and Matthew M Chingos. 2014. Is a Student Loan Crisis on the Horizon? Brown Center on Education Policy at the Brookings Institution Washington, DC.
- Albert, Christoph, and Joan Monras. 2017. "Immigrants' Residential Choices and Their Consequences," IZA Discussion Paper.
- Albouy, David, Walter Graf, Ryan Kellogg, and Hendrik Wolff. 2016. "Climate Amenities, Climate Change, and American Quality of Life." Journal of the Association of Environmental and Resource Economists 3 (1): 205–246.
- Avery, Christopher, and Sarah Turner. 2012. "Student Loans: Do College Students Borrow Too Much—Or Not Enough?" Journal of Economic Perspectives 26 (1): 165–92.
- Barrow, Lisa, and Ofer Malamud. 2015. "Is College a Worthwhile Investment?" Annual Review of Economics 7 (1): 519–555.
- Baum, Sandy, Jennifer Ma, Matea Pender, and CJ Libassi. 2019. "Trends in Student Aid 2019." New York, NY: College Board.
- Baum-Snow, Nathaniel, Matthew Freedman, and Ronni Pavan. 2018. "Why Has Urban Inequality Increased?" American Economic Journal: Applied Economics 10 (4): 1–42.

- Berry, Christopher R, and Edward L Glaeser. 2005. "The Divergence of Human Capital Levels Across Cities." Papers in Regional Science 84 (3): 407–444.
- Bettinger, Eric, Oded Gurantz, Laura Kawano, Bruce Sacerdote, and Michael Stevens. 2019. "The Long-Run Impacts of Financial Aid: Evidence from California's Cal Grant." American Economic Journal: Economic Policy 11 (1): 64–94.
- Black, Sandra E, Jeffrey T Denning, Lisa J Dettling, Sarena Goodman, and Lesley J Turner. 2020. "Taking It to the Limit: Effects of Increased Student Loan Availability on Attainment, Earnings, and Financial Well-Being," NBER Working Paper.
- Bleemer, Zachary, Meta Brown, Donghoon Lee, Katherine Strair, and Wilbert Van der Klaauw. 2021. "Echoes of Rising Tuition in Students' Borrowing, Educational Attainment, and Homeownership in Post-recession America." Journal of Urban Economics 122:103298.
- Bleemer, Zachary, Meta Brown, Donghoon Lee, Wilbert Van der Klaauw, et al. 2014. "Tuition, Jobs, or Housing: What's Keeping Millennials at Home." Staff Reports, no. 700.
- Bound, John, and Harry J Holzer. 2000. "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s." Journal of labor Economics 18 (1): 20–54.
- Burdman, Pamela. 2005. "The Student Debt Dilemma: Debt Aversion as a Barrier to College Access." The Institute for College Access and Success, Inc.

- Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro. 2019. "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock." Econometrica 87 (3): 741–835.
- Callender, Claire, and Jonathan Jackson. 2005. "Does the Fear of Debt Deter Students from Higher Education?" Journal of Social Policy 34 (4): 509–540.
- Card, David. 2009. "Immigration and inequality." American Economic Review 99 (2): 1–21.
- Chakrabarti, Rajashri, Nicole Gorton, and Michael F Lovenheim. 2020. "State Investment in Higher Education: Effects on Human Capital Formation, Student Debt, and Long-Term Financial Outcomes of Students." NBER Working Paper.
- Chapman, Stephanie. 2015. "Student Loans and the Labor Market: Evidence from Merit Aid Programs." Northwestern University Department of Economics, Chicago, IL.
- Conzelmann, Johnathan G, Steven W Hemelt, Brad Hershbein, Shawn M Martin, Andrew Simon, and Kevin M Stange. 2022. "Grads on the Go: Measuring College-Specific Labor Markets for Graduates." NBER Working Paper.
- Dettling, Lisa J, and Joanne W Hsu. 2018. "Returning to the Nest: Debt and Parental Co-residence among Young Adults." Labour Economics 54:225–236.
- Di Maggio, Marco, Ankit Kalda, and Vincent Yao. 2020. "Second Chance: Life Without Student Debt." NBER Working Paper.
- Diamond, Rebecca. 2016. "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000." American Economic Review 106 (3): 479–524.

- Durante, Alex, Jeff Larrimore, Christina Park, Anna Tranfaglia, et al. 2017. Report on the Economic Well-Being of US Households in 2016. Technical report. Board of Governors of the Federal Reserve System (US).
- Dynarski, Susan. 2008. “Building the Stock of College-Educated Labor.” Journal of Human Resources 43 (3): 576–610.
- Eeckhout, Jan, Roberto Pinheiro, and Kurt Schmidheiny. 2014. “Spatial Sorting.” Journal of Political Economy 122 (3): 554–620.
- Field, Erica. 2009. “Educational Debt Burden and Career Choice: Evidence from a Financial Aid Experiment at NYU Law School.” American Economic Journal: Applied Economics 1 (1): 1–21.
- Gervais, Martin, and Nicolas L Ziebarth. 2019. “Life After Debt: Postgraduation Consequences of Federal Student Loans.” Economic Inquiry 57 (3): 1342–1366.
- Giannone, Elisa. 2017. “Skilled-Biased Technical Change and Regional Convergence.” University of Chicago. Unpublished manuscript.
- Gicheva, Dora. 2016. “Student Loans or Marriage? A Look at the Highly Educated.” Economics of Education Review 53:207–216.
- Gicheva, Dora, and Jeffrey Thompson. 2015. “The Effects of Student Loans on Long-Term Household Financial Stability.” Student loans and the dynamics of debt, 287–316.
- Glaeser, Edward L. 2008. Cities, Agglomeration, and Spatial Equilibrium. OUP Oxford.
- Glaeser, Edward L, and Joshua D Gottlieb. 2008. “The Economics of Place-Making Policies.” NBER Working Paper.

- Greenwood, Michael J. 1975. "Research on internal migration in the United States: A survey." Journal of economic literature, 397–433.
- Hornbeck, Richard, and Enrico Moretti. 2018. "Who benefits from productivity growth? Direct and indirect effects of local TFP growth on wages, rents, and inequality."
- Houle, Jason N, and Lawrence Berger. 2015. "Is student loan debt discouraging homeownership among young adults?" Social Service Review 89 (4): 589–621.
- Howard, Greg, Jack Liebersohn, et al. 2018. "The geography channel of house price appreciation." Fisher College of Business Working Paper No.
- Hsieh, Chang-Tai, and Enrico Moretti. 2019. "Housing constraints and spatial misallocation." American Economic Journal: Macroeconomics 11 (2): 1–39.
- Kane, Thomas J. 2003.
A quasi-experimental estimate of the impact of financial aid on college-going.
- Katz, Lawrence F, and David H Autor. 1999. "Changes in the wage structure and earnings inequality." In Handbook of Labor Economics, vol. 3A, 1463–1555. Elsevier.
- Katz, Lawrence F, and Kevin M Murphy. 1992. "Changes in relative wages, 1963–1987: supply and demand factors." The quarterly journal of economics 107 (1): 35–78.
- Kline, Patrick, and Enrico Moretti. 2014. "People, places, and public policy: Some simple welfare economics of local economic development programs." Annu. Rev. Econ. 6 (1): 629–662.

- Krishnan, Karthik, and Pinshuo Wang. 2019. "The cost of financing education: can student debt hinder entrepreneurship?" Management Science 65 (10): 4522–4554.
- Ladinsky, Jack. 1967. "The geographic mobility of professional and technical manpower." Journal of Human Resources, 475–494.
- Liu, Shimeng, Weizeng Sun, and John V Winters. 2019. "Up in STEM, down in business: changing college major decisions with the great recession." Contemporary Economic Policy 37 (3): 476–491.
- Long, Bridget Terry. 2014. "The financial crisis and college enrollment: How have students and their families responded?" In How the financial crisis and Great Recession affected higher education, 209–233. University of Chicago Press.
- Looney, Adam, and Constantine Yannelis. 2018. "Borrowers with large balances: Rising student debt and falling repayment rates." The Brookings Institution: Washington, DC.
- Luo, Mi, and Simon Mongey. 2019. Assets and job choice: Student debt, wages and amenities. Technical report. National Bureau of Economic Research.
- Malamud, Ofer, and Abigail Wozniak. 2012. "The impact of college on migration evidence from the Vietnam generation." Journal of Human resources 47 (4): 913–950.
- Manson, Steven, Jonathan Schroeder, David Van Riper, Tracy Kugler, and Steven Ruggles. 2022. IPUMS National Historical Geographic Information System: Version 17.0. <http://doi.org/10.18128/D050.V17.0>.

- Marx, Benjamin M, and Lesley J Turner. 2019. “Student loan nudges: Experimental evidence on borrowing and educational attainment.” American Economic Journal: Economic Policy 11 (2): 108–41.
- Mezza, Alvaro, Daniel Ringo, Shane Sherlund, and Kamila Sommer. 2020. “Student loans and homeownership.” Journal of Labor Economics 38 (1): 215–260.
- Minicozzi, Alexandra. 2005. “The short term effect of educational debt on job decisions.” Economics of Education Review 24 (4): 417–430.
- Molloy, Raven, Christopher L Smith, and Abigail Wozniak. 2011. “Internal migration in the United States.” Journal of Economic perspectives 25 (3): 173–96.
- Monras, Joan. 2018. “Economic shocks and internal migration.”
- Monras, Joan, and Christoph Albert. 2022. “Immigration and Spatial Equilibrium: The Role of Expenditures in the Country of Origin.” The Center for Growth and Opportunity.
- Moretti, Enrico. 2011. “Local labor markets.” In Handbook of labor economics, 4:1237–1313. Elsevier.
- . 2013. “Real wage inequality.” American Economic Journal: Applied Economics 5 (1): 65–103.
- Neumark, David, and Helen Simpson. 2015. “Chapter 18 - Place-Based Policies.” In Handbook of Regional and Urban Economics, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, 5:1197–1287. Handbook of Regional and Urban Economics. Elsevier.
<https://doi.org/https://doi.org/10.1016/B978-0-444-59531-7.00018-1>.

- Oreopoulos, Philip, and Uros Petronijevic. 2013. "Making college worth it: A review of research on the returns to higher education."
- Roback, Jennifer. 1982. "Wages, rents, and the quality of life."
Journal of political Economy 90 (6): 1257–1278.
- Rosen, Sherwin. 1979. "Wage-based indexes of urban quality of life."
Current issues in urban economics, 74–104.
- Rothstein, Jesse, and Cecilia Elena Rouse. 2011. "Constrained after college: Student loans and early-career occupational choices." Journal of Public Economics 95 (1-2): 149–163.
- Ruggles, Steven, Sarah Flood, Ronald Goeken, Megan Schouweiler, and Matthew Sobek. 2022. IPUMS USA: Version 12.0 [dataset]. IPUMS.
<https://doi.org/10.18128/D010.V12.0>.
- Saiz, Albert. 2010. "The geographic determinants of housing supply."
The Quarterly Journal of Economics 125 (3): 1253–1296.
- Scala, Dante J, and Kenneth M Johnson. 2017. "Political polarization along the rural-urban continuum? The geography of the presidential vote, 2000–2016."
The ANNALS of the American Academy of Political and Social Science 672 (1): 162–184.
- Shao, Ling. 2014. "Debt, marriage, and children: The impact of student loans on marriage and fertility."
PhD diss., Department of Economics, University of Chicago, San Diego.

- Shapiro, Jesse M. 2006. “Smart cities: quality of life, productivity, and the growth effects of human capital.” The review of economics and statistics 88 (2): 324–335.
- Suárez, Serrato, Juan Carlos, and Owen Zidar. 2016. “Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms.” American Economic Review 106 (9): 2582–2624.
- Van Horn, Carl, Cliff Zukin, Mark Szeltner, and Charley Stone. 2012. “Left out. Forgotten? Recent high school graduates and the great recession.” John J. Heldrich Center for Workforce Development, Rutgers University.
- Weidner, Justin. 2016. “Does student debt reduce earnings?” Princeton University.
- Wozniak, Abigail. 2010. “Are college graduates more responsive to distant labor market opportunities?” Journal of Human Resources 45 (4): 944–970.

CHAPTER 2

Does Increased Union Power Cause Pension Under-Funding in the Public Sector?

2.1 Introduction

Unfunded public-sector pension obligations are the biggest fiscal challenge that many U.S. states (and cities) face in the coming decades: Novy-Marx and Rauh (2009) estimate that the aggregate funding gap for the largest state and municipal plans is above \$3 trillion. Where the funding gap is sufficiently large, it has culminated in municipal bankruptcy, and could conceivably culminate in a state's bankruptcy.¹ Pension under-funding is first and foremost a political economy problem that is best understood through the lens of models where politicians trade off the separate demands of Special Interest Groups (SIGs) and voters to maximize the probability of re-election (Lindbeck and Weibull 1987; Grossman and Helpman 1994; Persson and Tabellini 2000).² Voters value lower taxes and balanced budgets (holding the level of public good provision fixed). Public-sector unions, the SIGs, negotiate for rents, of

1. Unlike federal social security, municipal and state benefit payments are legally binding commitments, and state laws make it almost impossible to renegotiate them outside of bankruptcy proceedings (Burns 2011; Trusts 2013). No U.S. state has defaulted since the 1840s, but according to Anderson (2013), there have been twenty-eight municipal bankruptcies between 2007 and 2013. These bankruptcies are never mono-causal, but pension obligations are almost always a core issue, as in Detroit in 2014 (The Economist, March 1st 2014).

2. Lindbeck and Weibull (1987) includes bargaining explicitly while in Grossman and Helpman (1994), the politicians maximize a utilitarian social welfare function that assigns differential weights to SIGs and voters. Note that these weights can still be intuitively interpreted as the reduced form of a bargaining solution.

which there are two types: pensions benefits and wages.

From the politician’s perspective, wage increases need to be financed by either taxes or debt, both of which voters dislike. The costs of pension benefits, by contrast, can be made to largely disappear from view. Glaeser and Ponzetto (2014) refer to this as the ‘shroudedness’ of pension obligations. One source of shroudedness is that pensions can be what we refer to as *actuarially* under-funded: unrealistic actuarial assumptions can, for a while, lower the required contributions needed to fund tomorrow’s benefits, and thus hide the true cost of promises of future benefits (Mitchell and Smith 1994; Kelley 2014).³ In the words of a member of the Society of Actuaries Pension Financing Task Force, “consistent low-balling of pension costs over the past two decades made it easy for elected officials and union representatives to agree on very valuable benefits, for very much smaller current pay concessions” (Malanga 2016).

A second source of shroudedness is that pensions can be *materially* under-funded: the employer (government) can, without penalty, hold off on paying even the actuarially required contributions (Brinkman, Coen-Pirani, and Sieg 2018). The resulting actuarial funding gap is the officially reported one, but may still stay under voters’ radar because it is not considered public debt despite the fact that legally binding obligations are being accrued (Brown and Dye 2015; Munnell, Aubry, Cafarelli, et al. 2015).⁴

From the politician’s perspective, pension promises, therefore, hold greater appeal as concessions to public-sector unions compared to wage increases. This viewpoint

3. Because of unrealistic actuarial assumptions, researchers and rating agencies have concluded that the true aggregate funding gap is in excess of three times as large as the official actuarial funding gap (Novy-Marx and Rauh 2009).

4. Even if pensions obligations were listed as official debt, future taxes needed to finance them might be less salient than today’s taxes to myopic voters, so that Ricardian Equivalence might fail even without the ‘shroudedness’.

is supported by examples such as the former mayor of Houston, Lee P. Brown, who once justified a largely unfunded pension benefit increase under his administration by the fact that “it was budget neutral”, and that he did not “have the funds to give municipal employees the raises they deserved” (Boylan 2016). Consequently, increases in the bargaining power of public-sector unions – which models of SIG politics tell us should result in higher rents for that group – are likely to result in increased pension benefits.

In this paper, we test this hypothesis by exploiting variation in collective bargaining rights for public-sector workers, i.e., an institutional shifter of their bargaining power. Collective bargaining laws dictate the extent to which government employers are obligated to negotiate with unions on compensation and other conditions of employment. They vary significantly across states, over time, and across the five large public-sector employment groups: state employees, police, fire-fighters, non-college teachers, and other municipal workers. This breakdown corresponds well to the way public-sector pension systems are organized since these are always unique to a state and commonly separately organized by employment group.

Fortunately, there is also an empirical basis for such an investigation: the NBER Public Sector Collective Bargaining Law Data Set (PSCBLD) measures public-sector collective bargaining rights from 1955 to 1986 by state and year, as well as by the five employment groups (R. Freeman and R. Valletta 1988). Section 2.B describes these data in detail. Unfortunately, the dataset is quite dated.⁵ We therefore extended this dataset to the present day by first identifying the laws governing collective bargaining for each employment group in each state and, if the status quo differed from what it was at the end of the R. Freeman and R. Valletta (1988) data or the current law

5. Kim Reuben partially extended the data to 1996. The update appears to have been subsumed into R. Freeman and R. Valletta (1988) on the NBER website, without a separate citation. One other extension is Sanes, Schmitt, et al. (2014), who provide a qualitative cross-sectional snapshot of 2014 bargaining rights.

post-dates the end of that dataset, then we use Nexis Uni to identify and date the corresponding legal changes. We refer to the extended data as the Extended Public Sector Collective Bargaining Law Data Set, or the E-PSCBLD.⁶

From the perspective of causal identification, there is a concern that states which grant more legal rights for collective bargaining may also be states where voters are inclined towards more generous compensation for public-sector workers. We address this concern by identifying the source of changes to collective bargaining rights between two options: political and judicial. Political changes are the result of legislative or gubernatorial executive orders – both of which are more likely to reflect public sentiment. Alternatively, judicial changes are the result of court rulings (case law) or attorney general opinions, which are less likely to reflect voter preferences.⁷

By linking the E-PSCBLD to pension plan financial information, we show that collective bargaining rights for public-sector employees increase plan generosity historically and are associated with plan under-funding in the post-2000 period. We first exploit state, group, and time variation in the legal status of public-sector collective bargaining rights to show that historical changes resulted in higher benefit payments per retiree relative to contributions per member. This result is consistent for legal changes that are politically driven and judicially driven, indicating that it is unlikely to be explained by unobserved local sentiment towards public-sector employees.

In the post-2000 period, where more detailed financial information is available, we show that collective bargaining rights are significantly associated with higher levels of unfunded pension liabilities. We provide descriptive evidence that increased plan generosity could have resulted in lower funding ratios via two potential mechanisms. First, pension plans in areas with stronger collective bargaining rights experienced lower government contributions, measured as the share of the actuarially required con-

6. See Section 2.C for detailed documentation of this extension.

7. See Section 2.C for detailed documentation of how to draw these distinctions.

tributions – a pattern indicative of *material* under-funding. Second, pension plans in areas with stronger collective bargaining rights are more likely to have rosier expectations of investment performance, even when accounting for prior realized returns – a pattern that is indicative of *actuarial* under-funding.

Related Literature. A number of influential political economy theories model policies as being chosen by politicians who trade off the separate demands of special interest groups (SIGs) and voters in order to maximize the probability of re-election (Lindbeck and Weibull 1987; Grossman and Helpman 1994; Persson and Tabellini 2000). Instead of this binary distinction into SIGs and voters, one can also conceptualize all voters as interest groups that vary in their influence on the politician. Groups have more influence if they are better able to organize and ‘act collectively’ (Olson 2009); if they are more homogenous in their political preference and only chose whether to ‘come out to vote’ rather than whom to vote for (Glaeser, Ponzetto, and Shapiro 2005); if they care more intensely about an issue (Bouton et al. 2021); or if they are better informed (Besley and Burgess 2002; Adsera, Boix, and Payne 2003; Strömberg 2004; Ferraz and Finan 2008; Glaeser and Ponzetto 2014). We contribute to this literature by estimating the impact that stronger legal provisions for a group’s bargaining rights have on policies favoring that group.

We also contribute to the literature examining the effects of collective bargaining. There is substantial evidence that private- and public-sector unions are able to bargain for increased wages and higher employment levels relative to their non-unionized counterparts (Hirsch and Rufolo 1982; R. B. Freeman and R. G. Valletta 1987; Zax and Ichniowski 1988).⁸ However, there appears to be no previous work on the effect

8. The literature that studies the effect of private-sector unionization on wages and employment is perhaps more developed than that analyzing the public sector. Ashenfelter (1978) is one of the earliest seminal papers in the private-sector literature, while Freeman (1986), Reder (1988), and Freeman and Han (2012) examine public-sector unionization and its causes.

of collective bargaining rights on retirement benefits.⁹

Lastly, we add to the literature on the causes of public-sector pension underfunding (Grosskopf, Hayes, and Kennedy 1985; Inman 1985; Inman and Albright 1987; Mitchell and Smith 1994; Novy-Marx and Rauh 2009; Burns 2011; Mohan and Zhang 2014; Munnell, Aubry, Cafarelli, et al. 2015). The political economy problems of pension underfunding are acknowledged by most observers, but empirical work on this is scant.¹⁰ Anzia and Moe (2016) shows that the presence of union representatives on plans' boards correlates with lower funding ratios, and also provides an illustrative account of the bruising political battles surrounding efforts to make pension plans' actuarial assumptions more realistic in Rhode Islands in 2011 and California in 2015. Causally identified evidence comes from Dippel (2022) who employs a regression discontinuity design around close mayoral elections to show that municipal pension benefits grow disproportionately under Democratic Party mayors, and argues that Democratic Party mayoral candidates use pension benefits to 'bring out their base', as suggested by theory (Glaeser, Ponzetto, and Shapiro 2005).

Section 2.2 provides background information on public pension plan finances and the political processes that determine retirement benefits. Sections 2.3 and 2.4 describe our data tracking collective bargaining rights and pension plan finances, respectively. Section 2.5 outlines the empirical framework for the analysis, while Section 2.6 presents the results. Section 2.7 concludes the chapter.

2.2 Background on Public-Sector Union Bargaining

9. Munnell, Aubry, and Quinby (2011) documents that there is correlation between union coverage and benefits, but union coverage does not measure collective bargaining power, and the paper does not attempt to estimate causal effects.

10. Rigorous studies on pension benefit funding have tended to focus on pensions' fund management and on accounting practices since these have been more easily measurable contributors (Novy-Marx and Rauh 2009, 2014a, 2014b; Brown and Wilcox 2009).

2.2.1 Pension Plan Finances 101: Actuarial Accounting

To assess the funding status of a public pension plan, policymakers rely on the Actuarial Funding Ratio (AFR), a measure that captures the share of a plan's discounted future benefit obligations to its pensioners and active members that is funded by current financial assets.¹¹ Although the actuarial accounting that goes into calculating a plan's AFR is more complicated, it can be broadly summarized (at time τ) by the following expression

$$\text{AFR}_{i\tau} = \frac{\text{Assets}_{i\tau}}{\sum_{t>\tau}^{\infty} \frac{\text{Benefits}_{it}}{(1+\text{AAR}_i)^t}} \quad , \quad (2.1)$$

where AAR_i is the Actuarially Assumed Return on a plan's assets. With an increase in the AAR, future benefit obligations to pensioners are discounted further and the funding status appears more sound. For this reason, the AAR has become a contentious point in policy circles, a topic we return to in Section 2.2.2.

Note that both past and future contributions do not appear directly in Equation 2.1. Past contributions by employers and employees and any returns on them are subsumed into the plan's current asset base. Future contributions and obligations are assumed to directly offset each other (backed up by actuarial accountant calculations), and neither appears in Equation 2.1. In Defined Contribution (DC) plans, which dominate the private sector, asset returns directly determine benefit payments and Equation 2.1 always equals zero. In stark contrast, almost all municipal and state pension systems in the U.S. are Defined Benefit (DB) plans, where future benefit obligations are legally binding contracts independent of a plan's asset performance.

11. An equivalent measure that is often used is the Unfunded Actuarially Accrued Liabilities (UAAL), which measures the difference between a plan's assets and its discounted future benefit obligations to its pensioners and active members.

2.2.2 How Funding Gaps Emerge

A pension plan funding gap can occur for one of two separate but not independent reasons¹²: from unrealistic actuarial assumptions, which result in *actuarial* underfunding, and from unmet expected contributions, which lead to *material* underfunding. Unrealistic actuarial assumptions lead to underfunding by lowering the Actuarially Required Contributions (ARC), which is the level of contributions needed to cover future benefit obligations. In the short run, unrealistic actuarial assumptions do not substantially affect a plan's funding status because the AFR is calculated under the same erroneous assumptions; however, over time they lead to sustained growth in funding gaps and significant jumps when assumptions are corrected. To provide more context on what unrealistic actuarial assumptions might entail, the following discusses how benefit increases are linked to the calculation of the ARC.

Pension plan rules regarding benefits can be changed through any number of processes, including collective bargaining, changes in state statutes, and executive directives (see 2.2.3 for a full discussion). They can take many forms, but the simplest is an across-the-board increase in benefits, e.g. 10% higher benefits for all recipients. A more common form is 'formula enhancements', where changes are made to the formula used to calculate retirement benefits for an employee. These formulas usually have two components: (1) a minimum age of retirement, and (2) a percentage that is applied to an employee's years of service and salary to determine their benefit level upon retirement. For example, consider the retirement formula '2 at 50', which means a worker can retire starting at age 50, and draw a pension that equals 2% of their last annual salary for every year of service.¹³ Put another way, if a police officer is part of

12. In addition, funding gaps can also open up because of variation in fund management since $\text{Assets}_{i\tau}$, i.e. the plan's assets at time τ are determined by the returns on past contributions. However, there is so much smoothing in actuarial return calculations of pension funds that short-run variation in fund management performance is not a strong force in the data.

13. Some plans consider the reference salary to be an average of employee's highest-earning years.

a plan with this formula and has worked at the department since the age of 20, they could retire at age 50 and receive 60% of their last year's salary as a pension, or retire at age 65 and draw 90% of their last annual salary. Formula enhancements take the form of moving a '2 at 50' formula to a '3 at 50' formula, or a '3 at 55' formula to a '3 at 50' formula.

After benefit levels have been decided, the ARC are calculated by actuarial accountants and are, by definition, deemed adequate to cover obligations under the given actuarial assumptions. These assumptions are not determined by the accountants, however, and often fall to pension boards where they are heavily politicized choices (Anzia and Moe 2016; Greenhut 2009). One element of these assumptions is the expected rate of return on pension assets, the AAR. Most plans assume an AAR between 7 and 8 percent, but this is almost always higher than actual returns have been over the last decade (Martin, Kantchev, and Nairioka, November 13th 2016).¹⁴ Another important assumption regards employee incentives to retire. Formula enhancements, in particular, could lead to earlier retirement choices by employees, which both increase the expected years of benefit payments and decreases the member's years of contributions (and thus, the asset base of the plan as well). Modeling adjustments in these incentives can be complex and actuarial estimations of ARC often fail to account for them, leading to 'blind spots' in the calculations (Mitchell and Smith 1994).

Responsibility for the ARC is split between employers and employees – a setup that facilitates the additional mechanism leading to the emergence of funding gaps.

Typically, employees receive their highest salaries at the end of their career so this variation does not meaningfully change retirement incentives.

14. A related issue is whether the practice of discounting future obligations at the expected rate of return on assets is appropriate. Logically, it is inconsistent to discount a stream of effectively 'risk-free obligations' at the rate of return of a risky portfolio of assets (Novy-Marx and Rauh 2009, 2011, 2014a, 2014b; Brown and Wilcox 2009). Yet, state laws sanction public-sector plans to do precisely this (while simultaneously prohibiting private-sector 401(k) plans from doing the same).

The employee portion of the ARC is drawn directly from employee paychecks and cannot be shirked. Public employers, on the other hand, are under no obligation to pay their portion of the ARC (Brown and Dye 2015; Brinkman, Coen-Pirani, and Sieg 2018). Recent work has shown that a significant amount of pension under-funding may be caused by just such under-paying by employers, which is called *material* under-funding (Brown and Dye 2015; Munnell, Aubry, Cafarelli, et al. 2015). Importantly, although governments have flexibility in the short-run, pension obligations remain a legally binding commitment and future payments must be made when they are due.

2.2.3 The Role of Politics

Funding gaps, including those emerging from unrealistic actuarial assumptions, are tolerated and encouraged by government leaders because they are politically expedient. Although the associated pension obligations inevitably increase the government burden and are naturally unwelcome to politicians, they are often costs in the long term and government leaders may push for unrealistic actuarial assumptions to ease pressure in the short-run. For example, increasing the projected rate of return on assets (AAR) can mask growing funding gaps or make it appear to maintain a constant funding status despite benefit increases. Indeed, pension boards have in the past neutralized the transmission from benefit expansion to the ARC by simultaneously increasing their AAR (Mitchell and Smith 1994; Kelley 2014; Novy-Marx and Rauh 2011). Eventually, if the actual rate of return does not meet the higher AAR, the funding gap will grow.¹⁵

15. Efforts to lower plans' AAR have been the most acrimonious battleground in the pension field in recent years, typically fought out between union representatives and treasury representatives on a plan's board. Lowering the AAR is consequential because it immediately opens up a gap in Equation 2.1, which then immediately results in higher ARC for both employers and employees. This topic and the implications of these shocks for the local economy are explored in Chapter 3.

In addition to adjusting actuarial assumptions for more pleasant funding gaps, local government leaders can also shirk their share of the ARC. This is due to the fact that unfunded liabilities are not accurately labeled as debt, despite the legal obligation to pay out future benefits. They are instead seen as ‘budget neutral’. This feature, which results in *material* under-funding, along with the *actuarial* under-funding outlined in the previous paragraph, allow politicians to ease financial concerns regarding public-sector pension debt when politically motivated to do so.

Unlike general government spending that may be determined by the median voter theorem, the political pressures for expenditure on public-sector pensions are more appropriately captured by models of special interest group politics. Special interest groups arise in political systems allowing targeted, selective transfers to narrowly-defined groups of beneficiaries.¹⁶ Groups capitalize on the fact that the costs associated with these targeted transfers are spread out broadly over the population, minimizing individual costs and weakening opposition. In the standard framework provided by Lindbeck and Weibull (1987), political candidates maximize their probability of winning elections by making promises to various SIGs about redistribution policies. Politicians cater to groups that have the potential to swing elections in their favor, i.e. groups with ‘political clout’. Originally, this notion was defined as groups with beliefs or income levels near a voting threshold that can be swayed by promises of transfers. In recent work, ‘political clout’ has taken on a number of interpretations, including a group’s voter turnout rate (Glaeser, Ponzetto, and Shapiro 2005) or a group’s ability to understand complex policies, and thus, accurately vote on them (Strömberg 2004).

Public-sector employees fit well into this SIG politics framework and politicians are motivated to give them increased benefits due to unique elements of pension policies. As Section 2.2.2 makes clear, understanding the true burden and current

16. See Persson and Tabellini (2000) for a thorough discussion.

funding status for pension plans is incredibly complex. This great degree of complexity creates an information asymmetry between voting groups. Ordinary taxpayers are ill-equipped to calculate and fully understand the extent of agreements. Unions, on the other hand, elect leaders and hire experts that can decipher this knowledge for their group, giving them the information advantage discussed in Strömberg (2004).

While the information advantage is important, public-sector employees in many states also enjoy another advantage: institutional recognition and the right to negotiations with governments. Collective bargaining laws dictate the extent to which government employers are obligated, authorized, or prohibited from negotiating with public-employee unions. Legal rights to collectively bargain not only strengthen union representatives' hand in negotiations but also help unions grow.^{17, 18} As in any framework of SIG or pork-barrel politics, the strength of a group should increase the level of transfers that they are able to extract – a hypothesis we empirically examine below in the context of public-sector collective bargaining rights and pension plan benefits.

2.3 Data on Public Sector Collective Bargaining Laws

The legal environment governing public-sector collective bargaining has evolved unevenly across the United States. In our dataset, which we refer to as the Extended Public Sector Collective Bargaining Law Data Set (E-PSCBLD), we capture this varied evolution over three dimensions: time, space, and occupational group. Our main variable of interest measures the extent to which government employers are required

17. Previous literature has measured union strength with membership. Both measures are suitable for this purpose; however, they both also suffer from the issue of endogeneity to pension outcomes (an issue we will address below). We chose collective bargaining laws because they offer a more comparable measure across regions and because policy recommendations from research on legal environment are much clearer than those from research using union size.

18. We assume union representatives are neutral between fully-funded and under-funded benefit increases in negotiations – an assumption we discuss further in Section 2.A.

to negotiate in good faith with registered public-sector employee unions. We track the changing legal environment in this way from 1955 to 2018, encompassing the entire period during which the movement for public-sector employee union rights emerged and expanded.

The foundation for the E-PSCBLD comes from the NBER Public Sector Collective Bargaining Law Data Set (PSCBLD) (R. Freeman and R. Valletta 1988). The PSCBLD covers a dozen aspects of public-sector union laws, and varies by year, by state, and by five main public-sector employee groups: state employees, municipal police, municipal firefighters, non-college teachers, and other municipal employees. The laws covered in the PSCBLD fall broadly into five categories illustrating the rights afforded to public-sector unions, the central of which is ‘contract negotiation provisions’. This category contains our primary measurement of the legal environment, collective bargaining rights, which records the baseline ability of unions to bargain with government employers over wages and pensions. In some states, collective bargaining is prohibited outright, while for many states and occupational groups, collective bargaining is allowed and employers have an obligation to negotiate in good faith with union representatives.¹⁹

The main drawback of the PSCBLD is that it only covers the period 1955 to 1985, leaving it largely inadequate for analyzing public-sector pension outcomes through the present day.²⁰ To address this, we extend the main series on collective bargaining rights through 2018 for all states and occupational groups. In doing so, we rely on records of state statutes, court records, and state attorney general reports. In addition to extending the data, we construct a new measure detailing the source of the law changes (i.e. by legislative action, by court case decision, or by attorney

19. See Section 2.B additional information on the variables contained in the PSCBLD.

20. Since its publication, there has only been a modest update extending some aspects to 1996 by Kim Reuben. This update appears to have been subsumed into R. Freeman and R. Valletta (1988) on the NBER website, without a separate citation.

general opinion) at each instance throughout the entirety of the dataset. The process of extending the data as well full documentation of state-group-specific legal changes in collective bargaining rights are detailed in Section 2.C.

2.3.1 Variation in E-PSCBLD

Until the 1950s, almost no public-sector employees were unionized and the president of the American Federation of Labor and Congress of Industrial Organizations (AFL-CIO) stated at the time that it was “impossible to bargain collectively with the government” (Freeman 1986). This changed in 1962 when President Kennedy’s Executive Order 10988 recognized the right of federal workers to do so. In the period that followed, public-sector unionization expanded rapidly across the country; however, this expansion hinged critically on state laws paving the way and there has been significant variation in this across states and over time (Freeman 1986; Reder 1988; Freeman and Han 2012).

Figure 2.1 illustrates the geographic variation in the legal environment by aggregating across groups and across states within regions of the United States. In general, all regions moved towards increased bargaining rights over the years 1955-2018; however, states varied in the speed of law changes and in their final ‘plateau’ level. The Northeast was an early leader in expanding collective bargaining rights and continues to have the highest levels. Meanwhile, states in the South were slow to adopt collective bargaining rights for public-sector employees and continue to have lower levels than the rest of the country. The West and Midwest regions both had moderate rates of legal change and reached a more middle-ground of present-day collective bargaining rights. Although the legal environment in all regions has been relatively steady since the mid-1980s, there has been a recent push-back against labor unions in some states and this is captured in our extended dataset (see e.g. Wisconsin’s push for ‘Right-to-Work’ laws in 2010).

Although some states passed broad comprehensive laws changing collective action provisions for all public-sector occupational groups, the legal environment generally shows significant variation across occupation groups. Figure 2.2 illustrates this variation across groups using a model that absorbs state, year, and group-year fixed effects (group-year fixed effects displayed). Non-college teachers and municipal firefighters were groups that received expanded collective bargaining rights earlier than others, while state employees, municipal police, and other municipal workers were slow to receive expanded rights. These features of the data fit well-established narratives within the collective bargaining history. First, large and well-defined occupational groups such as teachers and firefighters are often framed as being the first that were able to coordinate an organized fight for union rights and likely held large political sway because of their size. Meanwhile, state employees and ‘other municipal workers’ are smaller and more loosely-defined groups that would have found it more difficult to organize and gather political power. Second, a common concern when expanding public-sector union rights has always been the interruption of vital public services. Due to this concern, police officers, serving perhaps the most vital of public employee roles, were frequently left out of early legislation. In our analysis, we will leverage the state, time, and group dimensions of variation in the legal environment to study their relation to the funding status of public-sector pension plan.

Finally, in the E-PSCBLD, we also document variation in the source of the legal changes between politically-driven and judicially-driven. Politically-driven changes encompasses legislative action and gubernatorial executive orders, both of which are likely to reflect changes in public opinion. Judicially-driven changes, which we define as court decisions and attorney general opinions, provide adjustments to the interpretation of statutes already in place that do not explicitly mention a specific public employee group. As a result, changes to collective bargaining rights resulting from court decisions and attorney general opinions are unlikely to be the result of shifts in

public opinion and we exploit this to address issues of endogeneity in our analysis. Of the 256 documented changes in collective bargaining rights, 182 are the result of changes in statutes, 33 are the result of a court decision, 32 are the result of an attorney general opinion, and 9 are the result of an executive order (by a governor).

2.4 Data on Public-Sector Pension Plans

There are two sources of financial information on public-sector pension plans. The first is the Center for Retirement Research at Boston College’s Public Plans Database (PPD). Based on pension plans’ official annual reports, it contains the headline measure of funding status, the AFR, and observes plans back to 2001. The second data source is a combination of the Annual Survey of Public Pensions (ASPP) and the Historical Database on Public Employee-Retirement Systems (PERS), which are both based on self-reporting by pension plans to the U.S. Census. In the ASPP-PERS, we can observe plans over a much longer historical period – back to 1957 – however, it does not report on funding status directly. Instead, it contains many of the factors that determine funding levels, including annual benefits paid, employee contributions, employer contributions, investment returns, and membership information. By mapping plans in both the PPD and the ASPP-PERS to their respective occupational group classification, we are able to link them to the state-, group-, and year-specific legal environment in which they operated, as coded in the E-PSCBLD.

2.4.1 Public Plans Data (PPD): 2001 – 2017

The PPD covers 177 of the largest public-sector pension plans in the United States: 87 state-run plans and 90 local-level plans.²¹ It includes over 100 detailed measures

21. PPD documentation claims that the plans in the database cover 95 percent of state/local pension assets and members in the US.

of pension plan financial status on an annual basis, including realized rates of return on assets, the Actuarially Assumed Return (AAR), and the Actuarial Funding Ratio (AFR) – the headline measure used by policymakers to understand a plan’s funding status.

Table 2.1 reports some descriptive financial statistics for pension plans in the PPD. Column 1 relates a plan’s funding ratio only to the actual returns on investments and the AAR, conditional on plan fixed effects. Both coefficients have the expected sign: actual returns improve a plan’s asset position, while a higher AAR means future obligations are more steeply discounted. Column 2 shows how funding status of plans evolved over time using a regression with only plan and year fixed effects. The associated yearly means (column 4) of funding status show that there is an almost monotonic decline over the 2001–2017 period: the average AFR in the Public Plans Data goes from 95% in 2001 to 73% in 2017. In column 3 we ask whether most of this decline is explained by variation in actual and projected returns. The answer is ‘no’. While both are significant, their contribution to the R-squared is easily trumped by the year fixed effects. The most likely explanation then for the patterns in columns 2–3 is that plans’ funding status have been eroding from 2001–2017 because of the essentially unstoppable object that is unfunded pension obligations moving closer to ‘maturity’ as the baby boomers are beginning to enter retirement.²²

2.4.2 ASPP-PERS Data: 1957 – 2017

The Annual Survey of Public Pensions (ASPP), and its predecessor, the Historical Database on Public Employee-Retirement Systems (PERS), provide an annual survey of state- and locally-administered pension plans, recording revenues, expenditures,

22. The baby boomers are commonly defined as the birth cohorts from 1945–1964. While some members of the baby boomers’ very first birth cohort (i.e. 1945) could have entered ‘early retirement’ (aged 55) as early as 2001, the majority of public employees still retire at age 65, and we are yet to hit the peak of baby-boomer retirement.

financial assets, and membership information from 1957 to 2017.²³ Pension programs included in the ASPP-PERS must be defined benefit plans that are sponsored by a unit of government recognized by the U.S. Census Bureau²⁴ and have a membership comprised of public employees compensated with public funds. For plans that meet these specifications, the ASPP-PERS provides a full census every five years (years ending in ‘2’ and ‘7’) and covers a subsample in the intervening years.²⁵ In total, ASPP-PERS contains over 3,000 unique public-sector pension plans, of which roughly 400 are state-administered and the rest are locally administered.

A couple of notable trends are apparent in the data. Beginning with membership information, we can illustrate the rise of public-sector pension programs. To address the inconsistent sampling of plans from year to year, we regress total active plan members (employees) and total benefit recipients (retirees) on plan and year fixed effects. The year fixed effects from this specification are shown in Figure 2.3. Measuring by both current employees and beneficiaries, pension plans have grown in size since the beginning of our sample. Current employee levels showed faster growth in the beginning of the sample, with a slower rate starting in the late-1970s. Beneficiaries, on the other hand, have continued to grow rapidly, and shows no sign of slowing to a plateau level as with employees. This difference likely reflects two factors: the growth in the size of local and state governments since 1960 as populations have grown, and the simultaneous increase in life expectancy. Both factors would lead to a faster-growing recipient base than employee base.²⁶

23. Survey conducted annually 1974-2017. Before 1972, it was every 5 years: 1957, 1962, 1967, and 1972.

24. This includes state governments as well as five types of local governments: county, municipal, township, school district, and special district.

25. Starting in 2004, a new sample is selected every five years. Prior to 2004, there was a non-probability sample.

26. Note that we do not say anything about relative members to beneficiaries. Current employees are not paying off current beneficiaries, so such a comparison would offer little insight into the

As for plan financial information, the ASPP-PERS data does not provide the headline measures of plan funding status (AFR); however, it does provide detailed information on financial flows that directly influence the AFR: benefits paid and contributions received. This distinction can be thought of as funding level stock (AFR) vs. flows that enter this stock (contributions and benefits). As shown in equation (2.1), there exists a theoretical mapping between benefits and contributions to the AFR; however, we are not privy to the inside accounting information (expected future benefit commitments, AAR, etc.) necessary to calculate AFR from the ASPP-PERS data. From expression (2.1), we know that AFR should be increasing in benefits and decreasing in contributions. It is this relationship that we will take advantage of in the historical panel by assuming under-funding to be associated with more generous benefits per recipient relative to contribution per employee.

As expected, financial data in the ASPP-PERS also illustrate an expansion of generosity in public-sector pension plans. Figure 2.4 shows the change in contributions per member and benefits per recipient from their 1957 levels over the sample period.²⁷ Both have increased consistently, at first in lockstep through the early 1990s, and then diverging, with benefits growing faster than contributions. Notably, employee contributions picked up around 2008-2009 when poor investment returns brought unfunded pensions to the forefront of national attention. In Figure 2.4, we also plot the ratio of these two components: a measure we call plan generosity. More specifically, we regress the ratio of benefits per retiree to contributions per member on year and plan fixed effects. We plot the associated year fixed effects and 95% confidence intervals. As shown, plan generosity was relatively steady through the

funding status captured by the AFR.

27. In this section, we focus on employee contributions because our data only shows actual payments made or received and, as noted in Section 2.2, employee contributions are mandatory while employer contributions reflect both their portion of the ARC and their level of shirking this funding requirement.

1980s, but then significantly picked up in the mid-1990s.

2.5 Empirical Strategy

Our empirical strategy starts with understanding the relationship between collective bargaining rights (CBR) and pension plan funding status in the PPD sample period (2001-2017). As noted, the headline measure of plan funding status – the Actuarial Funding Ratio (AFR) – is available during this period, but not in the period covered only by the ASPP-PERS data (1957-2000). In this 21st century period, we examine the association between collective bargaining rights laws and pension plan funding using the following pooled cross-sectional specification:

$$AFR_{it} = \alpha + \beta \mathbb{1}[\text{Collective Bargaining Rights}_{sgt}] + \mu_{s/g/t} + \epsilon_{it}, \quad (2.2)$$

where AFR is the Actuarial Funding Ratio of plan i in year t . The coefficient of interest, β captures the association between AFR and the state-, group-, and year-specific collective bargaining rights pertaining to plan i . We represent CBR with an indicator variable for whether the government employer has a duty to bargain with the plan's public employee union. In various regressions, we include state, group, and year fixed effects to consider different sources of variation in CBR. We cluster standard errors by pension plan.

To exploit the full extent of the state-, occupational group-, and year-specific variation in collective bargaining rights, we utilize the rich historical data available in the ASPP-PERS data. As noted, the ASPP-PERS data do not include headline measures of plan funding status (AFR), so we instead use our measure of plan generosity: the ratio of benefit payments per retiree to contributions per employee. We examine the relationship between collective bargaining rights and plan generosity using the

following Difference-in-Differences specification:

$$\log \left(\frac{\text{Benefits per Retiree}}{\text{Employee Contributions per Member}} \right)_{it} = \alpha + \beta \mathbf{1}[\text{CBR}]_{sgt} + \mu_t + \mu_i + \epsilon_{it}, \quad (2.3)$$

where the dependent variable is our measure of plan generosity and CBR takes the same meaning as in Equation 2.2. We include plan and year fixed effects and cluster standard errors by pension plan. The identifying assumption is that, in the absence of changes in the CBR, pension generosity would be similar between plans.

One concern when taking a causal interpretation of the coefficient of interest is that collective bargaining rights and public-sector pension plan generosity may both be correlated with local sentiment towards public employees. To probe this, we take advantage of the source of changes in collective bargaining rights. As noted in Section 2.3, changes in CBR can be politically or judicially driven. It would naturally follow that politically driven changes should have a higher correlation with local sentiment toward public employees. Thus, by interacting our indicator for CBR with the source of the change leading to that state-, group-, and year-specific status, we can assess the validity of this concern.

2.6 Results

We first document the robust relationship between public-sector employee collective bargaining rights and the headline measures of pension plan funding status – the Actuarial Funding Ratio (AFR) – in the 21st century. Results from the estimation of this pooled cross-sectional analysis (Equation 2.2) are reported in Table 2.2. As shown, in the 2001-2017 period, pension plans in states and for employee occupational groups with full collective bargaining rights (the government has a duty to bargain) have funding ratios that are 2-7 percentage points lower than plans for employees without these bargaining rights. Results in columns 1 through 4 explore this relationship in

more detail by taking advantage of different types of variation in CBR. Column 1 provides our baseline analysis with no fixed effects. The remaining columns include year fixed effects to control for macroeconomic developments that may impact all pension plans similarly from year to year. In columns 3 and 4, results report this association when exploiting variation within states across groups and within groups across states, respectively. Results are qualitatively consistent throughout.

Next, we utilize the historical sample provided by the ASPP-PERS data in a Difference-in-Difference analysis that exploits the full extent of the state, occupational group, and time variation in collective bargaining rights to understand how it impacts pension plan generosity. The results of the estimation of Equation 2.3 are shown in column 1 of Table 2.3. Consistent with our hypothesis, collective bargaining rights for public-sector unions increase pension plan generosity by about 5 percent. In column 2 of Table 2.3, we break this effect down by decade. Interestingly, collective bargaining rights increased plan generosity by up to 20 percent in the pre-1990 period; however, this generosity reversed in the post-1990 period. One potential explanation for this pattern is that CBR allowed unions to extract higher benefits initially, but then, as the financial status eroded, these plans had to restrict generosity to a higher extent.

As noted in Section 2.5, a threat to the identifying assumption in this framework is that collective bargaining rights and plan generosity could both be correlated with local sentiment towards public employees. In columns 3 and 4 of Table 2.3, we augment the DID analysis by interacting CBR with an indicator for whether the previous change in the legal status was made due to a court ruling or attorney general opinion, i.e., ‘judicially driven’. As shown in column 3, the estimated effect of CBR from each source are not statistically different. If anything, results appear to be stronger for judicially-driven changes in CBR, which are plausibly more exogenous to local sentiment towards public-sector employees. Column 4 repeats the breakdown of this effect by decade for CBR from each source. Results are qualitatively similar.

Given that employee unions granted collective bargaining rights are able to extract more generous retirement benefits, we next examine potential ways in which this could result in under-funded liabilities.²⁸ To do so, we estimate two variations of the model in Equation 2.2. First, we examine the pooled cross-sectional relationship between collective bargaining rights and the percentage of required government contributions that were actually paid in the post-2000 period. The estimated coefficients from this regression are reported in Table 2.4. As shown, there is some evidence that government employers who have a duty to bargain with public-sector employee unions paid a lower share of their expected contributions (roughly 4 percentage points less), although results vary and are insignificant when accounting for state or occupational group fixed effects.

In Table 2.5, we report results examining another potential mechanism by which plan generosity could lead to under-funding: government over-optimism on investment returns. These estimates reflect another variation of the model in Equation 2.2 where the dependent variable is now the Actuarially Assumed Return (AAR) in pension plan investments. As reflected in the table, collective bargaining rights has a positive and significant relationship with AAR. A higher AAR lowers both mandatory employee contributions and expected employer contributions but does not change the plan liabilities that eventually need to be paid out. In columns 5-8, we consider the possibility that AAR may reflect plan-specific performance by controlling for realized 5-year investment returns. The estimates are unchanged. Although descriptive in nature, the results in Table 2.4 and Table 2.5 suggest that governments, having increased pension generosity, could have avoided the associated costs in the short-term by shirking their required contributions or by being overly optimistic about investment returns – both of which would result in a rise of unfunded liabilities in the

28. As discussed in Section 2.2, actuarial accountants ensure that benefit increases are fully funded at the time of agreements.

medium- and long-term.

2.7 Conclusion

In this chapter, I have shown that collective bargaining rights for public-sector employee unions significantly increased pension plan generosity in the 20th century, and is associated with higher levels of unfunded liabilities in the 21st century. I have done so by exploiting variation in collective bargaining rights by state, year, and occupational group in a Difference-in-Differences framework. To facilitate this analysis, I extended an existing dataset tracking collective bargaining rights in the United States to cover the period 1955 to 2018²⁹ and documented the source of changes in the legal environment to aid in identification. Using more detailed financial data in the post-2000 period and a pooled cross-sectional analysis, I have provided descriptive evidence suggesting that increased plan generosity resulted in higher levels of unfunded liabilities because governments shirked their required contributions and made overly optimistic assumptions on investment returns.

This research empirically examines how institutional shifters of bargaining power for specific groups in the electorate can meaningfully affect policy outcomes in settings well-characterized by models of special interest group politics. These shifts lead to higher rents for the designated group and, in the case of public-sector pensions, unsustainable outcomes when information asymmetries exist. This phenomenon is further exaggerated when governments can obscure true costs from general voters in the short-term. When the associated costs are fully realized, however, they can have important implications for the local economy – a topic explored in Chapter 3.

29. Data previously only included information through 1986.

Table 2.1: Funding Status of Pension Plans: 2001–2017

Outcome:	(1) % Funded	(2) % Funded	(3) % Funded	(4) Yearly Mean
Return (5-yr)	0.4508*** (0.0584)		0.5870*** (0.0976)	
AAR	16.3039*** (2.4150)		4.6681** (1.8064)	
Year = 2002		-0.0744*** (0.0082)	-0.0454*** (0.0089)	0.9476
Year = 2003		-0.1297*** (0.0112)	-0.0876*** (0.0123)	0.8923
Year = 2004		-0.1561*** (0.0131)	-0.1159*** (0.0136)	0.8659
Year = 2005		-0.1792*** (0.0158)	-0.1409*** (0.0161)	0.8428
Year = 2006		-0.1840*** (0.0169)	-0.1644*** (0.0162)	0.8380
Year = 2007		-0.1677*** (0.0178)	-0.1747*** (0.0172)	0.8543
Year = 2008		-0.1962*** (0.0177)	-0.1808*** (0.0172)	0.8258
Year = 2009		-0.2505*** (0.0177)	-0.2004*** (0.0185)	0.7715
Year = 2010		-0.2722*** (0.0183)	-0.2244*** (0.0187)	0.7498
Year = 2011		-0.2844*** (0.0195)	-0.2379*** (0.0197)	0.7376
Year = 2012		-0.3052*** (0.0199)	-0.2438*** (0.0209)	0.7168
Year = 2013		-0.3044*** (0.0209)	-0.2680*** (0.0205)	0.7176
Year = 2014		-0.2925*** (0.0223)	-0.2878*** (0.0216)	0.7295
Year = 2015		-0.2875*** (0.0227)	-0.2667*** (0.0222)	0.7345
Year = 2016		-0.2973*** (0.0230)	-0.2561*** (0.0232)	0.7247
Year = 2017		-0.2912*** (0.0231)	-0.2547*** (0.0231)	0.7308
Constant	-0.5064*** (0.1900)	1.0220*** (0.0170)	0.5886*** (0.1451)	
Observations	2,650	2,650	2,650	
R-squared	0.7245	0.8324	0.8394	
# Plan FEs	177	177	177	

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Column 1 regresses plans' AFR on the actual return and the AAR, i.e. the assumed long term return that is used to discount obligations. Column 2 regresses the AFR on year fixed effects. Column 3 regresses plans' AFR on actual returns, AAR, and year fixed effects. Plan fixed effects are included in all regressions. Column 4's 'yearly mean' reports the AFR yearly mean for the model that is saturated in year and plan fixed effects, i.e. columns 2, deducting the year fixed effects from the constant, e.g. $0.9476 = 1.0220 - 0.0744$ in column 1 in 2003. Standard errors (clustered by plan) are reported in parentheses.

SOURCE: Public Plans Data. 2001-2022. Center for Retirement Research at Boston College, MissionSquare Research Institute, National Association of State Retirement Administrators, and the Government Finance Officers Association.

Table 2.2: Collective Bargaining Rights and Actuarial Funding Ratio

	(1)	(2)	(3)	(4)
	Actuarial Funding Ratio	Actuarial Funding Ratio	Actuarial Funding Ratio	Actuarial Funding Ratio
1[CBR]	-0.0157** (0.0068)	-0.0150** (0.0062)	-0.0653*** (0.0123)	-0.0170*** (0.0062)
Observations	2,961	2,961	2,961	2,961
R-squared	0.0015	0.1731	0.4604	0.1929
Dep. Mean	0.7955	0.7955	0.7955	0.7955
Dep. SD	0.1947	0.1947	0.1947	0.1947
Year FE	-	✓	✓	✓
State FE	-	-	✓	-
Group FE	-	-	-	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Standard errors (clustered by plan) are reported in parentheses.

SOURCE: Public Plans Data. 2001-2017. Center for Retirement Research at Boston College, MissionSquare Research Institute, National Association of State Retirement Administrators, and the Government Finance Officers Association.

Table 2.3: Collective Bargaining Rights and Increased Plan Generosity

	Baseline		Heterogeneity by Legal Path	
	(1) Benefits/Contr.	(2) Benefits/Contr.	(3) Benefits/Contr.	(4) Benefits/Contr.
$\mathbf{1}[\text{CBR}]$	0.0528** (0.0257)		0.0462 (0.0370)	
× 1960		0.1213**		0.1005**
× 1970		0.2046***		0.1718***
× 1980		0.1280***		0.0892**
× 1990		-0.0174		-0.0687
× 2000		-0.0830*		-0.1176**
× 2010		-0.0723		-0.1110*
$\mathbf{1}[\text{CBR}] \times \mathbf{1}[\text{Judicial}]$			0.0587* (0.0337)	
× 1960				-
× 1970				-
× 1980				0.1246***
× 1990				0.0492
× 2000				-0.0681
× 2010				-0.0412
Observations	36,492	36,492	36,492	36,492
R-squared	0.6280	0.6320	0.6280	0.6323
Dep. Mean	1.9148	1.9148	1.9148	1.9148
Dep. SD	0.5984	0.5984	0.5984	0.5984
Plan FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Standard errors (clustered by plan) are reported in parentheses.

SOURCE: SOURCE: ASPP-PERS.

Table 2.4: Collective Bargaining Rights and Government Contributions

	(1)	(2)	(3)	(4)
	Percent of Required Government Contribution Paid			
1[CBR]	-0.0381* (0.0203)	-0.0375* (0.0205)	0.0099 (0.0248)	-0.0339 (0.0218)
Observations	2,882	2,882	2,882	2,882
R-squared	0.0014	0.0178	0.0880	0.0202
Dep. Mean	0.9437	0.9437	0.9437	0.9437
Dep. SD	0.4919	0.4919	0.4919	0.4919
Year FE	-	✓	✓	✓
State FE	-	-	✓	-
Group FE	-	-	-	✓

*** p< 0.01, ** p< 0.05, * p< 0.1.

NOTE: Standard errors (clustered by plan) are reported in parentheses.

SOURCE: Public Plans Data. 2001-2017. Center for Retirement Research at Boston College, MissionSquare Research Institute, National Association of State Retirement Administrators, and the Government Finance Officers Association.

Table 2.5: Collective Bargaining Rights and Government Optimism


	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Actuarially Assumed Return (AAR) on Investments							
1[CBR]	0.0009*	0.0009***	0.0012***	0.0008***	0.0009***	0.0009***	0.0012***	0.0008***
	[0.0005]	[0.0001]	[0.0002]	[0.0001]	[0.0001]	[0.0001]	[0.0002]	[0.0001]
Investment Return (5-year)					-0.0097	0.0053	0.0077	0.0048
					[0.0090]	[0.0067]	[0.0073]	[0.0066]
Observations	2,650	2,650	2,650	2,650	2,650	2,650	2,650	2,650
R-squared	0.0104	0.1807	0.3818	0.1916	0.0177	0.1813	0.3829	0.1921
Dep. Mean	0.0780	0.0780	0.0780	0.0780	0.0780	0.0780	0.0780	0.0780
Dep. SD	0.0043	0.0043	0.0043	0.0043	0.0043	0.0043	0.0043	0.0043
Year FE	-	✓	✓	✓	-	✓	✓	✓
State FE	-	-	✓	-	-	-	✓	-
Group FE	-	-	-	✓	-	-	-	✓

*** p< 0.01, ** p< 0.05, * p< 0.1.

NOTE: Standard errors (clustered by plan) are reported in parentheses.


SOURCE: Public Plans Data. 2001-2017. Center for Retirement Research at Boston College, MissionSquare Research Institute, National Association of State Retirement Administrators, and the Government Finance Officers Association.

Figure 2.1: Historical Variation of Collective Bargaining Laws by Region




NOTE: This figure plots the raw data average of collective bargaining rights by four U.S. regions: Northeast, Midwest, South, West. Some states have rolled back union rights in the post-1985 period. Collective bargaining rights is divided by: prohibited (0), authorized (1), and required (2). Although most went after secondary elements of collective bargaining (e.g. union shop and strike provisions), some rolled back collective bargaining rights entirely. These include Indiana State Employees (2005), Kentucky State Employees (2003, 2015), Oklahoma Other Municipal Employees (2011), and Texas Non-College Teachers (1993).

Figure 2.2: Historical Variation of Collective Bargaining Laws by Occupational Group




NOTE: This figure plots estimated 'group-year' fixed effects from a regression of collective bargaining rights on year fixed effects, state fixed effects, and group-year fixed effects.

Figure 2.3: Trends in Pension Plan Membership



NOTE: This figure plots the year fixed effects and 95% confidence intervals from a regression of log total active plan members (employees) and log total pension recipients (retirees) on year fixed effects and plan fixed effects. Errors clustered by plan.
SOURCE: ASPP-PERS.

Figure 2.4: Trends in Pension Plan Generosity



NOTE: Each line plots the year fixed effects and 95% confidence intervals from separate regressions of the respective outcome (in logs of 2012 USD) on year fixed effects and plan fixed effects. Errors clustered by plan.
SOURCE: ASPP-PERS.

APPENDICES

2.A How Valuable are Pension Promises When They are Under-Funded?

This section provides an answer to the question of why both public-sector unions and politicians may favor unfunded pension benefit expansions.

The determination of pension benefits and contributions may be best characterized as a bargaining process between a politician and a public-sector union representative, in which the politician maximizes votes from core supporters (union members) and other voters, while the union representative can earn rents from union members for generating higher benefits, and from the politician for mobilizing political support. The politician can promise pension benefits to secure the political support of their core supporters. In practice, the blind spots in the actuarial calculations combined with the misleading budget neutrality of letting actual employer contributions fall behind actuarially required levels (both discussed in Section 2.2) make pension promises a ‘shrouded’ benefit from the politician’s point of view. They can bring out their core supporters while keeping a balanced budget in the eyes of other voters.³⁰

If the ability to under-fund pensions is key to the ‘shroudedness’ of pension benefits, it also raises the question how union representatives and union members discount under-funded pension benefits relative to fully funded ones in practice. It is possible that there is no discount at all because under-funded benefits are still legally binding commitments. One caveat to that view is that even if all obligations end up being paid in full, many union members may belong to the tax base from which they are

30. As noted in the main text, many of these features are incorporated in the theory in Glaeser and Ponzetto (2014) and Persson and Tabellini (2000).

paid.³¹ However, this should be equally true of funded benefits. If, as suggested by Inman 1982, retired pensioners are more likely to move out of the tax base, then under-funding may actually be preferred.

There is also a separate question of how union members view biased actuarial assumptions (such as an over-optimistic AAR), when these are likely to be eventually adjusted and lead to future increases in employee-paid actuarially required contributions. It seems probable that this scenario is not salient enough to impact the average union member's views of their benefits, although it is likely to be very salient to the union representatives on pension boards. The narrative evidence of union representatives pushing for and defending unrealistically high AARs supports this characterization (Greenhut 2009; Anzia and Moe 2016).

2.B Additional Background on the Public Sector Collective Bargaining Law Data Set

The PSCBLD contains a variety of legal measures to capture the legal standing towards public-sector collective bargaining. The main category is contract negotiation provisions, which includes collective bargaining rights. The various methodologies used to code CBR include that of R. Freeman and R. Valletta (1988) and Kim Reuben (subsumed into R. Freeman and R. Valletta (1988)). Table 2.6 reports these codings as well as our own binary measure of CBR (Dippel-Sauers). The other categories of legal measures included in the PSCBLD include union recognition provisions, union security provisions, impasse procedures, and strike policy. Although these may all play a role in shaping public employee bargaining power, we limit our analysis to the core measure of collective bargaining rights.

31. They may also be homeowners, and unfunded pension obligations may be capitalized into house prices. See Chapter 3 of this dissertation for empirical evidence, as well as Daly (1969), Glaeser and Ponzetto (2014), and Brinkman, Coen-Pirani, and Sieg (2018).

Table 2.6: Measuring Public-Sector Collective Bargaining Rights

Variable	Key
Collective Bargaining Rights (Freeman-Valletta)	0 = Collective bargaining prohibited 1 = No provision 2 = Employer authorized but not required to bargain with union 3 = Right to present proposals 4 = Right to meet and confer 5 = Duty to bargain I (implied) 6 = Duty to bargain II (explicit)
Condensed Collective Bargaining Rights (Reuben)	0 = No provision or collective bargaining prohibited (i.e. 0 or 1 from F-V) 1 = Employer authorized but not required to bargain with union or right to present proposals or right to meet and confer (i.e. 2, 3, or 4 from F-V) 2 = Duty to bargain I (implied) or duty to bargain II (explicit) (i.e. 5 or 6 from F-V)
Collective Bargaining Rights (Dippel-Sauers)	0 = No duty to bargain (i.e. 0-4 from F-V) 2 = Duty to bargain (implied or explicit) (i.e. 5 or 6 from F-V)

NOTE: Collective bargaining rights variable codings include those from R. Freeman and R. Valletta (1988), Kim Reuben, and Dippel-Sauers (this work).

2.C Extending the PSCBLD

To build our Extended Public Sector Collective Bargaining Law Data Set (E-PSCBLD), we take a two-step approach. The first step is to extend the collective bargaining rights measure in R. Freeman and R. Valletta (1988) to 2018. To do so, we compare the last observed record of the collective bargaining rights measure in the PSCBLD with the current version of the law for each state and occupational group. To aid us in identifying the present-day status of the legal environment, we reference a cross-sectional snapshot of present-day bargaining rights provided by Sanes, Schmitt, et al. (2014), along with state statutes, court records, and state attorney general reports.

In addition to extending our collective bargaining rights measure to 2018, we also identify the source of the documented legal environment change. Specifically, we distinguish changes brought about by political sources (statute changes and gubernatorial executive orders) and judicial sources (court rulings and attorney general opinions). To identify the source of legal changes, we rely heavily on the U.S. Department of Labor Summary of Public Sector Labor Relations Policies (DOL reports), which cover collective bargaining rights across the 50 states for the same 5 occupational groups, but only provide snapshots in the years when a report was published. This includes 1971, 1973, 1975, 1976, 1979, and 1981. To confirm the DOL report entries, we located as many of the source documents as possible. Court cases and modern-day statutes were accessed from Nexis Uni, while historical texts were examined in the UCLA law library.

Detailed case-by-case documentation on the history and source of law changes for the entirety of the E-PSCBLD can be found on www.zacharysauers.com or by contacting the author. For each change in the legal environment, we provide the date of the change, the cause of the change (statute change, court case decision, executive order, or attorney general opinion), the relevant text from the change (if available), and the Condensed Collective Bargaining Rights codification associated with that change (see Table 2.6). We also provide a concise summary of the code history for each subgroup. In some cases, we were unable to view the exact text of statutes before and after a prospective change in the collective bargaining status. For these instances, which are all changes early in the dataset, we defer to the DOL reports in accounting for the source of code changes.

References

- Adsera, Alicia, Carles Boix, and Mark Payne. 2003. "Are you being served? Political accountability and quality of government." The Journal of Law, Economics, and Organization 19 (2): 445–490.
- Anderson, Michelle Wilde. 2013. "The new minimal cities." Yale Law Review 123:1118.
- Anzia, Sarah F, and Terry M Moe. 2016. "Interest Groups on the Inside: The Governance of Public Pension Funds." Unpublished draft (August).
- Ashenfelter, Orley. 1978. "Union relative wage effects: new evidence and a survey of their implications for wage inflation." In Econometric contributions to public policy, 31–63. Springer.
- Besley, Timothy, and Robin Burgess. 2002. "The political economy of government responsiveness: Theory and evidence from India." The quarterly journal of economics 117 (4): 1415–1451.
- Bouton, Laurent, Paola Conconi, Francisco Pino, and Maurizio Zanardi. 2021. "The tyranny of the single-minded: Guns, environment, and abortion." Review of economics and statistics 103 (1): 48–59.
- Boylan, Richard T. 2016. "The Impact of the Voting Rights Act on Municipal Reform and Pensions."
- Brinkman, Jeffrey, Daniele Coen-Pirani, and Holger Sieg. 2018. "The political economy of municipal pension funding." American Economic Journal: Macroeconomics 10 (3): 215–46.

- Brown, Jeffrey R, and F Dye Richard. 2015. "Illinois Pensions in a Fiscal Context: A (Basket) Case Study."
National Bureau of Economic Research Working Paper 21293.
- Brown, Jeffrey R, and David W Wilcox. 2009. "Discounting state and local pension liabilities." American Economic Review Papers & Proceedings 99 (2): 538–42.
- Burns, Debra Brubaker. 2011. "Too big to fail and too big to pay: states, their public-pension bills, and the constitution." Hastings Const. LQ 39:253.
- Daly, George G. 1969. "The burden of the debt and future generations in local finance." Southern Economic Journal, 44–51.
- Dippel, Christian. 2022. "Political parties do matter in us cities... for their unfunded pensions." American Economic Journal: Economic Policy 14 (3): 33–54.
- Ferraz, Claudio, and Frederico Finan. 2008. "Exposing corrupt politicians: the effects of Brazil's publicly released audits on electoral outcomes."
The Quarterly journal of economics 123 (2): 703–745.
- Freeman, Richard, and Robert Valletta. 1988. "Appendix B: The NBER Public Sector Collective Bargaining Law Data Set."
When Public Sector Workers Unionize.
- Freeman, Richard B. 1986. "Unionism Comes to the Public Sector."
Journal of Economic Literature 24 (1): 41–86.
- Freeman, Richard B, and Eunice Han. 2012. "The war against public sector collective bargaining in the US." Journal of Industrial Relations 54 (3): 386–408.

- Freeman, Richard B, and Robert G Valletta. 1987. "The Effect of Public Sector Labor Laws on Collective Bargaining, Wages, and Employment." National Bureau of Economic Research Working Paper.
- Glaeser, Edward L, and Giacomo AM Ponzetto. 2014. "Shrouded costs of government: the political economy of state and local public pensions." Journal of Public Economics 116:89–105.
- Glaeser, Edward L, Giacomo AM Ponzetto, and Jesse M Shapiro. 2005. "Strategic extremism: Why Republicans and Democrats divide on religious values." The Quarterly Journal of Economics 120 (4): 1283–1330.
- Greenhut, Steven. 2009. "Plunder!: How Public Employee Unions are Raiding Treasuries Controlling Our Lives and Bankrupting the Nation."
- Grosskopf, Shawna, Kathy Hayes, and Thomas Kennedy. 1985. "Supply and Demand Effects of Underfunding of Pensions on Public Employee Wages." Southern Economic Journal, 745–753.
- Grossman, Gene, and Elhanan Helpman. 1994. "Protection for Sale." The American Economic Review 84 (4): 833–850.
- Hirsch, Werner Z, and Anthony M Rufolo. 1982. "Determinants of municipal wages: Some tests of the competitive wage hypothesis." Research in Urban Economics 2:309–327.
- Inman, Robert P. 1982. "Public employee pensions and the local labor budget." Journal of Public Economics 19 (1): 49–71.
- . 1985. "The funding status of teacher pensions: An econometric approach." NBER Working Paper 1727.

- Inman, Robert P, and David J Albright. 1987. "Central policies for local debt: The case of teacher pensions." NBER Working Paper 2166.
- Kelley, Dashle G. 2014. "The political economy of unfunded public pension liabilities." Public Choice 158 (1-2): 21–38.
- Lindbeck, Assar, and Jörgen W Weibull. 1987. "Balanced-budget redistribution as the outcome of political competition." Public choice 52 (3): 273–297.
- Malanga, Steve. 2016. "Covering Up the Pension Crisis. States and Actuaries Are Trying to Stifle Debate About the Growing Shortfall in Fund Assets." Wall Street Journal.
- Martin, Timothy W., Georgi Kantchev, and Kosaku Nairioka. November 13th 2016. "Era of Low Interest Rates Hammers Millions of Pensions Around the World." Wall Street Journal.
- Mitchell, Olivia S, and Robert S Smith. 1994. "Pension Funding in the Public Sector." The Review of Economics and Statistics.
- Mohan, Nancy, and Ting Zhang. 2014. "An analysis of risk-taking behavior for public defined benefit pension plans." Journal of Banking & Finance 40:403–419.
- Munnell, Alicia H, Jean-Pierre Aubry, Mark Cafarelli, et al. 2015. "How did statelocal plans become underfunded." Center for Retirement Research at Boston College, no. 42.
- Munnell, Alicia H, Jean-Pierre Aubry, and Laura Quinby. 2011. "Public pension funding in practice." Journal of Pension Economics & Finance 10 (2): 247–268.

- Novy-Marx, Robert, and Joshua D Rauh. 2009. "The liabilities and risks of state-sponsored pension plans." Journal of Economic Perspectives 23 (4): 191–210.
- . 2011. "Public pension promises: how big are they and what are they worth?" The Journal of Finance 66 (4): 1211–1249.
- . 2014a. "Linking benefits to investment performance in US public pension systems." Journal of Public Economics 116:47–61.
- . 2014b. "The revenue demands of public employee pension promises." American Economic Journal: Economic Policy 6 (1): 193–229.
- Olson, Mancur. 2009. The logic of collective action. Vol. 124. Harvard University Press.
- Persson, Torsten, and Guido Tabellini. 2000. "Political economics: explaining economic policy. Zeuthen lecture book series." The MIT Press 1 (8): 9.
- Reder, Melvin W. 1988. "The rise and fall of unions: The public sector and the private." Journal of Economic Perspectives 2 (2): 89–110.
- Sanes, Milla, John Schmitt, et al. 2014. Regulation of public sector collective bargaining in the states. Center for Economic / Policy Research Washington, DC.
- Strömberg, David. 2004. "Radio's impact on public spending." The Quarterly Journal of Economics 119 (1): 189–221.
- The Economist. March 1st 2014. "Detroit's Bankruptcy: Cram Down."

Trusts, Pew Charitable. 2013. “The state role in local government financial distress.” The Pew Charitable Trusts Report.

Zax, Jeffrey, and Casey Ichniowski. 1988. “The effects of public sector unionism on pay, employment, department budgets, and municipal expenditures.” In When public sector workers unionize, 323–364. University of Chicago Press.

CHAPTER 3

Shocking Debt Illusion: Actuarial Pension Plan Adjustments and House Price Capitalization

3.1 Introduction

The Ricardian Equivalence Theorem suggests that economic outcomes are independent of whether or not a public program is financed by current taxes or by bond issues, i.e., future taxes (Barro 1974; Seater 1993). At the local level, this equivalence of tax and debt financing arises as long as residents and home buyers are ‘fiscally aware’, so that future taxes are priced into local property values today (Daly 1969; Akai 1994). Potential homeowners want a compensating differential for anticipated future taxes. Tiebout sorting, or ‘voting with your feet,’ reinforces this prediction (Tiebout 1956). Debt illusion refers to a failure of Ricardian equivalence, whereby citizens may prefer debt financing over tax financing because they are not ‘fiscally aware’ (in the words of Banzhaf and Oates (2013)).

Because declining house prices today are more salient than higher taxes tomorrow, many view the local version of Ricardian Equivalence as a more compelling empirical proposition than its national equivalent (Daly 1969; Dollery and Worthington 1996; Fishel 2001; Banzhaf and Oates 2013; Brinkman, Coen-Pirani, and Sieg 2018). In other words, debt illusion is viewed as less likely at the local level. This view, however, requires believing that future taxes (property or other local taxes) are indeed capitalized into house prices today, and there is not a lot of evidence to support this

belief.

In fact, debt illusion may very well exist at the local level because the biggest source of local debt is unfunded public-employee pension obligations, which are what Glaeser and Ponzetto (2014) call ‘shrouded,’ i.e., difficult to understand and quantify for local homeowners and voters.¹ The main reason for this shroudedness is that the calculation of public-sector pension plan finances hinges critically on actuarial assumptions. Perhaps the most important of these assumptions is the actuarially assumed return (AAR) on a plan’s investments. Whether actuarial assumptions are overly optimistic is difficult for homeowners and voters to assess, although this issue is widely recognized as a big source of cities’ future funding problems (Mitchell and Smith 1994; Greenhut 2009; Novy-Marx and Rauh 2011; Anzia and Moe 2019; Kelley 2014). Novy-Marx and Rauh (2009) estimate that the ‘true’ level of unfunded liabilities is about three times higher than the official level.

Debt illusion due to the shrouded-ness of pension under-funding does not imply that homeowners are irrational or short-sighted. Instead, it is more nefarious because even informed homeowners and voters find it difficult to gauge their city’s true exposure to pension under-funding if it diverges from the officially reported actuarial level. This discussion implies that the biggest shocks to residents’ information about public debt occur when municipal plans have to adjust their actuarial assumptions and, in particular, their AAR. The National Association of State Retirement Administrators estimates that cutting the AAR by 0.25 percentage points immediately increases the required contributions into a pension plan by two to three percentage points, so that “it is in no one’s interest to make more realistic assumptions about returns” (The Economist, October 5th 2017). This game of ‘pretend and extend’, however, cannot

1. Unfunded pension obligations are the biggest source of local/municipal debt in the U.S. (Novy-Marx and Rauh 2014). There are many factors that have led to this phenomenon. See, for example, Chapter 2 of this dissertation. Inman (1982) also note that “in the United States the use of the municipal bond market is restricted to capital outlays”...“the only way for local governments to finance current expenditures with future taxes is through pension under-funding.”

be sustained forever. When the AAR exceeds the actual returns for long enough, the actuarial assumptions do eventually have to be brought in line with reality, sometimes by state law or gubernatorial executive order (Anzia and Moe 2019).

This paper investigates the effect of shocks to municipalities' official public debt levels on local house prices, using as the primary source of variation the 294 occurrences of pension plan adjustments to their AAR that have occurred in U.S. cities since 2000. When a plan's AAR is adjusted downward, investment assets cover less of the anticipated future obligations to pensioners and, thus, the reported funding status of the plan drops. In the raw data on pension plan finances, we find that a one-percent decrease in the AAR is associated with a concurrent three-percent decrease in the official funding ratio of pension obligations, i.e., the share of discounted future obligations that can be met with current assets.

We link these 294 Δ AAR events to a monthly panel of house prices at the city-level to investigate the speed with which (and extent to which) their implied shocks to cities' official pension under-funding obligations were capitalized into house prices. Our paper is focused on two particular questions: First, we ask how quickly information about funding shocks diffuses into house prices. We distinguish for each event between three different dates at which updated financial information is revealed or increases in salience: the date of the plan's internal Board Meeting Decision (BMD), when they agree to adjust the AAR; the plan's Actuarial Report's date of publication (ARD), when the transmission from Δ AAR to a plan's funding status is first publicized; and the date of a city's Fiscal Year End (FYE), when a city's budget first accounts for Δ AAR and the resulting changes in the city's expected contributions. We find that there is a significant decline in house prices around the BMD and ARD but not around the FYE. This result, empirically supporting local Ricardian Equivalence, suggests that information diffuses rather quickly and is fully capitalized into house prices by the FYE, about 9 months from the BMD on average.

Second, we explicitly consider that cities may be quite constrained in their ability to raise taxes, in which case increased pension obligations start to crowd out public good provision in cities' budgets. This constraint is particularly important in states like California, where Proposition 13 constrains local property taxes (Brunner and Sonstelie 2006). Cutting public goods and services may then be the only option. This, in turn, is likely to affect families with children much more than residents without children, because it directly affects the funding for public schools. We therefore consider separately the house prices for single-family homes and condominiums. Consistent with the hypothesis, we find that the effect is mainly driven by changes in prices for single-family homes, while there is no significant change in condo prices.

In the remainder of this Chapter, Section 3.2 provides background information on public pension plan finances and adjustments in the AAR. Section 3.3 describes our data on pension plan finances and local housing prices. Section 3.4 outlines the empirical framework for the analysis, while Section 3.5 presents the results. Section 3.6 concludes the chapter.

3.2 Background on Public-Sector Union Bargaining

3.2.1 Pension Plan Finances 101: Actuarial Accounting

To assess the funding status of a public pension plan, policymakers rely on the Actuarial Funding Ratio (AFR), a measure that captures the share of a plan's discounted future benefit obligations to its pensioners and active members that is funded by current financial assets.² Although the actuarial accounting that goes into calculating a plan's AFR is more complicated, it can be broadly summarized (at time τ) by the

2. An equivalent measure that is often used is the Unfunded Actuarially Accrued Liabilities (UAAL), which measures the difference between a plan's assets and its discounted future benefit obligations to its pensioners and active members.

following expression

$$\text{AFR}_{i\tau} = \frac{\text{Assets}_{i\tau}}{\sum_{t>\tau}^{\infty} \frac{\text{Benefits}_{it}}{(1+\text{AAR}_i)^t}} , \quad (3.1)$$

where AAR_i is the Actuarially Assumed Return on a plan's assets. Note that both past and future contributions do not appear directly in Equation 3.1. Past contributions by employers and employees and any returns on them are subsumed into the plan's current asset base. Future contributions and obligations are assumed to directly offset each other (backed up by actuarial accountant calculations), and neither appears in Equation 3.1. In Defined Contribution (DC) plans, which dominate the private sector, asset returns directly determine benefit payments and Equation 3.1 always equals zero. In stark contrast, almost all municipal and state pension systems in the U.S. are Defined Benefit (DB) plans, where future benefit obligations are legally binding contracts independent of a plan's asset performance.

3.2.2 Adjustments to the AAR

Unrealistic assumptions of the Actuarially Assumed Return (AAR) can lead to pension plan under-funding by lowering the Actuarially Required Contributions (ARC), which is the level of contributions needed to cover future benefit obligations. This results from the fact that, in Equation 3.1, an increase in the AAR means future benefit obligations to pensioners are discounted further and the funding status appears more sound. In the short run, unrealistically high AARs do not substantially affect a plan's funding status because the AFR is calculated under the same erroneous assumptions; however, over time they lead to sustained growth in funding gaps. When corrections are made to the AAR, there can be significant jumps in the level of unfunded liabilities. For this reason, local politicians often oppose adjustments to actuarial

assumptions.³ State and national governments, on the other hand, fear that local governments will put adjustments off for too long and that they will have to step in to help when unfunded liabilities spike drastically. To avoid such a scenario, they have recently put pressure on local governments to bring the AAR closer in line with realized returns.⁴

Adjustments to actuarial assumptions are not determined by politicians, however, and often fall to pension boards where they are heavily debated choices (Anzia and Moe 2016; Greenhut 2009). When adjustments to the AAR are agreed upon by pension boards, there are three important dates capturing this. The first is the date of the plan's internal Board Meeting Decision (BMD), when they agree to adjust the AAR. The second is the plan's Actuarial Report's date of publication (ARD), when the transmission from Δ AAR into a plan's funding status (AFR) is first publicized. The final point is the date of a city's Fiscal Year End (FYE), when a city's budget first accounts for Δ AAR and the resulting changes in the city's expected contributions. Each of these dates offers an opportunity for the local electorate to learn of the growth of unfunded liabilities and the potential impact on other public good provision.

3.3 Data

To identify AAR adjustments, we build on the Center for Retirement Research at Boston College's Public Plans Data (PPD). The PPD aggregates over 100 detailed measures of pension plan financial information on an annual basis from 2001 to 2019, including the AAR and key measures of funding status. It includes over 100 detailed

3. Funding gaps, including those emerging from unrealistic actuarial assumptions, are tolerated and encouraged by government leaders because they are politically expedient – a topic discussed in further detail in Section 2.2.

4. Anzia and Moe 2016 provide an illustrative account of the bruising political battles surrounding efforts to reduce the state pensions' AAR in Rhode Islands in 2011 and California in 2015.

measures of pension plan financial status on an annual basis, including realized rates of return on assets, the Actuarially Assumed Return (AAR), and the Actuarial Funding Ratio (AFR) – the headline measure used by policymakers to understand a plan’s funding status. The PPD covers 177 of the largest public-sector pension plans in the United States: 87 state-run plans and 90 local-level plans that are administered by 89 unique local governments.⁵ In our analysis, we focus on the local-level pension plans and omit state-level plans as AAR adjustments at that level may be more likely to correlate with other policy decisions.

Table 3.1 reports some descriptive financial statistics for pension plans in the PPD. Column 1 shows the evolution of actual five-year returns from 2001 (the constant term) to 2017. Next to it are the yearly mean returns in column 2. Because the model in column 1 is saturated in year and plan fixed effects, these means can be calculated by adding the year fixed effects to the constant, e.g. $0.0312 = 0.0995 - 0.0683$ in 2003. The table reports 5-year returns because the specific earnings accounting of pension plans means that shorter-term returns of 1 or 3 years are uncorrelated with the actuarial value of assets (AVA) in the data. The patterns in column 1 clearly reflect the impact of the financial crisis but also the actuarial smoothing of returns as is common in pension funds (Munnell, Aubry, and Quinby 2011).⁶

Column 3 shows the evolution of plans’ AAR, i.e. the assumed long-term return that is used to discount obligations. The yearly means in column 4 display a glacially slow decrease in expected returns from an average of 8% in 2001 to an average of 7.4% in 2017. As discussed Section 3.2, the choice of AAR is a hot-button topic and the general consensus is that AARs far exceed actual returns earned by pension fund managers. This consensus is not borne out in this data to the degree that might be

5. PPD documentation claims that the plans in the database cover 95 percent of state/local pension assets and members in the US.

6. As a result of this smoothing, the financial crisis only becomes visible in the returns in 2009, and depresses returns until 2012, when the stock market had already robustly picked up again.

expected, as the average actual returns were in line or even exceeded the average AAR from 2013–2017. It is likely that this partly reflects the fact that the PPD includes only the largest plans which, tend to attract the best fund managers and generate the highest returns.

Column 5 relates changes in a plan’s funding gap only to the actual return and the AAR, conditional on plan fixed effects. Both coefficients have the expected sign: actual returns improve a plan’s asset position, while a higher AAR means future obligations are more steeply discounted. Column 6 shows how funding status of plans evolved over time using a regression with only plan and year fixed effects. The associated yearly means (column 4) of funding status show that there is an almost monotonic decline over the 2001–2017 period: the average AFR in the Public Plans Data goes from 95% in 2001 to 73% in 2017. In column 7 we ask whether most of this decline is explained by variation in actual and projected returns. The answer is ‘no’. While both are significant, their contribution to the R-squared is easily trumped by the year fixed effects. The most likely explanation then for the patterns in columns 6–7 is that plans’ funding status have been eroding from 2001–2017 because of the essentially unstoppable object that is unfunded pension obligations moving closer to ‘maturity’ as the baby boomers are beginning to enter retirement.⁷ As this unfolds, unrealistic assumptions, like the AAR, have been pushed to the forefront of policy debates.

Of the 294 adjustments to the actuarially assumed return, the average adjustment was a decrease of 31 basis points. The geographic distribution of Δ AAR events is shown in Figure 3.1. They are not isolated to any particular regions or significantly correlated with political leaning, but there is variation by state – California, Penn-

7. The baby boomers are commonly defined as the birth cohorts from 1945–1964. While some members of the baby boomers’ very first birth cohort (i.e. 1945) could have entered ‘early retirement’ (aged 55) as early as 2001, the majority of public employees still retire at age 65, and we are yet to hit the peak of baby-boomer retirement.

sylvania, and Florida experience the most. Figure 3.2 shows a binned scatterplot of these ΔAAR events against the resulting (fiscal year to fiscal year) changes in the official funding ratio of the respective pension plans. The relationship is obvious: when a plan's AAR is adjusted downward, future obligations are discounted at a lower rate, and the plan's funding ratio drops. In the raw data, a one-percent decrease in the AAR is associated with a concurrent 2.4 percentage-point decrease in the official funding ratio.

From the observed annual plan financial information in the PPD, we collect more precise dates of AAR changes to observe them on a monthly frequency. The precise date of a pension board's AAR change approval, the BMD, was identified from a review of board meeting minutes available on individual plan websites. When meeting minutes were unavailable online, Freedom of Information Act (FOIA) requests were submitted to retrieve the relevant information. The publication date of the first actuarial valuation report incorporating an AAR change, the ARD, was obtained from a review of the annual financial reports collected by Center for Retirement Research at Boston College.⁸ To pinpoint the transmission of pension shocks to local government budgets, we link the PPD pension plans to the U.S. Census of Governments and find the start date of the fiscal year following the publication of the actuarial valuation report, i.e., the first time governments address changes to their expected contributions in a budget. On average, the ARD is 4.2 months after the BMD and the FYE is 5.2 months after the ARD. Using these three dates, we create alternative AAR paths that represent different scenarios of public information acquisition. In each alternative, a given AAR adjustment takes effect at a different date (BMD, ARD, or FYE), providing a way to test for housing market responses at different times depending on the salience of AAR changes.

We match pension plans to their local economy and observe house prices with

8. Reports originally sourced from plan-specific websites or via FOIA requests.

monthly frequency housing data from Zillow. More specifically, we use Zillow’s Home Value Index (ZHVI) at the city- and county-level to match pension system-specific administrative boundaries. The ZHVI captures the typical home value (35th to 65th percentile) for a region and is constructed based on data for more than 100 million U.S. homes. We rely on two versions of the ZHVI that capture house prices for single-family homes and condominiums.

3.4 Empirical Strategy

Our empirical strategy exploits the plausibly exogenous shocks to AAR that are plan- and time-specific in a Difference-in-Differences framework. More specifically, we estimate the following regression model:

$$\text{HPI}_{iht} = \alpha + \beta_1 \text{AAR}_{it}^{\text{BMD/ARD/FYEND}} + \beta_2 \mathbf{1}[\text{h}=\text{SFH}] + \epsilon_{iht} \quad , \quad (3.2)$$

where HPI is Zillow’s (log) housing price index for market i (which is synonymous with pension plan i ’s administrative boundaries), housing type h (single-family home or condominium), and month t .

As outlined in Section 3.3, there are three versions of the AAR series which capture plan i ’s assumed rate of return in month t where approved AAR adjustments take effect at one of three dates: the board meeting date (BMD), the actuarial valuation report publication date (ARD), or the end of the first fiscal year when the local government budget accounts for adjustments in expected contributions (FYE). These variations provide a way to test for housing market responses at different times depending on the salience of AAR changes. To account for different levels by housing type, we have an indicator that takes a value of 1 if it is for single-family homes and 0 otherwise (condominiums). In all regressions, we include city and state-year fixed effects. In some, we also include big city-year fixed effects to capture poten-

tially changing housing stock composition in larger cities that face supply constraints. We cluster standard errors by pension plan. The identifying assumption is that, in the absence of changes in the AAR, the evolution of house prices would be similar between plan administrative areas (conditional on respective fixed effects).

To understand heterogeneous effects by housing type, we decompose the coefficient of interest in Equation 3.2 by interacting the AAR series with housing type in the following specification:

$$\begin{aligned} \text{HPI}_{iht} = \alpha + \sum_{g=\text{SFH,Condo}} \beta_1^g \left[\text{AAR}_{it}^{\text{BMD/ARD/FYEND}} \times \mathbf{1}[\text{h}=\text{g}] \right] + \\ \beta_2 \mathbf{1}[\text{h}=\text{SFH}] + \epsilon_{iht} \quad , \end{aligned} \tag{3.3}$$

where the coefficients of interest, β_1^{SFH} and β_2^{Condo} capture the effect of AAR on house prices for single-family homes and condominiums, respectively. The rest of the specification is the same as for Equation 3.2.

3.5 Results

The baseline results from estimating Equation 3.2 are reported in Table 3.2. Column pairs 1–2, 3–4, and 5–6 consider AAR shocks taking place at the board meeting date (BMD), the actuarial valuation report publication date (ARD), or the end of the first fiscal year when the local government budget accounts for adjustments in expected contributions (FYE), respectively. Beginning with column 1, we find that AAR increases denoted at the BMD cause house prices to rise. Put another way, a 1 percentage point decrease in the Actuarially Assumed Return causes house prices to decrease by 0.68 percent. In column 2, we address the concern that the housing supply elasticity can vary significantly by city, which can in turn result in adjusting housing stock composition (more condominiums than single-family homes, for example). We

find that results are robust to controlling for this potential confounder by including big city-year fixed effects.

Comparing across column pairs, we can reflect on the speed with which the implied shocks to cities' official pension under-funding obligations were capitalized into house prices. First, note that estimated coefficients in columns 3 and 4, which consider AAR shocks at the actuarial valuation report publication date (ARD), closely match those in columns 1 and 2. In contrast, estimates in columns 5 and 6, which consider AAR shocks at the end of the first fiscal year when the local government budget accounts for adjustments in expected contributions (FYE), are not significantly different from 0. This suggests that information diffuses rather quickly to local homeowners and is fully capitalized into house prices by the FYE, which is an average of 9 months after the BMD and 5 months after the ARD.

We next consider whether shocks to reported levels of government debt have heterogeneous effects on the housing market. In columns 2, 4, and 6 of Table 3.3, we report estimated coefficients from the model in Equation 3.3. Columns 1, 3, and 5 repeat the baseline results from columns 2, 4, and 6 in Table 3.2. As shown, the main effect in the baseline estimates is driven by price responses of single-family homes, while there is no statistically significant effect on condominium prices. One potential explanation for this result is that cutting public goods and services may be the only option for some local governments to address spikes in pension spending. In other words, the increased ARC following a decrease in the AAR crowds out spending on other public goods. These cutbacks, which could include funding for public schools, may be more prominent to owners of single-family homes, who are more likely to be families with children.

3.6 Conclusion

In this chapter, I have shown that public-sector pension under-funding, which is the largest source of public debt at the local level in the United States, has meaningful implications for the local economy. I have done so by exploiting plausibly exogenous shocks to the reported levels of unfunded pension liabilities in a Difference-in-Differences framework. To facilitate this analysis, I have used various public records to document three important dates (BMD, ARD, FYE) when updated financial information is revealed to the public or increases in salience. I find that increases in public debt are capitalized into local house prices relatively quickly (within 9 months). Additionally, this effect is driven by responses in the price of single-family homes, owners of which may be more likely to rely on public goods that experience cuts following spikes in reported pension under-funding.

These findings provide rare empirical evidence of the local Ricardian Equivalence. More specifically, I have shown that homeowners expect a compensating differential for higher levels of local debt and the implied future tax increases. This can have important implications over the coming years, as “public pensions are the time bomb of government finance” (Mary 2020). What remains unclear is the extent to which ballooning pension debt crowds out other spending on public goods – a topic warranting future research. If so, this could create a downward spiral for municipalities throughout the country as pension debt, public goods cuts, and a shrinking tax base interact.

Table 3.1: Funding Status of Pension Plans: 2001–2017

Outcome:	Return _{it}		AAR _{it}		% Funded _{it}			
	(1)	(2) Yearly Mean	(3)	(4) Yearly Mean	(5)	(6)	(7)	(8) Yearly Mean
Return (5-yr)					0.4508*** (0.0584)		0.5870*** (0.0976)	
AAR					16.3039*** (2.4150)		4.6681** (1.8064)	
Year = 2002	-0.0480*** (0.0025)	0.0515	-0.0002 (0.0001)	0.0801		-0.0744*** (0.0082)	-0.0454*** (0.0089)	0.9476
Year = 2003	-0.0683*** (0.0025)	0.0312	-0.0004** (0.0002)	0.0799		-0.1297*** (0.0112)	-0.0876*** (0.0123)	0.8923
Year = 2004	-0.0624*** (0.0025)	0.0371	-0.0008*** (0.0002)	0.0795		-0.1561*** (0.0131)	-0.1159*** (0.0136)	0.8659
Year = 2005	-0.0591*** (0.0026)	0.0404	-0.0008*** (0.0002)	0.0795		-0.1792*** (0.0158)	-0.1409*** (0.0161)	0.8428
Year = 2006	-0.0264*** (0.0024)	0.0731	-0.0009*** (0.0003)	0.0794		-0.1840*** (0.0169)	-0.1644*** (0.0162)	0.8380
Year = 2007	0.0198*** (0.0023)	0.1193	-0.0010*** (0.0003)	0.0793		-0.1677*** (0.0178)	-0.1747*** (0.0172)	0.8543
Year = 2008	-0.0169*** (0.0030)	0.0826	-0.0012*** (0.0003)	0.0791		-0.1962*** (0.0177)	-0.1808*** (0.0172)	0.8258
Year = 2009	-0.0747*** (0.0020)	0.0248	-0.0013*** (0.0003)	0.0790		-0.2505*** (0.0177)	-0.2004*** (0.0185)	0.7715
Year = 2010	-0.0676*** (0.0021)	0.0319	-0.0018*** (0.0004)	0.0785		-0.2722*** (0.0183)	-0.2244*** (0.0187)	0.7498
Year = 2011	-0.0598*** (0.0019)	0.0397	-0.0024*** (0.0004)	0.0779		-0.2844*** (0.0195)	-0.2379*** (0.0197)	0.7376
Year = 2012	-0.0803*** (0.0021)	0.0192	-0.0031*** (0.0004)	0.0772		-0.3052*** (0.0199)	-0.2438*** (0.0209)	0.7168
Year = 2013	-0.0349*** (0.0033)	0.0646	-0.0034*** (0.0004)	0.0769		-0.3044*** (0.0209)	-0.2680*** (0.0205)	0.7176
Year = 2014	0.0225*** (0.0019)	0.1220	-0.0038*** (0.0004)	0.0765		-0.2925*** (0.0223)	-0.2878*** (0.0216)	0.7295
Year = 2015	-0.0008 (0.0020)	0.0987	-0.0044*** (0.0004)	0.0759		-0.2875*** (0.0227)	-0.2667*** (0.0222)	0.7345
Year = 2016	-0.0281*** (0.0021)	0.0714	-0.0053*** (0.0004)	0.0750		-0.2973*** (0.0230)	-0.2561*** (0.0232)	0.7247
Year = 2017	-0.0097*** (0.0019)	0.0898	-0.0066*** (0.0005)	0.0737		-0.2912*** (0.0231)	-0.2547*** (0.0231)	0.7308
Constant	0.0995*** (0.0017)		0.0803*** (0.0003)		-0.5064*** (0.1900)	1.0220*** (0.0170)	0.5886*** (0.1451)	
Observations	2,650		2,650		2,650	2,650	2,650	
R-squared	0.7699		0.7138		0.7245	0.8324	0.8394	
# Plan FEs	177		177		177	177	177	

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Columns 1, 3, and 6 regress plans' investment returns, AAR (the assumed long-term return that is used to discount obligations), and funding ratio (AFR) on year fixed effects, respectively. Column 5 regresses plans' AFR on the actual return and the AAR. Column 7 regresses plans' AFR on actual returns, AAR, and year fixed effects. Plan fixed effects are included in all regressions. Columns 2, 4, and 6 report 'yearly mean' of the respective variable for the model that is saturated in year and plan fixed effects, deducting the year fixed effects from the constant, e.g. 0.9476 = 1.0220 - 0.0744 in column 8 in 2002. Standard errors (clustered by plan) are reported in parentheses.

SOURCE: Public Plans Data. 2001-2022. Center for Retirement Research at Boston College, MissionSquare Research Institute, National Association of State Retirement Administrators, and the Government Finance Officers Association.

Table 3.2: AAR Shocks and House Price Capitalization: DID Baseline

	Board Meeting Date		Actuarial Report Publication Date		Fiscal Year End for Admin. Government	
	(1) House Price (log)	(2) House Price (log)	(3) House Price (log)	(4) House Price (log)	(5) House Price (log)	(6) House Price (log)
Actuarially Assumed Return (AAR)	0.678** (0.272)	0.704*** (0.260)	0.736*** (0.249)	0.734*** (0.241)	0.243 (0.255)	0.232 (0.249)
1[Single Family Home]	0.160*** (0.010)	0.160*** (0.010)	0.160*** (0.010)	0.160*** (0.010)	0.160*** (0.010)	0.160*** (0.010)
Number of Events	294	294	287	287	269	269
Number of Event Cities	85	85	82	82	82	82
Number of Total Cities	89	89	89	89	89	89
R-squared	0.911	0.913	0.911	0.913	0.910	0.912
Observations	36,762	36,762	36,150	36,150	36,150	36,150
City FE	✓	✓	✓	✓	✓	✓
State-Year FE	✓	✓	✓	✓	✓	✓
Big City-Year FE	-	✓	-	✓	-	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Standard errors (clustered by plan) are reported in parentheses.

SOURCE: Public Plans Data. 2001-2017. Center for Retirement Research at Boston College, MissionSquare Research Institute, National Association of State Retirement Administrators, and the Government Finance Officers Association.

Table 3.3: AAR Shocks and House Price Capitalization: Heterogeneity by Housing Type


	Board Meeting Date		Actuarial Report Publication Date		Fiscal Year End for Admin. Government	
	(1) House Price (log)	(2) House Price (log)	(3) House Price (log)	(4) House Price (log)	(5) House Price (log)	(6) House Price (log)
Actuarially Assumed Return (AAR)	0.704*** (0.260)		0.734*** (0.241)		0.232 (0.249)	
× 1[Condo]		0.529 (0.460)		0.535 (0.458)		0.072 (0.472)
× 1[Single Family Home]		0.878** (0.408)		0.937** (0.425)		0.393 (0.428)
1[Single Family Home]	0.160*** (0.010)	0.133** (0.053)	0.160*** (0.010)	0.129** (0.056)	0.160*** (0.010)	0.135** (0.057)
Number of Events	294	294	287	287	269	269
Number of Event Cities	85	85	82	82	82	82
Number of Total Cities	89	89	89	89	89	89
R-squared	0.913	0.913	0.913	0.913	0.912	0.912
Observations	36,762	36,762	36,150	36,150	36,150	36,150
City FE	✓	✓	✓	✓	✓	✓
State-Year FE	✓	✓	✓	✓	✓	✓
Big City-Year FE	✓	✓	✓	✓	✓	✓

*** p < 0.01, ** p < 0.05, * p < 0.1.

NOTE: Standard errors (clustered by plan) are reported in parentheses.


SOURCE: Public Plans Data. 2001-2017. Center for Retirement Research at Boston College, MissionSquare Research Institute, National Association of State Retirement Administrators, and the Government Finance Officers Association.

Figure 3.1: Geographic Distribution of AAR Adjustments



NOTE: Map plots number of adjustments to AARs by plans in each state over the 2001-2017 period.
SOURCE: Public Plans Data. 2001-2017. Center for Retirement Research at Boston College, MissionSquare Research Institute, National Association of State Retirement Administrators, and the Government Finance Officers Association.

Figure 3.2: How AAR Adjustments Affect Funding Ratio



NOTE: Figure is a binned scatterplot of the roughly 300 adjustments to AAR by public-sector pension plans in 2001-2017 against the reported change in the Actuarial Funding Ratio (AFR).
SOURCE: Public Plans Data. 2001-2017. Center for Retirement Research at Boston College, MissionSquare Research Institute, National Association of State Retirement Administrators, and the Government Finance Officers Association.

References

- Akai, Nobuo. 1994. "Ricardian equivalence for local government bonds: Budget constraint approach." Economics Letters 44 (1-2): 191–195.
- Anzia, Sarah F, and Terry M Moe. 2016. "Interest Groups on the Inside: The Governance of Public Pension Funds." Unpublished draft (August).
- . 2019. "Interest groups on the inside: The governance of public pension funds." Perspectives on Politics 17 (4): 1059–1078.
- Banzhaf, H Spencer, and Wallace E Oates. 2013. "On fiscal illusion in local public finance: Re-examining Ricardian equivalence and the renter effect." National Tax Journal 66 (3): 511.
- Barro, Robert J. 1974. "Are government bonds net wealth?" Journal of political economy 82 (6): 1095–1117.
- Brinkman, Jeffrey, Daniele Coen-Pirani, and Holger Sieg. 2018. "The political economy of municipal pension funding." American Economic Journal: Macroeconomics 10 (3): 215–46.
- Brunner, Eric J, and Jon Sonstelie. 2006. "California's School Finance Reform: An Experiment in Fiscal Federalism." In The Out Model at Fifty: Essays in Public Economics in Honor of Wallace Oates, edited by William Fishel. Cambridge, MA: Lincoln Institute of Land Policy.
- Daly, George G. 1969. "The burden of the debt and future generations in local finance." Southern Economic Journal, 44–51.
- Dollery, Brian E, and Andrew C Worthington. 1996. "The empirical analysis of fiscal illusion." Journal of Economic Surveys 10 (3): 261–297.

- Fishel, William. 2001. "The Homevoter Hypothesis: How Home Values Influence Local Government Taxation." School Finance, and Land-Use Policies.
- Glaeser, Edward L, and Giacomo AM Ponzetto. 2014. "Shrouded costs of government: the political economy of state and local public pensions." Journal of Public Economics 116:89–105.
- Greenhut, Steven. 2009. "Plunder!: How Public Employee Unions are Raiding Treasuries Controlling Our Lives and Bankrupting the Nation."
- Inman, Robert P. 1982. "Public employee pensions and the local labor budget." Journal of Public Economics 19 (1): 49–71.
- Kelley, Dashle G. 2014. "The political economy of unfunded public pension liabilities." Public Choice 158 (1-2): 21–38.
- Mary, Williams Walsh. 2020. "Coronavirus Is Making the Public Pension Crisis Even Worse." The New York Times. April 2.
- Mitchell, Olivia S, and Robert S Smith. 1994. "Pension Funding in the Public Sector." The Review of Economics and Statistics.
- Munnell, Alicia H, Jean-Pierre Aubry, and Laura Quinby. 2011. "Public pension funding in practice." Journal of Pension Economics & Finance 10 (2): 247–268.
- Novy-Marx, Robert, and Joshua D Rauh. 2009. "The liabilities and risks of state-sponsored pension plans." Journal of Economic Perspectives 23 (4): 191–210.
- . 2011. "Public pension promises: how big are they and what are they worth?" The Journal of Finance 66 (4): 1211–1249.

Novy-Marx, Robert, and Joshua D Rauh. 2014. "The revenue demands of public employee pension promises." American Economic Journal: Economic Policy 6 (1): 193–229.

Seater, John J. 1993. "Ricardian equivalence." Journal of economic literature 31 (1): 142–190.

The Economist. October 5th 2017. "The Big Squeeze American Public Pensions Suffer From a Gaping Hole."

Tiebout, Charles M. 1956. "A pure theory of local expenditures." Journal of political economy 64 (5): 416–424.