

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Beyond Flows: Essays on the Economic Effects of Migrant Composition

Permalink

<https://escholarship.org/uc/item/3fx6r9xv>

ISBN

9798180199072

Author

Gray, Andre

Publication Date

2026-07-24

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Beyond Flows: Essays on the Economic Effects of Migrant Composition

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

in

Economics

by

Andre Daniel Gray

Committee in Charge:

Professor Samuel Bazzi, Co-Chair

Professor Sara Lowes, Co-Chair

Professor Gaurav Khanna

Professor Tom Vogl

Copyright

Andre Daniel Gray, 2026

All rights reserved.

The Dissertation of Andre Daniel Gray is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2026

TABLE OF CONTENTS

Dissertation Approval Page	iii
Table of Contents	iv
List of Figures	viii
List of Tables	x
Acknowledgements	xiii
Vita	xv
Abstract of the Dissertation	xvi
1 Density and Diversity in African Cities	1
1.1 Introduction	1
1.2 Data	9
1.2.1 Proxied Origin-Destination Panels	10
1.2.2 Push Shocks to Outmigration	11
1.2.3 Pull Characteristics to Predict Immigration	16
1.2.4 Destination Productivity Outcomes	16
1.2.5 Validating a Proxy Origin-Destination Panel	20
1.3 Empirical Strategy	22
1.3.1 A Shift-Share Instrument for Migrant Birthplace Diversity	24
1.3.2 Exclusion Restriction	28
1.3.3 Predicting Outmigration Shifts with Push Shocks	29
1.3.4 Predicting Settlement Shares with Pull Characteristics	31
1.4 Results	34

1.4.1	Baseline Results: Standard Shift-Share with Nonlinearity	35
1.4.2	Instrumenting Outmigration Shifts with Push-Shocks	40
1.4.3	Instrumenting Settlement Shares with Pull Characteristics	44
1.4.4	Evidence from a Natural Experiment: Ending Apartheid	45
1.5	Mechanisms	48
1.5.1	Ethnic Conflict	49
1.5.2	Linguistic Distance	50
1.5.3	Skill-Level, Segregation and Industry Concentration	52
1.5.4	Productivity Evidence from Firms	53
1.5.5	Migration and Ethnic Attitudes	54
1.6	Urban Growth and Diversity in the Long-Run	55
1.7	Conclusion	59
1.8	Acknowledgements	60
1.9	Appendix 1: Simulations of Shift-Share Instruments with Nonlinearity	68
1.10	Appendix 2: Additional Tables & Figures	74
1.10.1	Supplementary Analysis	74
1.10.2	Birthplace-Level Tests Following Borusyak et al. (2022)	74
1.10.3	Robustness Checks	76
1.10.4	Pass-through of International Commodity Prices	76
1.10.5	Relationship of Local GDP to Night Light Density	79
2	Historical Immigration and the Market for Schooling in American Cities	92
2.1	Introduction	92
2.2	Data	97
2.2.1	Enumeration Districts and 1960 Tracts	98
2.2.2	Teacher-Level Outcomes	98
2.2.3	School-Level Outcomes	100
2.2.4	Linked Census Sample	100

2.3	Historical Background: European Migration and Catholic Schooling	102
2.4	Empirical Method	104
2.4.1	The 1920s National-Origin Quota System	107
2.4.2	Catholic Quota Exposure	111
2.4.3	Predicted Change in Catholic Exposure	114
2.4.4	Cohort-Level Exposure to Quota	114
2.5	Results	116
2.5.1	Impacts of Cohort Level Exposure to Catholic Migration	116
2.5.2	Impacts of Catholic Quota Exposure	119
2.5.3	Impacts of Relative Change in Catholic Migration	123
2.6	Robustness Checks	124
2.7	Conclusion	126
2.8	Acknowledgements	127
2.9	Appendix: Additional Figures and Tables	128
3	The Effect of Migrant Composition on Housing Prices	138
3.1	Introduction	138
3.2	Data	142
3.2.1	Changes in Composition by Origin	144
3.2.2	Predictors of Headship	147
3.3	Methodology	151
3.3.1	Theoretical Framework: Housing Demand and Immigrant Composition . .	151
3.3.2	Shift-Share Instruments for Migrant Composition	153
3.3.3	Shift Share Instruments by Age and Skill Cells	158
3.3.4	Estimation Strategy	159
3.3.5	Interpreting Migrant Inflows and Migrant Shares	161
3.4	Results	162
3.4.1	First Stage Tests of Instrument Relevance	162

3.4.2	Observational Relationship Between House Price and Migrant Inflows . . .	162
3.4.3	2SLS Estimates of Migrant Impacts on House Prices	164
3.4.4	Instrument Independence and Pretrend Diagnostics	167
3.5	Conclusion	169
3.6	Acknowledgements	170
3.7	Appendix 1: Additional Figures and Tables	171
3.8	Appendix 2: Simulations of Shift-Share Identification with Inflow and Composition Effects	180
3.8.1	Common Simulation Environment	180
3.8.2	Exercise 1: Two Endogenous Regressors (Inflow and Composition)	181
3.8.3	Exercise 2: Inflow-Only Estimation with Omitted Composition	181
3.8.4	Simulation Protocol	182

LIST OF FIGURES

1.1	Data Examples from Mozambique 1997-2007	15
1.2	Relationship between Diversity (HHI) and Labor Variables	25
1.3	Migrant Size and Composition on Light Density	28
1.4	Portage Sites, DRC Example	33
1.5	Comparing Migration Estimates to Other Samples	39
1.6	Predicting Outmigration with Individual Shocks Across Horizons	43
1.7	Ethnic Homelands under Apartheid, South Africa	46
1.8	Black Diversity in South Africa After Apartheid	47
1.9	Conflict and Population Size in the Cross-Section	49
1.10	Data Examples from Nigeria	58
1.11	Simulation of OLS and SSIV with Omitted Variable Bias	70
1.12	SSIV Diversity Beta Coefficients with Varying Instrument Correlation	71
1.13	SSIV Labor Beta Coefficients with Varying Instrument Correlation	72
1.14	F-Statistic by Varying Instrument Correlation	73
1.15	First Stage Scatterplots in a Birthplace-Level Regression	79
1.16	Reduced Form Scatterplots in a Birthplace-Level Regression	80
1.18	Standard Error Adjustments Following (Borusyak et al., 2022)	81
1.20	Balance at Baseline in Birthplace-Level Regressions	82
1.21	Conley Spatial Standard Errors for Labor Size Coefficient $\Delta\ell$	85
1.22	Conley Spatial Standard Errors for Migrant Diversity Coefficient Δdiv	86
1.23	Estimates with Individual Country Drops for Labor Size Coefficient $\Delta\ell$	87
1.24	Estimates with Individual Country Drops for Migrant Diversity Coefficient Δdiv	88
1.25	Placebo Shocks for Labor Size Coefficient $\Delta\ell$	89
1.26	Placebo Shocks for Migrant Diversity Coefficient Δdiv	90
1.27	IRF of International Prices on Local Urban/Rural Gap	90

1.28	Correlations of Port Activity and Night Light Density	91
2.1	Enumeration District Map Example	99
2.2	Change in Teacher Characteristics Over Time	101
2.3	Change in Catholic School Enrollment Share	105
2.4	Migrant Characteristics by Birth Country in 1880	108
2.5	Immigrant Compositional Change	110
2.6	Catholic Exposure by 1960 Tract	113
2.7	Event-Study of Quota Policies by Cohort	119
2.8	Event-Study of Predicted Catholic Change by Cohort	120
2.9	Catholic School Expansion Example	129
2.10	Event-Study of Catholic Quota Exposure by Cohort with csdid	130
2.11	Leave One City Out Composition and Teacher Trends	133
2.12	Leave One Origin Out Composition and Teacher Trends	134
2.13	Board of Education Archive	135
3.1	Migrant Flows by Birthplace	146
3.2	Migrant Age by Birthplace	147
3.3	Headship Rate by Immigrant Birthplace	148
3.4	Variance decomposition of headship rates	171
3.5	Pairwise Correlations Between HPI Indices	172
3.6	Tests of Instrument Independence in First Stage	176
3.7	Event Study Tests of Pre-Trends	177
3.8	Pairwise Correlations Between Migrant Inflows and Composition	178
3.9	Pairwise Correlations Between Shift-Share Instruments	179
3.10	Bias with Two Endogenous Regressors by Shock Correlation	183
3.11	Bias and Total Effect in Inflow-Only Shift-Share Simulation	184

LIST OF TABLES

1.1	Observational Returns by Migrant Status and Distance	23
1.2	Census Birthplace Shift-Share First Stage	36
1.3	Census Birthplace Shift-Share Second Stage, Migrant Flows	38
1.4	Predicting Outmigration with Contemporary and Lagged Shocks	42
1.5	First Stage Prediction of Migrant Size and Composition with Predicted Shifts and Shares	61
1.6	Second-Stage Effects of Migration on Productivity with Predicted Shifts and Shares	62
1.7	South Africa Shift-Share Second Stage, Migrant Flows	63
1.8	Shift-Share Second Stage, Conflict Outcomes	64
1.9	Change in Linguistic Distance and Productivity	65
1.10	First Stage Relationship of <i>div</i> and Linguistic Distance	65
1.11	Ethnicity and Firm Productivity in Ghana, Tanzania	66
1.12	2SLS of Historical Diversity and Productivity Instruments	66
1.13	2SLS of Historical Diversity and Productivity Instruments: Mineral Deposits and Portage	67
1.14	African Censuses with Geolocation, Birthplace or Ethnicity	75
1.15	Gravity Regressions of Migrant Flow using Proxy O-D Panels	76
1.16	Census Birthplace Shift-Share Second Stage, Migrant Flows	77
1.17	Census Birthplace Shift-Share Second Stage, Migrant Flows	78
1.18	Self-Reported Identity by Migrant Status and Distance	83
1.19	OLS Cross-Section of Historical Diversity and Urban Density	84
2.1	School Characteristics By Sector	102
2.2	Immigrant Characteristics in 1880 (Northeast Cities)	106
2.3	Historical Catholic Shares and Private School Share	117

2.4	1920 Quota Impacts on Child Educational Attainment	118
2.5	1920 Quota Impacts on Tract Characteristics	121
2.6	1920 Quota Impacts on Teacher Characteristics	121
2.7	1920 Quota Impacts on Schools	122
2.8	1920 Quota Impacts on Tract Characteristics	123
2.9	1920 Quota Impacts on Teacher Characteristics	124
2.10	City Coverage Across Data Sources	128
2.11	1920 Quota Impacts on Child Educational Attainment	128
2.12	Educational Attainment in Placebo Cohorts	130
2.13	Pre Period Tract Level Characteristics by Quota Exposure	131
2.14	Teacher Characteristics with Conley Standard Errors	131
2.15	Tract Level Characteristics with Conley Standard Errors	132
2.16	Individual Linking Probability by Quota Exposure	132
3.1	Correlation of Migrant Characteristics and Headship	150
3.2	Correlation of Migrant Characteristics and Migrant Inflows	155
3.3	First Stage Estimates of Shift-Share Instruments for Migrant Characteristics	163
3.4	OLS Estimates of House Price Growth and Migrant Characteristics	164
3.5	2SLS Estimates of House Price Growth and Migrant Characteristics	165
3.6	2SLS Estimates of House Price Growth with Migrant Flows and Composition	166
3.7	2SLS Estimates of House Price Volatility and Migrant Characteristics	167
3.8	2SLS Estimates of Dwelling and Migrant Characteristics	168
3.9	OLS House Price Volatility and Immigrant Composition	172
3.10	OLS Single Family House Price and Immigrant Composition	173
3.11	OLS Multi Family House Price and Immigrant Composition	173
3.12	OLS Housing Characteristics and Immigrant Composition	174
3.13	Decade Level House Price Growth And Immigration Composition	174
3.14	2SLS Estimates of House Price Growth and Individual Migrant Characteristics	175

3.15	2SLS Estimates of HPI Volatility and Migrant Characteristics	175
3.16	2SLS Estimates of House Price Volatility and Migrant Flows and Composition . .	176
3.17	Event-Study Lead Coefficients (Pretrend Check)	177
3.18	IV Correlations with Lagged MSA Trends	178

ACKNOWLEDGEMENTS

I am grateful to my advisors Samuel Bazzi and Sara Lowes. Their generous guidance, consistent encouragement and expert advice were essential to my PhD experience and to the development of this dissertation. I would also like to thank Gaurav Khanna, Tom Vogl, Katherine Meckel, Prashant Bharadwaj and Julie Cullen for their insights and tough questions across many seminars. Thank you to Anya Samek, my academic mentor for the last 10 years. She first hired me as an undergraduate research assistant, introduced me to the practice of research and has continued to be a resource and a well of encouragement.

I am grateful to my fellow graduate students who have become great friends. Kurtis Gilliat, Juan Hernandez, Vinny Armentano, Steven Brownstone, Ricardo Di Cato and Erica Chuang are brilliant scholars and excellent people. I am lucky to have met you all. Thanks to Francesco Amodeo for many nights spent at the movies, or over a drink. Thanks to Jordan Mosqueda and Paula Donaldson for sharing your intelligence and humor over lunches, dinners and coffees. A special thanks to my co-author Fulvia Budillon, who endured months with me in the field and has always been a source of energy.

Last I would like to thank my very patient family. My mother Montse Gray, from whom I inherited a love of learning, and who showered me with books from the start. My father Douglas Gray, a globetrotter, whose passion for travel and curiosity for other places and people have been an inspiration to my life and work. My sister Gabriella Gray, the more successful and capable sibling. And my grandparents, Candelaria and Raul Martin, for their unwavering love and support.

Chapter 1, “Density and Diversity in African Cities,” is currently being prepared for submission for publication of the material therein. The dissertation author, Andre Daniel Gray, is the sole author of this chapter.

Chapter 2, “Historical Immigration and the Market for Schooling in American Cities,” is currently being prepared for submission for publication of the material therein. The dissertation

author, Andre Daniel Gray, is the sole author of this chapter.

Chapter 3, “The Effect of Migrant Composition on Housing Prices,” is currently being prepared for submission for publication of the material therein. The dissertation author, Andre Daniel Gray, is the sole author of this chapter.

VITA

- 2017 Bachelor of Science, University of Southern California
- 2026 Doctor of Philosophy, University of California San Diego

ABSTRACT OF THE DISSERTATION

Beyond Flows: Essays on the Economic Effects of Migrant Composition

by

Andre Daniel Gray

Doctor of Philosophy in Economics

University of California San Diego, 2026

Professor Samuel Bazzi, Co-Chair

Professor Sara Lowes, Co-Chair

This dissertation consists of three chapters on the local economic effects of immigrant composition. Across settings, I study how the demographic, socioeconomic, and cultural characteristics of migrants shape urban growth, public-good provision, and housing markets.

Chapter one studies internal migration and urban growth in Africa. I examine how migrant inflow size and birthplace diversity jointly affect city-level economic activity in rapidly urbanizing countries. I construct a subnational panel of internal migration across Africa and develop a nonlinear shift-share instrument that separately identifies shocks to migrant levels and birthplace composition.

I find that cities receiving migrants from more diverse birthplaces experience lower short-run growth but long-run urbanization gains. Greater diversity is associated with higher ethnic conflict, but also with faster structural transformation. The chapter shows that migrant composition can generate effects distinct from total migration flows.

Chapter two examines how historical Catholic migration reshaped schooling markets in the United States. In the early twentieth century, Catholic migrants increased demand for private religious education, changing the allocation of children across public and private schools. To study the impact on native children, I build a panel of Catholic migration to Northeastern cities from 1880 to 1940 and construct census-tract exposure to the 1920s national-origin quota system using pre-policy settlement patterns and quota-induced missing inflows. Combining historical census microdata with newly digitized Catholic and public school records, I find that native children exposed to relatively larger Catholic migrant inflows attained more schooling.

Chapter three studies migrant composition, headship, and housing price growth. Standard shift-share measures aggregate migrant supply shocks linearly, potentially obscuring the role of age, skill, and household-formation patterns in shaping housing demand. Using metropolitan-level data from the American Community Survey, I construct measures of international migrant inflows and demographic composition. I show that nonlinear functions of origin-specific shifts and baseline settlement shares predict changes in migrant composition, then estimate how these features affect housing markets. I find that increases in the migrant share of the population are associated with lower house price growth, while younger migrant inflows reduce prices and increase price volatility.

Chapter 1

Density and Diversity in African Cities

1.1 Introduction

Measuring the impact of migrant workers is usually done in levels. The size of migrant inflows, or the average level of their skill, changes labor demand and supply, prices and traffic. Past work has argued that as the flow of migrants increases, they can raise the productivity of destinations through agglomeration effects. An increase in the raw number of people in a city stimulates competition and generates new ideas. If agglomeration effects exist, then the level of migrant flows can scale-up city populations and in turn productivity growth (Glaeser and Gottlieb, 2009). But to achieve these effects, migrants have to interact with each other and with the local population. These interactions might be smooth or difficult, and the particular skills, backgrounds and social networks that migrants bring with them may be more or less complementary. The complementarity of migrant workers is a function of their composition; the relative sizes of different groups at destination. In this paper I disentangle these two effects, density and diversity, in the context of African rural-urban migration. I find that both increasing migrant labor size and diversity produce lower per-capita city growth, indicative of short-run congestion effects from ethnic conflict. However, higher migrant diversity also predicts increases to non-agricultural labor share, suggesting that migrant diversity plays an important role in structural transformation.

The urban share of population in Africa rose from 31 to 54% between 1990 and 2020 (Christiaensen et al., 2025). Across African cities about a third of the urban labor force is composed of migrants from rural parts of the country (Christiaensen et al., 2023). In developed countries, the most productive cities are both dense and cosmopolitan. New York or London host a diversity of industries, amenities and services that benefit from a wide range of workers with a healthy mix of skills. Whether developing country cities grow under the same conditions is an open question. In the context of Africa, there are reasons to be skeptical that cities benefit from either density or diversity. First, African cities suffer from high congestion costs and poor urban infrastructure. Traffic, pollution, poor housing and small industrial sectors may all prevent cities from taking advantage of increased density (Castells-Quintana, 2017). Second, African countries are saddled with ethnic and religious conflict. A literature in political economy documents correlations between measures of ethnic and linguistic diversity, low GDP growth and high incidence of violent conflict (Arbatli et al., 2020; Alesina et al., 2003; Robinson, 2020). While the mechanisms are not well understood, microeconomic evidence has shown that ethnic divisions can directly lower the productivity of firms (Hjort, 2014).

This paper offers a framework to synthesize the literatures on agglomeration and ethnic conflict. Density and diversity are jointly determined by the size of the labor force, and the mix of worker types. As migrant flows increase or decrease, they increase or decrease the level of density in a city. If migrants are all from the same origin, then they lower the relative diversity of the workforce. If migrants come from many different origins, the relative diversity of the migrant pool increases. Capturing the effects of migrants on urban Africa requires an empirical strategy that can disentangle the effects of migrant levels and composition. To identify these parameters separately, I construct a shift-share instrument that predicts linear changes in migrant flows, and nonlinear changes in migrant diversity. Then, I consider the role of pull and push shocks in driving changes in migrant levels and composition. I estimate the impact of climate, commodity prices and conflict on the size and spatial pattern of migrant flows in Africa.

The paper begins by creating a proxied origin-destination panel of African migrant flows.

To overcome data gaps in the measurement of African migration, I use the universe of publicly available African censuses and leverage a mix of worker birthplace and language data to identify worker origins. While past work on African migration has relied on cross-sectional data or small household panels, this origin-destination panel captures aggregate changes in population size and composition by origin between censuses.

Destinations are exposed to migrants from different birthplaces according to a pre-period settlement share of workers from a given origin o living in destination d . Each origin birthplace has a country-wide outmigration “shift”, measured as the total number of people from that birthplace observed living outside of that homeland across census periods. The combination of the aggregate shifts with the destination-specific shares is an instrument for migrant labor size, or the “levels” of migration. To capture the effects of composition, I add a second shift-share instrument. This “composition” instrument captures how changes in the distribution of the outmigration shifts predictably alter a Herfindahl Index (HHI) of migrant shares. The result is a linear and nonlinear shift-share instrument to predict changes in both migrant levels and birthplace composition, as measured by an HHI of birthplace shares. The instruments capture predicted changes in migrant levels and composition that are plausibly exogenous to contemporaneous labor demand shocks.

To address concerns about destination productivity changes being endogenous to birthplace-level outmigration shifts, I build on the shift-share strategy by replacing the outmigration shifts with predicted outmigration rates. Following work on historical immigration to the US by Boustan et al. (2010), I predict outmigration from each origin using a set of plausibly exogenous shocks related to climate, conflict and international prices. In a zero stage I show that drought reduces outmigration, while conflict and high commodity prices boost outmigration rates. Leveraging these variables to predict our shifts, I re-estimate the 2SLS model for levels and composition.

While migrant settlement shares are often treated as an endogenous component in Bartik-style migration instruments, serial correlation between pre-period settlement shares and contemporary labor demand may bias results. In a last exercise I attempt to isolate exogenous variation in shares that comes from a zero stage regression of origin-destination distances and historical destination

characteristics. These “pull characteristics” predict a destination’s overall attractiveness to migrants based on historical productivity shocks that are plausibly unrelated to contemporary demand. “Pull characteristics” can be thought of as historical instruments for agglomeration. Examples include distance to colonial railroads, historical mineral deposits, and portage sites.

The outcomes of interest measure different aspects of a destination’s urban growth. I use different functions of changes in satellite night-light luminosity as proxies for changes in city size and GDP per capita. I also use census measured non-agricultural labor share as a proxy for urbanization, and a principal component of housing characteristics as a measure of local average household wealth. The results show that in-migration to a destination increases city light density, but not per-capita growth rates. Higher migrant levels scale up city sizes, but don’t yield higher productivity growth. A higher diversity of migrants coming from a wider set of origins decreases both light density and per-capita growth. Moving from a perfectly homogeneous to a perfectly heterogeneous migrant labor force reduces light density growth by 10%, conditional on migrant labor size. These effects are large. Relative to other developing countries, the negative effect of migrant labor size on light density growth is 4 times larger in an African context, and the congestion effects of diversity are twice as large. These results suggest that congestion forces created by migrants may constrain the potential benefits of increasing density.

When examining non-light based measures of productivity the picture becomes more complicated. More diverse migrant pools increase the non-agricultural labor share, an indicator of urbanization and structural transformation. Again the size of the effects is economically meaningful. Increasing migrant diversity creates a more than 1 for 1 increase in non-agricultural labor share, conditional on migrant labor size. Migrant labor size by itself has a modest and negative effect on non-agricultural labor share. I find that the result is driven mainly by an increasing share of services. This heterogeneity points to the importance of separately identifying size and composition effects of migration. I interpret these results as evidence that migrant diversity may have negative short-run effects on city growth, but may yield long-run benefits in the form of industry mix and structural transformation. As a further exploration of heterogeneity, I consider evidence from the end of

Apartheid in South Africa, which created a sudden drop in migration costs for black South Africans. Leveraging the abolition of Pass Laws as a natural experiment, I study the impact of the end of Apartheid on changes to migrant labor size and diversity. I find that while the size of the native black migrant labor force still produces negative effects on urban productivity growth consistent with the baseline results, there are no congestion effects from black migrant diversity. This result shows that particular country experiences and history may alter the implications of migrant diversity for urban growth.

I consider several mechanisms to explain my results, including ethnic conflict, linguistic diversity, firm size and skill composition. I find evidence consistent with the congestion effects of ethnic conflict described in the cross-sectional political economy literature. Both increasing migrant labor size and diversity predict higher rates of conflict in receiving cities. Unlike past experimental work, I don't find evidence that increasing worker diversity lowers wages or output per worker in a panel of firms, or that migrant diversity affects the average skill level at destination. In a final section of the paper, I explore the possible long-run benefits of migrant diversity suggested by the positive impact on non-agricultural labor share. I leverage my constructed historical agglomeration pull factors in a long-run version of the empirical strategy to explore the cross-sectional relationship of diversity, population and productivity across African cities. I find that cities that happened to be located in more diverse areas due to exogenous factors benefitted more from historical productivity shocks in the long-run. I conclude that migrant diversity brings short-term costs, but may play a beneficial role in long-run structural change.

This paper makes a contribution to the identification of linear and nonlinear effects of migration. The literature on developed country migration regularly quantifies the effects of migrant flows on growth, often parsing results by skill level. Recently, researchers have begun to consider the importance of broader categories of worker type in explaining the heterogeneous effects of immigration on destinations. Alesina et al. (2016) directly studies the role of birth country diversity on destination outcomes using a gravity model to predict both the share and diversity of immigrants. They find that immigrant diversity is positively associated with measures of TFP and patent intensity.

However, their estimation strategy does not consider the simultaneous role of migrant labor size, and risks conflating the effect of birthplace diversity with migrant labor size. A more recent literature has sought to microfound the role of particular ethnic or national groups in productivity growth through knowledge spillovers or home-country connections (Boberg-Fazlić and Sharp, 2024; Choi et al., 2024; Burchardi et al., 2019; Imbert et al., 2022). Another strand of work has studied the impact of ethnic enclaves and different assimilation rates on the migrant labor market (Albert et al., 2021). Assimilation rates, ethnic enclaves and knowledge-spillovers are all a function of the composition of migrants. My paper develops a framework to aggregate these ideas into a shift-share design that can disentangle the effects of pure increases in labor size from compositional changes in migrant worker types.

In developing contexts, a few papers have considered the direct effect of birthplace-specific human capital on destination productivity (Bazzi et al., 2016). Most work on ethnicity and migration in development has been interested in estimating linguistic or cultural distance as a migration cost, rather than evaluating the aggregate effects on destinations (Wang, 2024). In political economy, studies of African diversity have been mostly correlational or leveraging long-run instruments (Alesina and Ferrara, 2005; Alesina et al., 2003; Arbatli et al., 2020; Ashraf and Galor, 2013). This work has generally found that more linguistically or culturally diverse countries see higher conflict, lower public goods and lower GDP growth. At subnational levels, Montalvo and Reynal-Querol (2021) find this correlation flips direction, hinting at the idea that agglomeration benefits and diversity play joint roles in city-level outcomes. This correlational work has suffered from two key empirical concerns: (1) historical measures of local diversity are correlated with local productivity factors and geography. The distribution of ethnic groups across space is related to many unobservable productivity fundamentals (Michalopoulos, 2012). (2) Diversity is intimately related to population size – historically diverse places tend to also have relatively more people. Montalvo and Reynal-Querol (2021) note that many historical trade centers, now larger cities, were founded at the intersection of ethnic boundaries. My approach addresses these concerns first by estimating a differenced equation that leverages changes in population and diversity within destinations over

time. Second, I design two shift-share instruments to separately identify the effects of migrant labor size and birthplace diversity.

This paper also speaks directly to the literature on the returns to migration. While this literature is large, work on Africa in particular is relatively sparse due to data gaps. Few surveys capture internal migrant flows, or even area of previous residence. Most results in African settings have relied on household panels that track individuals over several waves, or cross-sectional data with simple markers of migrant status. Young (2013) and Gollin et al. (2021) use Demographic Health Surveys (DHS) to quantify large differences in consumption and amenities between rural and urban areas. Hamory et al. (2021) and Lagakos et al. (2020) use individual fixed effects in household panels to compare migrant returns before and after migrating from rural to urban areas, and find relatively lower estimates of urban premia. Related work looks for evidence of agglomeration forces, showing increasing returns to density across African cities (Henderson et al., 2021; Castells-Quintana, 2017).

Across these papers migrants are identified in coarse categories as coming from a “city”, “town” or “rural area”. The data is either not large enough or not detailed enough to consider origin-destination pairs, making it hard to estimate migration costs. Relatedly, while these papers consider heterogeneous migration costs across origins, which may include linguistic differences, there’s no conception of complementarities between migrants of different origins. This paper contributes in two ways. First, I use a combination of birthplace data and ethnic linking methods to tie workers at destination to specific regional origins. This novel proxied origin-destination panel is able to answer more specific questions about the role of particular origins on migration returns. I show that my proxied origin-destination migration panel can replicate key patterns in standard migration gravity models and estimates of migration returns. Second I bring shift-share methods to an African setting, allowing me to calculate a density premium by leveraging exogenous shifts in migration levels, rather than individual returns.

Finally, I contribute to a separate literature that considers the role of shocks on migration. My paper uses climate, commodity price and conflict shocks as plausibly exogenous instruments for shifts in outmigration. Kamukeni and Naito (2024) and Henderson et al. (2017) find drier

conditions in African countries increase the rate of internal migration and urbanization respectively. Henderson et al. hypothesize that cities that specialize in non-agricultural production provide a safe-haven for rural migrants. McGuirk and Nunn (2024) links drying conditions with increased conflict in Africa between pastoralists and farmers. Few papers study the effects of price shocks on migration levels or composition in African countries (Brückner, 2012). Gollin et al. (2016) considers the role of commodity exports on urbanization, finding that higher value exports also drive urbanization at the country level. My work most relates to a separate thread in the political economy of conflict that identifies the association of granular price shocks of agriculture and minerals with conflict (Bazzi and Blattman, 2014; McGuirk and Burke, 2020). This paper is a more granular study of push shock dynamics, using origin-destination decisions and subnational exposure to shocks as a first-stage to predict migrant flows.

This paper brings a novel perspective to the impact of shocks on migration. Suppose that where migrants come from is an important variable in the production function of a destination. If both levels and composition of migrants matter, then the level and composition of shocks matters as well. An intense climate or price shock that affects a singular region of a country may drive many migrants that are relatively homogeneous in terms of language or skill. A shock that is broader, affecting many locations at different intensities will create a more diverse flow of outward migrants. I show that different push-shocks contribute differently to migrant flows, affecting both levels and composition simultaneously.

A key innovation of this paper is the use of nonlinear functions of shift-share instruments to capture nonlinear effects of migration. The methods used to disentangle size and composition of migrant flows can be applied to a number of other contexts where the impact of flows have linear and nonlinear components, such as in the study of skill-complementarity, assimilation, segregation or industry concentration (Lewis, 2011). Immigration shocks do not only cause changes in population levels or average skill level. The relative sizes of immigrant populations can affect the rate at which group assimilate, the formation of enclaves, rates of crime and conflict and the organization of firms.

1.2 Data

An ideal dataset to study the aggregate impacts of migration would capture changes in migrant flows by destination and origin over time. Few African data sources contain granular migration flows within or across countries. To measure changes in migrant composition, I construct a proxy for origin-destination panel data by linking people at a destination to their origin birthplace or ethnic homeland. The main analysis uses the subset of available African censuses that include subnational information on household location and either birthplace, ethnicity or mother tongue. The spatial granularity is at the second administrative level, which corresponds to counties or districts. While the main tables use censuses to capture changes in flows over time, I supplement this evidence in different parts of the paper with geolocated Demographic Health Surveys (DHS) and Afrobarometer Surveys. These cross-sectional household surveys collect data on fewer households but contain more variables and years. Both DHS and Afrobarometer contain information on workers' ethnicity, mother tongue or tribe. Using the ethnic linking methods described below, I can identify the spatial origins of individuals in DHS and Afrobarometer, allowing me to construct a panel at the sample cluster to ethnic homeland pair level.

In the section below, I discuss the construction of the origin-destination panel data. Next I discuss the data inputs used in producing exogeneity in the outmigration shifts and settlement shares. "Push shocks" are used to predict changes in outmigration rates for different origins. The variables used for outmigration shocks include drought events, agricultural and mineral price shocks, and geolocated conflict events. "Pull characteristics" are used to predict settlement shares. These are characteristics that are correlated with early settlement in a destination, but are unrelated to contemporary labor demand shocks. I use colonial rail locations, mineral deposits, and river characteristics to predict which destinations attracted migrants historically.

Next I discuss the productivity outcomes used to evaluate changes at migrant destinations. I use a variety of functions of night light density that have been used in past work on urban productivity in developing countries (Chiovelli et al., 2023). I also leverage census data that captures measures

of industry-specific labor shares, and individual's housing characteristics as a wealth proxy. The last part of this section explores observational returns to migration using the proxy origin-destination dataset, and finds patterns that are consistent with more standard origin-destination panels of migration. A concern with using ethnic or language data to infer migrant origins is that changes over time may capture changes in fertility at destination, rather than real migration flows from the origin to destination. The observational returns serve as a sanity test that provides evidence that my approach is capturing migration changes rather than relative growth rates.

1.2.1 Proxied Origin-Destination Panels

1.2.1.1 Census data

When available, African censuses are taken from IPUMS. In a few cases I supplement with census data from the World Bank microdata portal or government statistical websites. Censuses are typically spaced about 10 years apart. The majority of the censuses are 10% random samples. The census years cover the period between 1970 and 2020. The main analysis will focus on the period 1990 - 2020, which is the window of time in which satellite data of nighttime light density is available.

A subset of censuses record a worker's birthplace at the second administrative level. I link each individual at a given destination and year to their birthplace, which I call their "origin". Because subnational birthplace is reported for migrants moving within country, this paper focuses on the effects of internal migrants. Aggregating the data to the origin-destination-year level, I can observe decade level shifts in the number of people from a given birthplace, at a given destination. This panel is unbalanced, as each country completed censuses in different years. For a given country, the unit of analysis is changes in migrant labor and composition across census years for a given destination-year.

I supplement the birthplace data using an ethnic linking strategy. In some censuses workers report their ethnicity or native language, but not their birthplace. Because ethnic groups were

historically organized into regional territories, I can use this information to connect them to their ethnic homeland of origin. I use the Linking Ethnographic Database (LEDA) by Müller-Crepon et al. (2022) to map each reported ethnicity to an ethnic homeland region on the Murdock Map, a common ethnographic resource used in African political economy. This linking procedure allows me to match a worker to a spatial origin or “homeland”. Figure 1.10 shows an example of the groups plotted in the Murdock map, most of which resemble counties in size. I can run the same analysis using these ethnic homelands as our “origins”, while destinations are still reported at the second administrative level. See Table 1.14 for a list of the censuses used in the analysis, including both samples with birthplace or ethnicity-linked origin data.

1.2.1.2 DHS and Afrobarometer Panel data

The shift-share results are based on the constructed panel of censuses. To supplement some analyses, I use standard geolocated household surveys. The Demographic Health Surveys (DHS) and Afrobarometer are household surveys with detailed information on assets, housing quality, identity and the geolocation of households (BenYishay et al., 2017; Boyle et al., 2024). Afrobarometer includes various questions about ethnic and economic attitudes. These variables are useful in understanding city-level and individual-level outcomes over shorter time horizons. What the surveys lack is detailed information on migration behavior ¹. Since most of the surveys ask about ethnicity or home language, I can use the same LEDA linking procedure as described above for censuses. Leveraging the ethnic information in DHS and Afrobarometer as an indicator of migrant origin allows me to study migration with a detailed origin subnational unit.

1.2.2 Push Shocks to Outmigration

In an exercise to isolate exogeneity of shifts in our shift-share design, I consider a variety of shocks that may push migrants from origins. Past empirical work on migration has used instruments

¹A subset of DHS surveys ask individuals about their region of previous residence. Typically these regions are listed at the first administrative level such as provinces, or may be as broad as identifying North vs. North-West of the country. Ethnic homelands are a more granular and consistent spatial unit.

for income shocks that affect an origin region's outmigration rate but are unrelated to productivity changes at destination. I explore several potential candidates for push shocks at origin based on past work. These include climate shocks, international commodity prices and conflict events. While all these subnational shocks have been used in other contexts, my paper is among the first to study these shocks in the context of an origin-destination migration panel for Africa.

1.2.2.1 Climate

A number of recent papers have studied climate change effects on migration in developing countries (Desmet and Rossi-Hansberg, 2024; McGuirk and Nunn, 2024; Kamuikeni and Naito, 2024). The hypothesis guiding this work is that drier conditions create negative income shocks in agricultural areas, which then affects migration decisions. For instance, a major drought in East Africa in 2011 forced a reported 920,000 people from Somalia to flee to neighboring countries, many of whom settled in cities like Nairobi. The main climatic variable of interest is a measure of drought conditions. I use data from the Standardised Precipitation-Evapotranspiration Index (SPEI), which measures drought intensity monthly by combining temperature and precipitation data (Vicente-Serrano et al., 2010). The data provides monthly estimates of drought intensity at a 0.5x0.5 degree cell resolution from 1900-2022, which I aggregate into yearly estimates. For each subnational unit, I calculate the average drought experience over time. A month-cell is under normal conditions when the SPEI index is around 0, which means there is balance between the precipitation rate and the potential evapotranspiration. I code the month as "in drought" if the index is a standard deviation lower than zero. This is the threshold for extreme dryness suggested by the index (Vicente-Serrano et al., 2010). For each region, I calculate a yearly drought index from monthly dummy variables that signal whether the month was a drought based on the SPEI (ie. the share of year in drought). As secondary climate variables, I also consider rainfall and temperature separately.

1.2.2.2 Agricultural and Mineral Commodities

Many rural areas in Africa produce commodities for export, including cash crops like maize and minerals like gold or diamonds. As international prices fluctuate, producers in different regions are differentially exposed to these potential income shocks. Migration may then respond positively or negatively to exogenous changes in international commodity prices, weighted by local commodity-specific exposure. To construct a subnational measure of price exposure for key commodities, I use data on both the share of production in commodities and international prices. I construct a yearly time series of the global market prices of key commodities spanning the years 1960-2024. I start from a price list assembled by (Bazzi and Blattman, 2014) which tracks major agricultural and mineral commodity prices from 1960-2009. I then manually extend this list using available data from the IMF Primary Commodities price system, and the World Bank Pink Sheet (WB, 2025; IMF, 2025). For mineral prices, I mainly rely on the US Geological Survey's (USGS) "Historical Statistics for Mineral and Material Commodities", which covers US prices of major minerals back to 1900 (Kelly et al., 2010).

I create a measure of subnational exposure to various agricultural commodities using FAO Global Agro-Ecological Zones (GAEZ) production maps, which estimate the cell-level average hectares dedicated to a set of major commodity crops across the continent (Berman and Couttenier, 2015). These production areas are estimated in the year 2000, and are used as baseline exposure. Product prices are normalized to 100 in 2000 and summed with weights by the hectare area. Exposure is weighted by the share of productive hectares dedicated to that crop. For each crop product p grown in region o :

$$AgriculturalPriceExposure_{ot} = \sum_p Price_t * HectareShare_{op} \quad (1.1)$$

To account for price shocks over longer time horizons, I construct different measures of

lagged price exposure. A 10 year lagged exposure is an average of price exposure between censuses:

$$PriceShock_{ot} = \frac{1}{10} \sum_{t-10}^t PriceExposure_{ot} \quad (1.2)$$

For mineral commodities, I weight exposure to given minerals using the “USGS Compilation of Geospatial Data (GIS) for the Mineral Industries and Related Infrastructure of Africa” (Kelly et al., 2010). This dataset contains geolocated mineral facilities and their estimated capacity. For each mineral, I measure the total capacity across countries. Then each mineral producing region is exposed to a given mineral price according to its relative share of total production capacity in that mineral. The total mineral price exposure for a region is then:

$$MineralPriceExposure_{ot} = \sum_m Price_t * RelativeCapacity_{om} \quad (1.3)$$


Where relative capacity in mineral m for region o is given by:

$$RelativeCapacity_{om} = \frac{Capacity_{om}}{Capacity_m} \quad (1.4)$$

1.2.2.3 Conflict

The Armed Conflict Location & Event Data (ACLED) provides geolocated conflict events across Africa² The ACLED data covers events between 1997 and 2025. The data is based on the manual coding of news articles, and provides basic information on each event including approximate dates, estimated death toll, and the types of actors involved. For each region, I aggregate all conflict events that appear in that region to produce an annual count of conflict events. Conflict intensity in a given year is the average number of conflict events per month that year. As a secondary outcome, I produce an index of average conflict fatalities in a given region-year.

²Different events have different degrees of confidence in the precise geolocation. I only include events where the researchers have marked having confidence in the geolocation at the second administrative level.



Notes: This figure shows examples of the data from one census pair, specifically Mozambique 1997 and 2007. Panel A shows the change in migrant flows to Maputo city from each birthplace between 1997 and 2007. Panel B shows the change in the weighted agricultural price exposure between 1997 and 2007. Panel C shows the average drought intensity by region between 1997 and 2007. Panel D shows the average annual conflict events in each region over this time period. Data for the change in migrant in-flows comes from the publicly available African censuses on IPUMS International. Data for agricultural price exposure is taken from IMF Primary Commodities and the World Bank Pink Sheets, combined with FAO production exposure. Data for drought exposure is calculated from an SPEI index (Vicente-Serrano et al., 2010), and Conflict Exposure data is from ACLED.

Figure 1.1: Data Examples from Mozambique 1997-2007

1.2.3 Pull Characteristics to Predict Immigration

Traditional migration shift-share instruments rely on endogenous settlement shares of different immigrant groups from past years. If settlement shares are serially correlated over time, it's possible that the settlement shares are endogenously related to contemporary labor demand shocks. For robustness, I'll also consider a group-specific settlement share that is instrumented by a collection of destination and origin-destination characteristics. In particular I instrument for past settlement using the interaction of the distance between origin-destination pairs and a set of "pull characteristics". These pull characteristics are designed to predict the attractiveness of destinations but are plausibly uncorrelated with current period demand shocks. I consider several historical shocks that predict urban formation, including the presence of mineral deposits, colonial railroads and portage sites. Data for mineral deposits is taken from the "USGS Compilation of Geospatial Data (GIS) for the Mineral Industries and Related Infrastructure of Africa". Data on colonial railroad projects is taken from the universe of colonial rail projects collected by Jedwab and Moradi (2016). Last, I construct an instrument for portage leveraging data on river flow from the HydroSHEDS database (Lehner and Grill, 2013). The details of these instrument constructions are described in the empirical strategy.

Last, I leverage a set of common geographical characteristics as controls in cross-sectional regressions. These include the ruggedness index from Nunn and Puga (2012), malaria suitability from Kiszewski et al. (2004), Tse Tse suitability from Alsan (2015) and soil suitability from Ramankutty et al. (2002).

1.2.4 Destination Productivity Outcomes

I use a variety of subnational proxies for regional productivity. I consider two functions of changes in light density as measures of urban extent and GDP per-capita growth. I also leverage census-level measures that capture structural transformation and urbanization characteristics including non-agricultural and services labor shares. Most censuses also include information on

housing quality, such as the material of the roof and floor of the surveyed household. I construct a principal component measure of these material characteristics as a proxy for household wealth. All outcomes are measured in terms of differences between census years.

1.2.4.1 Light Density

Nighttime light density is a commonly used proxy for economic development, especially at subnational levels. A variety of papers have shown correlations between measures of luminosity and other measures of wealth from either surveys of household income or administrative level human capital and non-agricultural labor share (Michalopoulos and Papaioannou, 2013; Chiovelli et al., 2023). There are three issues that must be dealt with to use luminosity as a proxy for wages or productivity. First, the light density data produced before 2013 and after 2013 are at different resolutions, which complicates comparisons over time³. I use a harmonized dataset produced by Li et al. (2020) that performs an inter-calibration to combine the datasets and generates a DN-style output at 30 arc-seconds which ranges from 0 to 63.

The second issue is dealing with blooming, where high light density in a given cell may bleed into neighboring cells producing unwanted spillovers, say from a large city to a small neighboring town. The harmonized dataset partially adjusts for this, and administrative regions are large enough to somewhat mitigate this concern. Additionally, in the robustness analysis I use spatially correlated standard errors to address contamination across regions.

The third issue is how to construct an appropriate measure of economic growth from the pixel-level light data. Different papers use different constructions. Chiovelli et al. (2023) aggregate pixels by summing over their regions of interest, and then including controls in their regression for log population and region area. Montalvo and Reynal-Querol (2021) use changes in log light density divided by a gridded population estimate as a measure of local economic growth. I consider

³Before 2013, satellite imaging comes from the Defense Meteorological Satellite Program (DMSP)/Operational Linescan System (OLS), which provides digital number (DN) values at a resolution of 30 arc-seconds. After 2013, the Visible Infrared Imaging Radiometer Suite (VIIRS) on the Suomi National Polar-orbiting Partnership satellite provides light density data at 15 arc-seconds, and at a higher radiometric resolution. This means it can detect smaller differences in light density relative to the past technology.

two light density constructions in differences to capture the size and growth of destinations. The first is the standardized level change in average light density in a given administrative region. This measure captures a change in light density levels, and can be thought of as a proxy for city size growth or extent. The second measure is the change in the log of light density over population. This measure captures the proportional change in light density per capita, and is used as a proxy of GDP per capita. As cities approach the maximum value of measured light density at the pixel level, there may be frontier effects. Large productive cities may show low light density growth because there is no further improvement to be made. I try to mitigate these concerns by controlling for pre-period light density in the differenced regressions. I also consider separate measures of productivity from asset data and industry shares.

$$\Delta \text{Log}(\text{Lights}/\text{Capita})_t = \text{Log}(\text{Lights}/\text{Capita})_t - \text{Log}(\text{Lights}/\text{Capita})_{t-1} \quad (1.5)$$

While past work has shown a correlation of night light density with household level proxies for subnational GDP in Africa, less is known about the relationship between light density and industrial or service sector activity. In the Appendix I make an effort to provide evidence of the relationship between city-level economic growth, proxied by port activity, and night light density. Figure 1.28 in the Appendix documents a positive relationship between shipping traffic levels estimated at African ports and monthly and annual aggregations of light density. The figure also shows suggestive evidence of possible frontier effects, as particularly high traffic port-year cells have an attenuating relationship with marginal increases in light density.

1.2.4.2 Assets and House Quality

I follow the spirit of Young (2013) using the ownership of durables and housing conditions as proxies for household wealth. DHS surveys include a rich set of variables that describe ownership of different durable assets like a TV, bicycle, car or microwave. I use these dummy indicators for

assets to construct a wealth score as a first principal component of these assets. This procedure is similar to the DHS's own measured wealth score, which leverages a first principal aggregation of asset categories. I only include the subset of assets available in all surveys, to ensure comparability across survey samples. In the census data, assets are more sparse. However, most censuses include information on the house, wall and roof materials. Leveraging this fact, I replicate the procedure for DHS estimations of wealth by taking a first principal component of a set of house quality measures for each census where it's available. Averaging these estimates across individuals in a district gives us a subnational measure of housing quality. Table 1.1 shows that these principal component measures are positively correlated with education and age, suggesting that the method accurately captures relative wealth.

1.2.4.3 Industry and Urbanization

As Gollin et al. (2016) highlights, growth in cities can take place in both tradable and non-tradable sectors. A common refrain about African cities is that they seem to grow without an expansion in the manufacturing sector; so-called "consumption cities" (Jedwab et al., 2025). I use two measures of urbanization growth as outcomes: the total non-agricultural labor share and the non-tradable services labor share. Each outcome is derived from individual census data that report general industry classifications such as agriculture, mining, wholesale trade, manufacturing, financial services, etc. These industries are aggregated into measures designed to reflect growing urbanization – as cities grow, we expect them to increase their non-agricultural share. For cities experiencing service-led growth rather than manufacturing growth, we expect "services share" to increase.

These outcomes are considered in terms of differences across census years. This is particularly important in an African data context because informal work is often underreported in census occupational data, even though it makes up a significant fraction of the labor. The industry measures used capture changes in the industry shares among the subset of individuals who report formal employment. The benefit of our differenced equation design is that baseline differences in

formalization rates are differenced out.

1.2.5 Validating a Proxy Origin-Destination Panel

Using origin-destination panel data from ethnicity and birthplace differs in significant ways from typical origin-destination panels. The data does not capture flows of migrants over time, but rather the stock of individuals from particular birthplaces or homelands for a set of unbalanced years. Observing an increase in people from a given birthplace o in a destination d may reflect increased migrant flow into the destination, or decreased return migration to o . Observing an increase in people from a given ethnicity may also reflect differential changes in the population growth of that ethnicity, rather than changes in migration from an ethnic homeland.

As a first exercise I estimate observational returns to migration, and compare the results to other cross-sectional estimates as a validation exercise. I use DHS panel data, where individuals are linked to ethnic homelands. I use DHS panel for this exercise because of the detailed catalogue of individual assets, which can be used to construct a proxy for migrant wealth as an outcome. Individuals are identified as “migrants” if they are residing in an administrative region outside of their ethnic homeland. Leveraging the DHS’s survey of individual assets, I construct a “durable assets” and “house quality” score. The procedure for doing this follows closely the DHS’s own method for estimating wealth, transforming each asset category into a dummy variable and calculating a first principal component measure. The outcome “Durables” is the first principal component of several binary variables for different assets, including electricity, phone, car, fridge, television and bicycle. The outcome “House Quality” is the first principal component of reported wall, roof and floor material of the house. Using the constructed principal component measures of durable assets and housing quality, I compare consumption between migrants to non-migrants as:

$$\begin{aligned} Wealth_{it} = & \beta_1 \text{Migrant}_i + \beta_2 \text{Migrant}_i \cdot \text{Distance}_{od} \\ & + \beta_3 \text{Migrant}_i \cdot \text{CoethnicShare}_{od} + Z_i + X_{od} + W_d + v_{st} + \gamma_o + \epsilon_{it} \end{aligned} \quad (1.6)$$

Where $Wealth_{it}$ is the wealth score for individual i in survey year t . $Dist_{od}$ is the log distance between the ethnic homeland and the destination, $Migrant_i$ is a dummy for migrant status, and $CoethShare_{od}$ reflects the fraction of individuals in d that are from ethnic homeland o . I include fixed effects for country-year ν_{st} , and γ_o to isolate variation within an ethnic group. The controls for the individual Z_i include age and schooling, while X_{od} includes the level of o-d distance and coethnic share, and W_d includes destination log population. Wealth is a measure of either durable assets or housing quality. Subsequently, I compare outcomes of migrants to natives within a destination by replacing the ethnicity fixed effect with a destination fixed effect. Finally, I isolate the sample to only migrants, and compare the outcomes of migrants within a destination from different homelands.

Table 1.1 shows the results for different variations of equation 1.6 with different fixed effects. While the average consumption benefit to migrants is weakly positive in terms of assets, columns 1-3 suggest that migrants that travel further have higher returns. This is consistent with a model of migration with heterogeneous costs in which the migrants that choose to pay high migration costs are wealthier ex-ante, or have a higher productivity draw for a particular location (Lagakos et al., 2020). Coethnic share is consistently negative across specifications, suggesting that migrants gain better returns in more diverse destinations. This is consistent with the findings of Wang (2024), which finds that migrants that move to more culturally distant destinations have higher returns in an Indonesian sample. As large productive cities are more diverse, this finding is consistent with the idea that migrants move to productive, cosmopolitan destinations. I also find that returns to migration are decreasing in destination population size, conditional on travel distance and coethnic share. Henderson et al. (2021) find similar negative effects of population scale on household incomes in their cross-sectional African data, consistent with a story about the negative effects of sprawl and slum formation on migrant returns. The modest to negative average effects of migrant status on wealth offer a preview of the paper's baseline results, where I find no per-capita growth premium from increased migrant labor size.

As an additional check, I also run a standard gravity model relating the share of migrants from origin o in d using a Poisson pseudo-likelihood (PPML) model. Table 1.15 shows the relation

of the share of migrants arriving in d to their travel distance and coethnic share. Consistent with past work on migration, I find that migration between origin-destination pairs in my panel is decreasing in distance and coethnic share, as individuals sort towards areas that are close by and majority coethnic. This finding is consistent with standard models of migration with linear migration costs in distance.

1.3 Empirical Strategy

The goal of the paper is to estimate the effects of changes in migration levels and composition on destination outcomes. Both the size of the migrant labor force and the diversity of migrants are likely endogenously determined by a destination's productivity. I use a shift-share instrumental variable strategy to overcome this identification challenge. The IV strategy adds a nonlinear component to a standard migration shift-share to simultaneously instrument for migrant size and composition. In the main analysis, the *shifts* are birthplace level outmigration rates, while the *shares* are origin-specific settlement shares at destination. Changes in migrant size and diversity are predicted as linear and nonlinear combinations of these shifts and shares. In a secondary analysis, I then isolate exogenous variation in the outmigration shifts using birthplace-level push shocks that move migrants in or out of their birthplace over time. Third I use destination-specific pull characteristics to isolate exogenous variation in settlement shares.

The strategy starts from an equation of interest that relates a destination's productivity to both migrant labor size and diversity. For a given destination d in year t , the equation of interest is:

$$y_{d,t} = \beta_1 l_{d,t} + \beta_2 div_{d,t} + v_t + \gamma_d + \epsilon_{d,t} \quad (1.7)$$

where $y_{d,t}$ is a proxy for log wages, or another city-level outcome. The logged migrant labor force size is $l_{d,t}$, and the diversity of the migrant labor force is represented by $div_{d,t}$. The diversity of the labor force is a measure of the distribution of groups within the labor force. I use a Herfindahl measure of the concentration of groups. For exposition, I specifically use the negative of a Herfindahl

Table 1.1: Observational Returns by Migrant Status and Distance

	(1) Within Ethnicity		(2) Within Ethnicity		(3) Within Destination		(4) House Quality		(5) Durables		(6) Migrants Only	
	Durables	House Quality	Durables	House Quality	Durables	House Quality	Durables	House Quality	Durables	House Quality	Durables	House Quality
Migrant==1	0.152 [0.079]*	-0.121 [0.083]	0.059 [0.080]	0.125 [0.082]								
Migrant*Population	-0.050 [0.008]***	-0.032 [0.008]***	-0.015 [0.008]*	0.000 [0.008]								
Migrant*Distance	0.064 [0.011]***	0.098 [0.012]***	0.024 [0.012]**	-0.009 [0.013]								
Migrant*CoethnicShare	-0.251 [0.033]***	-0.208 [0.035]***	-0.221 [0.038]***	-0.270 [0.042]***								
ln(O-D Distance km)	0.024 [0.008]***	-0.026 [0.008]***	0.030 [0.010]***	0.027 [0.011]**			0.064 [0.008]***	0.028 [0.009]***				
ln(Population)	0.227 [0.008]***	0.203 [0.008]***										
Age	0.007 [0.000]***	0.005 [0.000]***	0.006 [0.000]***	0.004 [0.000]***			0.006 [0.000]***	0.004 [0.000]***				
School Years	0.101 [0.001]***	0.075 [0.001]***	0.081 [0.001]***	0.058 [0.001]***			0.082 [0.001]***	0.057 [0.001]***				
Coethnic Share	-0.163 [0.029]***	-0.115 [0.031]***	-0.018 [0.033]	0.041 [0.035]			-0.263 [0.022]***	-0.254 [0.026]***				
Mean Dep. Var	-0.017	-0.017	-0.017	-0.017			0.021	0.007				
Observations	359,411	395,543	359,401	395,529			183,244	199,415				
Destination FE	N	N	Y	Y			Y	Y				
Ethnicity FE	Y	Y	N	N			N	N				
Migrant Only	N	N	N	N			Y	Y				

Notes: This table estimates observational returns to migrations from different model specifications. The outcome "Durables" is the first principal component of several binary variables for different assets, including electricity, phone, car, fridge, television and bicycle. The outcome "House Quality" is the first principal component of reported wall, roof and floor material of the house. Origins are ethnic homelands, and destinations are administrative units at the second level. The first two columns include ethnic group fixed effects to isolate variation within ethnic group across destination and migrant status. Columns 3-6 include destination fixed effects to isolate variation within destination across migrant status. All regressions include country-year fixed effects. Standard errors are clustered at the DHS sampling cluster level. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

index as the measure of “migrant diversity”, which is the inverse of the concentration of birthplaces. For workers coming from a set of birthplaces $o \in O$ and arriving in destination d , I calculate div_{dt} as:


$$-div_{dt} = \sum_o \pi_{odt}^2 \quad (1.8)$$

where π_{odt} is share of people in d at t from o $\frac{count_{od}}{count_d}$. As this number moves from 0 to 1, the composition of migrants moves from totally homogeneous to perfectly heterogeneous in terms of birthplaces. Scholars have used a variety of other indices for diversity, including fractionalization, which is the inverse of a Herfindahl Index, and polarization, which is more sensitive to large, equally sized groups. I choose to use the HHI because it is the simplest functional form of diversity to integrate into a shift-share design, and is highly correlated with other popular measures.

Theoretically we expect workers to sort into cities that offer higher wages (Combes et al., 2010). Therefore productive cities should attract a higher number of workers from a greater set of origins, creating a correlation between higher y_{dt} , higher l_{dt} and higher div_{dt} . I address the endogeneity of migrant flows and destination outcomes in the sections that follow. Section 3.1 describes a shift-share instrument that can incorporate nonlinear functions of migration like the diversity HHI. Then, I describe the exclusion restriction of the instrument in section 3.2. Sections 3.3 and 3.4 discuss extensions using pull and push shocks to predict outmigration shifts and pre-period settlement shares.

1.3.1 A Shift-Share Instrument for Migrant Birthplace Diversity

The goal is to develop a shift-share instrument for both the level of migration and its diversity, as measured by a birthplace HHI. I take inspiration from Schubert et al. (2024), which studies the effect of firm concentration on wages using a firm-level Bartik shock as an instrument for the change in employer labor market concentration. Similar to my setting, Schubert and co-authors instrument for an HHI variable as a function of Bartik style shifts and shares. I adapt this approach to the construction of two simultaneous instruments in a migration setting. In the case of migrant labor, the



(a) Worker Diversity vs. Log Total Labor



(b) Migrant Diversity vs. Log Migrant Labor

Notes: This figure shows scatterplots relating diversity and labor at the destination-year level. The left panel plots total worker diversity against total log labor size, while the right panel looks at migrant diversity against migrant log labor size. The blue line is a line of best fit across all points, while the red line is a local polynomial fit with confidence bands.

Figure 1.2: Relationship between Diversity (HHI) and Labor Variables

instrument is the aggregation of birthplace outmigration shifts and destination settlement shares. For example a given destination like Kigali Rwanda has a certain fraction of all migrants from birthplace Kirehe, in the East. The change in migrant labor from Kirehe is a combination of pre-period settlement shares of Kirehe migrants in Kigali, combined with shifts in the aggregate migration rate from Kirehe to all destinations. The change in the total number of people from Kirehe living outside Kirehe across census years constitutes a shift. To instrument for diversity, the destination settlement shares are weighted by the squared relative growth rates of each birthplace's migrant flow. The key is that this aggregation is a nonlinear combination of the individual shift-shares. The nonlinearity of the function mapping the shift-shares to aggregate diversity allows for the same set of shift-shares to predict both endogenous variables.

I begin by differencing equation 1.7 between censuses, which are typically 10 years apart in our data:

$$\Delta y_{d,t-(t-10)} = \beta_1 \Delta \ell_{d,t-(t-10)} + \beta_2 \Delta div_{d,t-(t-10)} + \epsilon_{d,t} \quad (1.9)$$

I instrument for $\Delta \ell_{dt}$ and Δdiv_{dt} using decade level shifts from different ethnic homelands and baseline ethnic shares.

Each origin o has a number of out-of-homeland migrants $count_{o,t}$, which I define as the

aggregate number of people from o observed outside of that homeland in a given census year. Our shifts are the growth rates in the number of people from o observed outside of their origin, which we define as $g_{o,t} = \frac{count_{o,t} - count_{o,t-10}}{count_{o,t-10}}$.

Each destination d is exposed to these shifts weighted by their baseline settlement share of people from that origin in the previous census period. For a given year these settlement shares are the fraction of people from o in d relative to the full migrant population from o : $Shares_{odt} = \frac{count_{odt}}{\sum_d count_{odt}}$

We can predict the change in labor demand in the region d as the sum of the shifts and shares:

$$\Delta \hat{L}_{dt} = \sum_o \frac{count_{od,t-10}}{\sum_d count_{od,t-10}} * g_{o,t} * count_{o,t-10} \quad (1.10)$$

In the analysis, we take a log of this predicted value as our measure $\Delta \hat{\ell}$.

To predict a change in diversity, I start from a definition of the change in birthplace concentration. Following the HHI formulation above, change in the relative concentration of migrants from different o living in destination d is:

$$-\Delta div_{d,t} = \sum_o \pi_{odt}^2 - \sum_o \pi_{od,t-10}^2 \quad (1.11)$$

This equation can be written as a function of the initial concentration in the base period, and the respective growth rates of migrants from o to d , and the total change in the migrant labor force at d . The equation becomes:

$$-\Delta div_{d,t} = \sum_o \pi_{od,t-10}^2 \left(\frac{(1 + g_{odt})^2}{(1 + g_{dt})^2} - 1 \right) \quad (1.12)$$

Where mechanically the change is a function of the growth in the number of people o in d , g_{odt} relative to the total growth of population in d , called g_{dt} . The negative term at the front of the equation reverses the order, such that increasing levels represent lower Herfindahl concentration and higher birthplace diversity.

Both g_{odt} and g_{dt} are likely correlated with contemporary productivity in the destination. An instrument replaces the endogenous current period div_{dt} by substituting aggregate shifts for the

two growth rates in equation 1.12. In particular, our instrument is:


$$-\widehat{\Delta div_{dt}} = \sum_o \pi_{od,t-10}^2 * \left(\frac{(1 + g_{ot})^2}{(1 + \widetilde{g_{odt}})^2} - 1 \right) \quad (1.13)$$

Where g_{ot} is again the aggregate growth of migrants out of an origin and $\widetilde{g_{odt}}$ is a predicted growth of migrant labor in destination d defined as $\sum_o \pi_{od,t-10} * g_{ot}$.

1.3.1.1 Functional Relationship of Composition, Population and Productivity

The estimation strategy leverages two functional forms for an instrument that aggregates shifts g_{ot} and shares π_{odt} across origins. This is possible because of the nonlinear relationship between the migration levels and their composition. Two potential issues emerge from this. First, the linear correlation between migrant levels and composition reduces the power of the estimates. Intuitively, the more related migrant levels and composition are, the less added variation is introduced by the second instrument. I show this in a simulation exercise in Appendix A, in which I calculate a distribution of beta coefficients as the linear correlation between migrant labor size and migrant diversity varies. In Figure 1.2 I plot the cross-sectional relationship between migrant diversity and migrant labor size. In general while the variables are related, the relationship is not strongly linear. This fact allows us to be more confident that migrant level changes do not linearly predict composition changes, and that both instruments provide identifying variation to the empirical estimates.

Second, we don't know ex-ante the functional relationship between city diversity and productivity. Panel B of Figure 1.3 presents the cross-sectional relationship between city diversity and light density in destinations. In the cross-section, I find a positive and linear relationship between worker diversity and light density. Cities that are bigger and less homogeneous are higher income. This cross-sectional result is consistent with the subnational associations studied by Montalvo and Reynal-Querol (2021). In their paper, the authors argue that, while country-level regressions show a negative relationship between higher diversity and light density, at the subnational level this



(a) Migrant Labor Size and Light Density

(b) Migrant Diversity and Light Density

Notes: This figure shows scatterplots relating migrant diversity, migrant labor size and light density at the destination-year level. The left panel plots migrant labor size against contemporary logged light density, while the right panel looks at migrant diversity against contemporaneous log light density. The blue line is a line of best fit across all points, while the red line is a local polynomial fit with confidence bands.

Figure 1.3: Migrant Size and Composition on Light Density

relationship is reversed. Their work leverages historical diversity measures from anthropological maps, rather than census-reported ethnic mix. However, the mechanism that drives their correlation is presumably the same – certain places that are high productivity attract more people, from more diverse sources. The goal of this paper is to move beyond these cross-sectional associations, and study a causally motivated estimate in first differences.

1.3.2 Exclusion Restriction

The exclusion restriction requires that the drivers of aggregate shifts out of a birthplace are not correlated with productivity shifts at particular destinations. For example, this would be violated if productivity gains in certain major cities like Kigali drive aggregate migration trends out of certain birthplaces. Borusyak et al. (2022) describes the exclusion restriction assumption in shift-share designs in which exogeneity comes from the individual shocks. In our setting the shocks are the estimated g_{ot} outmigration rates from origins o in period t . Borusyak and co-authors show that the exclusion restriction in this case is equivalent to an orthogonality condition at the birthplace

level. In particular, given regional exposure weights s_o and a distribution of unobservables ϕ_o :

$$\mathbb{E}[g_{ot}|\phi_o] = \mu \quad (1.14)$$

A second requirement is that the shocks are independent and exposure is dispersed across birthplaces. In particular:

$$\mathbb{E}[(g_{1t} - \mu)(g_{2t} - \mu)|\phi_1, \phi_2] = 0 \quad (1.15)$$

The dispersion condition can be defined as a Herfindahl index that goes to zero as the number of birthplaces increases: $\sum_o s_o^2 \rightarrow 0$.

While I can't directly test these conditions, I perform several tests at the birthplace level. First, I derive shift-share standard errors that address the possible covariance between g_o and ϕ_o following Borusyak et al. (2022). These estimates are presented in Figure 1.18. Second, I check pre-trends by running birthplace-level regressions of the exposure-weighted residuals of destination characteristics on our growth rate shocks. This analysis is produced in Figure 1.20. The destination characteristics I include are geographical characteristics such as disease ecology and soil suitability, as well as past-period light density. I show that my birthplace shocks are uncorrelated with pre-period destination characteristics. Another way to address endogeneity concerns is to directly leverage plausibly exogenous variation in shifts and shares. I describe two potential strategies in the following sections that utilize migration push shocks and pull characteristics.

1.3.3 Predicting Outmigration Shifts with Push Shocks

In the standard migration shift-share, the shift terms g_o are the real change in the number of people leaving origin o across years. The argument for exogeneity comes from an assumption that past shares are unrelated to current period changes in outcomes, except through the labor and composition channels. Since our specification is in differences, this is comparable to a parallel trends assumption in a difference-in-differences design, but with many individual treatment exposures to the aggregate shift treatment (Borusyak et al., 2025; Goldsmith-Pinkham et al., 2020). This may be

violated in cases where changes in particularly large labor markets drive g_{ot} or when country-wide shocks affect both the total outmigration rate and particular destinations (Jaeger et al., 2018). Many migration papers have considered instrumenting for shifts in the outmigration rates from origins using push shocks (Mullins and Bharadwaj, 2021; Boustan et al., 2010; Bazzi et al., 2023; Kamuikeni and Naito, 2024). If the push shocks are plausibly uncorrelated with destination outcomes, then a zero-stage regression of outmigration on push shocks can isolate exogenous variation in migration shifts. I consider the total outmigration from a given origin o and year t as a function of a set of past shocks $shock_{o,t-y}$, where the shocks are measured at some lag y from the current period:

$$g_{ot} = \omega_1 \log(Shock)_{o,t-y} + \nu_o + \delta_s t + \epsilon_{ot} \quad (1.16)$$

I include a linear time trend $\delta_s t$ and origin fixed effects ν_o ⁴. I substitute these predicted outmigration rates $\widehat{g_{ot}}$ for the raw shift g_{ot} in equations 1.10 and 1.13.

Past work on the historic US has leveraged a variety of local economic conditions and weather variables (Boustan et al., 2010; Bazzi et al., 2023). In developing countries, estimates of migration elasticities to income or other shocks vary by context. For example, income shocks have been shown to move migration both positively and negatively depending on the country and time period (Marchiori et al., 2012; Bazzi, 2017; Shrestha, 2017). My strategy is to consider a set of plausible local shocks that may move outmigration. In a zero stage, I consider each shock and its lags separately, and then perform a joint estimation of the relevant shocks to predict the outmigration rate. I consider climate variables, conflict and international commodity price changes, weighted by local exposure.

⁴Another option would be to use year fixed effects in this regression. As noted by Mullins and Bharadwaj (2021), this isolates variation in the shock between counties in the same year. But this specification considers only short term adjustments to weather. Because my migration adjustments are across decades, I substitute the year fixed effects for state specific linear time trends. This allows for detrended variation within origins over time to predict the migration rate.

1.3.4 Predicting Settlement Shares with Pull Characteristics

Even with a predicted flow estimated for g_{ot} , another concern may be that previous period settlement patterns are correlated with current period labor demand shocks in destinations. This may be the case if a given settlement of migrants was driven by a previous labor demand shock, and demand shocks are serially correlated in destination (Jaeger et al., 2018). Instead of using past settlement patterns as exposure shares, I consider replacing the shares with a gravity-model prediction of which destinations migrants are likely to choose. In particular I construct instruments for settlement shares of origins in destinations as a function of the o-d distance between regions and the interaction of distance with a set of “pull characteristics”.

I estimate migrant shares of o in d as

$$\sigma_{odt} = \omega_1 Pull_d * \log(Dist)_{od} + \omega_2 \log(Dist)_{od} + \omega_3 \log(Dist)_{od} * Z_d + \mu_d + \nu_o + \gamma_t + \epsilon_{odt} \quad (1.17)$$

Where σ_{odt} represents the migrant shares $\sigma_{odt} = \frac{count_{odt}}{\sum_d count_{odt}}$. The vector Z_d represents a set of historical destination characteristics that predict the attractiveness of destinations. We can think of these as historical shocks that predict which locations are more predisposed to urban growth. The key is that these pull characteristics Z are not correlated with contemporary shocks to demand. In particular, we use distance to colonial rail lines, mineral deposits and portage sites as predictors of destination attractiveness. We run the same regression to separately estimate π_{odt} for our prediction of diversity change.


1.3.4.1 Historical Pull Characteristics as Predictors of Agglomeration

My pull characteristics draw on historical sources of agglomeration unrelated to contemporary productivity. Colonial railroads were often built to connect coastlines to a particular resource in the interior of the country. An example is the British Uganda railway, which connected Mombasa on the coast to Lake Victoria for geopolitical reasons. This railway incidentally also increased the productivity of regions that lay along the least-cost path between these points. Human settlements


grew everywhere along the railway, and the railway's path within Kenya predicts the location of contemporary Kenyan cities (Jedwab et al., 2017). In Jedwab and Moradi (2016) the authors show that colonial rails continue to predict urban agglomeration, long after the rail lines fell into disuse. I calculate a distance to the nearest colonial rail line as a predictor of a destination's attractiveness, or "agglomeration potential".

I next consider a destination's propensity to be a portage site (Bleakley and Lin, 2012). Maritime trade often requires ships to move inland from the coast along navigable rivers. Sharp changes in elevation along rivers create rapids and waterfalls, preventing large ships from traveling further. It becomes necessary to create infrastructure at the point at which a river is no longer navigable to transfer goods from ships to land transport. Prior work in the US has shown that many US cities developed along the Atlantic Seaboard Fall Line, which creates a point of elevation change at which inland rivers are no longer navigable on the east coast (Bleakley and Lin, 2012). Using the same logic, I create a prediction of rapids and waterfall locations along African rivers. Non-navigable river segments are predicted using HydroSHEDS data on river discharge. In particular I measure the change in river discharge (cubic meters per second) along each river network. I define possible portage sites as places where the size class of a river, which is a logarithmic function of discharge, changes. Each potential portage site is then saved as a point in space, and joined to the administrative census regions. Figure 1.4 shows an example of predicted portage sites along the Congo River, along with a map of the river network and associated cities. We see for example that the mouth of the river in the West shows a high portage propensity near where the Congo River has large rapids that precede the city of Kinshasa.

Last I consider historical mineral deposits as a predictor of past urban growth. This intuition follows Combes et al. (2010) in using geological factors as an instrument for contemporary agglomeration. While deposits should be related to historical roads and infrastructure, they should not be related to contemporary demand shocks. Using interactions of the pull characteristics and origin-destination distance, I instrument for settlement shares for each origin-birthplace pair in a baseline period. I then merge these predicted shares with the outmigration shifts from the baseline



(a) River Network Map



(b) Portage Propensity

Notes: This figure shows an example of the portage score estimation for the portage pull characteristic. The left panel shows a map of the Congo river network, with locations of cities and rapids. On the right I plot my estimated portage points, which are located in bright yellow. All colors reflect nearest distance to these estimated portage sites, which is our portage score.

Figure 1.4: Portage Sites, DRC Example

analysis.

In an additional section, I explore a second use for these historical pull characteristics in estimating long-run dynamics. The main results of the paper consider contemporary changes in diversity and migrant labor across 10-year census periods. To understand the longer term relationship between migrant diversity and urban growth, I need a strategy that can be estimated in a cross-section, and doesn't rely on data that captures differences over census years. The ideal experiment would place urban centers at random locations across Africa, and expose them to different levels of migrant worker diversity. Then I would compare the growth prospects of cities located in more or less dense areas, with more or less worker ethnic diversity. To proxy for this idealized experiment, I run a cross-sectional regression that interacts my historical instruments for agglomeration with historical measures of regional diversity. I use my agglomeration instruments as predictors of city location, and evaluate outcomes in cells that were historically likely to agglomerate. The interaction between the agglomeration instrument and regional diversity measures allows me to compare two cells that were both predicted to agglomerate, but were exposed to different ex-ante levels of worker diversity. The empirical strategy is discussed in greater detail below.

1.4 Results

In this section I will first discuss the results of the nonlinear migration shift-share instrument using standard outmigration shifts and pre-period settlement shares. I will then discuss results from different predictions for shifts and shares using pull and push shocks. I show zero-stage regressions of push shocks predicting outmigration rates at the birthplace level, and then show the results for the first and second-stage SSIV.

In Section 4.1, Table 1.2 and Table 1.3 show the first and second stage results of a regression of outcome in destination d and time t , instrumenting for $\Delta\ell$ and Δdiv :

$$\Delta y_{d,t} = \alpha \widehat{\Delta \ell}_{d,t} + \gamma \widehat{\Delta div}_{d,t} + \epsilon_{d,t} \quad (1.18)$$

I find that level increases in migrant labor size predict increases in city size, as measured by average light density change, but do not correspond to changes in per-capita light density growth. Increased migrant labor causes cities to grow, but does not map to clear effects on either productivity growth or urbanization. As the migrant labor pool becomes more diverse, cities grow slower and have lower productivity in terms of light density changes. However, increasing diversity of migrants predicts higher rates of urbanization, as captured by an increasing non-agricultural labor share. This suggests that migrant composition may have a positive benefit on industry mix, and yield long-term gains even if there are short-run costs. This argument is substantiated in the long-run estimates of diversity and city growth produced in Table 1.12. In the long-run, I find evidence of positive interactions between diverse migrant labor and historical productivity shocks, suggesting that benefits of migrant composition manifest over longer periods.

In section 4.2 I first consider the effect of push shocks as an instrument for the measured birthplace-level shifts g_{ot} . I find that drought predicts lower outmigration, while conflict predicts higher outmigration from birthplace. I find heterogeneous effects of prices across commodity type and time horizon. Leveraging a combination of push shocks to predict g_{ot} , I re-run our 2SLS

regression and find that the first stage succeeds in predicting changes in labor size and composition. I find that the effects of migrant labor size are consistent, and that migrant diversity continues to predict higher urbanization rates.

In section 4.3 I instrument for settlement shares using the interaction of distance with the measured destination-level pull characteristics. I show consistent results of migrant labor size causing lower productivity growth in terms of light density per capita. However, the joint F-statistics in this final exercise are too weak to draw strong conclusions.

In section 4.4 I consider heterogeneity in the returns of migrant diversity by evaluating a natural experiment. I leverage the fact that South Africa's Apartheid regime generated high migration barriers, which were suddenly dropped in the 1990s. The end of migration restrictions for black workers provides a plausibly exogenous shift in the outmigration rates (g_{ot}) from South Africa's native homelands. Following the baseline shift-share model, I instrument for changes in migrant labor size and diversity at destination. In this exercise I isolate variation in the composition of the black migrant composition, rather than the relationship between the black and white population. While the estimates are weaker due to low sample sizes, I find that the impact of migrant diversity in this context is inverted. Black migrant diversity in South Africa contributes to higher city size growth and city productivity. I interpret this as evidence that some African states are able to overcome linguistic and ethnic differences and benefit from diverse birthplace composition.

1.4.1 Baseline Results: Standard Shift-Share with Nonlinearity

Table 1.2 presents the results for the first stage of the aggregate shift-share, as described in equation 1.9. The predicted variables are the aggregations of baseline shares in the previous census years, and aggregate shifts across census years. The linear and nonlinear shift-share instruments are strong predictors of the real differences in migrant labor and migrant diversity. Surprisingly, predicted l also predicts lower diversity (higher HHI), which means larger total migration shifts are typically also more homogeneous. We'd normally expect larger shifts in migration to require a broader composition, as migrants come from further and further locations to meet a destination's

Table 1.2: Census Birthplace Shift-Share First Stage

	(1)	(2)
	<i>Migrant Flows</i>	
	$\Delta \ell$	Δdiv
Predicted $\Delta \ell$	0.452	-0.004
	[0.027]***	[0.001]***
Predicted Δdiv	1.847	0.558
	[0.447]***	[0.110]***
Mean Dep.	7.58	-0.01
Observations	829	829

Note: This table estimates the first-stage regression of predicted on real changes in migrant labor and composition $real_{d,t} = \alpha \Delta l_{d,t} + \gamma \Delta div_{d,t} + v_t + \mu_c$. Δl is the logged difference in total migrant labor in destination d between census years. The predicted value is the shift-share instrumented change in migrant labor. Δdiv is the difference in the negative HHI between census years for all non-native residents in location d . The predicted value is the estimated change in HHI based on the relative aggregate shifts and shares across origins. Regressions include fixed effects for country and year. Data for this table comes from an origin-destination panel of workers in African IPUMS Census samples. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

labor demand. The inverted correlations between diversity and migrant labor provide further evidence that the relationship between the two instruments is nonlinear. The jointly calculated F-statistics and Sanderson-Windmeijer values are presented in Table 1.3.

Panel A in Table 1.3 gives the OLS estimates of a regression of migrant labor and diversity on light outcomes. Column 1 measures the standardized change in average light density, while column 2 measured the change in the log of per-capita light density. These measures capture different aspects of light density change at destination. The first is related to city size and extent, as expanding cities increase average light density. The second has historically been compared to GDP/capita growth. As other measures of urban productivity, I include changes in the non-agricultural labor share and changes to the housing quality index, both measured from census data. I find that a 1%

increase in migrant labor size corresponds to a 0.19 SD increase in average light density, and a .76% decrease in light density per capita growth. The diversity term is measured as changes in a 0 to 1 measure of diversity. Moving from a completely homogeneous to a perfectly heterogeneous migrant labor pool is associated with a 1.3 standard deviation drop in average light density. A 0.01 increase in the HHI is associated with a 3.4% reduction in light density per capita growth. I find that an increase in migrant diversity is also related to a substantial increase in non-agricultural labor share. Going from 0 to 0.01 on a diversity measure is associated with 0.86 percent point increase in the non-agricultural labor share.


Panel B in Table 1.3 shows the second stage regression using the aggregate shift-share instruments for migrant labor and diversity. The OLS and IV results are broadly consistent. Total migrant labor size is positively related to city size growth, but negatively related to per-capita productivity growth. The size of the coefficients are larger. A 1% increase in migrant labor size yields a .3 SD increase in average light density, and a -2% drop in per-capita light growth. A 0 to 0.01 increase in migrant diversity reduces average light density by 0.02 SD, and per-capita growth by 10%. The change in the non-agricultural share is more than a 1 for 1 increase.

To better interpret these measures, I take a set of censuses from non-African developing countries and replicate the shift-share exercise using available birthplace data from these countries. I then measure light density changes across census years for these countries. The strength of the linear and nonlinear shift-share instrument varies across these countries. I take the estimated coefficient on changes in migrant labor and migrant diversity, and plot them in two dimensions to compare with the African sample. I find that my estimated impact of migrant labor size corresponds to higher city size growth and lower productivity growth relative to other developing countries. This suggest that level changes in migrant labor increase African city sprawl and extent without translating to faster per-capita growth. Across both light density measures, the African estimated impacts of migrant diversity are much lower than the non-African sample. The impact of migrant diversity on lower city night light density and per-capita growth is nearly double the estimates of non-African developing countries.


Table 1.3: Census Birthplace Shift-Share Second Stage, Migrant Flows

<i>Panel A: OLS Results</i>				
	(1) Δ Lights	(2) Δ Log(Lights/Capita)	(3) Δ Non-Agriculture Share	(4) Δ Housing Quality
Δ ℓ	0.189 [0.043]***	-0.763 [0.079]***	-0.030 [0.016]*	0.039 [0.041]
Δ div	-1.375 [0.360]***	-3.397 [1.080]***	0.860 [0.278]***	-0.453 [0.512]
Mean Dep.	-0.00	3.31	0.09	0.01
Observations	829	829	634	671
<i>Panel B: Shift-Share IV</i>				
	Δ Lights	Δ Log(Lights/Capita)	Δ Non-Agriculture Share	Δ Housing Quality
Δ ℓ	0.296 [0.100]***	-2.081 [0.173]***	-0.060 [0.036]	0.126 [0.092]
Δ div	-2.091 [0.703]***	-10.304 [2.757]***	1.310 [0.586]**	-0.135 [0.695]
Mean Dep.	-0.00	3.31	0.09	0.01
Observations	829	829	634	671
Kleibergen-Paap Fstat	16.462	16.462	13.716	26.970
Sanderson-Windmeijer ℓ	245.593	245.593	141.084	106.168
Sanderson-Windmeijer div	27.686	27.686	23.386	19.570

Note: This table presents the results of second-stage regressions of instrumented changes in migrant labor size and composition on destination productivity outcomes. Panel A presents the OLS estimation of changes in migrant labor and diversity on light density outcomes. $y_{d,t} = \alpha \Delta l_{d,t} + \gamma \Delta div_{d,t} + v_{tc}$. Panel B presents the IV 2SLS estimation of changes in migrant labor and diversity on light density outcomes. The original harmonized range of light density is from 0 to 63. The first column outcome is the standardized change in average light density, while the second column is the log change in light density per capita. The third column outcome is the change in labor share in non-agricultural industries, and the fourth column measures changes in housing quality, measured by a principal component of housing characteristics. All regressions include country-year fixed effects. Data for this table comes from an origin-destination panel of workers in African IPUMS Census samples. * p<0.01, ** p<0.05, *** p<0.01.



(a) Δ Lights



(b) Δ Log(Lights/capita)

Notes: This figure compares the shift-share migration estimates for $\Delta \ell$ and Δdiv to a sample of other countries available from IPUMS. The coefficients of the SSIV are plotted in 2-d space. The same shift-share estimate represented in Table 1.3 is estimated at the second administrative level for Cambodia, Chile, Indonesia, Mexico, Peru and Thailand. For each country, the available birthplace data is used to calculate migrant labor size and diversity. Birthplace is at the state level for Mexico, province for Peru, Thailand and Indonesia, commune for Chile, municipality for Colombia, and district for Cambodia. IV regressions include country-year fixed effects. Gray bars plot the standard errors from the SSIV. Horizontal bars represent the SEs on the diversity coefficient, while the vertical bars represent the SEs on the labor coefficient. Coefficients for labor and diversity in Table 1.3 are presented in red. The original harmonized range of light density is from 0 to 63. The first panel outcome is the standardized change in average light density, while the second panel is the log change in light density per capita. Data for this table comes from an origin-destination panel of workers in African IPUMS Census samples.

Figure 1.5: Comparing Migration Estimates to Other Samples

Table 1.16 shows the same regressions but for alternative destination outcomes. I include the levels of log light density, the change in services labor share, and the levels of housing wealth. share of the native workforce dedicated to services. I find that higher migrant diversity predicts higher services share, lower average light density in levels, and higher average levels of housing wealth. These level relationships between migrant diversity and light and housing levels suggest that migrant diversity is associated with destinations that may be smaller but wealthier on average.

1.4.1.1 Robustness Checks of Shift-Share

I perform a variety of robustness checks for our results. To address concerns about the correlation of outcomes like light density across administrative regions, Figure 1.21 and Figure 1.22 show robustness to Conley spatial standard errors across a variety of distance bandwidths. Next I test the sensitivity of our results to changes in our sample. Figure 1.23 and 1.24 show how the estimates

vary when dropping individual countries from the sample. Last I check the results when replacing the real estimated birthplace shocks with randomly distributed placebo shocks. Figures 1.25 and 1.26 show how the IV results appear in response to placebo shocks.

In these robustness results, I find some sensitivity to country drops, due to the short-panel structure of our data. Results are robust to spatially correlated errors across bandwidths. The placebo shocks help us understand how our shift-share results leverage variation from the shifts and shares portions of our instrument construction. While placebo shocks produce zero-results for our diversity estimator, the logged light density per capita outcome in the labor shift-share estimate is biased upward under placebo shocks. This suggests that some identifying variation is coming from the settlement shares. The prediction of settlement shares leveraging distance and pull characteristics will allow us to isolate exogenous variation in settlement shares.

1.4.2 Instrumenting Outmigration Shifts with Push-Shocks

Next I look at the estimation results when instrumenting for aggregate shifts using push-shocks. As described in the empirical strategy, I predict the aggregate shifts in migrant labor using shocks, or the changes between origin-destination pairs. We present the results of this analysis in three parts. First, we estimate zero-stage relationships between the raw number of people leaving birthplaces and different potential migration shocks. We find that climate and price variables do predict population movements at different lags from the contemporary period. Leveraging this result, we then estimate a first-stage regression that leverages predicted g_{ot} variables in our shift-share design, producing push-shock motivated instruments for migrant labor size and migrant composition. We find that the push shock version of our first stage is still able to predict changes in both migrant labor and migrant composition. Next, we leverage this first stage to look at our destination productivity outcomes. While the results for migrant labor effects remain consistent, our second stage for diversity is too weak to consistently estimate an effect. We conclude that the benefits of migrant labor density outweigh potential costs of migrant diversity, at least among migrants motivated by birthplace shocks.

1.4.2.1 Zero-Stage Regressions of Migrant Flows on Shocks

Table 1.4 shows the results of regressions of the log population from o in destination d on shocks at different lags from the current period. The first column includes country fixed effects, the second column includes birthplace fixed effects, and the third column includes 10-year lags of the given variables. The outcome in columns 1-3 is the number of people living outside of the given birthplace in a given year. I find that conflict consistently predicts greater outmigration. Contemporary drought conditions prevent migration, which is consistent with a model where negative wealth shocks prevent households from migrating (Bazzi, 2017). High agricultural prices in the contemporary period also encourage migration, which is also consistent with a model where positive wealth shocks increase migration. Different shocks to birthplaces may not only drive different levels of outmigration, but also may change the skill-level composition of the marginal migrant. If the type of workers who leave in response to high commodity prices are different from workers that leave in response to conflict, then changes in migrant outflows also reflect average skill level changes. Columns 4-6 show the results when the population counts are weighted by human capital, defined as schooling years for workers 18 and above. These labor efficiency units for a given origin-destination pair are an aggregation of these weighted population counts, defined as $le_{od} = \sum_i^N h_i$, where h_i are schooling years for an individual i from o living in d . The results are fairly consistent when using labor efficiency units rather than log counts, although the standard errors are larger as the weighting introduces more dramatic variation. In column 6 there is evidence that lagged conflict has a negative impact on labor efficiency units, which is a flip in the sign relative to column 3. While past period conflict events move more individuals out of origins, the marginal migrant has lower schooling years. As expected, different types of shocks have different implications for the human capital of the marginal migrant. In the section on mechanisms, I directly consider the impact of outmigration shifts on changes to the average skill level at destination.


To better understand the dynamic impact of different shocks on outmigration, Figure 1.6 plots the beta coefficients of individual regressions of migration counts against particular shocks at different lagged periods. Each dot represents a particular shock variable at a particular time horizon

Table 1.4: Predicting Outmigration with Contemporary and Lagged Shocks

	(1)	(2)	(3)	(4)	(5)	(6)
	<i>Migrant Flow</i>			<i>Labor Efficiency Units</i>		
	log(N)	log(N)	log(N)	log(le)	log(le)	log(le)
Drought	-1.308 [0.393]***	-0.288 [0.263]	0.002 [0.259]	-14.324 [0.984]***	-23.666 [1.393]***	-23.926 [1.423]***
Agricultural Price	0.168 [0.067]**	0.287 [0.064]***	0.676 [0.077]***	-0.010 [0.186]	-0.043 [0.288]	1.002 [0.398]**
Mineral Price	-0.010 [0.062]	-0.048 [0.068]	-0.131 [0.067]*	0.012 [0.144]	0.040 [0.283]	-0.229 [0.282]
Conflict	0.018 [0.004]***	0.016 [0.004]***	0.002 [0.004]	0.000 [0.010]	-0.016 [0.015]	0.019 [0.018]
Year	7.068 [0.944]***	6.815 [0.578]***	5.503 [0.604]***	4.483 [2.458]*	7.173 [2.916]**	2.533 [2.958]
Year ²	-0.002 [0.000]***	-0.002 [0.000]***	-0.001 [0.000]***	-0.001 [0.001]*	-0.002 [0.001]**	-0.001 [0.001]
Lagged Drought			0.472 [0.694]			13.785 [3.625]***
Lagged Agricultural Price			-0.652 [0.087]***			-1.840 [0.438]***
Lagged Conflict			0.029 [0.006]***			-0.070 [0.024]***
Mean Dep. Var	7.264	7.447	7.447	8.534	8.492	8.492
Observations	2,252	1,944	1,944	1,473	1,423	1,423
Country FE	Y	Y	Y	Y	Y	Y
Origin FE	N	Y	Y	N	Y	Y

Notes: Each column is a joint regression of migrant flows in terms of population or labor efficiency units on a set of contemporary and lagged shocks. The drought index is the SPEI indicator for drought intensity at an annual level. The agricultural prices are a weighted regional price exposure aggregated from individual crop shares and international prices normalized to 2000. Mineral prices are weighted by mineral and mine capacity in the origin location, and conflict represents the number of ACLED events. Lagged shocks represent average changes in the variable in the last 10 years. All regressions include country fixed effects. Year is included as a linear and squared value.

relative to the migration count year. We find that, when estimated separately, the effects of drought and conflict are consistent across time horizons, but that commodity prices show heterogeneity. Recent price spikes increase migration, while historical price spikes decrease migration. This suggests that long-run trends in high prices for key commodities convince more people to stay in their birthplace.



Notes: This column shows beta coefficients from individual regressions of birthplace-level outmigration rates on individual shocks at different horizons. The lags are at the 1 and 10-year horizon.

Figure 1.6: Predicting Outmigration with Individual Shocks Across Horizons

1.4.2.2 First and Second Stage Estimates using Predicted Shifts

Panel A in Table 1.5 shows a first stage regression of equation 1.9 where the g_{ot} growth rates are instrumented using the predicted outmigration shifts. The prediction model used leverages contemporary, 1 and 10-year lags of the selected push shock variables drought, conflict and commodity prices. I also include a linear time trend in our prediction model. I find that the shift-share with predicted shifts successfully predicts changes in migrant labor size and composition. This result suggests that push-shocks have an effect on both migrant flows and the mix of migrants that arrive. Panel A of Table 1.6 gives the second stage regression of equation 1.9 using the predicted migration shifts. I find that migrant labor size is positively related to growth in city size, but negatively related to per-capita growth in light density. I also find that increases to migrant labor size predict decreasing levels of housing quality at destination. This is consistent with a story of increasing migration driving higher housing prices and slum formation at destination. The estimates for diversity impacts on light density are too noisy to make an inference. However, I find consistent evidence that increasing migrant diversity predicts higher non-agricultural labor

share. This lends further evidence that migrant diversity has consistent impacts on urbanization and structural transformation in destinations.

1.4.3 Instrumenting Settlement Shares with Pull Characteristics

The placebo shocks exercise from the headline shift-share estimation showed that some variation in the instrument for labor size changes may be coming from endogenous shares. In this section I leverage pull characteristics interacted with origin-destination distances to isolate an exogenous component of settlement shares. The implicit assumption is that the historical pull characteristics, including distance to colonial rail, mineral deposits and portage sites are unrelated to contemporary labor demand shocks. In a zero stage regression, I predict the following settlement share variables:

$$\sigma_{odt} = \omega_1 Pull_d * \log(Dist)_{od} + \omega_2 \log(Dist)_{od} + \omega_3 \log(Dist)_{od} * Z_d + \mu_d + \nu_o + \gamma_t + \epsilon_{odt} \quad (1.19)$$

$$\pi_{odt} = \omega_1 Pull_d * \log(Dist)_{od} + \omega_2 \log(Dist)_{od} + \omega_3 \log(Dist)_{od} * Z_d + \mu_d + \nu_o + \gamma_t + \epsilon_{odt} \quad (1.20)$$

Where σ_{odt} represents the migrant shares $\sigma_{odt} = \frac{count_{odt}}{\sum_d count_{odt}}$. π_{odt} is the baseline Herfindahl estimate, which behaves as the settlement share for the diversity instrument. I use the same growth rates g_{ot} as in the baseline shift-share strategy above. Panel B of Table 1.5 shows the results of the first-stage, where the predicted shifts are used to produce the instrument aggregates. I find that the diversity predictor is weaker than in the main shift-share, but still statistically significant at the 5% level.

Panel B of Table 1.6 gives the second stage regression of equation 1.9 using the predicted shifts. I find that the labor size estimate remains negative and significant for per-capita light density growth. I interpret this as strong evidence that level changes in migrant labor size slow productivity growth, even as they increase city size and extent. The estimates of diversity are noisy. This is partially a result of a relatively weak instrument. I also consider this suggestive evidence that the effects of migrant labor size outweigh the effects of diversity.


Across estimation strategies, I find evidence that increasing migrant labor size causes cities to

grow in light density, but lowers per-capita GDP growth. Increasing migrant diversity decreases city growth, both in level and per-capita terms, but increases the non-agricultural labor share. Migrant labor size does not have consistent effects on industry mix, reflecting the importance of identifying both level and composition parameters to understand the aggregate impacts of migration. My results are broadly consistent with cross-sectional work that finds limited evidence of agglomeration returns in African cities, and negative impacts of ethnic diversity. My findings on the effects of migrant diversity on labor allocation suggest long-run urbanization benefits of migrant diversity, which I explore further in my long-run estimation strategy below.

Using a battery of shocks to birthplaces and historical characteristics of destinations, I predict both shifts and shares. I find that migration shocks do successfully predict both the size and composition of migrant labor. While the effects of labor size remain consistent, the effects of migrant diversity on light density outcomes are not statistically significant. I find consistent evidence that migrant diversity increases non-agricultural labor share. The exercise predicting settlement shares yields a weak first stage. However, I do find that the negative effects of migrant labor size on per-capita growth rates remains consistent and statistically significant.

1.4.4 Evidence from a Natural Experiment: Ending Apartheid

South Africa represents a particular case of internal sorting. The end of Apartheid brought a sudden lifting of migration restrictions that had forced the black population to remain in “ethnic homelands”. These homelands, established by the Native Land Act in 1913, were loosely tied to distinct historical tribes. Transkei, for example, was a homeland attached to the Xhosa people. While KwaZulu, around Durban, was the homeland for the Zulu people. The Pass Laws required black South Africans to carry internal passports to regulate their movement outside of native lands (Amodio and Chiovelli, 2018). While not all black South Africans lived in these ethnic homelands, millions are believed to have been forcibly resettled in the homelands between 1960 and 1991 (Lochmann et al., 2023). The Pass Laws were repealed in 1986, followed by the Native Land Act in 1991. These events amount to a sudden reduction in migration costs for black South Africans from




Notes: These figures show the location of South Africa's ethnic homelands, where the majority of the black population were forced to settle. The figure on the right shows the distribution of major language families, according to the Ethnologue dataset of languages.

Figure 1.7: Ethnic Homelands under Apartheid, South Africa


different ethnic origins, who could now move freely to new destinations. Leveraging municipality level data from censuses between 1991-2022, we can study the long-term labor market effects of this migration shock.

Figure 1.7 shows the distribution of African ethnic homelands during the Apartheid era in South Africa. The panel on the right shows the distribution of language families, as measured in the Ethnologue dataset of languages. These figures give a sense of the correspondence between the homelands and particular black ethnic groups or linguistic families. The regions outlined in black are local municipalities as of 2022. These regions will serve as our units of analysis when studying regional outcomes after Apartheid.

Between 1986 and 1991, many of the restrictions on internal migration were lifted for black South Africans. Using census data from 1991 - 2022, we can trace the trajectory of diversity and population size over this period. While the South African censuses don't contain consistent information on subnational birthplace, they do identify native language or mother tongue of respondents. Using this language information, we can link black respondents to their ethnic homeland using the same language-ethnic group linking procedure leveraged in the main analysis. For each destination, we can calculate an ethnic HHI index as the relative diversity of the black population. We don't include English or Afrikaans speakers in this index.



(a) Ethnic -HHI by Homeland Distance



(b) Δ Ethnic -HHI in Major Cities

Notes: These figures show change in the homeland composition of the black population in destinations over time. The left figure shows the HHI score for destinations at different average distances from the native lands in the eastern part of the country. The right figure shows the change in this HHI measure for major cities, from the end of Apartheid onward.

Figure 1.8: Black Diversity in South Africa After Apartheid

Figure 1.8 shows the change in diversity over time in destinations. The left figure splits destinations in terms of average distance to ethnic homelands, with the furthest quantile representing destinations on the west coast of the country. The y-axes capture the negative Herfindahl index of black population diversity by native homeland, with higher values representing increased diversity. The right figure considers major urban destinations, and plots the change in the HHI index, rather than the levels. I find heterogeneous fluctuations in black ethnic diversity across destinations. In levels, the closest quantile of municipalities to ethnic homelands is the most diverse in terms of the black population's ethnic mix. After the abolishment of pass laws, this group experiences a decrease in diversity, as South Africans depart for further off destinations. The furthest quantile of municipalities experienced a long increase in diversity from 2001 to 2011, nearing the levels of the native homelands. Cape Town, one of the major towns in the furthest quantile group, saw a strong increase in diversity from 1991-1996. This is shown in Panel B, which captures changes to the negative HHI over time. Other major cities saw clear increases in diversity, represented as points above 0 in Panel B. Though the rates are less dramatic in terms of proportional changes.


I replicate the baseline shift-share strategy from Table 1.3, applied to the South Africa case. The shifts g_{ot} are the post-Apartheid growth in migrants from origin o , as defined by their native

language, living outside of o . The settlement shares are the Apartheid shares of each migrant group living in a particular municipality. One might instead think to use “distance to homeland” as an instrument for the settlement shares. However, as seen in the trend figures, migrants did not necessarily travel to cities close to homelands. Many major cities like Cape Town, far from the native homelands, experienced sharp shifts in the migrant population.

Table 1.7 shows the results of a second-stage shift-share. The linear effects of migrant labor flows are consistent with the headline results. Migrants increase city sizes, but migrant flows do not produce per-capita light density growth. However, the compositional effects are inverted. Increasing black migrant diversity causes larger cities and higher per-capita growth in light density. These results show that the effects of migrant diversity are heterogeneous across locations. South Africa’s particular history of oppression and black-white conflict may have taken precedence over ethnic tensions between black ethnic groups.

1.5 Mechanisms

In this section I consider a few possible mechanisms for the headline result of the paper. Migrant flows increase city sizes but slow productivity growth. Migrant birthplace diversity causes lower city sizes and lower productivity growth, but increases urbanization. The negative impact of diversity is indicative of a nonlinear congestion force that dampens agglomeration benefits in the short-run, but may yield positive long-run benefits in terms of structural transformation. The “birthplace composition” measure, or Δdiv , is a broad index of diversity. Because it is derived from workers’ differences in birthplaces, there could be many types of diversity captured by this measure. Diversity in birthplace may be correlated with linguistic diversity, ethnic or religious diversity. To the extent that different origins may have different skills or experiences, it also encompasses diversity of skill. I consider a few examples of how birthplace diversity may create short-run negative externalities through ethnic conflict, difficulties in cooperation related to linguistic distance, or differences in skill complementarity.



Notes: This figure shows the cross-sectional relationship between the number of ACLED battle-events in a region and population size, ordered into deciles. The regions defined in this analysis are grid-cells of the African continent of approximately $1000km^2$. Population is calculated from the Worldpop estimates for 2013. The y-axis measures the average probability of observing a conflict event in that grid-cell in a given year between 1997 and 2025.

Figure 1.9: Conflict and Population Size in the Cross-Section

1.5.1 Ethnic Conflict

The short-run costs of migrant diversity may be a result of urban conflict generated between disparate groups. A standard result in political economy links ethnic diversity in Africa to ethnic conflict (Arbatli et al., 2020)⁵. Most of the evidence of this relationship comes from cross-sectional associations of diversity measures on geolocated battle events from ACLED or UCDP.

Beyond the normal causality concerns, another challenge for this conflict literature is understanding the role of population size on conflict. Indeed, in the cross-section population size is strongly positively correlated with conflict. To the extent that diversity and population move together, it's difficult to disentangle the role that diversity plays independent of population size. If larger population centers are naturally more diverse, increased conflict or crime could be due to either force. The construction of my empirical strategy is able to disentangle these effects.

⁵There are debates about the particular functional form this relationship takes, with scholars arguing for different ways to index diversity and weight relative group sizes as a metric for polarization, fractionalization, etc. The relationship between relative group size and conflict is not obvious. Many ethnic conflicts, like the Rwandan genocide, seem to emerge without the need for many groups, or equally sized groups. I continue to rely on changes in a negative Herfindahl index, for consistency with the baseline specification

Following the Bartik style shift-share strategy elaborated in the empirical methods section, I study the effect of changes in migrant labor size and migrant diversity on conflict outcomes. Table 1.8 shows the results of this analysis. Conflict is measured as the number of battle events that occur in a given destination-year, according to the ACLED dataset. I also weight the conflict events by estimated average fatalities, which is represented as “deaths” in the table columns.

Column 3 presents results for the differenced change in conflict at destination, in response to changes in migrant labor size and diversity. A 1% increase in migrant labor size adds 0.347 average conflict events per year. A .1 point increase in migrant diversity increases the number of conflict events by 0.7 per year. This is strong evidence of a congestion force that affects destinations as larger and more heterogeneous migrants arrive in destinations. Ethnic conflict may be one mechanism that reduces light density growth in increasingly diverse destinations.

1.5.2 Linguistic Distance

A natural problem for migrants moving to new destinations may be learning a new language. If migrants arriving in a destination speak many different languages, these linguistic differences could pose challenges for coordination in firms, households, or neighborhoods. A mostly US literature on assimilation of migrants suggests that closing cultural gaps is an important part of unlocking returns to migration (Abramitzky et al., 2020a). So far none of our measures of diversity have tried to weight differences across birthplaces in terms of cultural or linguistic distance. Two points for why it is not ex-ante obvious that cultural distance is the important mechanism driving our results: (1) This paper looks at internal migration. While many states have an abundance of local indigenous languages, it is also true that most states have a condensed set of national languages. Take Mozambique for example. While the country hosts a wide variety of local languages and dialects, most people in urban centers speak Portuguese fluently. (2) There are many examples of conflict or animosity between groups that are on paper culturally similar (Posner, 2004). Differences in groups over space are often the result of historical, geographical and political factors that may arise even when groups speak the same language.

To consider the role of linguistic distance, I match each administrative region to its majority spoken language, as according to the Ethnologue, a large geographically coded dataset of world languages. A useful feature of the Ethnologue is that each language is coded within a nested tree of language families and sub-families. Oko, for example, is a Nigerian language spoken in Edo state and part of the Niger-Congo family. Within the Niger-Congo linguistic family, its part of the Volta-Niger linguistic subgroup. Between any two languages, we can then create a simple measure of linguistic distance based on how many linguistic families and sub-families they have in common. Oko is related to Ibo, another Volta-Niger language, but distant from the northern Nigerian Hausa, an Afro-Asiatic family language.

For each migrant from an origin o in destination d , I calculate the linguistic distance between the majority language in that given origin and destination pair. Weighting by the total migrants from each origin into a destination d , I calculate an average linguistic distance index that measures the average cultural distance between migrants and destination natives. Table 1.9 presents results from an OLS regression in first differences, relating destination productivity to changes in migrant labor size and changes in the average linguistic distance of migrants. As above, these differences are calculated across census periods. I find that increasing linguistic distance is only weakly related to changes in light density, but strongly related to an increasing share of non-agricultural labor.

Is increasing linguistic distance captured up by the diversity shift-share instrument used in the analysis? In Table 1.10, I regress the measure of average linguistic distance against my first-stage predictors from Table 1.2. In levels, higher ethnic concentration is in fact associated with a lower level of linguistic distance. However, in differences I find the opposite of the expected relationship. Increasing diversity of birthplace composition predicts decreases in average linguistic distance. That is, our predictor of changes in birthplace diversity does not simultaneously predict changes in linguistic distance. Therefore, while we do observe an association between changes in linguistic distance and higher agricultural labor share, my shift-share strategy does not directly map to a levels change in average linguistic distance.

1.5.3 Skill-Level, Segregation and Industry Concentration

The changes in migration discussed in this paper are in terms of the number of people over 18 that have settled in a destination. While the empirical model is in differences, it's possible that the effects are being partially driven by differential trends in the human capital of the migrants arriving in destination. Changes in migrant diversity may be related to changes in the average skill level of the workforce. Another possible channel is through other features of the labor market. If migrants arriving in destinations bring many diverse skills, they may reduce the industrial concentration of cities and in-turn diversify production at destination. Lower industry concentration may be a channel through which migrant diversity creates long-run structural transformation.

To study these mechanisms, I consider the impact of instrumented migrant levels and composition on changes to average worker skill level, industry concentration and industry segregation. Industry concentration is measured as the HHI of industry labor-share across industries within a region. Industry segregation is a measure of the extent to which migrants from different origins sort into particular industries. The measure captures the extent to which industry labor shares for a given migrant birthplace deviate the population average in that region (Alesina and Zhuravskaya, 2011). For a given migrant origin group $o \in O$ living in destination d and working in industry $i \in I$, segregation is defined as:

$$Segregation_d = \frac{1}{O-1} \sum_{o=1}^O \sum_{i=1}^I \frac{N_o}{N_d} \frac{(\pi_{io} - \pi_{od})^2}{\pi_{od}} \quad (1.21)$$

Where π_{od} is the fraction of group o in destination d , and π_{io} is the fraction of group o in industry i of destination d . N_o is the total population of group o in destination d and N_d is the total population in destination d . Higher values of this segregation index correspond to greater segregation of ethnic groups across industries within a destination.

Table 1.17 presents the OLS and SSIV results for changes in average worker skill level, industry concentration, and industry segregation by birthplace. The model is the same baseline shift-share examined in Table 1.3. I find that both increasing migrant labor size and diversity lowers

industry concentration, as measured by an HHI of labor-shares. This suggests that migrant flows and composition introduce increase diversity to the industrial mix of a destination. I also find evidence that increasing migrant labor size reduces average worker skill level. This is a possible mechanism by which increasing migrant labor size does not translate to higher per-capita productivity growth. Low skilled migrants bring down the average skill level of workers, compete for low-skilled wages, but may not generate the ideas, businesses and networks necessary to create agglomeration benefits. I find no evidence that increasing migrant diversity affects the average skill level of workers or the segregation of groups across industries. While increasing migrant diversity does introduce diversity to the industry mix, we don't see evidence that particular migrants specialize within certain industries. The industry categories are broad, therefore it's possible that given more granular occupational categories different specialization patterns would emerge.

1.5.4 Productivity Evidence from Firms

Past work has suggested that ethnic divisions within firms can lower productivity in team-based production (Hjort, 2014). This suggests a mechanism via poor coordination or discrimination within firms, rather than outright conflict. In this section, I consider evidence on changes in diversity within firms, and the subsequent impact on productivity. An ideal microeconomic dataset to study the role of migrants in productivity would include firm-level information on input-output, as well as detailed characteristics of the firm's labor force. While some enterprise level surveys exist for African countries, few provide information on workers, including ethnic identity or birthplace. Two exceptions come from the "Regional Programme on Enterprise Development (RPED)", led by the World Bank in collaboration with the Centre for the Study of African Economies (CSAE). In the 90s and early 2000s a panel of manufacturing firms was collected for Ghana and Tanzania which, for some waves, include data on the ethnic composition from a sample of workers.

These panels were collected between 1992 and 2003, recording basic characteristics of the firm such as wages and labor size. In addition, a worker supplement is collected for each firm in which up to 10 workers are interviewed and asked about their experience and background. I leverage

this data to examine associations between firm-level productivity, number of workers and worker ethnic concentration. Table 1.11 shows the results of firm-level regressions of productivity measures on firm labor size and worker ethnic concentration. Outcomes include log wages per worker, and log manufacturing output value per worker. Columns 1 and 3 include fixed effects for country and year, while columns 2 and 4 look at within-firm variation over time using firm fixed effects. I find some evidence that in the cross-section, firm productivity is increasing in labor size and decreasing in ethnic concentration. This is the opposite association of what we'd expect to see if ethnic diversity hinders firm productivity. Looking within firms over time, we see evidence that firms are becoming less productive as they grow in size, with no evidence of an effect by changing ethnic HHI.

These samples are small, and the analysis is not causal. However we don't see evidence of an ethnic diversity penalty within or across firms that has been posited by microeconomic papers.

1.5.5 Migration and Ethnic Attitudes

Are migrants more or less prone to ethnic conflict relative to natives or non-migrants? It may be that the act of migrating itself is associated with increasing or decreasing tribalism among ethnic groups. For example, a migrant from a minority group may choose to move to a city, and begin to experience economic disenfranchisement or discrimination. In turn, the migrant becomes more allied with their ethnic group, and deepens their coethnic preference. The inverse of this process is also possible. Work with panel data in Kenya has shown reduced tribalism and increasing national identity in individuals after they migrate to cities (Kramon et al., 2022). This suggests that urban centers play a diffusing role in ethnic tensions. Using Afrobarometer data, I explore observational differences between migrants and non-migrants in terms of their reported national identity and future economic expectations. I use the same model of observational returns leveraged in Table 1.1. I replace our outcome of interest with a measure of ethnic allegiance or economic

disenfranchisement:

$$\begin{aligned} Attitudes_{it} = & \beta_1 Migrant_i + \beta_2 Migrant_i \cdot Dist_{od} \\ & + \beta_3 Migrant_i \cdot CoethShare_{od} + Z_i + X_{od} + W_d + \nu_{st} + \gamma_o + \epsilon_{it} \end{aligned} \quad (1.22)$$

Where $Dist_{od}$ is the log distance between the ethnic homeland and the destination, $Migrant_i$ is a dummy for migrant status, and $CoethShare_{od}$ reflects the fraction of individuals in d that are from ethnic homeland o . We include fixed effects for country-year ν_{st} , and γ_o to isolate variation within an ethnic group. The controls for the individual Z_i include age and schooling, while X_{od} includes the level of o-d distance and coethnic share, and W_d includes destination log population. $Attitudes_{it}$ is a measure of ethnic or economic attitudes, “Nationalism” or “Expectations”. National identity, or “nationalism” is an individual’s response to a question about how much the individual identifies with the nation relative to their ethnic group. We use it here as a measure of ethnic allegiance. Future expectations is a question styled after the Michigan Consumer Confidence Survey, which asks individuals if they expect economic conditions to improve in the next year. We use this question as a measure of economic disenchantment or distress. Table 1.18 shows these results, with no significant differences between migrants and natives in terms of national identity or economic expectations.

1.6 Urban Growth and Diversity in the Long-Run

The analysis in this paper has focused on decade level changes in migrant labor size, composition and productivity. The positive effects of migrant diversity on non-agricultural labor share suggests that migrant diversity may affect fundamental characteristics of the economy at destination. It may be that the benefits of such effects manifest over longer time horizons. The problem with studying longer-run effects of labor size and diversity on growth is a lack of data. Most early African censuses, where they exist, don’t include ethnic or birthplace identifiers. In addition, few measures of subnational productivity exist in earlier periods.

In this section I leverage the historical pull shocks constructed for the shift-share instrument

to estimate a long-run effect of labor size and diversity on urban development. The historical productivity shocks, including colonial rail lines, portage sites and mineral deposits serve as agglomeration instruments, which can be used to predict contemporary city locations. The empirical strategy takes inspiration from the literature on estimating labor demand curves from shocks to labor demand (Diamond, 2016a; Notowidigdo, 2020) and housing supply elasticities (Saiz, 2010; Guedes et al., 2023). In this work inverse demand and supply elasticities are estimated using an interaction between a labor demand shock and a housing supply constraint.

I start from the following equation. For each region i , the long-term relationship between diversity, population and productivity is described as:

$$y_i = \beta_0 + \beta_1 \ell_i + \beta_2 \text{div}_i + X_i + \epsilon_i \quad (1.23)$$

where ℓ_i and div_i capture a destination's size and diversity in the long-run, and X_i is a vector of geographic controls. This cross-sectional regression is similar to the associations estimated by Montalvo and Reynal-Querol (2021). What I add to this framework is a causal inference strategy. Since there is no long-run panel data on diversity div_i over time, I use a historic fractionalization index that captures each region's exposure to historical ethnic groups defined by anthropological maps. I use the Murdock Map, and define each region's historical diversity exposure as the distribution of land occupied by different ethnic homelands. A region's contemporary diversity is proxied by the interaction of this historical exposure to different groups, and historical labor demand. Intuitively, areas become more diverse as labor demand shocks compel them to draw in labor from surrounding areas. The more inherently diverse this potential pool of labor (due to the historic spatial distribution of groups), the more diverse the city will be. Substituting historic diversity HistDiv_i for div_i , I estimate:

$$y_i = \beta_0 + \beta_1 \ell_i + \beta_3 \text{HistDiv}_i + \beta_4 \text{HistDiv}_i * \ell_i + X_i + \epsilon_i \quad (1.24)$$

where HistDiv_i captures a region's potential exposure to diverse migrants based on the Murdock

Map, and ℓ_i captures a measure of employment density. The interaction of fractionalization and labor is the variable of interest. The interaction captures how historic fractionalization affects the labor demand elasticity, and in turn output and productivity. $HistDiv$ represents a fixed regional quality, and β_4 represents an elasticity of urban growth with respect to this quality.


Population density is an endogenous variable. In addition, $HistDiv_i$ by itself will be related to a variety of geographic fundamentals that governed the distribution of groups over space (Michalopoulos, 2012). Our IV strategy will predict population density and its interaction by exploiting temporary shocks to regional productivity ΔA_i that drove labor demand historically, but that are no longer correlated with unobserved productivity fundamentals today. I use these historic productivity shocks as agglomeration instruments, and predict both labor size $\hat{\ell}$ and the interaction of labor size and diversity $\ell * \widehat{HistDiv}$:

$$\hat{\ell}_i = \alpha + \beta_1 \Delta A_i + \beta_2 \Delta A_i * HistDiv_i + \omega_s + \epsilon_i \quad (1.25)$$

$$\ell_i * \widehat{HistDiv}_i = \alpha + \beta_1 \Delta A_i + \beta_2 \Delta A_i * HistDiv_i + \omega_s + \epsilon_i \quad (1.26)$$

ΔA_i represents one of the three historical pull shocks, either distance to colonial rail, distance to portage site, or distance to a mineral deposit. ω_s are state fixed effects. Using these regressions as two first-stages, I will then estimate the second stage effect on contemporary development outcomes for each region in the cross-section. The exclusion restriction requires that a historical productivity shock ΔA_i is uncorrelated with unobserved other factors in ϵ_i that drive outcomes of light density, wealth, lights/capita, and conflict in the contemporary period. Note we do not need that $HistDiv_i$ is uncorrelated with ΔA_i in this estimation.

I split the African continent into equally sized hexagonal grids of approximately $1000km^2$, which I use as my regions i . For each grid I aggregate data on conflict, light density, population across years and the geographic variables including malaria suitability, ruggedness and soil suitability. Figure 1.10 shows an example of how the Murdock Map is used to calculate a historical diversity index for each region. This is a fractionalization index of the relative share of each grid taken up by



Notes: This figure shows examples of the grid-level data constructed from a historical map of diversity. Fractionalization is calculated as the dispersion in the share of each grid taken up by different ethnic homelands.

Figure 1.10: Data Examples from Nigeria

a particular Murdock group. Higher values imply more diversity.

I use a dummy for labor size ℓ that marks regions as cities (or high density) if they overlap with a city of over 20k inhabitants in 2015, as defined by the urban database Africapolis. Table 1.19 shows the results of OLS regressions for different development outcomes y following equation 1.24. Panel A shows results using a city dummy for ℓ , while panel B uses an interpolated population measure at the grid level, taken from the Worldpop dataset. Urban agglomerations have higher light density, and higher conflict. Due to the high relative population size, per capita measures in these regressions are negative relative to less populated regions. In the cross-section, I find that cities located in more diverse areas have a higher lights/capita value, but also higher incidence of conflict as measured by a summation of ACLED conflict events.

Table 1.12 shows results for our three instrumental variable strategies following the historical pull characteristics described in the Empirical Strategy section. I find that across strategies, high diversity places that receive a productivity shock have higher values of our preferred development measure, lights/capita. The baseline negative coefficient on diversity is in line with the correlations studied in past political economy work on the relationship of diversity and development. While baseline diversity has a negative coefficient, the elasticity of urban growth with respect to diversity is positive, suggesting that in the long-run diversity of ethnic groups can be a boon for urban growth.

1.7 Conclusion

This paper disentangles the effects of migration levels and composition on urban African productivity. I build a panel that proxies for origin-destination migration flows using census data on reported birthplace and ethnicity. Then I implement a shift-share instrument that simultaneously identifies the linear change in migration labor size and nonlinear birthplace composition. I find that destinations that receive more migration labor grow in light density, but are not more productive in per-capita terms. This is consistent with a story of growing African cities with limited productivity returns or urban structural transformation (Jedwab et al., 2025). Migrants that are more diverse cause lower growth in levels and per-capita terms. However, increasing migrant diversity also yields higher non-agricultural labor force shares in destinations. Migrant diversity creates short-run costs, but seems to generate long-run benefits to destinations in terms of urbanization. The results are broadly consistent when I instrument for shifts using plausibly exogenous push shocks to outmigration, including international commodity prices, conflict events and drought conditions. In zero stage regressions, I find evidence that these channels move outmigration from origins, and have predictable effects on migrant labor size and composition. This fact has implications for how we study the impact of migration shocks, as the spatial dispersion of shocks changes both the size and composition of migrants.

I explore several mechanisms to better understand the effect of migrant labor size and diversity on destinations. I find evidence that increased migrant diversity and migrant labor size increase urban conflict. I don't find evidence that diversity hurts productivity at the firm level in a panel of firms. The mixed results for diversity across different outcomes lead me to pursue two further exercises. In a study of the Apartheid period of South Africa, I leverage the repeal of the Pass Laws as a migration shock event that generates high outmigration shifts from native homelands. Measuring the size and composition of the black migrant labor force and estimating the same baseline shift-share yields an inverted effect of diversity. Productivity of South African municipalities is increasing in birthplace diversity, suggesting that South African cities benefit from

the diversity of the black population. The particular history of South Africa's conflict between the black and white populations may have diffused tensions across different black African ethnic groups.

Lastly, I consider the long-run effects of diversity. I leverage my pull characteristics as "agglomeration predictors", which forecast optimal locations for long-run urban growth. By interacting these predictors with historical measures of diversity from anthropological maps, I estimate an elasticity of regional diversity with respect to these historical productivity shocks. I find that in a long-run cross-section, diverse areas that experienced a historical productivity shock fared better than less diverse areas who experienced comparable shocks. This pattern is consistent across different historical pull characteristics. I conclude that migrant labor diversity is a long-term boon for African cities, and may outweigh congestion costs over time.

1.8 Acknowledgements

Chapter 1, "Density and Diversity in African Cities," is currently being prepared for submission for publication of the material therein. The dissertation author, Andre Daniel Gray, is the sole author of this chapter.

Table 1.5: First Stage Prediction of Migrant Size and Composition with Predicted Shifts and Shares

<i>Panel A: Predicted Shifts</i>		
	(1)	(2)
	<i>Migrant Flows</i>	
	$\Delta \ell$	Δdiv
Predicted $\Delta \ell$	0.463 [0.034]***	-0.006 [0.003]**
Predicted Δdiv	2.213 [2.629]	1.607 [0.345]***
Mean Dep.	7.88	-0.01
Observations	661	661
Fstat	93.081	12.106

<i>Panel B: Predicted Shares</i>		
	$\Delta \ell$	Δdiv
Predicted $\Delta \ell$	0.214 [0.042]***	0.005 [0.002]*
Predicted Δdiv	-22.876 [12.033]*	3.721 [0.901]***
Mean Dep.	7.77	-0.01
Observations	710	710
Fstat	28.672	9.009

Note: This table estimates the first-stage regression of predicted on real changes in migrant labor and composition $real_{d,t} = \alpha \Delta l_{d,t} + \gamma \Delta div_{d,t} + v_t + \mu_c$. Δl is the logged difference in total migrant labor in destination d between census years. In Panel A, shifts are taken from predicted changes in outmigration across birthplaces, while shares are pre-period settlement shares from each birthplace. Shifts are instrumented by drought events, conflict events and commodity price shocks. In Panel B shifts are taken from the baseline changes in outmigration across birthplaces, while shares are predicted settlement shares from each birthplace, leveraging historical productivity shocks and origin-destination distance. Historical productivity shocks include distance to colonial rail, distance to mineral deposits, and distance to portage sites. Δdiv is the difference in HHI between census years for all non-native residents in location d . Regressions include fixed effects for country and year. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.6: Second-Stage Effects of Migration on Productivity with Predicted Shifts and Shares

<i>Panel A: Predicted Shifts</i>				
	(1) Δ Lights	(2) Δ Lights/Capita	(3) Δ Non-Agriculture Share	(4) Δ Housing Quality
Δ l	0.256 [0.076]***	-1.387 [0.157]***	0.012 [0.070]	-0.299 [0.130]**
Δ div	-1.849 [1.521]	0.406 [4.300]	2.104 [1.265]*	-2.012 [2.034]
Mean Dep.	-0.18	2.43	0.35	0.03
Observations	661	661	465	472
Kleibergen-Paap Fstat	12.788	12.788	13.289	11.121
Sanderson-Windmeijer ℓ	71.251	71.251	43.754	43.862
Sanderson-Windmeijer div	12.810	12.810	13.497	11.174
<i>Panel B: Predicted Shares</i>				
	Δ Lights	Δ Lights/Capita	Δ Non-Agriculture Share	Δ Housing Quality
Δ l	-0.283 [0.264]	-3.150 [0.717]***	0.612 [0.185]***	0.223 [0.347]
Δ div	4.546 [11.405]	6.553 [8.927]	-1.331 [2.099]	0.068 [2.731]
Mean Dep.	-0.13	2.70	0.23	0.01
Observations	710	710	517	553
Kleibergen-Paap Fstat	9.051	9.051	10.804	4.734
Sanderson-Windmeijer ℓ	11.744	11.744	7.938	8.339
Sanderson-Windmeijer div	9.329	9.329	10.576	5.742

Note: This table presents the 2SLS estimation of changes in migrant labor and diversity on light density outcomes. $y_{d,t} = \alpha \Delta l_{d,t} + \gamma \Delta div_{d,t} + \nu_{tc}$. In Panel A, shifts are taken from predicted changes in outmigration across birthplaces, while shares are pre-period settlement shares from each birthplace. In Panel B shifts are taken from the baseline changes in outmigration across birthplaces, while shares are predicted settlement shares from each birthplace, leveraging historical productivity shocks and origin-destination distance. The original harmonized range of light density is from 0 to 63. The first column outcome is the standardized change in average light density, while the second column is the log change in light density per capita. The third column outcome is the change in labor share in non-agricultural industries, and the fourth column measures changes in housing quality, measured by a principal component of housing characteristics. All regressions include country-year fixed effects. Data for this table comes from an origin-destination panel of workers in African IPUMS Census samples. All regressions include country-year fixed effects. * p<0.01, ** p<0.05, *** p<0.01.

Table 1.7: South Africa Shift-Share Second Stage, Migrant Flows

<i>Panel A: OLS Results</i>		
	(1) Δ Lights	(2) Δ Lights/Capita
Δl	0.047 [0.016]***	-0.283 [0.033]***
Δdiv	1.646 [0.301]***	1.582 [0.662]**
Mean Dep.	0.07	-0.15
Observations	798	798
<i>Panel B: Shift-Share IV</i>		
	Δ Lights	Δ Lights/Capita
Δl	0.083 [0.026]***	-0.215 [0.046]***
Δdiv	4.524 [1.640]***	10.525 [3.161]***
Mean Dep.	0.07	-0.15
Observations	798	798
Kleibergen-Paap Fstat	9.037	9.037
Sanderson-Windmeijer L	374.608	374.608
Sanderson-Windmeijer Div	22.360	22.360

Note: Panel A presents the OLS estimation of changes in migrant labor and diversity on light density outcomes. $y_{d,t} = \alpha \Delta l_{d,t} + \gamma \Delta \text{div}_{d,t} + v_{tc}$. Panel B presents the IV 2SLS estimation of changes in migrant labor and diversity on light density outcomes. Light density is included with various specifications. The original harmonized range is from 0 to 63. The first column gives the log level in the contemporary period, the second is the logged difference across census years, the third is the non-logged distance and the fourth divides the light density value by the population of the administrative region in that period. All regressions include country-year fixed effects. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.8: Shift-Share Second Stage, Conflict Outcomes

<i>Panel A: OLS Results</i>				
	(1) Conflict	(2) Deaths	(3) Δ Conflict	(4) Δ Deaths
$\Delta \ell$	0.410 [0.268]	-0.241 [0.296]	0.347 [0.140]**	-0.023 [0.445]
Δdiv	2.573 [1.081]**	-3.249 [4.983]	7.122 [2.074]***	1.159 [3.681]
Mean Dep.	1.19	1.26	0.63	0.14
Observations	829	829	581	581
<i>Panel B: Shift-Share IV</i>				
	Conflict	Deaths	Δ Conflict	Δ Deaths
$\Delta \ell$	0.719 [0.588]	-1.214 [0.747]	0.738 [0.257]***	0.543 [0.882]
Δdiv	4.041 [2.022]**	-7.763 [6.253]	10.991 [3.349]***	-4.733 [8.157]
Mean Dep.	1.19	1.26	0.63	0.14
Observations	829	829	581	581
Kleibergen-Paap Fstat	16.462	16.462	12.030	12.030
Sanderson-Windmeijer ℓ	245.593	245.593	156.499	156.499
Sanderson-Windmeijer div	27.686	27.686	22.170	22.170

Note: Panel A presents the OLS estimation of changes in migrant labor and diversity on conflict at destination. $y_{d,t} = \alpha \Delta l_{d,t} + \gamma \Delta \text{div}_{d,t} + v_{tc}$. Panel B presents the IV 2SLS estimation of changes in migrant labor and diversity on conflict at destination. Outcomes include the number of ACLED battle events, which we call “conflict”, as well as the estimated number of conflict deaths in a destination across years. We include these outcome both as levels and in differences. The differenced terms are in raw changes in the number of conflict events or deaths over time. All regressions include country-year fixed effects. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.9: Change in Linguistic Distance and Productivity

<i>OLS Results</i>				
	(1) Δ Lights	(2) Δ Log(Lights/Capita)	(3) Δ Non-Agriculture Share	(4) Δ Housing Quality
$\Delta \ell$	0.217 [0.049]***	-0.779 [0.077]***	-0.027 [0.016]*	0.027 [0.036]
Δ Linguistic Distance	-2.812 [1.442]*	3.582 [5.019]	4.600 [1.505]***	1.701 [2.344]
Mean Dep.	0.01	3.32	0.13	0.03
Observations	881	881	685	722

Note: This table presents the OLS estimation of changes in migrant labor and average linguistic distance on productivity at destination. Average linguistic distance is calculated as the inverse of the similarity between a migrant's language and the majority language at destination, weighted by flow size across all migrant groups. Similarity is defined according to Ethnologue linguistic families. This measure is then differenced across census periods. All regressions include country-year fixed effects. Clustering is at the second administrative level. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.10: First Stage Relationship of *div* and Linguistic Distance

<i>First Stage Correlation</i>		
	Linguistic Distance	Δ Linguistic Distance
Predicted Δl	0.015 [0.001]***	-0.002 [0.000]***
Predicted Δdiv	0.094 [0.022]***	-0.027 [0.010]***
Mean Dep.	-0.00	-0.00
Observations	944	944

Note: This table presents a first-stage regression relating the predicted migrant labor and predicted diversity change estimates from our shift-share design to average changes in linguistic distance at the destination. Average linguistic distance is calculated as the inverse of the similarity between a migrant's language and the majority language at destination, weighted by flow size across all migrant groups. Similarity is defined according to Ethnologue linguistic families. The first column outcome is the average linguistic distance of migrants to destination natives, weighted by the group size of migrants. The second column outcome is this measure differenced across census periods. All regressions include country-year fixed effects. Clustering is at the second administrative level. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.11: Ethnicity and Firm Productivity in Ghana, Tanzania

	Log(Wages/Worker)	Log(Wages/Worker)	Log(Output/Worker)	Log(Output/Worker)
Log(labor)	0.299 [0.053]***	-0.903 [0.136]***	0.361 [0.104]***	-1.326 [0.511]**
Ethnic HHI	-1.067 [0.269]***	0.126 [0.224]	-0.457 [0.530]	0.868 [0.843]
Mean Dep.	15.19	15.78	12.82	13.44
Observations	511	332	496	330
Country FE	Y	N	Y	N
Wave FE	Y	N	Y	N
Firm FE	N	Y	N	Y

Notes: This table shows the result of a regression of log wages or output per worker on firm size and worker ethnic concentration $\log(\frac{y}{\ell})_{ist} = \beta_1 \ell + \beta_2 HHI + \epsilon_{ist}$. Where i is firm, s country, and t is survey wave. The first column in each pair includes country and survey wave fixed effects, while the second includes firm fixed effects. Output represents the total monetary value of manufacturing output, while wages reflect the total wage bill. Each regression also controls for the number of workers that data was collected on for the HHI variable, up to 10 workers. * p<0.01, ** p<0.05, *** p<0.01.

Table 1.12: 2SLS of Historical Diversity and Productivity Instruments

Panel A: Colonial Rail				
	(1) Light Density	(2) Lights/Capita	(3) Growth Lights/Capita	(4) Conflict
City	4.537 [0.209]***	-2.447 [0.254]***	0.294 [0.257]	34.455 [10.677]***
Diversity	0.154 [0.127]	-0.713 [0.155]***	0.116 [0.157]	21.367 [6.509]***
City*Diversity	-1.265 [0.910]	3.300 [1.106]***	0.265 [1.121]	-136.214 [46.520]***
Mean Dep.	-1.92	-10.90	-1.71	8.50
Observations	16,322	16,321	16,321	16,322
Cragg-Donald F-Stat	179.838	179.830	179.830	179.838
Sanderson-Windmeijer City	201.228	201.276	201.276	201.228
Sanderson-Windmeijer City*Div	186.097	186.088	186.088	186.097

Notes: This table presents cross-sectional regressions of instrumented urban growth and the interaction of instrumented urban growth and historical diversity on contemporary productivity outcomes. All regressions include state fixed effects. Light density outcomes are calculated in 2013, for comparison to Montalvo and Reynal-Querol (2021). Column 1 measures log light density, column 2 is a log measure of light density over a WorldPop estimated population for the grid in 2010, column 3 measures the change in log lights per capita from 1992 to 2013. Column 4 measures the number of conflict events in the grid since 1997, measured in ACLED battle events. “City” is an indicator equal to 1 if an Africapolis city is located within the grid and the population is above 20,000. “Diversity” is a historical measure of diversity calculated as the fractionalization of land share of different Murdock ethnic groups in the grid cell. All regressions control for distance to coast, malaria and TseTse suitability, ruggedness, distance to a major river, agricultural land productivity, and a historical estimate of population size in 1800. * p<0.10, ** p<0.05, *** p<0.01.

Table 1.13: 2SLS of Historical Diversity and Productivity Instruments: Mineral Deposits and Portage

Panel A: Mineral Deposit				
	(1) Light Density	(2) Lights/Capita	(3) Growth Lights/Capita	(4) Conflict
City	7.526 [0.271]***	-1.325 [0.271]***	0.538 [0.340]	105.514 [13.407]***
Diversity	0.212 [0.148]	-1.058 [0.148]***	-0.002 [0.186]	16.038 [7.319]**
City*Diversity	-2.012 [1.008]**	3.390 [1.009]***	0.359 [1.268]	-116.340 [49.963]**
Mean Dep.	-1.78	-9.88	-1.80	8.07
Observations	22,682	22,680	22,680	22,682
Cragg-Donald F-Stat	306.111	306.087	306.087	306.111
Sanderson-Windmeijer City	318.667	318.634	318.634	318.667
Sanderson-Windmeijer City*Div	421.934	421.894	421.894	421.934

Panel B: Portage Propensity				
	(1) Light Density	(2) Lights/Capita	(3) Growth Lights/Capita	(4) Conflict
City	-1.424 [1.350]	-25.776 [7.386]***	-24.468 [7.549]***	-11.948 [90.547]
Diversity	0.539 [0.154]***	-2.138 [0.845]**	-0.102 [0.863]	1.241 [10.345]
City*Diversity	-5.528 [1.478]***	15.076 [8.097]*	1.820 [8.276]	24.295 [99.161]
Mean Dep.	-1.97	-9.34	-1.76	6.89
Observations	32,933	32,931	32,931	32,933
Cragg-Donald F-Stat	5.441	5.451	5.451	5.441
Sanderson-Windmeijer City	5.897	5.908	5.908	5.897
Sanderson-Windmeijer City*Div	91.601	91.596	91.596	91.601

Notes: This table presents cross-sectional regressions of instrumented urban growth and the interaction of instrumented urban growth and historical diversity on contemporary productivity outcomes. All regressions include state fixed effects. Light density outcomes are calculated in 2013, for comparison to Montalvo and Reynal-Querol (2021). Column 1 measures log light density, column 2 is a log measure of light density over a WorldPop estimated population for the grid in 2010, column 3 measures the change in log lights per capita from 1992 to 2013. Column 4 measures the number of conflict events in the grid since 1997, measured in ACLED battle events. “City” is an indicator equal to 1 if an Africapolis city is located within the grid and the population is above 20,000. “Diversity” is a historical measure of diversity calculated as the fractionalization of land share of different Murdock ethnic groups in the grid cell. All regressions control for distance to coast, malaria and TseTse suitability, ruggedness, distance to a major river, agricultural land productivity, and a historical estimate of population size in 1800. * p<0.10, ** p<0.05, *** p<0.01.

1.9 Appendix 1: Simulations of Shift-Share Instruments with Nonlinearity

In this section, I run a simulation of the instrumental variable strategy to observe how the double instrumenting procedure performs when the true coefficients for labor size and diversity are known. The set-up of the simulation is as follows. Suppose there are 500 regions, sourcing population from 50 homelands or origins. Each destination d receives a flow of migrants from homeland o , according to:

$$X_{od} = \alpha_1 Z_{od} + \alpha_2 U_d + \theta V_d + \psi W * V_d + \epsilon_{od} \quad (1.27)$$

Where Z_{od} is an exogenous component of origin flows that is normally distributed, but origins vary in size and volatility $Z_{od} \sim N(10 * o, 3 * o)$. As o is assigned numerical values 1 to 50, higher values have higher means and standard deviations. U_d is an unobserved omitted variable that is specific to a given region, $U_d \sim N(10, 5)$, and epsilon is $\epsilon_{od} \sim N(0, 1)$. We can think of U_d as a bias that captures the universal "attractiveness" of particular regions for people migrating from any homeland. In addition, I include two parameters that govern the relationship between total labor and total diversity in region d . V_d is distributed normally $V_d \sim N(0, 1)$, and W is a weighting vector that assigns 1 to the first homeland, and 0 to the rest, creating a skew in flows towards homeland 1. The strenght of these forces are governed by θ and ψ parameters. The relative strength of Z_{od} and the unobserved variable U_d are governed by α_1 and α_2 .

Given these numbers, the real labor supply and diversity for a destination city d can be calculated as:

$$L_d = \sum_o^O X_{od} \quad (1.28)$$

As before, diversity is calculated as the herfindahl index (HHI), which is a nonlinear function

of X_{od} .

$$Div_d = \sum_o \left(\frac{X_{od}}{L_d} \right)^2 \quad (1.29)$$

These aggregate components map into a city-level outcome Y_d following:


$$Y_d = \beta_1 \log(L_d) + \beta_2 \log(Div_d) + \beta_3 \log(U_d) + \epsilon_{od} \quad (1.30)$$

Each of these aggregates will introduce bias into the second stage equation via the unobserved feature or "attractiveness" of destinations d captured by U_d and V_d . Using 2SLS I predict L_d and Div_d using just the Z_{od} components, which are used for instruments in a zero stage as follows:

$$\widehat{X_{od}} = \alpha_3 Z_{od} + \epsilon_{od} \quad (1.31)$$

Using the predicted $\widehat{X_{od}}$ from this equation, I then instrument for the aggregates L_d and Div_d by replacing the real X_{od} flows with the predicted values. Figure 1.11 shows the results comparing OLS and SSIV beta coefficients in a model with 500 regions, 50 homelands, and a true β_1 of 5, and a true β_2 of -1. I also set α_1 and α_2 to 0.5, which sets the strength of the instrument Z_d relative to bias U_d . The correlation between L_d and Div_d is controlled by parameters $\theta = 1, \psi = 2$.


Using two instruments requires that the instruments should not be linearly correlated with each other. In practice, L_d and Div_d may be correlated, as more attractive destinations also draw in a wider array of migrant groups. To explore how the beta estimator changes as the relationship between the instruments becomes stronger, I alter the model to create variation in the amount of correlation between L_d and Div_d . I draw θ and ψ from a gamma distribution such that $\theta \sim \text{Gamma}(60, 10)$ and $\psi \sim \text{Gamma}(200, 50)$. I also alter the weighting vector such that W is $[1, 2, 5, 0, 0, 0 \dots]$. Figure 1.12 shows the distribution of beta coefficients for the SSIV model of the diversity variable. The boxplots are separated into quantiles of the absolute value of the correlation between L_d and Div_d . As the correlation between the two aggregate variables moves towards 1 the accuracy of the SSIV estimator is reduced. Figure 1.13 shows the same result for the labor beta coefficient. Figure 1.14 plots how the F-statistic for the diversity instrument declines as the correlation between L_d and Div_d increases.



Notes: This figure shows the distribution of beta coefficients for migrant labor size L_d and migrant diversity Div_d , using both SSIV and OLS. The simulation is run 500 times, and the dotted lines show the real values for the beta coefficients, which are 5 for labor and -1 for diversity. The OLS model regresses $Y_d = \beta_1 \log(L_d) + \beta_2 \log(Div_d)$, while the SSIV uses aggregates calculated from predicted X_{od} flows. 500 destinations and 50 origins are included. The parameters used are $\alpha_1 = \alpha_2 = 0.5$, $\theta = 1$, $\psi = 2$. The weighting vector W is $[1, 0, 0, \dots]$.


Figure 1.11: Simulation of OLS and SSIV with Ommitted Variable Bias

Moving the linear correlation between the instruments from 0.5 to 0.7 reduces the F-statistic by nearly half.




Notes: This figure shows the distribution of beta coefficients for migrant diversity Div_d , using SSIV as the correlation between the linear and nonlinear instrument varies. Box plots are grouped by quantiles of the absolute correlation coefficient between L_d and Div_d . The average of the absolute value of the correlation coefficient is shown on the x-axis. The simulation is run 500 times, and the dotted lines show the real values for the beta coefficient. The SSIV model regresses $Y_d = \beta_1 \log(L_d) + \beta_2 \log(Div_d)$, where the aggregates calculated from predicted X_{od} flows. 500 destinations and 50 origins are included. The parameters used are $\alpha_1 = \alpha_2 = 0.5$. The correlation between L_d and Div_d is governed by $\theta \sim \text{Gamma}(60, 10)$ and $\psi \sim \text{Gamma}(200, 50)$. The weighting vector W is $[1, 2, 5, 0, 0, 0, \dots]$.

Figure 1.12: SSIV Diversity Beta Coefficients with Varying Instrument Correlation



Notes: This figure shows the distribution of beta coefficients for migrant labor size L_d , using SSIV as the correlation between the linear and nonlinear instrument varies. Box plots are grouped by quantiles of the absolute correlation coefficient between L_d and Div_d . The average of the absolute value of the correlation coefficient is shown on the x-axis. The simulation is run 500 times, and the dotted lines show the real values for the beta coefficient. The SSIV model regresses $Y_d = \beta_1 \log(L_d) + \beta_2 \log(Div_d)$, where the aggregates calculated from predicted X_{od} flows. 500 destinations and 50 origins are included. The parameters used are $\alpha_1 = \alpha_2 = 0.5$. The correlation between L_d and Div_d is governed by $\theta \sim \text{Gamma}(60, 10)$ and $\psi \sim \text{Gamma}(200, 50)$. The weighting vector W is $[1, 2, 5, 0, 0, 0, \dots]$.

Figure 1.13: SSIV Labor Beta Coefficients with Varying Instrument Correlation



Notes: This figure shows the distribution of F-statistics for the instrument of migrant diversity Div_d , using SSIV as the correlation between the linear and nonlinear instrument varies. Box plots are grouped by quantiles of the absolute correlation coefficient between L_d and Div_d . The average of the absolute value of the correlation coefficient is shown on the x-axis. The simulation is run 500 times, and the dotted lines show the real values for the beta coefficient. The SSIV model regresses $Y_d = \beta_1 \log(L_d) + \beta_2 \log(Div_d)$, where the aggregates calculated from predicted X_{od} flows. 500 destinations and 50 origins are included. The parameters used are $\alpha_1 = \alpha_2 = 0.5$. The correlation between L_d and Div_d is governed by $\theta \sim \text{Gamma}(60, 10)$ and $\psi \sim \text{Gamma}(200, 50)$. The weighting vector W is $[1, 2, 5, 0, 0, 0, \dots]$.

Figure 1.14: F-Statistic by Varying Instrument Correlation

1.10 Appendix 2: Additional Tables & Figures

1.10.1 Supplementary Analysis

The following section presents additional tables and figures that are referenced in the text.

1.10.2 Birthplace-Level Tests Following Borusyak et al. (2022)

In this section, I conduct the robustness checks of the shift-share model suggested in Borusyak et al. (2022). The authors recommend estimating an IV coefficient from a birthplace-level regression, relating exposure-weighted outcome residuals on exposure-weighted treatment residuals. The procedure accounts for a single instrument model, so I proceed by estimating a labor-only version of the model. The purpose of this exercise is to analyze the identifying shock-level variation, so while the coefficients do not match our headline estimates due to the exclusion of the diversity parameter, they are informative about our shocks g_{ot} .

When generating the birthplace level data, I consider the standardized change in migrant labor size as the treatment, while the shocks are the growth rates at the birthplace level g_{ot} , and the instrument Z is the predicted labor change according to the shift-share. I first plot our residualized average treatment residual across destinations against the birthplace growth shocks in a binned scatterplot. Figure 1.15 shows a strong first stage relationship between the growth shocks and the average treatment residual. Next I plot the average outcome residual against the same binned birthplace shocks. These figures, by outcome, are shown in Figure 1.16. Again, I see a clear positive relationship, as higher level shocks relate to higher residualized outcomes. The benefit of these scatterplots is that we can visualize the distribution of the treatment effect across the distribution of birthplace growth shocks. We see evidence of nonlinearity, with particularly high growth shocks driving a disproportional amount of the variation in our treatment.

Next we calculate adjusted standard errors, as recommended in Borusyak et al. (2022), which are standard errors in a birth-place IV regression of outcome on treatment, weighted by exposure

Table 1.14: African Censuses with Geolocation, Birthplace or Ethnicity

Country	Available Census (ADM2)	Birthplace	Ethnicity	Mother Tongue
Benin	1979 1992 2002 2013	1979 1992 2002 2013	1979 1992 2002 2013	2013
Botswana	1991 2001 2011			
Burkina Faso	1985 1996 2006	1985 1996 2006		2006
Cameroon	1976 1987 2005	1976 1987 2005		
Côte d'Ivoire	1988 1998	1988 1998	1988 1998	
Ethiopia	1984 1994 2007			1994 2007
Guinea	1983 1996 2014	1983 1996 2014		
Ghana	1984 2000 2010		2000 2010	
Kenya	1969 1979 1989 1999 2009 2019	1969 1979 1989 1999 2009 2019		
Malawi	1987 1998 2008 2018		2008 2018	
Mali	1987 1998 2009	1987 1998 2009		1987 1998 2009
Mozambique	1997 2007 2017	1997 2007 2017		2007 2017
Rwanda	1991 2002 2012	2002 2012		
Senegal	1988 2002 2013		1988 2002 2013	1988 2002 2013
Sierra Leone	2004 2015	2004 2015	2004 2015	2004 2015
South Africa	1996 2001 2011 2016		1996 2001 2011 2016	
Tanzania	1988 2002 2012	1988 2002 2012		
Togo	1960 1970 2010		1960 1970 2010	
Uganda	1991 2002 2014	1991 2002 2014	1991 2002 2014	
Zambia	1990 2000 2010	1990 2000 2010	1990 2000 2010	

shares. Figure 1.18 plots the normal clustered robust standard errors from the IV regression, along with the adjusted standard errors. The gray bars represent confidence intervals. As expected, the standard errors are higher in the adjusted version, but the coefficients remain significant.

The last use of the birthplace-level dataset is to test correlations between the shocks and baseline characteristics of destinations, weighted by exposure. Our exclusion restriction requires that there is no significant correlation between pre-period characteristics and the birthplace shocks. To test this, we regress a set of baseline characteristics, averaged across destinations and weighted by exposure shares, against the birthplace-level shocks. The characteristics include pre-period light density measured in 1993, as well as geographic characteristics including soil suitability, TseTse fly suitability, and malaria suitability. The data sources for these variables are enumerated in the data section of the paper. Figure 1.20 plots the beta coefficients from this regression, run separately for each outcome of interest related to light density.

Table 1.15: Gravity Regressions of Migrant Flow using Proxy O-D Panels

	N	N	N
$\ln(\text{distance})$	-1.359 [0.017]***	0.163 [0.093]*	0.312 [0.070]***
$\ln(\text{distance}^2)$		-0.155 [0.009]***	-0.144 [0.007]***
Coethnic Share			8.634 [0.511]***
Mean Dep. Var	65.753	65.753	65.753
Observations	83,661	83,661	83,661
Destination FE	Y	Y	Y
Origin FE	Y	Y	Y
Year FE	Y	Y	Y

Notes: This table shows results from a gravity regression $\pi_{odt} = \exp(\mu_{dt} + \gamma_{ot} + \beta_1 \text{CoethShare}_{odt} + \beta_2 \text{Dist}_{od} + \beta_3 \text{Home}_{od}) + \epsilon_{odt}$. Where π_{odt} is an estimated probability defined as the fraction of individuals from o that appear in destination d at time t . That is, $\pi_{odt} = \frac{M_{odt}}{L_{ot}}$ where odt is the number of people from o in d . The regression includes destination-year and origin-year fixed effects, μ_{dt} and γ_{ot} respectively. Regressions are estimated using Poisson pseudo-likelihood (PPML). Distance is calculated as the log kilometer distance between the centroids of origin and destination administrative regions.

1.10.3 Robustness Checks

In this section we present a series of figures that show robustness for our main shift-share specification. We perform a variety of robustness checks for our results. Figure 1.21 and Figure 1.22 show robustness to Conley spatial standard errors across a variety of distance bandwidths. Figure 1.23 and 1.24 show how our estimates vary when dropping individual countries from the sample. Last we simulate a series of random outmigration shocks and apply them in our shift-share design. Figures 1.25 and 1.26 show how the IV results appear in response to placebo shocks.

1.10.4 Pass-through of International Commodity Prices

Our push shock instrument leverages origin-specific shocks to predict outmigration rates. These predicted outmigration shifts are then used as an instrument to estimate the impact of migration

Table 1.16: Census Birthplace Shift-Share Second Stage, Migrant Flows

<i>Panel A: OLS Results</i>			
	Log(Lights)	Δ Services Share	Housing Quality
$\Delta \ell$	1.036 [0.113]***	1.021 [1.218]	0.049 [0.039]
Δdiv	-9.999 [1.632]***	28.617 [17.441]	0.553 [0.360]
Mean Dep.	-2.49	5.01	-0.17
Observations	829	595	822
<i>Panel B: Shift-Share IV</i>			
	Log(Lights)	Δ Services Share	Housing Quality
$\Delta \ell$	2.529 [0.228]***	1.305 [1.892]	0.377 [0.096]***
Δdiv	-10.425 [2.979]***	10.640 [4.442]**	2.612 [0.850]***
Mean Dep.	-2.49	5.01	-0.17
Observations	829	595	822
Kleibergen-Paap Fstat	16.462	12.581	53.053
Sanderson-Windmeijer ℓ	245.593	161.907	183.079
Sanderson-Windmeijer div	27.686	23.125	27.438

Note: Panel A presents the OLS estimation of changes in migrant labor and diversity on wealth and urbanization outcomes. $y_{d,t} = \alpha \Delta l_{d,t} + \gamma \Delta \text{div}_{d,t} + v_{tc}$. Panel B presents the IV 2SLS estimation of changes in migrant labor and diversity on wealth and urbanization outcomes. Outcomes include the levels of logged light density, change in the labor share in services, and the level of housing quality. All regressions include country-year fixed effects. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

on destination outcomes. We might be concerned that these push shocks are not only correlated across space, but also affect productivity in destinations through other channels. For instance, if a price shock hits an origin area and affects the trade of crops to a port destination, this may be realized in light or wealth growth measures, unrelated to the price shock effects on migration. International price shocks may also drive up the cost of food in destination areas, changing wealth and labor supply in destination unrelated to the migration channel.

In this section, we study the impact of plausibly exogenous push shocks on local price behavior in order to test how shocks may create differences in the relative attractiveness of urban and rural locations. In particular we estimate impulse response functions from local projections that show the impact of international price changes on local urban/rural price dispersion. If international

Table 1.17: Census Birthplace Shift-Share Second Stage, Migrant Flows

<i>Shift-Share IV</i>			
	Δ Industry HHI	Δ Avg. Human Capital	Δ Industry Segregation
$\Delta \ell$	-0.038 [0.016]**	-1.297 [0.290]***	-1.919 [1.682]
Δdiv	-0.524 [0.175]***	1.817 [3.346]	-1.329 [6.515]
Mean Dep.	-0.06	1.59	1.36
Observations	506	529	506
Kleibergen-Paap Fstat	12.380	4.334	12.380
Sanderson-Windmeijer ℓ	122.771	94.947	122.771
Sanderson-Windmeijer div	22.630	4.739	22.630


Note: This table presents the SSIV estimation of changes in migrant labor and diversity on skill and industry concentration outcomes.. $y_{d,t} = \alpha \Delta l_{d,t} + \gamma \Delta \text{div}_{d,t} + v_{tc}$. Industry concentration is measured as the change in the labor share HHI across general industry categories such as construction, mineral production, agriculture, manufacturing and retail trade. Segregation is a measure of the deviation from random of the distribution of migrants from particular origins to particular industries. Human capital changes are measured as the change in the average skill level of workers in a given destination over time, measured in terms of schooling years. All regressions include country-year fixed effects. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

price shocks create wedges in urban/rural prices for affected commodities, this may be evidence of direct effects on destinations and a violation of the exclusion restriction.


Our data on local prices is gathered at the crop-month level from the Famine Early Warning system (FEWS NET) and World Food Program (VAM) (Porteous, 2019). These datasets provide crop-month price observations at a set of geolocated markets across African countries. We code markets as urban or rural based on the local population density measured by Worldpop. We then create a urban-rural price gap as the difference in prices for a crop-month across urban and rural markets in the same country. We then estimate the following local projection. For crop c in state s at month t we estimate:

$$Urban - Rural_{cs,t+h} = \omega International_{ct} + \mu_s + \gamma_c + \epsilon_{cs,t+h} \quad (1.32)$$


Where $Urban - Rural_{cs,t+h}$ is the MoM change in the urban-rural price difference in crop c at horizon h months from t . The variable $International_{ct}$ captures the MoM change in the international




(a) Δ Lights



(b) $\Delta \text{Log(Lights/Capita)}$



(c) Δ Non-Agricultural Share



(d) Δ House Quality


Notes: This figure shows binned scatterplots of birthplace-level treatment residuals against the birthplace-level shocks, organized in 50 bins. The OLS lines of best fit is shown in red. The residualized procedure is accomplished using the method explained in (Borusyak et al., 2022).

Figure 1.15: First Stage Scatterplots in a Birthplace-Level Regression


price of crop c at month t , and crop and state fixed effects are included. Figure 1.27 shows the result of this projection over a 12 month horizon. We don't see evidence of a systematic effect of shocks to MoM international price changes on crop specific urban-rural price differences.

1.10.5 Relationship of Local GDP to Night Light Density


While light density has frequently been used to measure local economic growth in developing countries, the functional relationship between light density and GDP may vary across countries, cities, and production types. A set of papers have sought to verify the relationship between local




(a) Δ Lights



(b) Δ Log(Lights/Capita)



(c) Δ Non-Agricultural Share




(d) Δ House Quality

Notes: This figure shows binned scatterplots of birthplace-level treatment residuals against the birthplace-level shocks, organized in 50 bins. The OLS lines of best fit is shown in red. The residualized procedure is accomplished using the method explained in (Borusyak et al., 2022).

Figure 1.16: Reduced Form Scatterplots in a Birthplace-Level Regression

economic change and light density using either independent sources of local GDP data (McCord and Rodriguez-Heredia, 2022) or other GDP proxies such as household wealth (Chiovelli et al., 2023). In African cities subnational GDP estimates are particularly scarce. Past attempts to verify the relationship of subnational GDP to light density in Africa have generally relied on household surveys. In this section I instead use data on port activity as an independent proxy for city-level GDP, and study how accurately differences in port activity are captured by light density.

The IMF hosts a database of aggregate port activity, measured by tracking ship-level satellite signals as they enter and leave ports. Using methods described in (Arslanalp et al., 2021), the




Notes: This figure shows beta coefficients and confidence intervals for both the normal IV regression used in the main tables, and the adjusted standard errors recommended by Borusyak and co-authors. The procedure is done separately for each light outcome of interest. The treatment are in standardized units of the raw changes in migrant labor size.

Figure 1.18: Standard Error Adjustments Following (Borusyak et al., 2022)

authors use features of a particular ship's size, capacity and draft to estimate the size of a ship's payload, and changes to the payload as they leave particular ports. This method allows the authors to estimate a daily index of estimated trade flows at a particular port. While the index produced is dimensionless it allows for comparisons of shipping traffic within and across different port locations.

Figure 1.28 documents the relationship between city-level estimates of VIIRS night light density and port activity for all recorded ports in Africa. I show associations of monthly and annual aggregations of avg night light density with export and import sizes measured by the IMF Port Watch index. The figure shows a positive relationship between relative port activity and light density. Ports that are more active in terms of shipping traffic also have higher light density. The relationship appears to attenuate for larger cities, suggest that light density as a proxy for economic activity becomes less accurate at the high end, as marginal changes in GDP produce less observable change in light density.



Notes: This figure shows beta coefficients from regressions of average residualized destination characteristics on birthplace shocks. The baseline destination characteristics used are pre-period light density (measured in 1993, prior to treatment for all groups), soil suitability, TseTse fly suitability, malaria suitability. Sources for these variables can be found in the data section on “Pull Characteristics”. The shocks are enumerated in raw population counts.

Figure 1.20: Balance at Baseline in Birthplace-Level Regressions

Table 1.18: Self-Reported Identity by Migrant Status and Distance

	Within Ethnicity		Within Destination		Migrants Only	
	Nationalism	Expectations	Nationalism	Expectations	Nationalism	Expectations
Migrant==1	0.166 [0.238]	0.188 [0.131]	0.201 [0.237]	0.089 [0.096]		
Migrant*Population	-0.038 [0.034]	-0.016 [0.016]	-0.034 [0.036]	-0.018 [0.016]		
Migrant*Distance	0.011 [0.033]	-0.028 [0.020]	0.011 [0.027]	-0.006 [0.020]		
Migrant*CoethnicShare	-0.054 [0.086]	0.051 [0.070]	-0.131 [0.087]	0.066 [0.061]		
ln(O-D Distance km)	0.038 [0.026]	0.014 [0.016]	0.026 [0.021]	0.002 [0.019]	0.046 [0.028]	-0.012 [0.018]
ln(Population)	0.040 [0.026]	-0.008 [0.025]				
Coethnic Share	0.048 [0.079]	-0.074 [0.089]	0.104 [0.078]	-0.079 [0.053]	0.067 [0.066]	-0.014 [0.052]
Mean Dep. Var	3.539	2.809	3.539	2.809	3.625	2.827
Observations	41,033	41,022	41,030	41,021	18,210	18,241
Destination FE	N	N	Y	Y	Y	Y
Ethnicity FE	Y	Y	N	N	N	N
Migrant Only	N	N	N	N	Y	Y


Notes: The data comes from Afrobarometer surveys, linking individuals to their origin based on reported ethnicity. The outcomes of interest are national identity and economic expectations. National identity, or “nationalism” is an individual’s response to a question about how much the individual identifies with the nation relative to their ethnic group. Future expectations is a question styled after the Michigan Consumer Confidence Survey, which asks individuals if they expect economic conditions to improve in the next year. All regressions include country-year fixed effects. Standard errors are clustered at the Afrobarometer sampling cluster level. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.19: OLS Cross-Section of Historical Diversity and Urban Density

<i>Panel A: Urban Dummy</i>				
	(1) Light Density	(2) Lights/Capita	(3) Growth Lights/Capita	(4) Conflict
City	1.548 [0.017]***	-0.465 [0.047]***	-0.383 [0.056]***	45.330 [2.153]***
Diversity	0.043 [0.023]*	-0.542 [0.065]***	0.200 [0.079]**	3.718 [3.006]
City*Diversity	-0.720 [0.071]***	0.460 [0.198]**	0.030 [0.240]	3.560 [9.159]
Mean Dep.	-2.00	-9.96	-1.89	7.81
Observations	28,656	28,654	28,654	28,656


<i>Panel B: Urban Population</i>				
	Light Density	Lights/Capita	Growth Lights/Capita	Conflict
2010 Population	0.136 [0.002]***	-0.864 [0.002]***	-0.528 [0.006]***	3.071 [0.265]***
Diversity	-0.478 [0.092]***	-0.478 [0.092]***	0.119 [0.250]	-33.902 [10.852]***
Population*Diversity	0.043 [0.010]***	0.043 [0.010]***	0.037 [0.028]	4.164 [1.196]***
Mean Dep.	-2.00	-9.96	-1.89	7.81
Observations	28,654	28,654	28,654	28,654

Notes: This table presents cross-sectional regressions of urban growth and the interaction of urban growth and historical diversity on contemporary productivity outcomes. All regressions include state fixed effects. Light density outcomes are calculated in 2013, for comparison to Montalvo and Reynal-Querol (2021). Column 1 measures log light density, column 2 is a log measure of light density over a Worldpop estimated population for the grid in 2010, column 3 measures the change in log lights/capita from 1992 to 2013. Column 4 measures the number of conflict events in the grid since 1997, measured in ACLED battle events. "City" in Panel A is an indicator marked as 1 if an Africapolis city is located within the grid and the population is above 20,000. In Panel B, this term is replaced by a log population estimate for 2010 from Worldpop. "Diversity" is a historical measure of diversity calculated as the fractionalization of land share of different Murdock ethnic groups in the grid cell. All regressions control for distance to coast, malaria and TseTse suitability, ruggedness, distance to a major river, agricultural land productivity and a historical estimate of population size in 1800. * p<0.01, ** p<0.05, *** p<0.01.




Notes: This figure plots the results of IV regressions of the main estimating equation where standard errors account for spatial autocorrelation. In particular the figure plots conley standard errors for the migrant labor size coefficient with different bounds on the decay of spatial autocorrelation.

Figure 1.21: Conley Spatial Standard Errors for Labor Size Coefficient $\Delta \ell$




Notes: This figure plots the results of IV regressions of the main estimating equation where standard errors account for spatial autocorrelation. In particular the figure plots conley standard errors for the diversity coefficient with different bounds on the decay of spatial autocorrelation.

Figure 1.22: Conley Spatial Standard Errors for Migrant Diversity Coefficient Δ_{div}




Notes: This figure plots the results of IV regressions of the main estimating equation where individual countries are dropped from the sample. This figure shows results for the migrant labor size coefficient.

Figure 1.23: Estimates with Individual Country Drops for Labor Size Coefficient $\Delta \ell$




Notes: This figure plots the results of IV regressions of the main estimating equation where individual countries are dropped from the sample. This figure shows results for the migrant diversity coefficient.

Figure 1.24: Estimates with Individual Country Drops for Migrant Diversity Coefficient Δdiv




Notes: This figure plots the results of IV regressions of the main estimating equation where g_{ot} is replaced with randomly drawn shocks. This figure shows results for the migrant labor size coefficient. The x-axis tracks individual draws of the placebo shocks.

Figure 1.25: Placebo Shocks for Labor Size Coefficient $\Delta \ell$




Notes: This figure plots the results of IV regressions of the main estimating equation where g_{ot} is replaced with randomly drawn shocks. This figure shows results for the diversity coefficient. The x-axis tracks individual draws of the placebo shocks.

Figure 1.26: Placebo Shocks for Migrant Diversity Coefficient Δdiv



Notes: This figure shows the results of a local project estimating the effect of an international price shock on the urban-rural price gap for the same commodity across African markets.

Figure 1.27: IRF of International Prices on Local Urban/Rural Gap



Notes: This figure shows the relationship of logged average measures of night light density and log port activity at an annual and monthly level. Data for import and export sizes is taken from IMF Port Watch. Port traffic is estimated using satellite-based measures of ship traffic as described in (Arslanalp et al., 2021). Night light density comes from daily snapshots of VIIRS light intensity.

Figure 1.28: Correlations of Port Activity and Night Light Density

Chapter 2

Historical Immigration and the Market for Schooling in American Cities

2.1 Introduction

Migrants represent an influx of both labor supply and demand for goods. When new migrants arrive, they can demand unique products or amenities, opening new markets and changing the competitive landscape of destinations. The emergence of public-private competition in US education is a clear example of migrants affecting the market structure. In the mid 19th century public elementary and secondary schools were an increasingly common feature of urban American life (Goldin, 1999). These schools were not secular and often featured Protestant values and Bible study. As the age of mass migration began in the 1880s (Sequeira et al., 2017), immigrants from new ethnic backgrounds began to populate the northeast. Catholic immigrants objected to the Protestant curriculum taught in the existing public school system. Founded by the Catholic Church, which was rapidly expanding its presence in major American cities, new parochial private schools were opened to serve Catholic students (Sanders, 1977). For the next 100 years, private schooling in the US was dominated by Catholic institutions, many of which continued to operate in urban areas long after the first and second generation of Catholic migrants had moved out of downtown (McCarthy, 1973).

This paper studies the effects of international migration on US public-private schooling competition in the early 20th century. I find that increased migration from Catholic countries improved native educational attainment relative to areas that saw migration from non-Catholic countries. Using data on school and teacher characteristics from northeastern US cities, I explore different channels through which migration affects local schooling.

Catholic immigrants may have impacted local schooling for native children in a few ways. (1) New immigrants may have encouraged “native flight” away from certain neighborhoods, impoverishing local public schools and depriving those who stayed from high ability peers (Betts and Fairlie, 2003b). (2) Catholic immigrants flocked to private schools, reducing the burden of immigration on the public school system and allowing public schools to thrive with smaller class sizes. (3) The growth in private school demand produced competition in the local schooling market, causing public schools to invest in school performance and teacher quality (Hoxby, 2003). The directional effect in the case of (2) and (3) predicts better public school performance in areas more exposed to Catholic migration, relative to non-Catholic migration. In case (1), the effects of native flight may outweigh potential benefits through competitive channels.

The central empirical question is whether migration-driven demand for religious private schooling changed the long-run market structure of education and the quality of nearby public schooling. I study this question at a neighborhood level, where both parish formation and school entry decisions were made, and where competition between school sectors was most immediate. This local focus is important because city-level averages can mask large within-city differences in migrant settlement and schooling access. The empirical strategy identifies shocks to Catholic migration by leveraging quota policies from the 1920s that differentially affected migration from different origin countries. Neighborhoods were relatively more or less exposed to a drop in Catholic inflows depending on past period settlement shares. Country-level quotas are treated like independent shocks that produce “missing migrants” relative to past flows, and are then used in a difference-in-differences framework to examine the effects on census tract trends after the policy is implemented in 1920. In addition to this cross-sectional variation, I leverage variation in cohort level exposure by comparing

individuals within a neighborhood that were of schooling age, versus students who had already aged out of the schooling market. Using a combination of cohort-level comparisons with varied Catholic quota exposure, I evaluate the educational attainment of native children more or less exposed to the drop in Catholic migrants. Then, I explore possible mechanisms for the change in native outcomes, looking at how exposure to missing Catholic migrants affected public school size, composition and teacher quality.

I make three main contributions. First, while past work on historic migration has mostly focused on the role of assimilation and effects on labor market and productivity outcomes, I focus on the structure of a specific market where migrants played a crucial role in changing the competitive landscape. Second, I combine tract-level shift-share exposure to the 1920s quotas with within-city variation in historical settlement, allowing identification at a neighborhood margin where school competition operated directly. Third, I link these migration shocks to downstream educational outcomes through teacher composition, school staffing intensity, and eventual child attainment. Using linked census records, this paper is able to overcome empirical concerns around post-treatment sorting of individuals across space.

This project contributes to a recent literature on the impact of historic US migration on local economic growth, labor markets, city mortality and other outcomes (Abramitzky et al., 2023; Ager et al., 2023; Abramitzky and Boustan, 2017; Abramitzky et al., 2025). Past work often leverages regional exposure to major immigration shocks such as the quota policies of the 1920s, which limited migration from particular countries. This work has generally found positive benefits of migrants in the age of mass migration (1860-1920) in terms of county incomes and urbanization. At the same time, areas that received migrants also experienced disproportionate rates of mortality and disease (Ager et al., 2023). A related literature considers the impact of migrating on the migrants themselves, especially on propensity to assimilate and labor market returns (Abramitzky et al., 2020b,c; Gagliarducci and Tabellini, 2022; Abramitzky et al., 2025). Abramitzky et al. (2025) and Gagliarducci and Tabellini (2022) are two recent papers that focus on the role of ethnic enclaves and the Catholic Church in particular in preventing assimilation for Catholic ethnic groups who

were able to maintain cultural identities from their home country. I make a few key contributions to the literature. My paper looks at the impact of immigrants on the market for schooling; while the impact of immigrants on growth is well documented, we know less about how immigrants shape demand and provision of public goods like education. Next, my paper focuses on the role of migrant composition as an important dimension of migrant effects. Catholic and non-Catholic migrants had different effects on the neighborhoods they settled in, plausibly due to differences in their skills and preferences. By comparing migrants from different birthplaces, I'm able to study questions related to migrant composition rather than purely migrant shares. Lastly, my paper leverages within-city variation, which is more granular than past papers that typically consider county-level effects. The enumeration districts utilized in this paper are on average smaller than contemporary census tracts. This allows for a more detailed picture of the impact of immigration at a neighborhood level.

Second this project contributes to work on the contemporary impacts of immigrants on public and private schooling (Boustan et al., 2023; Betts and Fairlie, 2003a; Murray, 2016). Work in education economics has looked at the effects of “native flight” away from public schools and towards private schools in response to new migrants. This literature finds evidence that native students leave public schools as new migrants enter, increasing demand for private schooling, with potentially long lasting effects on the quality of local public education. I contribute to this literature by studying a historical example in which immigrants increased demand for private schooling. I can then examine the spillover effects on public schools over a longer time horizon.

Third this paper contributes to a literature on the relative impact of private schools on public schools (Urquiola, 2016; Bagde et al., 2022; Neal, 1997). Past work on US private-public competition already acknowledges the significant role of Catholic religious institutions in the competitive schooling market. Despite a significant decrease in Catholic school shares of enrollment after the 1960s, many private schools in the US continue to have a religious mandate. Some work in this space has used historical Catholic shares as an instrument for contemporary private school presence in different regions (Cohen-Zada, 2009). I add to this past work by leveraging plausibly exogenous immigration shocks that shifted demand for private schools, allowing me to study both

contemporaneous and long term effects of new private schools on local public school quality and size.

Lastly this paper contributes to an economic history literature on the emergence of the public school system in the US (Goldin and Katz, 1999; Gross, 2017). While much has been written about the explosion of mass schooling in the late 19th century US and its economic impact, less is known about how the public school system emerged, and the factors contributing to its success, especially in urban areas. Immigrants play an important part in the early political economy of mass schooling, and may have had a long term impact on the relative performance of public schools in US cities.

Previewing the results, I first find that neighborhoods with higher historical Catholic concentration have persistently lower public-school enrollment shares by 1960, consistent with long-term persistence of private schooling demand. Second, native children residing in neighborhoods that were exposed to more Catholic migrants in 1920 had relatively higher educational attainment compared to neighborhoods exposed to non-Catholic migrants. This suggests that the Catholic composition of migrant inflows had a meaningful impact on local schooling outcomes. I then explore possible mechanisms for these results, including the effects of Catholic migrants on local teacher characteristics and schooling quality in public and private elementary schools. I find that Catholic migration increased the fraction of local Catholic teachers, but did not meaningfully affect measures of teacher quality. I find that local public schools in areas with more Catholic migrants saw improved teacher/pupil ratios, suggesting that smaller class sizes was a plausible channel through which migration and public school openings affected local public schools.

The rest of the paper proceeds as follows. Section 2 describes the historical census, school, and linked individual data. Section 3 provides institutional background on Catholic migration and parochial school expansion. Section 4 presents the identification strategy based on quota-induced variation in Catholic inflows. Section 5 reports main results and robustness checks. Section 6 concludes with implications for migration, education markets, and the political economy of public goods.

2.2 Data

This paper combines data from the census and school directories to produce a decade-level panel of census tract level characteristics over time, and the characteristics of schools that opened in those tracts. Data on native workers, teachers and immigrant shares are taken from the 100% and 1% IPUMS USA Census files for the period 1880-1960. The universe of individuals is available for the period 1880-1940. Until 1940, I leverage the census designated enumeration districts as the geographic unit of interest. Enumeration districts are historical delineations capturing an area that a census taker could survey within two weeks. These districts are usually smaller than contemporary census tracts. In large cities enumeration districts are comparable to neighborhoods or city blocks. The historical census records contain limited information on individuals. For this project, I focus on socioeconomic information available including occupation, birthplace and literacy, as well as educational attainment which is available in 1940.

Using the enumeration districts to trace locations over time requires maps of where these districts were located in a given census sample year and city. For this I use data from the Urban Transition Historical GIS project, which has digitized a number of enumeration district maps for northern cities and linked them to the census files (Shertzer et al., 2016; Logan et al., 2011). This project has been used by historians and social scientists to study segregation and racial patterns in the US at a granular level (Shertzer, 2016a).

Data on school locations and characteristics come from a variety of archival and digital sources. Catholic schools are digitized from the Official Catholic Directory for each year. The Catholic Church publishes an annual list of all established churches, including the street address. As part of these records, the Church includes associated parochial schools, which are typically attached to a particular Church parish, though they are often located in a separate building with a separate address. The directories also include information on the number of sisters or fathers that attend the school, as well as lay (secular) teachers. The records also typically include the number of boys and girls that attend the school in that year.

Public school data is more difficult to find for the period. Each city keeps directories of schools in different formats, or provides information on school addresses in annual education reports. In this paper I use data from the directory of public schools of Baltimore Maryland, which were regularly published between 1910 and 1960. I also use the New York City Board of Education directories of teachers, which are available for select years in the period from the New York City Municipal Archives. These sources provide information on each public school's address in that year. They also include a list of all teacher names, and the schools in which they work.


For both the Catholic and public schools, I geocode the addresses to extract longitude and latitude points of the school's location. In cases where a Catholic school address is not provided, I use the address of the associated Catholic Church.

2.2.1 Enumeration Districts and 1960 Tracts

Each city has a unique enumeration district map for each census year (1880–1940). A subset of these maps have been digitized and linked to historical census records by the Urban Transition Historical GIS Project for major cities including NYC, Baltimore, Pittsburgh, Philadelphia, Boston and Chicago (Shertzer, 2016b; Logan et al., 2011). Figure 2.1 To trace outcomes over time, I match enumeration districts to the census tracts from the 1960 census. Data at the enumeration district level is aggregated to the tract level, weighted by the share of each tract overlapped by a particular enumeration district. Several city-years of enumeration districts are still remaining to be digitized from available ED maps, including Manhattan 1880. Because enumeration districts are on average smaller than tracts, another strategy will be to match all districts to the 1940 ED map for a more granular panel of locations in the city. In what follows, observations are aggregated at the 1960 tract level unless otherwise specified.

2.2.2 Teacher-Level Outcomes

I am interested in the impact of Catholic migration and private school openings on public school quality. One measure of public school quality is the quality of public school teachers. I use



Notes: This figure shows an enumeration district map for the borough of Brooklyn in 1950. Numbers correspond to the tracts in the census files. Maps were digitized by the Urban Transition Historical GIS project for NYC, Baltimore, Pittsburgh, Philadelphia, Boston and Chicago (Shertzer, 2016b; Logan et al., 2011).

Figure 2.1: Enumeration District Map Example

observable characteristics of teachers residing in a particular enumeration district as a proxy for local school quality. This proxy is imperfect, as teachers do not necessarily teach in the enumeration district they reside in.

Typically teacher quality can be measured using educational outcomes for the teacher. Educational attainment is available for the 1940 100% census sample, but not for prior years. To get a continuous measure over census years, I use literacy and occupational score measures from linked spouses, mothers and fathers. This restricts the teacher sample to those who live with a spouse or parents in a given census year. The intuition is that by assortative matching, a teacher's education and socioeconomic status will be correlated with their spouse. It's also natural to expect that higher income or higher literacy parents are associated with teachers of higher educational and socioeconomic status. Other outcomes of interest include a teacher's birthplace and age. I consider the impact of Catholic migration on the composition of teachers within an enumeration district, including the share of teachers from Catholic countries.

Figure 2.2 shows the change in teacher Catholic share and teacher spousal occupational score over the time period of interest (1880–1940). As public schools became increasingly common


during this period, the number of Catholic teachers declined and then grew in key cities such as New York and Chicago. Over time, spouses of teachers rose in socioeconomic status, with the exception of Boston and Pittsburgh. This suggests that teacher socioeconomic status changed differentially across cities.

2.2.3 School-Level Outcomes


In addition to teacher specific outcomes, I am also interested in understanding overall changes to the schooling market in a given tract or district. To capture this I use geocoded public and Catholic schools to calculate distance to the nearest public and private school for a given tract. I also calculate the number of private and public schools within a radius of each tract centroid and the size of individual public private schools, estimated by personnel and classroom numbers for public schools, and pupil totals for Catholic schools. For the public schooling results, my sample is restricted to records from Baltimore, where clean public school directories allow me to capture changes to teacher/pupil ratios over time. Table 2.1 provides summary statistics of the collected school-level data by sector. Catholic schools report 3 types of teachers, including priests (brothers), nuns (sisters) and non-religious faculty, or “lay teachers”. As expected, private Catholic schools are relatively smaller than public schools. However, the teacher/pupil ratios are generally comparable in the cross-section.

2.2.4 Linked Census Sample

To understand impacts on children, I leverage linked data from the Census Tree project (Buckles et al., 2025) that connects individuals across censuses over time. This project leverages crowd-sourced data from the archival “Family Search” platform where users trace their ancestry and link census records to relatives. Leveraging this data, Census Tree has achieved match rates across census between 69 and 86% for men, and between 58 and 79% for women. The linked records allow me to capture an observed set of adults in 1940 that were children at the time of major Catholic immigration policy changes in 1920, and resided in Northeastern cities. For these individuals I



(a) Catholic Share



(b) Spouse Occupational Score

Notes: This figure shows the change in teacher characteristics over census years. The spousal occupational score is a measure of the socioeconomic status of a spouse's reported occupation. The share Catholic is defined as the share of teachers that report a foreign birthplace that is a majority Catholic country. Each line represents a large city recipient of immigrants where enumeration district level data is available. Cities included are NYC, Baltimore, Pittsburgh, Philadelphia, Boston and Chicago.

Figure 2.2: Change in Teacher Characteristics Over Time

Table 2.1: School Characteristics By Sector

	Catholic	Public	Difference	SE
Panel A: Catholic–Public Differences				
Pupils	555.286	702.187	146.901***	(26.663)
Teachers	13.854	18.878	5.025***	(0.867)
Teachers per Pupil	0.026	0.028	0.003	(0.002)
Panel B: Catholic School Characteristics				
	Mean	SD	N	
Sisters	11.730	6.881	404	
Brothers	4.761	3.303	67	
Lay Teachers	5.872	5.985	172	

Notes: This table reports average characteristics of Catholic and public elementary schools in 1910–1940. Catholic school data come from the Official Catholic Directory and include records from New York, Baltimore, Philadelphia, and Boston. Public school data come from Baltimore Board of Education directories. T-test comparisons of school size and teacher-pupil ratios are reported in the final two columns of Panel A. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

observe their educational attainment in 1940, measured in terms of schooling years, along with their age in 1940 and their tract of residence in 1920. An advantage of the linked census is that it captures individuals who moved away from particular districts as adults. I’m therefore able to evaluate effects on children regardless of post-treatment sorting towards other neighborhoods and cities. Therefore results at the cohort level partially offset concerns about effects through local compositional changes among native children.

2.3 Historical Background: European Migration and Catholic Schooling

In the late 19th century, Europeans began to arrive in the United States in large numbers, often settling in northeastern urban centers like New York, Boston and Baltimore (Abramitzky and Boustian, 2017). Many Europeans came from countries with Catholic majorities like Poland,

Ireland and Italy. In response to the influx, the Catholic Church began to open new churches in urban centers. The Church also coordinated the opening of “ethnic churches” for particular migrant groups. Ethnic parishes serviced Polish and Italian migrant neighborhoods, providing sermons in their respective languages (Abramitzky et al., 2025).

Public schooling in the United States functioned as a locally funded and quasi-monopolistic system that was not religiously neutral. Public schools in many cities incorporated Protestant practices, including Bible reading and prayer, creating a mismatch between the preferences of Catholic immigrants and the existing education system. Catholic immigrants were deeply concerned about their children receiving Protestant instruction in the growing public school sector. In response, parochial Catholic schools emerged as an alternative type of private school where immigrants could send their children (Sanders, 1977).

The emergence of Catholic schooling was a distinct attempt by the Church and families to preserve their particular cultural heritage, while also supporting upward mobility for immigrant children. The Catholic Church was able to rapidly expand schooling capacity due to the availability of low-cost religious labor. Classes were taught primarily by nuns, who typically received little or no formal compensation. This substantially lowered the marginal cost of schooling provision relative to public schools and enabled the expansion of parochial schools even in relatively low-income immigrant communities.

Because both Catholic parishes and public school systems were organized at the local level, the expansion of Catholic schooling generated highly localized variation in educational supply. Parishes operated schools at the neighborhood level, while public schools were funded and governed by local jurisdictions. As a result, exposure to Catholic schooling varied sharply across cities and within cities, generating substantial heterogeneity in the degree of public–private competition.

While parochial schools were initially fragmented in the late 1800s, by the early 20th century they became increasingly centralized and standardized under diocesan authority. Curricula were formalized, governance structures were consolidated, and training for religious teachers improved. These changes facilitated the transition of Catholic schools from loosely organized parish institutions

to a coordinated system capable of serving a large urban population.


The expansion of Catholic schooling also occurred alongside political efforts to standardize public education and limit the role of private providers. Compulsory schooling laws and broader reform movements sought to reinforce the dominance of public schools. However, legal challenges ultimately constrained the ability of the state to suppress private education, most notably in *Pierce v. Society of Sisters* (1925), which affirmed the right of private schools to operate. These conflicts contributed to the emergence of a regulated but persistent mixed system of public and private education.

Catholic school enrollment continued to increase through the early 1900s. Figure 2.3 plots the increase in the number of students attending Catholic elementary schools, which peaked in the early 1960s. By 1920, Catholic schools represented a significant alternative to public schools, breaking the prior monopoly of the public school system and introducing one of the first large-scale systems of school choice in the United States (Gross, 2017). In this sense, Catholic immigration generated an endogenous shift from a largely unified public system to a mixed public–private equilibrium driven by heterogeneity in preferences and institutional responses.

The growth of Catholic schooling potentially affected public education through multiple channels, including enrollment substitution, changes in public school resource allocation, and competitive responses in quality. These mechanisms form the basis for the empirical analysis that follows.

2.4 Empirical Method

In this section I outline a strategy to capture the causal effect of increasing Catholic migration on the opening of Catholic schools, and the downstream effects on public school quality and native child human capital. I define a measure of tract-level exposure to Catholic migrants. I then compare tracts differentially exposed to a 1920 migration restriction policy in a difference-in-differences (DiD) framework.



Notes: This figure shows the change in the share of elementary school children attending Catholic schools from 1920 to 2020. The figure illustrates the steady rise in enrollment through the early and mid-20th century, followed by a sharp decline after the 1960s. Data are from the Digest of Education Statistics at the National Center for Education Statistics.

Figure 2.3: Change in Catholic School Enrollment Share

The key identifying assumption is that, absent the quota policy, tracts with different values of predicted Catholic quota exposure would have followed similar trends in schooling market outcomes and child attainment. The policy-induced national shock provides plausibly exogenous time variation, while pre-period settlement shares provide cross-sectional exposure. Because settlement patterns may capture persistent neighborhood differences, all specifications include rich fixed effects and baseline controls to absorb time-invariant tract heterogeneity and common city-wide shocks.

The empirical challenge to studying the relationship of local Catholic share and schooling is that Catholics may have selected into destinations where public schools and human capital are on a different trend relative to areas with less Catholic migrants due to preexisting local economic factors. For example, Catholic migrants were often poorer than native workers, and tended to reside in affordable, lower income neighborhoods within Northeastern cities. Table 2.2 provides comparisons of Catholic and Non-Catholic immigrants to native workers. Catholic workers are defined as all immigrants from Ireland, Italy, Poland, Mexico, and the Austro-Hungarian Empire.

The comparisons include fixed effects for city, and a control for age. Relative to native workers, Catholic immigrants were older and less educated. Compared to Non-Catholic migrants, Catholics are less literate and less wealthy, as measured by Census occupational scores. The occupational scores present a stark difference in socioeconomic status between average non-Catholic migrants and Catholic migrants. Non-Catholic migrants were closer in occupational score to native workers. This is partly because non-Catholic migrants tended to have settled in the US before Catholic migrants. Some migrants, such as those from England, were also more readily assimilated.

Table 2.2: Immigrant Characteristics in 1880 (Northeast Cities)

	Literacy	Occscore	School	Age
Catholic	-0.144 [0.000]***	-1.072 [0.022]***	-0.100 [0.001]***	19.321 [0.023]***
Non-Catholic Imm	-0.009 [0.000]***	0.095 [0.021]***	-0.061 [0.001]***	18.966 [0.021]***
Mean Dep.	0.934	16.779	0.227	25.130
Observations	2,828,279	2,258,812	2,374,399	3,789,455

Notes: All regressions have fixed effects for city and gender. The reference category is a native-born individual. Cities included are Baltimore, Boston, Chicago, Cincinnati, Cleveland, Detroit, New York, Pittsburgh. Catholic and Non-Catholic designations are based on whether the reported birthplace is a Catholic majority country. Literacy and Occscore regressions keep only the samples aged 18-65. Literacy is a binary variable, while Occscore is a Census derived measure of socioeconomic status based on linking occupational titles to income categories. "School" measures whether the individual is currently attending a school, and only retains individuals aged 10-29. Age is controlled for everywhere except the final column. Standard errors are in brackets. * p<0.10, ** p<0.05, *** p<0.01.

Figure 2.4 shows comparisons of migrant populations in 1880 by country of origin. Across birthplaces, Catholic origin migrants tend to have lower occupational scores. Literacy is particularly low among Italian and Polish migrants, who made up a large share of the Catholic influx into the country between 1880 and 1920.

In terms of composition, these features suggest that areas with more Catholics were likely to be poorer, and Catholics themselves were not reaching the educational attainment of native students. Instead, what separated Catholic migrants from other groups were the institutions that came with them. Figure 2.9 shows the geocoded expansion of Catholic schools that took place from


1900-1920 in New York. Catholic churches pushed outward, beginning in the central downtown where Catholics first settled and agglomerated, and eventually reaching Brooklyn and Queens.

An ideal experiment would generate exogenous shocks to Catholic migrant share across different neighborhoods or cities, unrelated to local fundamentals. I adapt the strategy of Ager et al. (2023) by leveraging the national-origin quota system introduced in the 1920s. These laws sharply reduced immigration from countries whose migrant stocks were underrepresented in earlier census base years, especially Southern and Eastern European origins. Because Catholic immigrants were concentrated in several of these highly restricted origin groups, the quota system generated differential shocks to the predicted Catholic share of immigrant inflows across neighborhoods with different pre-policy birthplace compositions.


Two threats are especially important. First, quota exposure might proxy for total immigrant decline rather than Catholic-specific composition shifts. I address this by estimating alternative specifications that net out non-Catholic quota exposure and by constructing a relative Catholic change measure. Second, differential pre-trends could bias the DiD estimates if high-exposure neighborhoods were already on distinct trajectories before 1920. I therefore report pre-trend and placebo-style checks and interpret post-1920 estimates in light of these diagnostics.

2.4.1 The 1920s National-Origin Quota System

Driven by rising nativist sentiment after World War I, the United States adopted a sequence of laws that sharply restricted immigration during the 1920s. The first major restriction was the Emergency Quota Act of 1921, which capped annual immigration from each country at 3 percent of the number of foreign-born residents from that country living in the United States in the 1910 Census. The 1924 Immigration Act, also known as the Johnson-Reed Act, made the system more restrictive by lowering the quota to 2 percent and changing the base year from 1910 to 1890. This change was consequential because the immigrant stock in 1890 was much more heavily weighted toward older sending countries in Northern and Western Europe. As a result, the law disproportionately reduced immigration from newer sending countries in Southern and Eastern Europe, including Italy, Poland,



(a) Occupational Score by Birth Country



(b) Literacy by Birth Country

Notes: This figure shows migrant characteristics for immigrants age 18-65 in 1880, residing in major Northeastern cities. The analysis includes migrants residing in NYC, Chicago, Boston, Pittsburgh, Detroit, Baltimore, Cleveland, Cincinnati. Panel A measures the occupational score of migrants on average, while Panel B measures the proportion of literate individuals. Error bars represent 95% confidence intervals.


Figure 2.4: Migrant Characteristics by Birth Country in 1880

Russia, and other origins with large Catholic, Jewish, and Orthodox populations.


The identifying variation comes from the fact that these national-origin quotas were based on historical settlement patterns rather than contemporaneous local economic conditions in receiving neighborhoods. Countries whose U.S. migrant stocks had grown rapidly between 1890 and 1910 faced especially severe quota reductions once the 1890 base was adopted. Italy provides a useful example. Italian immigration expanded substantially in the late nineteenth and early twentieth centuries, but Italians represented a relatively small share of the foreign-born population in 1890. Consequently, the 1924 quota sharply reduced new Italian arrivals. By contrast, countries with longer-established migration streams, such as England, Ireland, or Germany, were less mechanically constrained because their immigrant stocks were already well represented in the 1890 Census.

This policy shock can therefore be used to predict changes in the local composition of immigration across neighborhoods or cities. Places that had previously relied more heavily on origins subject to severe quota cuts experienced larger predicted declines in new immigrant arrivals from those groups after the policy. Because Catholic immigrants were disproportionately drawn from several of the newly restricted Southern and Eastern European origins, the quota system generated plausibly exogenous variation in Catholic migrant exposure across local labor and schooling markets. Importantly, this variation operates through the interaction between national-origin quota intensity and pre-existing local immigrant settlement patterns: two neighborhoods in the same city could face different exposure depending on the birthplace composition of their pre-1920 foreign-born population.

Figure 2.5 documents changes to the foreign-born population in major Northeastern cities from 1880–1940. The figure highlights the rapid growth of newer immigrant origins before the quota regime and the subsequent slowdown after the 1920s border closure. This timing is useful for the empirical design because it separates pre-policy immigrant settlement patterns from the post-policy contraction in flows. The empirical strategy therefore compares neighborhoods with higher versus lower predicted exposure to the quota-induced decline in Catholic immigrant inflows before and after the border closure.



(a) Immigrant Share



(b) Foreign Born Population Size

Notes: This figure shows the change in immigrant composition in the analysis sample, which includes all cities with linked enumeration district maps for the period 1880 to 1940. Cities include NYC, Chicago, Boston, Pittsburgh, Detroit, Baltimore, Cleveland, Cincinnati. Panel A measures the foreign born population in terms of immigrant shares, while Panel B takes raw numbers from the 100% Census files.

Figure 2.5: Immigrant Compositional Change

2.4.2 Catholic Quota Exposure

The quotas instituted in 1921 produced sudden missing flows in immigration patterns. Potential Italian immigrants who would have arrived in 1922 were no longer able to migrate. Leveraging these missing flows, I create a census tract level measure of exposure to the quota policy and the effect on total Catholic inflows. The exposure measure is constructed as a Bartik instrument, combining pre-period settlement shares from particular immigrant groups with changes in inflows as a result of the quota policy. Following Ager et al. (2023), I construct the following measure of Catholic immigrant quota exposure at the tract level d :

$$\text{Quota Exposure}_d = \frac{100}{P_{d,1910}} \sum_{n=1}^N \max(\hat{M}_{n,1922-1930} - Q_{n,1922-1930}, 0) \frac{FB_{nd,1910}}{FB_{n,1910}} \quad (2.1)$$

The term $\hat{M}_{n,1922-1930}$ is a linear prediction of the inflow of migrants from each Catholic country without the quota. This represents a linear extrapolation of the pre period trend in immigrant inflow. In particular, the prediction is estimated for each origin country as:

$$M_{nt} = \beta_1 \ln(t) + \beta_2 \ln(t^2) + \epsilon_{nt} \quad (2.2)$$

Where M_{nt} represents the real flows from country n in a pre-period year t between 1900-1914.

The term $Q_{n,1922-1930}$ is the total quota for a given Catholic country of origin n . This term captures the number of permitted migrants from a particular birthplace, according to their 1890 immigrant share. The difference between $\hat{M}_{n,1922-1930}$ and $Q_{n,1922-1930}$ captures the predicted “missing” flow of immigrants from an origin country, relative to what would have arrived without the quota in place. The term $\frac{FB_{nd,1910}}{FB_{n,1910}}$ captures the fraction of immigrants from Catholic country n that are located in tract d . The construction is comparable to a Bartik instrument that combines pre-period exposure by settlement shares and plausibly exogenous shifts by origin.

The scaling in equation 2.1 is per 100 tract residents in 1910. A one-unit increase in QE_d

therefore corresponds to one additional predicted missing Catholic immigrant per 100 baseline residents over 1922–1930. This scaling is used throughout the tract-level DiD tables, so coefficients can be interpreted as outcome changes associated with this standardized exposure increment.

Given the tract level quota exposure, we can then estimate a DiD model for an outcome y_{dt} as:


$$y_{dt} = \alpha_d + \gamma_t + \beta(QE_d * \text{post}_t) + \Gamma(FB_{d,1910} * \text{post}_t) + \epsilon_{dt} \quad (2.3)$$

In this model, $FB_{d,1910}$ denotes the total foreign born population of tract d , and QE_d is the enumeration district quota exposure. The year t marks the census year of interest at which the outcome is measured. Post_t is a dummy variable marking whether the observation is before or after the 1920 quota policy. α_d and γ_t capture tract and year level fixed effects respectively.

Across the main tract-level specifications, the identifying variation comes from differences in pre-1920 origin composition across tracts within the same city, interacted with a common national policy shock. The key assumption is that, conditional on tract fixed effects, year fixed effects, and baseline controls, high- and low-exposure tracts would not have experienced systematically different post-1920 schooling trends absent the quota shock. In this design, country-level quota assignments are policy-driven and national, while tract-specific exposure is pinned down by pre-period settlement shares measured before implementation.

The coefficient of interest, β , measures whether more exposed tracts differentially changed after quota implementation. For interpretation, a negative β for Catholic share outcomes implies that tracts with larger predicted missing Catholic inflows experienced larger post-1920 declines in Catholic presence. For schooling outcomes, the sign of β indicates whether reduced Catholic inflows shifted public-private market shares, school staffing intensity, or teacher composition.

Figure 2.6 shows an example of the quota exposure measure for tracts in New York City. Catholics were centrally located near downtown New York, in historically low income neighborhoods. Across tracts, there are substantial differences in Catholic exposure, with a diminishing gradient moving towards the upper west side of the city.



Notes: This figure shows the calculation of quota exposure by 1960 census tract. Predicted Catholic inflows are taken from a linear extrapolation of past immigrant inflows. Settlement shares are defined in the 1910 census at the enumeration district level, and then aggregated as a weighted sum to the 1960 tract level. Quota levels for each immigrant group are taken directly from Ager et al. (2023).

Figure 2.6: Catholic Exposure by 1960 Tract

2.4.3 Predicted Change in Catholic Exposure

The proposed measure of quota exposure captures plausibly exogenous negative shocks to Catholic immigration at the tract level. A concern with this approach may be that the measure is correlated with total immigration exposure, or total immigrant share. Therefore the measure is partially capturing a negative shock to immigrant population, rather than a feature of Catholic migration in particular. An alternative construction would be to capture the relative change in Catholic migration, compared to non-Catholic migration.

I construct a measure of relative Catholic exposure as a difference in quota exposure to Catholic vs. non-Catholic migration. Quota exposure for non-Catholic migration is defined equivalently to the method in 2.1, substituting non-Catholic origin countries for the predicted missing flows. The relative change in exposure to Catholic exposure is then:

$$\Delta QE_d = QE_{Non-Catholic} - QE_{Catholic} \quad (2.4)$$

Intuitively, the quota exposure measure captures a local area's exposure to missing migrants. If a location is losing more Non-Catholic migrants relative to Catholic migrants, then they are differentially receiving a more Catholic composition of migrants compared to the pre-quota period.

2.4.4 Cohort-Level Exposure to Quota

In order to understand the impact of Catholic migration on native child human capital, I use linked records by the Census Tree Project in a cohort design to compare educational attainment across children more or less exposed to changes in Catholic migration during the quota policy of 1920. I observe total educational attainment defined as total schooling years in 1940. I then link these records to identify individuals who resided in Northeastern cities in 1920 as children.

The cohort comparison defines treated individuals as those aged 0-10 in the year 1920, as they are most likely to be in elementary school during significant changes to local Catholic composition. The control group is defined as individuals aged 10-20 in 1920, as these individuals

were plausibly out of elementary school by the time that changes to local Catholic immigration took place.

Utilizing 1940 educational attainment for individual i in cohort c as my outcome of interest y_{idc} , I estimate the cohort level DiD as:

$$y_{idc} = \alpha_d + \gamma_c + \mu_{j \times c} + \beta (QE_d \times TreatedCohort_i) + X'_i \delta + \epsilon_{idc} \quad (2.5)$$

Where treatment is defined by cohort c , and QE_d represents the relative district-level exposure to immigration quota. I include district and cohort fixed effects, α_d and γ_c . I also include city level effects interacted with cohorts $\mu_{j \times c}$. Controls $X'_i \delta$ include individual characteristics such as a linear control for age and a binary gender dummy, as well as baseline foreign share, and interactions of baseline foreign share with cohorts.

This cohort design compares children differentially exposed to post-quota neighborhood composition changes during schooling ages while holding fixed persistent district-level differences. The interpretation is that β_1 captures exposure-intensity effects for younger cohorts relative to older cohorts that were less likely to be in elementary school when quota-driven composition changes occurred.

To understand the dynamic effects of quota exposure by cohort, I next test an event-study version of the cohort level analysis. In particular, I estimate:

$$y_{idc} = \alpha_d + \gamma_c + \mu_{j \times c} + \sum_{a \neq a_0} \beta_a (QE_d \times \mathbf{1}\{Cohort_i = a\}) + X'_i \delta + \epsilon_{idc} \quad (2.6)$$

Where each 5-year cohort between age 0 in 1920 and age 35 in 1920 is interacted separately with the treatment exposure. I estimate this model for various versions of quota exposure, including the main Catholic quota exposure measure, a non-Catholic immigrant quota, and the predicted change in Catholic composition. The age bin 20-25 is used as the reference category, as individuals aged above 20 were plausibly unaffected by Catholic-related changes to the elementary schooling market. The control variables and fixed effects are defined the same as above.

The parallel trends assumption for this strategy requires no differential pre-1920 trend in outcomes by quota exposure. The strategy also requires exogeneity of national quota shocks with respect to tract-level schooling shocks, conditional on fixed effects and baseline controls. To assess these assumptions, I report event-study style pre-trend estimates for tract outcomes (Table 2.13), placebo cohort tests for child attainment (Table 2.12), and robustness to spatially correlated errors using Conley standard errors (Tables 2.14 and 2.15).

2.5 Results

In this section I report effects of exposure to shocks in Catholic migration on teacher outcomes, school outcomes, and child attainment. I first use the baseline Catholic quota exposure measure, then the relative Catholic change measure, and finally cohort-level specifications. Table 2.3 reports the long-run relationship between historical Catholic concentration and the 1960 public-school enrollment share. Catholic share in both 1900 and 1930 predicts lower public enrollment in 1960, and quota exposure remains negative even after controlling for baseline Catholic share. In magnitude terms, a 10 percentage-point increase in Catholic share in 1900 corresponds to about a 3.2 to 4.8 percentage-point lower public enrollment share in 1960, depending on specification.

Across columns, the coefficients are stable in sign and statistically precise, which supports the view that Catholic settlement generated persistent local demand for private schooling. In magnitude terms, these estimates imply that historically Catholic neighborhoods remained substantially less public-school intensive even by 1960, long after the peak migration period. This persistence is consistent with durable institutional effects, such as parish school networks and intergenerational schooling preferences.

2.5.1 Impacts of Cohort Level Exposure to Catholic Migration

The OLS correlation of quota exposure and local educational attainment of native children is shown in Table 2.11. This table shows the relationship between children born in a given enumeration

Table 2.3: Historical Catholic Shares and Private School Share

	Public Share	Public Share	Public Share
Catholic Share 1900	-0.480 [0.056]***		-0.323 [0.068]***
Catholic Share 1930		-0.319 [0.045]***	
1920 Catholic Quota Exposure			-0.002 [0.001]***
Mean Dep.	0.697	0.695	0.697
Observations	2,537	3,062	2,533

Notes: Unit of observation is the 1960 census tract. The dependent variable is the public-school enrollment share among elementary-age children in 1960. “Catholic Share 1900” and “Catholic Share 1930” are tract-level population shares from those census years. “1920 Catholic Quota Exposure” is the Bartik-style predicted missing Catholic inflow due to quota policy. All columns include city fixed effects. Standard errors are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

district with a given level of quota exposure, between 1900 and 1920. Catholic quota exposure is negatively associated with later attainment, while non-Catholic exposure and relative Catholic change are positively associated with attainment. This pattern suggests that native children in tracts exposed to more missing Catholic migrants had on average lower educational attainment.

Table 2.4 reports cohort-based DiD estimates for child educational attainment, where total attainment is measured in schooling years in census year 1940. Younger exposed cohorts have lower attainment where Catholic quota exposure is higher, while relative Catholic change is associated with higher attainment.

Lower Catholic exposure during school ages predicts lower completed schooling in adulthood for native children. Magnitudes are economically meaningful: the estimated one-standard-deviation effect is -0.092 years for Catholic quota exposure (about -1.11 months), compared with $+0.087$ years ($+1.05$ months) for non-Catholic exposure and $+0.098$ years ($+1.17$ months) for relative Catholic change. Relative to the sample mean of 13.681 years, these correspond to roughly -0.68% , $+0.64\%$, and $+0.71\%$ changes, respectively. The opposite sign for non-Catholic exposure and the positive coefficient on relative Catholic change are consistent with a composition channel, where local migration mix altered schooling environments faced by children.

Table 2.4: 1920 Quota Impacts on Child Educational Attainment

	Ed Attainment	Ed Attainment	Ed Attainment
Treat * Non-Catholic QE	0.003 [0.000]***		
Treat * Catholic QE		-0.011 [0.001]***	
Treat * Relative Catholic Change			0.004 [0.000]***
Mean Dep.	13.681	13.681	13.681
SD Regressor	26.946	8.054	27.026
Effect of 1 SD (Years)	0.087	-0.092	0.098
Effect of 1 SD (Months)	1.047	-1.109	1.172
Effect of 1 SD (% of Mean)	0.64	-0.68	0.71
Observations	1,319,932	1,319,932	1,319,932


Notes: Unit of observation is a linked individual observed in 1920 and 1940. Outcome is years of schooling in 1940. “Treat” is an indicator for being ages 0–10 in 1920; controls are ages 10–20 in 1920. Reported coefficients are interactions between Treat and tract-level exposure measures. Regressions include tract, city, and sex fixed effects with baseline controls as in the text. Standard errors are in brackets. * p<0.10, ** p<0.05, *** p<0.01.

2.5.1.1 Event-Study Specification of Cohort Effects

To understand dynamics, I estimate cohort-by-exposure interactions in 5-year age bins, using ages 20–25 in 1920 as the omitted reference group. Coefficients for younger bins can be interpreted as differential attainment gradients with respect to quota exposure, relative to this older cohort that was less likely to be in elementary school during the policy window. Figure 2.7 shows that coefficients become more negative for Catholic exposure in younger cohorts and more positive for non-Catholic exposure, consistent with larger schooling-market effects for children more exposed during school ages.

Figure 2.8 shows the same dynamic pattern using the relative Catholic-change measure. Coefficients are larger for the youngest cohorts, while older cohorts are closer to zero, which is consistent with the intended timing logic of the cohort design. Appendix Figure 2.10 provides a complementary group-time ATT implementation based on (Callaway and Sant’Anna, 2021); the qualitative direction of the cohort pattern is similar, though magnitudes differ because the estimands

and weighting schemes are not identical.




Notes: This figure shows the results of a dynamic DiD specification where age cohort bins are interacted with different quota exposure metrics, including catholic and non-catholic migrant quota exposure. The reference category is Age 20-25 in 1920, who were plausibly unaffected by changes that mostly affected children of elementary school age.

Figure 2.7: Event-Study of Quota Policies by Cohort

2.5.2 Impacts of Catholic Quota Exposure

How did tracts respond to missing migrants? In work by Abramitzky et al. (2023) the authors find that missing migrants tended to be replaced over time by migrants from other origins less



Notes: This figure shows the results of a dynamic DiD specification where age cohort bins are interacted with a quota-based prediction of increased Catholic composition. The reference category is Age 20-25 in 1920, who were plausibly unaffected by changes that mostly affected children of elementary school age.

Figure 2.8: Event-Study of Predicted Catholic Change by Cohort

affected by quota restrictions. It's possible that missing Catholics drove changes in the local migrant share that affected schooling outcomes through compositional channels. Table 2.5 reports the first-stage relationship between quota exposure and tract composition outcomes using equation 2.3. As expected, higher exposure predicts lower Catholic share after 1920. Exposure to missing Catholics also predicts higher shares of both native-born Americans and non-Catholic immigrants. This pattern shows that quota exposure captures future declines in Catholic migrant inflows. It also suggests that declining Catholic numbers were not 1 for 1 replaced by migrants from other non-Catholic groups. Instead, the gap was filled by both natives and non-Catholic migrants.

The estimates also show that the decrease in Catholic share is mirrored by increases in non-Catholic immigrant and native shares, rather than by a collapse in total neighborhood population. Quantitatively, a one-unit increase in quota exposure is associated with a 0.193-point lower Catholic share, a 0.117-point higher non-Catholic immigrant share, and a 0.076-point higher native share. Relative to pre-period means (9.572, 17.376, and 73.051), these are about -2.0% , $+0.7\%$, and $+0.1\%$ shifts. This composition pattern is important for interpretation: the policy shock appears to have

Table 2.5: 1920 Quota Impacts on Tract Characteristics

	Catholic Share	Non-Catholic Imm Share	US Native Share
Quota Exposure * Post	-0.193 [0.013]***	0.117 [0.017]***	0.076 [0.016]***
Mean Dep.	9.572	17.376	73.051
Observations	11,990	11,990	11,990

Notes: Unit of observation is tract-year, with all variables aggregated to 1960 tract boundaries. The reported coefficient is the interaction “Quota Exposure $\times$ Post” from equation 2.3. “Post” equals 1 in 1930 and 0 in 1900, 1910, and 1920. Outcomes are tract shares of Catholic immigrants, non-Catholic immigrants, and U.S.-born natives. Regressions include the fixed effects and controls in the baseline DiD specification. Standard errors are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

shifted who lived in exposed tracts, which is the channel needed for subsequent schooling-market effects.

Table 2.6 reports the same design for tract-level teacher outcomes. More exposed tracts have lower teacher Catholic share, while most other teacher background measures are not statistically distinguishable from zero. This indicates a compositional shift in teacher origin without strong evidence of broad changes in observable teacher-quality proxies.

Table 2.6: 1920 Quota Impacts on Teacher Characteristics

	Teacher Age	Teacher Spouse Lit	Teacher Father Occ	Teacher Catholic Share
QE * Post	-0.027 [0.022]	-0.000 [0.001]	-0.034 [0.025]	-0.098 [0.023]***
Mean Dep.	27.685	0.511	16.497	4.018
Observations	12,001	12,001	12,001	11,676

Notes: Unit of observation is tract-year, mapped to 1960 tract boundaries. Teacher outcomes are tract-level averages or shares for teachers residing in the tract. “Teacher Catholic Share” is the share of teachers from Catholic-origin groups. “Post” equals 1 in 1930 and 0 in 1900, 1910, and 1920. Regressions include the fixed effects and controls in the baseline DiD specification. Standard errors are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Substantively, this table indicates that the teacher workforce adjusted primarily along identity/composition margins. The Catholic-share coefficient is -0.098 (s.e. 0.023), compared with a mean of 4.018, implying about a 2.4% decline relative to baseline levels. By contrast, point estimates for teacher age (-0.027 on a mean of 27.685), spouse literacy (approximately zero on a mean of 0.511), and father occupation score (-0.034 on a mean of 16.497) are small and statistically

weak. This suggests that reduced Catholic inflows did not produce broad observable declines in teacher socioeconomic background, but did change the Catholic representation of teachers in local areas.

Table 2.7 examines school-level outcomes using the Baltimore public-school panel. The main outcomes are enrollment counts and teacher-per-pupil ratios in Catholic and public schools. The enrollment coefficients are imprecise, but the public teacher-per-pupil coefficient is negative and precisely estimated, implying lower staffing intensity in more exposed locations after the quota shock.

Table 2.7: 1920 Quota Impacts on Schools

	Catholic Pupils	Cath Teach/Pupil	Public Pupils	Public Teach/Pupil
QE * Post	-10.058 [6.467]	0.002 [0.008]	-5.146 [10.644]	-0.042 [0.012]***
Mean Dep.	801.780	-3.788	795.268	-3.698
Observations	132	104	209	167

Notes: Unit of observation is school-year (Baltimore sample). Exposure is assigned from the school's tract-level quota exposure. "Post" equals 1 in 1930 and 0 in 1900, 1910, and 1920. "Cath/Public Teach-Pupil" are the staffing-intensity measures used in the estimation sample. Regressions include the fixed effects and controls used in the baseline DiD setup. Standard errors are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The key effect in this table is on the public staffing margin. Interpreted directly, exposed locations saw lower teacher-per-pupil intensity in public schools after the quota shock, with an estimated coefficient of -0.042 (s.e. 0.012). In contrast, enrollment effects are imprecise: -10.1 for Catholic pupils (s.e. 6.5) and -5.1 for public pupils (s.e. 10.6). This pattern indicates that the clearest adjustment appears on staffing intensity rather than enrollment counts in this Baltimore panel.

Taken together, Tables 2.5–2.7 support a mechanism centered on local composition and supply adjustments rather than broad teacher-quality changes. First, quota exposure shifts tract composition away from Catholic immigrants. Second, Catholic representation among teachers declines, but age and family-background proxies for teacher quality move little. Third, the clearest school-level response appears on public staffing intensity. This combination is consistent with

a channel in which migration composition changed the demand mix across schooling sectors and altered how local systems staffed schools, rather than generating large immediate changes in observable teacher quality.

2.5.3 Impacts of Relative Change in Catholic Migration

I next use a relative quota-exposure measure to compare regions that experienced a relative increase in Catholic composition versus non-Catholic composition. Table 2.8 confirms the intended first stage under this alternative treatment: higher relative Catholic exposure increases Catholic share and decreases non-Catholic immigrant share in the post period.

Table 2.8: 1920 Quota Impacts on Tract Characteristics

	Catholic Share	Non-Catholic Imm Share	US Native Share
Catholic Change * Post	0.104 [0.005]***	-0.185 [0.006]***	0.081 [0.006]***
Mean Dep.	9.572	17.376	73.051
Observations	11,990	11,990	11,990

Notes: Unit of observation is tract-year. “Catholic Change” is the relative exposure measure $QE_{Non-Catholic} - QE_{Catholic}$, so larger values imply a more Catholic composition after quota restrictions. “Post” equals 1 in 1930 and 0 in 1900, 1910, and 1920. Outcomes are tract shares. Regressions include the fixed effects and controls in the baseline DiD specification. Standard errors are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

This table validates the alternative treatment definition by showing the expected directional shifts in composition. A one-unit increase in relative Catholic change is associated with a 0.104-point rise in Catholic share, a 0.185-point decline in non-Catholic immigrant share, and a 0.081-point rise in native share. Relative to means, these correspond to roughly +1.1%, −1.1%, and +0.1% changes. That pattern indicates that the “relative change” measure is capturing composition reweighting rather than random demographic noise.

Table 2.9 shows a similar pattern to the baseline design: relative Catholic exposure increases teacher Catholic share, while most other teacher background outcomes move little. The consistency between Table 2.6 and Table 2.9 strengthens the interpretation that this is not an artifact of one

Table 2.9: 1920 Quota Impacts on Teacher Characteristics

	Age	Spouse Occ	Father Lit	Catholic Share
Catholic Change * Post	0.006 [0.008]	-0.024 [0.007]***	-0.000 [0.000]	0.065 [0.009]***
Mean Dep.	27.685	5.690	0.656	4.018
Observations	12,001	12,001	12,001	11,676

Notes: Unit of observation is tract-year, mapped to 1960 tract boundaries. Reported coefficients are from “Catholic Change $\times$ Post” interactions. Outcomes are teacher averages or shares. Regressions include the fixed effects and controls in the baseline DiD specification. Standard errors are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

exposure construction. Magnitudes are also similar across specifications: Catholic teacher share rises by 0.065 (s.e. 0.009) under the relative-change measure, while age is near zero (0.006 on a mean of 27.685) and father literacy is effectively unchanged. The one additional movement is spouse occupation score (-0.024 , s.e. 0.007), which is statistically precise but modest relative to its mean (5.690). Across both approaches, teacher composition responds in the expected direction, while broader teacher background proxies remain comparatively stable.

2.6 Robustness Checks

The empirical strategy requires a parallel-trends assumption. Educational trends for treated cohorts in more exposed tracts should be parallel to educational attainment trends in less exposed tracts, absent the quota policy. For robustness, I test pre-trend balance between groups that were more or less exposed to Catholic quota implementation. Table 2.12 reports a placebo cohort comparison among individuals aged 20–30 at the time of the policy. The ages 20–25 cohort is assigned as the placebo treatment group and compared with the ages 25–30 cohort. Estimated placebo effects are small and statistically insignificant across all three exposures (0.001, -0.006 , and 0.001, with standard errors 0.003, 0.008, and 0.002), which reduces concern that the main cohort results are driven by generic age-gradient differences unrelated to schooling-age exposure.

The tract-level event-study estimates in Table 2.13 document the pre-period immigrant composition that generates variation in quota exposure; they are not a test of the parallel-trends

assumption used in the cohort attainment design. More-exposed tracts had higher Catholic shares before quota implementation by construction, because the exposure measure assigns greater weight to tracts with larger pre-period populations from Catholic origins that were subsequently restricted by the quotas. The sign on Catholic share flips once quota exposure is interacted with the year 1920, consistent with the quota policy reversing the prior relationship between exposure and Catholic immigrant presence. More-exposed areas also had lower US-native shares before the policy, but do not appear to have had differentially lower native shares relative to non-Catholic immigrant shares.

Another empirical challenge is the spatial correlation of outcomes at the tract level. In the appendix I show robustness of results to adjustments for Conley standard errors. Table 2.14 replicates the effects of quota exposure on local teacher characteristics using standard errors adjusted for spatial correlation. Table 2.15 shows the results for the first stage relationship between Catholic quota exposure and tract level immigration characteristics with adjusted standard errors.

The Conley results preserve the main directional findings. For teacher outcomes, the Catholic-share effect remains negative and statistically precise (-0.121 , s.e. 0.029), while other teacher outcomes remain small and imprecise. For tract composition, the Catholic-share first-stage remains negative and significant (-0.241 , s.e. 0.054), while non-Catholic and native-share responses are less precise under spatial correction. Overall, allowing for spatial dependence weakens precision for some margins but does not overturn the central composition result.

The cohort-level results rely on a linked subsample of individuals who are observed in both the 1920 census and the 1940 census. This creates selection concerns, as it may be that individuals that stayed locally are more likely to be observed, relative to individuals who left. The results on educational attainment may therefore partially reflect selection differences in the composition of linked individuals across treatment. Table 2.16 shows the results of regressions relating linking probability to quota exposure and cohort exposure. Individuals residing in tracts with higher Catholic quota exposure in 1920 are slightly less likely to be linked in the 1940 sample. Individuals in younger cohorts in more exposed tracts are more likely to be linked. This suggests that there may have been sorting out of tracts exposed to missing Catholic migrants.

2.7 Conclusion

This paper studies how migration shocks can reshape the market for a central local public good: schooling. Using newly assembled neighborhood-level data on immigration composition, school locations, school characteristics, and linked child outcomes, I show that Catholic immigration was not only a demographic shock but also a demand shock that altered the long-run public-private education equilibrium in Northeastern US cities.

Historical Catholic concentration predicts lower public-school enrollment shares decades later, consistent with durable private-school demand. Quota-driven variation indicates that the 1920s border closure changed local Catholic composition in the expected direction. Tracts exposed to missing Catholic migrants became relatively more non-Catholic, with a mix of native and non-Catholic migrants changing the local composition. Native children who grew up in tracts with missing Catholic migrants have worse educational performance. The effects of quota-driven changes in Catholic inflows are the opposite of those for non-Catholic inflows, suggesting that features of the two migrant groups had an important and differing effect on local outcomes. I find evidence of downstream adjustments in school staffing intensity and teacher composition. Teachers were less likely to be Catholic in tracts that lost Catholic migrants. Public schools in Baltimore exposed to higher levels of missing Catholic inflows saw increased class sizes. The effect on class sizes is consistent with a story of crowding. Private schools eased the population attending public schools, possibly by removing Catholic children from the public school classrooms. In areas with less private schools, the public schools had to accommodate larger class sizes, incorporating migrants from non-Catholic backgrounds into the classroom.

Taken together, these findings are consistent with a broader view in which migration affects destinations through preferences and institutions in addition to labor supply. When incoming populations demand specific goods—in this case religious schooling—local providers and public systems may adjust, and those adjustments can persist long after the initial migration wave.

2.8 Acknowledgements

Chapter 2, “Historical Immigration and the Market for Schooling in American Cities,” is currently being prepared for submission for publication of the material therein. The dissertation author, Andre Daniel Gray, is the sole author of this chapter.

2.9 Appendix: Additional Figures and Tables

Table 2.10: City Coverage Across Data Sources


City	ED maps	Public school data	Catholic school data
Baltimore	Yes	Yes	Yes
Boston	Yes	No	Yes
Chicago	Yes	No	Yes
Cincinnati	Yes	No	No
Cleveland	Yes	No	No
Detroit	Yes	No	No
New York	Yes	Yes	Yes
Philadelphia	Yes	No	Yes
Pittsburgh	Yes	No	No
St. Louis	Yes	No	No

Notes: ED maps are available in `data/ED_Wards` for the cities listed in Table 2.10. Public school data are currently available for Baltimore and New York. Catholic school data are currently available for Boston, Baltimore, New York, Philadelphia, and Chicago.


Table 2.11: 1920 Quota Impacts on Child Educational Attainment

	Ed Attainment	Ed Attainment	Ed Attainment
1920 Non-Catholic QE		0.016 [0.000]***	
1920 Catholic QE	-0.023 [0.000]***		
Relative Catholic Change			0.009 [0.000]***
Mean Dep.	13.026	13.026	13.026
SD Regressor	7.819	21.485	24.543
Effect of 1 SD (Years)	-0.178	0.338	0.228
Effect of 1 SD (Months)	-2.140	4.056	2.742
Effect of 1 SD (% of Mean)	-1.37	2.60	1.75
Observations	2,257,231	2,257,231	2,257,231

Notes: Unit of observation is a linked individual. Outcome is years of schooling in 1940. Coefficients are cross-sectional associations between tract exposure in 1920 and later attainment. Regressions include city and sex fixed effects with baseline controls as in the main text. Standard errors are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.




(a) Catholic Schools Before 1900



(b) Catholic Schools in 1920

Notes: This figure shows the expansion of Catholic schools in the early 20th century. Data is taken from the Official Catholic Directory, and geocoded by street intersections. The heat map captures census estimates of local Catholic shares in 1900. Schools include elementary and secondary institutions.

Figure 2.9: Catholic School Expansion Example



Notes: This figure shows the results of a dynamic DiD specification where age cohort bins are interacted with a quota-based prediction of Catholic quota exposure. Periods to treatment are annual, and given in terms of student age in 1920. Students are born later as the graph moves to the right. Period zero in this graph captures students aged 15 in 1920. The ATT estimates use the specification proposed by (Callaway and Sant'Anna, 2021).

Figure 2.10: Event-Study of Catholic Quota Exposure by Cohort with csdid

Table 2.12: Educational Attainment in Placebo Cohorts

	Ed Attainment	Ed Attainment	Ed Attainment
Placebo * Non-Catholic QE	0.001 [0.003]		
Placebo * Catholic QE		-0.006 [0.008]	
Placebo * Relative Catholic Change			0.001 [0.002]
Mean Dep.	12.167	12.167	12.167
Observations	32,076	32,076	32,076

Notes: Unit of observation is a linked individual. Outcome is years of schooling in 1940. This is a placebo cohort DiD among older cohorts (ages 20–30 in 1920), where ages 20–25 are placebo treated and ages 25–30 are placebo controls. Reported coefficients are interactions between placebo treatment and tract-level exposure measures. Regressions include the same fixed effects and controls used in the main cohort specification. Standard errors are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.13: Pre Period Tract Level Characteristics by Quota Exposure

	Catholic Share	Non-Catholic Imm Share	US Native Share	Teacher Catholic Share
QE x 1880	-0.021 [0.027]	-0.136 [0.036]***	0.157 [0.028]***	-0.007 [0.054]
QE x 1900	0.129 [0.024]***	-0.065 [0.023]***	-0.064 [0.025]***	0.128 [0.056]**
QE x 1910	0.429 [0.016]***	-0.217 [0.025]***	-0.211 [0.021]***	0.205 [0.054]***
Mean Dep.	9.572	17.376	73.051	4.018
Joint p-value	0.000	0.000	0.000	0.001
Observations	11,990	11,990	11,990	11,676

Notes: Unit of observation is tract-year, mapped to 1960 tract boundaries. Columns report event-study interactions between quota exposure and year indicators, with 1920 omitted as the reference year. Outcomes are tract shares (Catholic immigrant share, non-Catholic immigrant share, U.S.-born share, and teacher Catholic share). The reported joint pre-trend statistic tests whether 1900 and 1910 interactions are jointly zero. Regressions include tract and year fixed effects. Standard errors are in brackets. * p<0.10, ** p<0.05, *** p<0.01.

Table 2.14: Teacher Characteristics with Conley Standard Errors

	Age	Spouse Lit	Father Occ	Catholic Share
QE * Post	-0.010 [0.040]	-0.000 [0.001]	-0.024 [0.050]	-0.121 [0.029]***
Mean Dep.	27.608	0.516	15.972	3.755
Observations	10,648	10,648	10,648	10,368

Notes: Unit of observation is tract-year, mapped to 1960 tract boundaries. Outcomes are teacher characteristics (age, spouse literacy, father's occupation score, and teacher Catholic share). Coefficients report the interaction "Quota Exposure $\times$ Post" from the baseline DiD specification. "Post" equals 1 in 1930 and 0 in 1900, 1910, and 1920. Regressions include tract and year fixed effects, and standard errors are Conley spatial HAC standard errors. * p<0.10, ** p<0.05, *** p<0.01.

Table 2.15: Tract Level Characteristics with Conley Standard Errors


	Catholic Share	Non-Catholic Imm Share	US Native Share
Quota Exposure * Post	-0.241 [0.054]***	0.113 [0.073]	0.128 [0.079]
Mean Dep.	8.920	17.090	73.989
Observations	10,638	10,638	10,638

Notes: Unit of observation is tract-year, mapped to 1960 tract boundaries. Outcomes are tract population shares (Catholic immigrant share, non-Catholic immigrant share, and U.S.-born share). Coefficients report “Quota Exposure $\times$ Post” from the baseline DiD specification. “Post” equals 1 in 1930 and 0 in 1900, 1910, and 1920. Regressions include tract and year fixed effects, and standard errors are Conley spatial HAC standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.16: Individual Linking Probability by Quota Exposure


	Linked	Linked	Linked	Linked
Catholic Quota Exposure	-0.0031 [0.0003]***		-0.0029 [0.0004]***	
Relative Catholic Change		-0.0003 [0.0001]***		0.0016 [0.0001]***
Treat * Catholic Quota Exposure			0.0009 [0.0003]***	
Treat * Relative Catholic Change				-0.0011 [0.0001]***
Baseline Foreign-Born Share			-0.2114 [0.0164]***	-0.3787 [0.0224]***
Mean Dep.	0.517	0.517	0.517	0.517
Observations	4,122,975	4,122,975	4,122,975	4,122,975

Notes: Unit of observation is an individual observed in 1920 aged 0-35. Linking probability captures a binary indicator for whether the individual appears in the linked sample with an observation for educational attainment. Linking probability is regressed on quota exposure and interactions of quota exposure with a treatment indicator for whether the individual was 0-15 at the time of the quota policy. All models include city and age level fixed effects * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.




Notes: This figure shows results for coefficients on tract level and teacher characteristics from running equation 2.3, leaving a particular city out of the data. The first three graphs show tract level characteristics by left out city, while the last three graphs show teacher spouse literacy, parent occupational score, and teacher age. The reported coefficients are on the *QuotaExposure* * *Post* term in the equation.

Figure 2.11: Leave One City Out Composition and Teacher Trends




Notes: This figure shows results for coefficients on tract level and teacher characteristics from running equation 2.3, leaving a particular immigrant origin out of the data. The first three graphs show tract level characteristics by left out birthplace, while the last three graphs show teacher spouse literacy, parent occupational score, and teacher age. The reported coefficients are on the $QuotaExposure \times Post$ term in the equation.

Figure 2.12: Leave One Origin Out Composition and Teacher Trends




(a) Example of Public School Directory, NYC




(b) Example of Teacher Directory, NYC

Notes: This figure shows examples of the evidence of public schooling characteristics in the New York Board of Education archives. Directories were published annually listing all schools in the New York Area, as well as a list of the teacher names for each school. Each school also includes a street address, which can be geocoded. Some directories also include information on number of pupils, and a gender breakdown.

Figure 2.13: Board of Education Archive



Notes: This figure shows a heatmap of teacher father's occupational score in 1910 New York. Father's are linked to teacher's that live in the household, through the census's family interrelationship data. Occupational score is a proxy for income, mapping occupation categories to average salaries and educational requirements. The sample only includes teachers who live with their parents in the same household.



Notes: This figure shows a heatmap of teacher spouse literacy in 1910 New York. While teacher's have 100% literacy, the literacy of their spouses can be taken as a proxy for socioeconomic status.

Chapter 3

The Effect of Migrant Composition on Housing Prices

3.1 Introduction

As new migrants enter a housing market, they shift out demand for housing units. In empirical work shocks to the number of new individuals arriving is used to predict local changes in house prices. But different types of migrants interact differently with the housing market. A group of single, high-income engineers are looking for different kinds of housing than a middle-income family of four. The number of units or rooms demanded per unit of population varies widely based on a particular migrant group's preferences and socioeconomic composition. These effects are absorbed or lumped in standard shift-share methods that relate linear changes in migrant inflows to changes in house price growth. In this paper, I study alternative shift-share constructions to decompose the effect of migrants in terms of key characteristics such as the foreign born population share, and the distribution of ages, skills and cultural backgrounds. I show that different functions of shift and share variables can predict changes to compositional measures of migrants, and separate the effects of migrant population size from migrant composition in studies of the housing market.

Migrant characteristics matter for housing markets first because they affect the total number

of units demanded. The headship rate, defined as the number of household heads per unit of adult population, is a key metric of housing demand (Paciorek*, 2016). Headship is a function of household formation decisions made by individuals who choose to live with parents, partners, roommates or alone. Headship is known to swing widely during recessions or periods of high house price growth (Farnham et al., 2011; Kaplan, 2012). Household formation is also subject to cultural norms that vary by birthplace (Huber and Schmidt, 2019). Conditional on income, occupation and age, different groups have preferences over the kind of housing and living arrangements they demand.

Beyond headship, migrant characteristics affect the type of housing demanded. Certain migrants increase demand for multi-family housing, while others may strictly prefer single-family homes. Migrant characteristics can therefore affect the composition of housing demand at different densities. Differences in preferences and headship are what make a supply shock of increasing migration different from a supply shock of increasing native workers on local housing.

The workhorse empirical framework for estimating the effects of migration is the shift-share instrument. Locations are exposed to new migrants based on origin-level shifts in the number of migrants leaving particular countries, and destination-specific settlement shares. Settlement shares, the share of migrants from a particular origin settling in a particular destination, are taken from a pre-period and treated as unrelated to the contemporary shocks driving origin-level out-migration. In this paper, I decompose the effects of migration in terms of population inflows and compositional characteristics by leveraging shift-share logic to instrument for changes to migrant characteristics. I identify shifts to migrant age, skill and headship at origin, and combine these shifts with destination exposure shares to simultaneously instrument for migrant inflow size and composition.

I show that functions of different origin-level shifts in composition can predict linear and non-linear moments that capture compositional changes to migrant inflows in destinations. Instruments for composition provide independent variation when controlling for shifts in immigrant inflows, and predict changes to house price growth. I find that increasing migrants as a share of the population consistently reduces house price growth at annual and decade level frequencies. This is

consistent with work by Saiz and Wachter (2011a) that finds a relationship between foreign born share and lower house prices at the neighborhood level. Increasing migrant shares also reduce house price volatility in the aggregate. Increasing youth share of migrants predicts lower house price growth but increasing higher price volatility. Increasing high skilled share of recent migrants increases the average number of rooms demanded, and the average size of buildings. My results suggest that features of migrant composition are an important dimension for understanding how population changes map to housing market effects.

This is not the first paper to study migrant compositional changes in shift-share designs. Prior work has largely focused on differences in migrant impact by skill level. Many papers have sought to decompose the effect of migrant inflows in terms of low-skilled and high-skilled labor. Caiumi and Peri (2024) update and extend the foundational work by Borjas (2003) in considering the impact of immigrants on wages decomposed by skill-experience cells. Diamond (2016b) create a model in which incoming population flows affect destinations in terms of amenities, which are an endogenous function of the ratio of high to low skilled labor.

This paper adds to a literature on the impact of immigration on housing prices. The seminal paper is Saiz (2007), which uses a shift-share instrument to show that a 1% increase in immigrant population increases MSA house prices by approximately 1%. Subsequent papers using similar methods have found heterogeneity in the effects on rents versus prices (Unal et al., 2026; Helfer et al., 2023), by initial destination city characteristics (Sharpe, 2019), and by differences between raw inflows and migrant shares (Saiz and Wachter, 2011a). An important strand of extended work has noted that shift-share designs don't adequately capture the general equilibrium effects of natives responding to migrant inflows, especially as far as they relocate or resort within the city (Sanchis-Guarner, 2023). Another entry from Ang et al. (2026) looks specifically at the impact of Asian migration on housing prices, decomposing the effects of immigration in terms of schooling quality channels and demand pressures. This paper extends this work by offering a general framework for decomposing the effects of immigration composition on house prices. Rather than focus on a particular case, I show that headship responds to various features of migrant flows, including age and

skill shares. I then show how shift-share instruments can be adapted to accommodate these features.

This paper also contributes to the literature on headship and housing prices. The organization of workers into households naturally drives the relationship between total population and housing demand. Sudden shocks to housing preferences, especially in response to crises, have been shown to explain periods of high house price volatility (García and Paciorek, 2022). Whether workers choose to live with parents, spouses, grandparents or roommates is endogenous to house price levels (Kaplan, 2012; Farnham et al., 2011). It is also a function of cultural and age-skill specific preferences (Huber and Schmidt, 2019; Paciorek*, 2016). In this paper I show how shocks to the supply of migrants with particular age and skill composition can have differential impacts on housing prices, due to origin-level differences in headship.

Last, this paper engages with prior work on shift-share designs and formula-based instruments. A number of papers have considered the potential endogeneity of both shifts and shares in Bartik style designs (Jaeger et al., 2018; Goldsmith-Pinkham et al., 2020). Two major concerns include the potential serial correlation of pre-period settlement shares, and the potential correlation between origin-level shocks and local destination economic trends (Borusyak et al., 2022). A related literature has begun to study formula-based instruments, which increase power by leveraging heterogeneous shock exposure in different ways. This paper is an applied example of a similar idea. I show that different functions of shifts and shares can generate variation in immigrant composition that is only weakly related to total immigrant inflows.

The paper proceeds first by discussing the datasets produced from ACS and census records to capture changes in migrant composition in cities. I then discuss predictors of headship by age, skill and birthplace. The following section discusses the methodology, beginning with a theoretical framework and then detailing different functions of shift-share instruments. I then discuss the estimation strategy of interest, and the results of regressions of housing market characteristics on migrant features.

3.2 Data

The paper traces demographic changes at the Metropolitan Statistical Area (MSA) level¹ using annual data from the American Community Survey (ACS) 2005-2024², and 1% samples from the decennial census for 1970 - 2020. The sample includes working age native and foreign born individuals, aged 18-65. I measure detailed changes in the demographic composition of MSAs, including foreign born shares, birthplace composition and age-skill cells. Age is defined in terms of quantiles, while skill is identified using worker reported educational attainment, where high-skilled is denoted for workers with a bachelor's degree or higher. These fine-grained demographic measures can be problematic in the annual ACS samples, which capture about 1% of the population and cover areas with more than 65,000 people. The more fine grained the demographic cells I study, the more likely that the low sample sizes produce noise when comparing across years. I therefore test robustness of all the results leveraging non overlapping 5-year aggregated ACS estimates, and decade level changes using the census files.

The ACS files also include several variables on features of an individual's dwelling, which I use to capture the housing decisions of recent migrants. In particular I use categorical variables capturing the number of rooms in a household's unit, the building size in terms of number of units and the building age. The number of rooms is numerical, and aggregated as an average at the MSA level. Building age is an index, with higher values representing buildings built in more recent years. Building size is a 1-10 number, with 1 representing buildings of only one unit (single family homes), and 10 representing buildings of 50 or more units.

House prices are taken from several sources. The main panel dataset for annual house price changes at the county and MSA level is the Federal Housing Finance Agency (FHFA) House Price Index. The FHFA index is a weighted repeat-sales index that measures average price changes for the same single-family properties over time, using sales and refinancing transactions on properties

¹Due to inconsistencies in metropolitan codes over longer time-horizons, the ACS and Census records are mapped to contemporary CBSA IDs from the county FIPS codes. This reduces the sample size as county fips are only available for a subset of the public records.

²The ACS sample begins in 2005 because that is the first ACS year in which county FIPS codes are available.

with mortgages purchased or securitized by Fannie Mae or Freddie Mac. Because it follows repeat transactions for the same properties, the FHFA index is designed to abstract from changes in the composition of homes sold across periods. I use the annual county and MSA FHFA series as the primary measure of single-family house price growth. This is supplemented with data from Redfin, which provides MSA-level measures separately for all homes, single-family homes, and condo/co-op or multifamily-type property categories. Redfin's standard market tracker measures are based on MLS and public-records transaction data and report median sale prices for homes sold in a given geography and period; these measures therefore capture actual transaction prices but may reflect changes in the composition of homes sold. My outcomes of interest are annual house price growth, defined as the percent change in the relevant house price index across years, and changes in house price volatility. House price volatility is calculated as the three-year rolling standard deviation of annual house price growth. Due to the differences in the construction of the redfin and FHFA indices, I examine the correlation between these time series, including across the sub-components for single-family and multi-family homes. Figure 3.5 shows correlations between the FHFA measure, Redfin measures and other common metrics for local HPI changes at the metro level from Zillow and Bogin et al. (2019). While the FHFA index closely tracks measures from Zillow and Bogin et al. (2019), the Redfin metrics are less strongly correlated. In particular, the multi-family index appears to have dynamics independent of the other indices reported.

The paper constructs shift-share measures compiled from changes in total immigration by birthplace and settlement share exposure to inflows at the MSA level. Immigrants are identified by a foreign birthplace, and by their reported year of immigration. A recent immigrant is defined as a foreign born worker who reports arriving within 5 years of the survey year. The endogenous objects of interest in this paper are the migrant share, defined in terms of changes to the foreign born share over time, and the composition of recent migrants, in terms of age, skill and headship. Changes in the migrant share do not purely reflect changes in migrant inflows. Instead, they capture several dynamics including growth in the migrant population, and decreases in the native population. One benefit to using changes in the migrant share as opposed to raw inflows is that it reduces

measurement error and improves consistency across census samples. However, this study will not distinguish migrant inflows from potential native flight at destination.

Initial settlement shares are taken from the 1980 1% sample of the Census and are defined for a given immigrant origin g and an MSA destination d as $s_{gd} = \frac{N_{gd}}{\sum_{g=1}^N N_{gd}}$. Where N_{gd} is the number of migrant workers from g residing in d .

The headship rate by immigrant group is a key variable of interest. Headship is defined as the ratio of households to total living adults. This is calculated for a given immigrant group g and destination d as the number of reported household heads over the total number of immigrants from g residing in d . All numbers are weighted by the ACS and Census provided person weights.

A number of MSA-level local controls are used in the estimation, including total population. I also include measures of housing supply, taken from Gyourko et al. (2008). In particular, I use the Wharton Residential Land Use Regulation Index (WRLURI), which captures how difficult it is to construct new housing in the MSA. The index takes into account regulatory stringency, as well as local density.

3.2.1 Changes in Composition by Origin


A large literature attributes cross-city differences in house price growth to differences in housing supply constraints, including land-use regulation, zoning, and geographic barriers to construction (Glaeser and Gyourko, 2025). This supply-side view predicts that a given increase in local housing demand should generate larger price increases and smaller quantity responses in more supply-constrained metropolitan areas. Recent evidence from Louie et al. (2026), however, challenges this interpretation. Using U.S. city data from 2000 to 2020, they show that higher income growth predicts similar growth in house prices, housing quantities, and population across cities with different measured housing supply elasticities. They interpret this as evidence that measured supply constraints are not the primary explanation for differences in house price and quantity growth across U.S. cities. Instead, their findings point toward differences in the nature of local housing demand. In particular, population growth and average income growth need not capture the same demand shock.

Some metropolitan areas experience income growth because they attract high-income workers with strong demand for high-quality housing, while others grow through broader employment expansion that increases population without equivalently increasing willingness to pay for scarce housing attributes. Thus, the composition of local labor demand and the composition of households entering a housing market can generate a disconnect between population growth, income growth, and house price growth.

This paper builds on that demand-side perspective by studying whether immigrant inflows affect house prices not only through the total number of new residents, but also through the composition of immigrant housing demand. Standard immigration-housing specifications typically treat immigration as a scalar population shock, measuring exposure with immigrant inflows or changes in immigrant shares. Yet immigrant groups differ systematically in age structure, household formation, marital status, fertility, income, occupation, and propensity to rent or own. These characteristics may affect the type, timing, and intensity of housing demand in ways that are only weakly related, or nonlinearly related, to total immigrant population growth. I therefore examine whether changes in immigrant worker composition help explain house price growth and volatility across metropolitan areas, beyond what is captured by aggregate immigrant population growth alone.

In a standard estimation strategy for capturing changes to immigrant populations, settlement-based shift-share instruments leverage variation in immigrant flows from specific birthplaces over time. The intuition is that as the total immigrant composition changes (in terms of how many people are coming from different origins), regions are affected differentially according to past settlement shares. Figure 3.1 uses ACS data to document year over year changes in the recent immigrant composition by region of birth. Particular origins go through periods of relative growth and decline. Central America and Mexico for example saw declines in the flow of new immigrants from 2010-2015, while immigrants from Asia increased during the same time period.

Embedded in these flows is additional variation in the composition of new migrants. Figure 3.2 shows proportional changes in mean age for recent immigrants. In some periods changes in mean age are correlated across birthplaces. New immigrants were older on average in 2010 across




Notes: This figure shows the year over year change in total recent immigrants by birthplace. The data is taken from the American Community Survey (ACS) 1-year records. “Recent immigrants” are defined as immigrants who report arriving within 5 years of the survey year. Birthplaces are grouped into categories according to the primary regions that send migrants to the US.

Figure 3.1: Migrant Flows by Birthplace

major origins. In other cases trends differ by origin. In 2011, Asian and European immigrants to the US were relatively younger, while the age of recent Central American and Mexican migrants remained elevated compared to previous years.

These changes in the type of immigrants that are arriving are typically left out of population-based instruments. However, to the extent to which changes in immigrant type are similar to the origin-level shocks that drive outmigration rates, they satisfy the same exclusion restriction required for standard shift-share instruments.

To assess how much headship dispersion can be attributed to observable demographic composition versus persistent birthplace-level differences, Figure 3.4 shows a variance decomposition from nested models at the MSA-by-birthplace-by-year level. Observable age and education measures explain only part of total headship variation, while birthplace-level differences account for an additional share.




Notes: This figure shows the year over year change in the average age of recent immigrants by birthplace. The data is taken from the American Community Survey (ACS) 1-year records. “Recent immigrants” are defined as immigrants who report arriving within 5 years of the survey year. Birthplaces are grouped into categories according to the primary regions that send migrants to the US.

Figure 3.2: Migrant Age by Birthplace

3.2.2 Predictors of Headship

The propensity of migrants to form households varies widely by birthplace. Some migrants may move in with established households, creating zero housing market pressure. Others may be single and high income, and boost demand for new apartments. Figure 3.3 shows the headship rates for recent immigrants across major regional birthplaces. The first thing to notice about this graph is the substantial distance in average headship between different migrant groups. European and Asian immigrants demand many more houses per person than Central American and Mexican immigrants. The second thing to notice about the graph is that the headship rates are relatively stable over time, with a gradual secular increasing trend across all groups. Over time, migrants are demanding more houses per person. While there are some irregularities such as the decline in headship during the financial crisis in 2008, headship rates are not starkly cyclical.

Migrants choose whether they live alone, with family or cohabit with unrelated individuals. These decisions may be a function of a host of socioeconomic variables. To understand which variables predict differences in headship across migrants, I consider the correlation between changes



Notes: This figure shows the headship rate of recent immigrants by birthplace. The data is taken from the American Community Survey (ACS) 1-year records. “Recent immigrants” are defined as immigrants who report arriving within 5 years of the survey year. “Headship rates” are defined as the ratio of self-identified household heads to total recent immigrants for a given birthplace. Birthplaces are grouped into categories according to the primary regions that send migrants to the US.

Figure 3.3: Headship Rate by Immigrant Birthplace

in migrant characteristics and changes in headship. Table 3.1 shows correlations between population characteristics and headship rates using a variety of fixed effects specifications. I use 1-year ACS data, aggregated to the migrant group origin and destination pair.

I examine a regression of age, education and birthplace composition on headship at the immigrant group level. Columns 1 and 2 consider the estimating equation:

$$y_{gdt} = \beta_1 Age_{gdt} + \beta_2 SchoolingYears_{gdt} + \gamma_d + v_t + \epsilon_{gdt} \quad (3.1)$$

where y_{gdt} measures the headship rate of an immigrant group from g residing in MSA d in year t . Age and schooling years are measured as the mean adult age and schooling years for the immigrant group at time t . Column 2 of the table includes fixed effects for MSA γ_d and year, v_t . As expected, immigrant groups that are on average older and have more education tend to demand more houses per capita.

In columns 3-4, I aggregate groups to an MSA level, and consider within MSA changes in

immigrant characteristics. I consider a first-differenced region of MSA headship rates year over year on changes to immigrant age and skill level, with a control for population change.

$$\Delta y_{dt} = \beta_1 \Delta Age_{dt} + \beta_2 \Delta SchoolingYears_{dt} + \beta_3 \Delta Population_{dt} + \gamma_d + v_t + \epsilon_{dt} \quad (3.2)$$

Added to these variables are two nonlinear functions of age and population. Changes in age distribution are captured as the change in a herfindahl index across binned age groups, while changes in birthplace diversity capture changes in the concentration of origins that immigrants arrive from. Birthplace diversity at time t is captured as 1 minus the sum of squared shares of immigrant groups to total immigrant population $M_{dt} = \sum_g N_{gdt}$:

$$\pi_{dt} = \sum_g \left(\frac{N_{gdt}}{M_{dt}} \right)^2 \quad (3.3)$$

The age distribution herfindahl is calculated similarly, but with the relative shares of immigrants from a set of age quantiles. Increases in average skill and age at the MSA level predict increases in headship. Increasing birthplace diversity also predicts higher headship. As the sourcing of immigrants comes from more dispersed origins, headship is higher at destination. This may be explained by migrants from different areas having starkly different housing preferences and cohabiting less.

For a given migrant group or MSA, increasing average age predicts higher headship rates. Older migrants demand more homes per unit of population. Higher education also predicts higher headship, as higher income individuals demand more housing and are more likely to live alone. Higher birthplace diversity also predicts higher headship. The more diverse the migrant workforce, the more housing units are required per unit of population.

Table 3.1: Correlation of Migrant Characteristics and Headship

	(1) Group	(2) Group	(3) MSA Δ	(4) MSA Δ
Mean Age	0.01*** (0.00)	0.01*** (0.00)		
Mean Education (years)	0.02*** (0.00)	0.02*** (0.00)		
Δ Mean Skill level			0.02*** (0.00)	0.01*** (0.00)
Δ Mean Age			0.01*** (0.00)	0.00*** (0.00)
Age distribution change			-0.00 (0.00)	0.00 (0.00)
Δ Birthplace Diversity (1-HHI)			0.08*** (0.02)	0.05*** (0.01)
$\Delta \log(1 + \text{Immigrant Pop})$			-0.01* (0.00)	-0.00 (0.00)
Region FE	No	Yes	No	No
Year FE	No	Yes	No	Yes
MSA-level differenced outcome	No	No	Yes	Yes
Observations	285801	285801	4541	4541

Note: This table presents OLS estimates of a regression of headship rates on migrant composition at the origin-destination pair level and the MSA level. The first two columns show the correlation of migrant inflow characteristics within a given origin-destination pair on headship rates. The second column controls for year and destination level fixed effects. Columns 3 and 4 consider a regression of headship changes on immigration composition changes at the MSA level. The 4th column includes year effects. Data is taken from the yearly cross-sectional ACS. * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$.

3.3 Methodology

The goal of the empirical method is to relate changes in house prices to changes in different components of the migrant labor force at the MSA level. Because migrants endogenously sort to productive locations that are plausibly experiencing steeper house price trends, the standard approach isolates migration supply shocks with a shift-share instrument. Past work has sought to decompose the labor supply shock induced by migration in terms of high and low wage workers. Work by (Caiumi and Peri, 2024) and (Borjas, 2003) starts from a decomposition of total labor demand in terms of high and low-skilled labor shares, mediated by an elasticity of substitution σ .

$$L_t = \left[\theta_H L_{Ht}^{\frac{\sigma-1}{\sigma}} + \theta_L L_{Lt}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad (3.4)$$

From this formulation the authors derive a log-linear relationship between the relative wages of immigrants to natives, and the relative labor supply of immigrants and natives for a given skill level, subject to an elasticity of substitution. Shocks to high and low skilled labor are produced with an instrument that combines origin-specific changes in total outmigration with skill-specific country shares to capture the relative increase in high and low skilled labor arriving at a destination.

I set out to use a similar logic to construct compositional instruments that capture various dimensions of immigrant groups, leveraging variation both in the compositional shares and the changing composition of origin-level outmigration. Crucially, I am interested in estimating equations that simultaneously instrument for flows and composition, rather than separating a population inflow into skill components.

3.3.1 Theoretical Framework: Housing Demand and Immigrant Composition

This subsection provides a demand-side framework that motivates why immigrant composition should enter the estimating equation separately from total immigrant inflows. Consider destination d in year t , and let immigrants be partitioned into groups $g \in \mathcal{G}$ (for example, by origin,

age, and skill category). Let N_{gdt} denote the population of group g , h_{gdt} denote the headship rate (households per person), and $q_{gdt}(p_{dt})$ denote housing services demanded per household at price p_{dt} .

Total housing demand can be written as:

$$H_{dt}^D(p_{dt}) = \sum_{g \in \mathcal{G}} N_{gdt} h_{gdt} q_{gdt}(p_{dt}). \quad (3.5)$$

Define total immigrant population $M_{dt} = \sum_g N_{gdt}$ and composition shares $s_{gdt} = N_{gdt}/M_{dt}$.

Then

$$H_{dt}^D(p_{dt}) = M_{dt} \sum_{g \in \mathcal{G}} s_{gdt} h_{gdt} q_{gdt}(p_{dt}). \quad (3.6)$$

Equation (3.6) shows two distinct channels:

1. a *scale channel* through total immigrant inflows M_{dt} ;
2. a *composition channel* through $(s_{gdt}, h_{gdt}, q_{gdt})$.

Holding prices fixed, a first-order decomposition of demand changes is:

$$\Delta H_{dt}^D \approx \underbrace{\bar{\eta}_{dt} \Delta M_{dt}}_{\text{scale}} + \underbrace{\bar{M}_{dt} \sum_g (\bar{h}_{gdt} \bar{q}_{gdt} \Delta s_{gdt} + \bar{s}_{gdt} \bar{q}_{gdt} \Delta h_{gdt} + \bar{s}_{gdt} \bar{h}_{gdt} \Delta q_{gdt})}_{\text{composition and within-group behavior}}, \quad (3.7)$$

where bars denote local averages and $\bar{\eta}_{dt} = \sum_g \bar{s}_{gdt} \bar{h}_{gdt} \bar{q}_{gdt}$. Equation (3.7) implies that equal immigrant inflows can generate different housing-demand shocks when arriving groups differ in household formation and housing consumption intensity.

Let housing supply satisfy

$$\Delta H_{dt}^S = \kappa_d \Delta p_{dt}, \quad (3.8)$$

with $\kappa_d > 0$ capturing local supply elasticity. Market clearing ($\Delta H_{dt}^D = \Delta H_{dt}^S$) implies

$$\Delta p_{dt} = \frac{1}{\kappa_d} (\beta_M \Delta M_{dt} + \beta_X^\top \Delta X_{dt}) + u_{dt}, \quad (3.9)$$

where ΔX_{dt} collects composition moments such as changes in youth share, high-skill share, and headship propensity. This directly motivates estimating specifications with both flow and composition terms.

The framework delivers three testable implications used in the empirical section. First, after conditioning on immigrant flow growth, composition variables should retain explanatory power for headship and prices if groups differ in h_{gdt} or q_{gdt} . Second, omitting composition terms biases the interpretation of flow coefficients by combining scale and composition effects. Third, price responses to comparable inflow shocks should be stronger in lower-elasticity markets (smaller κ_d), creating heterogeneity across MSAs even with similar immigrant growth.

3.3.2 Shift-Share Instruments for Migrant Composition

An OLS estimation of equation 3.9 will be biased by the selective sorting of migrants to locations. If migrants choose to settle in higher productivity areas where wages and prices are on an upward trend, coefficients for M_{dt} and X_{dt} may be biased upward. I will instrument for each endogenous regressor separately using instruments.

Each MSA destination d has a baseline exposure to immigrants from country g according to a settlement share s_{gd} defined as the number of immigrants that are from g relative to all other origins residing in the destination MSA:

$$s_{gd} = \frac{N_{gd}}{\sum_{g=1}^G N_{gd}} \quad (3.10)$$

where N_{gd} is the number of working age immigrants from g in d in the baseline year. These shares are fixed in a baseline period, which in our case is the ACS year 2005. Year to year, the growth rate of immigrants from origin g residing in the US is defined as:

$$j_{gt} = \frac{N_{gt} - N_{g,t-1}}{N_{g,t-1}} \quad (3.11)$$

The instrument for the predicted total inflow of immigrants into MSA $\Delta \widehat{M}_{dt}$ is then the combination

of these elements, amounting to a weighted sum of growth rates, scaled by the baseline period immigrant share:

$$\Delta \widehat{M}_{dt} = Z_{dt}^M = ImmigrantShare_{d0} * \sum_g s_{gd} * j_{gt} \quad (3.12)$$

Now suppose that immigrants arrive at destinations with a set of characteristics that also matter for the outcome of interest. Consider a covariate X of the immigrant flow such as average age or education that may vary over time. Let x_{gdt} denote the mean value of characteristic X among immigrants from origin g in destination d at time t . Then the immigrant-weighted mean characteristic in destination d is:

$$X_{dt} = \sum_g s_{gdt} x_{gdt}, \quad s_{gdt} = \frac{N_{gdt}}{\sum_g N_{gdt}}. \quad (3.13)$$

Just as in the case of population inflows X_{dt} is a weighted average of origin-specific characteristics, where the weights are the origin shares within the destination.

The change in destination-level composition can be decomposed as

$$\Delta X_{dt} = \sum_g (s_{gdt} - s_{gd,t-1}) x_{gd,t-1} + \sum_g s_{gdt} (x_{gdt} - x_{gd,t-1}). \quad (3.14)$$

The first term captures *between-origin reallocation*: composition changes because the relative importance of different origins changes within the destination. The second term captures *within-origin change*: composition changes because the average characteristic of a given origin group evolves over time.

Changes in destination-level composition need not be proportional to changes in total inflows. Even if total immigrant inflows ΔM_{dt} are large, ΔX_{dt} may be small if all origins expand proportionally and within-origin characteristics remain constant. Conversely, composition may shift substantially even when total inflows are modest if inflows are concentrated among origins with systematically different values of X .

Table 3.2 reports correlations between MSA level measures of migrant inflows, defined as migrants that arrived in the last 5 years, migrant population shares and recent migrant characteristics.

Each column is a regression of inflows or shares on the average MSA level migrant age, years of education, headship rate and birthplace diversity. The specification includes MSA and year level fixed effects, and all variables represent log values. Increased migrant inflows are associated with younger recent migrants and lower average education, but increasing diversity. The association with lower education years is nearly 1 to 1, while a 1% drop in average age is associated with a 3% increase in migrant inflows. This is compatible with a story of higher inflows of migrants being drawn from a wider range of origins, and potentially a pool of less productive workers on the margin. While migrant inflows are positively associated with migrant shares, increasing migrant shares are associated with higher headship, while inflows are not. This suggests that changes in the rate of households per capita are more associated with the share of migrants and native workers, rather than pure increases in inflows.

Table 3.2: Correlation of Migrant Characteristics and Migrant Inflows

	Inflows	Share	Share
Log Mean Age	-3.412 [0.244]***	-0.756 [0.103]***	-0.171 [0.077]**
Log Mean Ed Years	-0.859 [0.260]***	-0.910 [0.096]***	-0.698 [0.080]***
Log Diversity	0.506 [0.133]***	0.055 [0.052]	-0.069 [0.039]*
Log Headship	-0.029 [0.108]	0.130 [0.035]***	0.094 [0.021]***
Log Migrant Inflows			0.166 [0.011]***
Mean Dep.	8.20	-2.37	-2.35
Observations	5,144	5,226	5,144

Note: This table presents OLS estimates of a regression of migrant inflows and immigrant shares on recent migrant mean age, skill level headship and birthplace diversity. All variables are in terms of logs, and the OLS regressions include MSA and year level fixed effects. Recent migrants are defined as foreign born individuals age 18-65 who report arriving in the country within 5-years of the survey. Data is taken from the yearly cross-sectional ACS. * p<0.01, ** p<0.05, *** p<0.01.

3.3.2.1 Implications of Composition for Shift-Share Estimation.

Suppose as in equation 3.9 that housing prices are a function of both total migrant inflows and a distribution of immigrant characteristics. For example, if older immigrants demand more homes than younger immigrants (as shown in our study of headship demand above), then the age distribution of migrants is an important feature of housing demand, and in turn housing prices. If the researcher ignores the composition margin, they might instead estimate

$$\Delta p_{dt} = b\Delta M_{dt} + \varepsilon_{dt}, \quad (3.15)$$

where Δp_{dt} is the change in house prices in destination d and period t , ΔM_{dt} is the total immigrant inflow. The change in composition ΔX_{dt} now appears in the error term:

$$\varepsilon_{dt} = \theta\Delta X_{dt} + u_{dt}. \quad (3.16)$$

The interpretation of b depends on the relationship between ΔX_{dt} and ΔM_{dt} . Suppose changes in composition are linearly related to inflows:

$$\Delta X_{dt} = a + \rho\Delta M_{dt} + v_{dt}, \quad (3.17)$$

where v_{dt} is orthogonal to ΔM_{dt} . Substituting (3.17) into equation 3.9 gives

$$\Delta p_{dt} = \beta\Delta M_{dt} + \theta(a + \rho\Delta M_{dt} + v_{dt}) + u_{dt} = \theta a + (\beta + \theta\rho)\Delta M_{dt} + (\theta v_{dt} + u_{dt}). \quad (3.18)$$

In this case, the conditional expectation of Δp_{dt} remains linear in ΔM_{dt} . A regression that omits composition is therefore not misspecified in functional form, but the coefficient on inflows no longer has a pure scale interpretation. Instead,

$$b = \beta + \theta\rho. \quad (3.19)$$

That is, the inflow coefficient conflates the direct effect of more migrants with the indirect effect operating through the composition response induced by those inflows.

If the researcher's estimand is the total reduced-form effect of inflows "as they arrive," this may be acceptable. But if the goal is to isolate the effect of the number of migrants separately from the effect of who those migrants are, then the specification is incomplete even under linearity.

Now suppose composition is a nonlinear function of inflows:

$$\Delta X_{dt} = f(\Delta M_{dt}) + v_{dt}, \quad (3.20)$$

for some nonlinear function $f(\cdot)$. Then the true outcome equation becomes

$$\Delta p_{dt} = \beta \Delta M_{dt} + \theta f(\Delta M_{dt}) + \theta v_{dt} + u_{dt}. \quad (3.21)$$

In this case, the conditional expectation of house prices is no longer linear in total inflows. Unless $f(\cdot)$ is well approximated by a linear function over the support of the data, estimating (3.15) imposes the wrong functional form. This case is particularly relevant when the omitted composition variable measures an object such as mean age, family structure, headship propensity, or skill diversity. For such variables, changes in composition may be negligible for small inflows but may accelerate once inflows become large enough to alter family reunification patterns, household formation, or the distribution of housing demand across unit types. In this case total inflows are not a sufficient statistic for migrant demand pressure in housing markets.

These distinctions matter directly for the interpretation of the shift-share instrument. The standard instrument

$$Z_{dt}^M = \sum_g s_{gd} * j_{gt} \quad (3.22)$$

is designed to generate exogenous variation in total inflows. But if origin-level shocks also induce predictable changes in destination-level composition, then the same instrument shifts both ΔM_{dt} and ΔX_{dt} . In that environment, a second-stage regression that includes only total inflows attributes

all instrument-induced variation in house prices to the scale of immigration, even though part of that variation may operate through composition.

What the instrument captures in this case varies depending on the correlation of inflow and composition related shocks. In the Appendix, I simulate an example where an inflow shift-share instrument is used to address endogeneity in the relationship between migrant inflows and an outcome y_{dt} . Figure 3.11 plots the results of an exercise where the true model is a function of both inflow and composition, but only inflows are included. The plots show the relative bias when composition and inflow are more or less correlated. If composition and inflow related shocks are uncorrelated, then the shift-share recovers something close to the structural β effect of increasing migrant population, without the related compositional effects. As composition and inflow shocks become more correlated, then the shift-share instrument captures an estimated total effect on y_{dt} through both inflow and compositional channels.

3.3.3 Shift Share Instruments by Age and Skill Cells

There are various strategies for constructing composition related shift-share instruments to add to the estimating equation. Caiumi and Peri (2024) propose a method for isolating migration flows by age and skill by splitting their measure of settlement shares into age-skill cells, defined in a pre-period. Rather than a share of Mexican migrants settling in MSA d in 1980, they consider the proportion of Mexican migrants of a given age and skill cohort arriving at d . This approach still uses the standard growth rate shifts, which is only a part of the origin-level variation available to the researcher.

An alternative strategy is to use age-skill specific settlement shares, and directly capture shifts in the demographic of interest at the origin level. Define a demographic-specific settlement share and growth rate as s_{gd}^X and j_{gdt}^X , where

$$s_{gd}^X = \frac{N_{gd}^X}{\sum_g N_{gd}^X}. \quad (3.23)$$

And N_{gd}^X represents the number of immigrants from g and demographic bin X that reside in d in the baseline period. The growth rates are calculated at a time t for each demographic bin or quantile.

$$j_{gt}^X = \frac{N_{gt}^X - N_{g,t-1}^X}{N_{g,t-1}^X} \quad (3.24)$$

Rather than combining these elements directly, I calculate deviations between the population level growth rates j_{gt} , and changes in the bin level growth rates j_{gdt}^X . Focusing on a given share, such as the share of immigrants in the youngest quantile, a predicted composition change instrument at time t is:

$$Z_{dt}^X = \sum_g s_{gd}^X * (j_{gdt}^X - j_{gt}) \quad (3.25)$$

By differencing out average growth rates, the constructed demographic instruments capture an idiosyncratic component of the origin level change in composition, plausibly independent from the change in total outmigration. Figure 3.9 plots pairwise correlations of the shift-share instruments against each other, and documents low correlations across compositional and inflow instruments. The demographic bins of interest in this study will be age quantiles and skill shares, as well as average headship rates. In particular I will predict changes in the youth share of recent immigrants, the high skill share of recent immigrants, and the average headship rate of recent immigrants.

3.3.4 Estimation Strategy

The estimating equation 3.9 relates house price growth to two endogenous regressors, immigration share growth and migrant compositional change. I instrument for ΔX_{dt} with Z_{dt}^X and ΔM_{dt} with Z_{dt}^M as described in equation 3.25 and equation 3.12 respectively.

In what follows, I test versions of the estimation including only one of the regressors, as well as both. Using both instruments simultaneously requires that both instruments have enough independent variation. In the appendix I show a simulation of the dual instrument strategy. Figure 3.10 shows the results of an exercise that mirrors the identification strategy with 2 shift-share

instruments and 2 endogenous regressors. I plot how bias varies for the inflow and composition coefficients as inflow and composition are more or less correlated. At high levels of correlation, the added variation provided by the second instrument is reduced, inflating the simulated confidence intervals. Figure 3.8 and Figure 3.9 show the pairwise correlations between the endogenous variables of interest, and between the constructed shift-share instruments respectively. In general I find low correlations between the variables of interest in terms of annual changes. For the shift-share instruments, this is by design as the youth and skilled share instruments are a function of deviations from the growth rates that appear in standard shift-share predictor for inflows. Headship does not use this construction, and is therefore naturally correlated with migrant population inflows.

I focus on a few key elements of migrant composition, and evaluate them separately. Immigrant youth share captures changes in the age distribution. Youth is defined as immigrant workers aged 18-33, or the first quantile. Immigrant high skill share captures changes in the average education level of migrants. A worker is labelled high-skilled if they attended four years of college. Immigrant headship rate captures changes in average headship, predicted by origin-level headship propensity in past periods.

For each composition margin, I construct an origin-level shift and interact it with baseline destination exposure shares. The age-composition instrument uses, for each origin g , the annual change in that origin's young-migrant share (e.g., ages 18–33 among recent immigrants), and aggregates these shocks to destination d using pre-period origin settlement shares $s_{gd,0}$; this yields a predicted change in destination age composition that is separate from predicted total inflows. The skill-composition instrument is constructed analogously, replacing the age share with the origin-level change in the high-skill share (college or above) among recent immigrants and aggregating with the same baseline settlement weights. The headship-composition instrument uses origin-level changes in immigrant headship rates (household heads per migrant) and applies the same shift-share aggregation to obtain predicted destination-level changes in household-formation intensity. In all three cases, the composition instrument is entered jointly with the standard inflow shift-share instrument so that identification comes from variation in compositional shocks beyond overall

migrant inflow growth.

The outcome y of interest is the percent change in the FHFA house price index, as well as the volatility of house price growth.

3.3.5 Interpreting Migrant Inflows and Migrant Shares

The empirical analysis considers both recent migrant inflows and changes in the migrant share of the population. These measures capture related but distinct margins of immigration. Scaled migrant inflows measure recent arrivals relative to the total population and therefore more directly capture the immediate increase in local population associated with new migrants. Holding other factors fixed, larger inflows increase the number of people demanding housing in an MSA.

Changes in the migrant share capture changes in the relative composition of the local population rather than new arrivals alone. The migrant share can increase because the immigrant population grows, because the native population declines, or through a combination of both. It may therefore reflect longer-run adjustments such as native outmigration, differential population growth, and changes in the relative housing demand of immigrant and native residents. Results using migrant shares should consequently not be interpreted as estimating only the effect of newly arriving migrants.

Because inflows and shares capture different population adjustments, their estimated relationships with house prices need not have the same sign. Indeed, in the results I find that a positive inflow effect on house prices, consistent with new arrivals increasing aggregate housing demand. I also find a consistent negative migrant-share effect, suggesting that on average migrants exert less housing demand per person than native residents. These results are consistent with prior work by (Saiz, 2007) and (Saiz and Wachter, 2011b). On top of these regressors, I then layer additional migrant features as described above.

3.4 Results

In what follows I will consider the effects of changes in migrant flow composition in terms of migrant inflows, shares, age, skill and headship propensity. The data is aggregated to the MSA level, and analyzed in terms of year over year percent changes in the variables of interest. I first discuss first stage tests of instrument relevance for the two endogenous regressors, and then consider OLS and 2SLS models relating immigrant composition to house price growth and volatility.

3.4.1 First Stage Tests of Instrument Relevance

Table 3.3 shows the results of the first-stage correlation of different instrument products on the endogenous variables of interest. The variables considered are the immigrant share, the share of youth defined as age bin 18-33, the share of high skilled immigrants defined as possessing a bachelor's degree, and the average headship rate. Each composition instrument is tested separately, controlling for the immigrant share in order to ensure that the instrument provides independent variation for the endogenous variable of interest. I find that the share, age and education composition instruments constructed are able to accurately predict changes in the relevant endogenous measures. However, the direct predictor of headship change is substantially weaker, and does not provide additional variation when immigrant share is controlled for.

3.4.2 Observational Relationship Between House Price and Migrant Inflows

Before running the second stage regression, I first run an OLS specification that regresses changes in house prices on changes in the immigrant composition. Table 3.4 shows the results from the following specification:

$$\Delta y_{gdt} = \beta_1 \Delta ImmigrantShare_{gdt} + \beta_2 \Delta ImmigrantComposition_{gdt} + \nu_t + \epsilon_{gdt} \quad (3.26)$$

Table 3.3: First Stage Estimates of Shift-Share Instruments for Migrant Characteristics

	Δ Immigrant Share	Δ Age	Δ Education	Δ Headship
Pred Immigrant Share	0.065 [0.021]***	-0.009 [0.020]	0.007 [0.013]	0.020 [0.011]*
Pred Immigrant Youth Share		0.014 [0.003]***		
Pred Immigrant High Skill Share			0.022 [0.006]***	
Pred Immigrant Headship				0.010 [0.008]
Mean Dep.	0.03	-0.00	0.00	0.01
Observations	2,938	2,671	3,377	3,454

Notes: First-stage regressions for endogenous MSA-level immigration variables. Dependent variables include immigrant share growth and composition changes (youth share, high-skill share, and headship rate). Instruments are shift-share predictors for total immigrant inflows and composition-specific shocks. Composition first stages include immigrant share growth controls to test independent predictive content. Specifications include year fixed effects with standard errors clustered by MSA. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The variable y_{gdt} is the percent change in HPI index constructed by the FHFA. The Immigrant Composition variable represents that vector of immigrant composition objects, including the youth share, high skilled share and headship rate. Immigrant share is consistently associated with increased house price growth. In the OLS I find limited impact of age and skill distribution on house prices, as measured by the aggregate FHFA house price index. These estimates capture within-year comparisons of changes in immigrant share and changes in house prices, and therefore suffer from the noise of low ACS observation counts as well as a lack of sharp changes in migration over the short time horizon.

However, some clear variation emerges when instead I use the Redfin single-family and multi-family house price indexes. Table 3.10 shows a weakly positive correlation between immigrant share growth and single-family house price growth. Conversely, Table 3.11 estimates equation 3.26 using a multi-family house price index, and finds that higher immigrant share growth predicts lower house prices. This correlational evidence suggests that immigrants may have a heterogeneous impact on the housing market, as different parts of the migrant distribution demand different types of housing.

In Table 3.9 I run the same specification for the outcome of housing price volatility, which

considers the deviation in house price using the 3-year rolling standard deviation in HPI growth. I find suggestive evidence that increasing migrant share growth is associated with reduced volatility in house prices.

Table 3.4: OLS Estimates of House Price Growth and Migrant Characteristics

	Δ HPI	Δ HPI	Δ HPI	Δ HPI
Δ Immigrant Share	0.055 [0.036]	0.067 [0.032]**	0.064 [0.032]**	0.058 [0.033]*
Δ Immigrant Youth Share		0.000 [0.002]	0.000 [0.002]	-0.000 [0.002]
Δ Immigrant High Skill Share			-0.002 [0.002]	-0.002 [0.002]
Δ Immigrant Headship Rate				-0.003 [0.003]
Mean Dep.	0.03	0.04	0.04	0.04
Observations	2,938	3,243	3,243	3,233

Notes: OLS estimates at the MSA-year level. The dependent variable is annual house price growth (FHFA). Key regressors include immigrant share growth and changes in immigrant composition variables (youth share, high-skill share, and headship rate). Regressions include year fixed effects and controls for housing supply and population size. Standard errors are clustered by MSA. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.4.3 2SLS Estimates of Migrant Impacts on House Prices

Table 3.5 estimates the second stage instrument regression of immigrant share growth and composition change on house price growth. The table reports Kleibergen-Paap F statistics and Sanderson-Windmeijer statistics for each model. I find that increasing migrant share predicts lower house price growth. Simultaneously instrumenting for migrant share growth and migrant composition yields underpowered results, as suggested by low F-statistics. The various instruments do not provide enough joint variation to identify the model.

In Table 3.13 I test the same model in the decade level census data, using changes across census years with 1990 as the baseline year. Migrant shares consistently predict lower house prices. This suggests that increased migrant shares affect demand preferences in destinations. As the local composition becomes more foreign born, there is lower pressure on house prices as migrants have lower headship or demand lower quality housing.

Table 3.5: 2SLS Estimates of House Price Growth and Migrant Characteristics

	Δ HPI	Δ HPI	Δ HPI	Δ HPI
Δ Immigrant Share	-0.045 [0.020]**	-0.101 [0.113]	-0.053 [0.028]*	-0.035 [0.076]
Δ Immigrant Youth Share		-0.102 [0.106]		
Δ Immigrant High Skill Share			-0.021 [0.036]	
Δ Immigrant Headship				-0.026 [0.113]
Mean Dep.	0.03	0.04	0.04	0.03
Observations	2,938	2,671	3,377	3,454
Kleibergen-Paap Fstat	9.104	0.540	5.050	1.469
Sanderson-Windmeijer Flows	9.104	1.472	6.879	3.813
Sanderson-Windmeijer Composition		8.293	13.468	3.026

Notes: Two-stage least squares estimates at the MSA-year level. The dependent variable is annual house price growth. Endogenous regressors are immigrant share growth and immigrant composition changes (youth share, high-skill share, and headship rate), instrumented with shift-share flow and composition instruments. All specifications include year fixed effects with standard errors clustered by MSA. Reported first-stage diagnostics include Sanderson-Windmeijer and Kleibergen-Paap statistics. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Again, the composition instruments do not provide independent variation and yield weak F-statistics. This is improved in table 3.14, where each instrument is estimated separately. The F-statistics are larger in the independent instrument case. I find that increasing youth share of migrants is associated with lower house price growth.

I next consider the effects using migrant inflow size instead of shares. Table 3.6 shows the same model replacing the migrant share variable with a measure of scaled migrant inflows defined as $\Delta \text{Immigrant Inflows} = \text{Recent Immigrants} / \text{Total Population}$. I find that increasing migrant inflows, unlike shares, are associated with increases in house price indices. This relationship is consistent whether or not I control for endogenous components of youth and high skilled share. I also find evidence that increasing youth share lowers house price growth, which is consistent with the evidence in table 3.14. A plausible explanation is that younger migrants are more likely to cohabit and less likely to demand single-family homes, putting less pressure on aggregate housing demand.

Table 3.7 takes each instrument separately, and estimates the effect of migrant characteristics on house price volatility. I find that increasing migrant share reduces house price volatility.

Table 3.6: 2SLS Estimates of House Price Growth with Migrant Flows and Composition

	Δ HPI	Δ HPI	Δ HPI	Δ HPI
Δ Immigrant Inflows	0.588 [0.169]***	0.798 [0.268]***	0.522 [0.194]***	1.318 [0.969]
Δ Immigrant Youth Share		-0.044 [0.026]*		
Δ Immigrant High Skill Share			-0.008 [0.036]	
Δ Immigrant Headship				0.220 [0.271]
Mean Dep.	0.04	0.04	0.04	0.04
Observations	3,085	2,671	3,243	3,422
Kleibergen-Paap Fstat	51.856	9.230	8.819	0.428
Sanderson-Windmeijer Flows	51.856	13.043	40.517	0.988
Sanderson-Windmeijer Composition		13.635	22.579	0.877

Notes: Two-stage least squares estimates at the MSA-year level. The dependent variable is annual house price growth. Endogenous regressors are immigrant inflows and immigrant composition changes (youth share, high-skill share, and headship rate), instrumented with shift-share flow and composition instruments. All specifications include year fixed effects with standard errors clustered by MSA. Reported first-stage diagnostics include Sanderson-Windmeijer and Kleibergen-Paap statistics. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Interestingly, increasing youth share increases house price volatility. This effect is maintained when also instrumenting for migrant inflows, as seen in Table 3.16. There is more churn in housing markets with younger tenants, likely due to the fact that younger households move more often, and change their housing formation decisions.

Last, I examine the impact of recent migrant composition on the type of dwellings chosen. Table 3.12 shows OLS regressions of migrant composition on MSA level changes in the average number of rooms of a migrant's dwelling, the building size and the building age. Higher skill shares of migrants are associated with larger buildings and younger buildings. Higher youth share is associated with fewer rooms and older buildings. These dynamics suggest that, controlling for migrant shares and inflows, the age and skill mix of migrants changes the composition of the types of buildings demanded within metros.

Table 3.8 shows the instrumented results for 2SLS regressions of changes in average number of rooms and building size. Positive shocks to the number of high-skilled migrants entering raise the number of rooms an average immigrant demands, as well as the building size. These results

Table 3.7: 2SLS Estimates of House Price Volatility and Migrant Characteristics

	HPI Vol	HPI Vol	HPI Vol	HPI Vol
Δ Immigrant Share	-2.707 [0.965]***			
Δ Immigrant Youth Share		1.992 [0.944]**		
Δ Immigrant High Skill Share			0.476 [0.422]	
Δ Immigrant Headship				21.758 [42.218]
Mean Dep.	0.03	0.04	0.04	0.04
Observations	3,465	3,243	3,377	3,615
Kleibergen-Paap	9.649	10.013	13.076	0.252

Notes: Two-stage least squares estimates at the MSA-year level. The dependent variable is the rolling standard deviation of house price growth. Endogenous regressors are immigrant share growth and immigrant composition changes (youth share, high-skill share, and headship rate), instrumented with shift-share flow and composition instruments. All specifications include year fixed effects with standard errors clustered by MSA. Reported first-stage diagnostics include Kleibergen-Paap statistics. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

suggest that high skilled migrants tend to settle in denser parts of a city, which might plausibly have a lower footprint on housing demand relative to single-family homes. At the same time, migrants with higher skill demand larger units, which changes the types of units that experience the most demand pressure.

3.4.4 Instrument Independence and Pretrend Diagnostics

This subsection reports three additional diagnostics to evaluate whether composition instruments capture variation beyond total immigrant-flow changes, and whether large composition shifts are preceded by differential trends in outcomes.

Figure 3.6 presents conditional first-stage estimates for composition outcomes. The first set of specifications regresses changes in immigrant mean age and mean education on composition instruments only. The second set adds immigrant share growth directly as a control. Comparing these specifications shows whether composition instruments retain predictive content once level-flow growth is partialled out. I find that the age-share IV and the high-skilled share IV are not affected by the inclusion of immigrant growth.

Table 3.8: 2SLS Estimates of Dwelling and Migrant Characteristics

	Rooms	Rooms	Rooms	Bldg Size	Bldg Size	Bldg Size
Δ Immigrant Inflows	17.877 [12.469]			-0.881 [13.413]		
Δ Immigrant Youth Share		0.629 [2.218]			0.758 [1.516]	
Δ Immigrant High Skill Share			3.002 [1.307]**			1.927 [1.152]*
Mean Dep.	-0.00	-0.01	-0.01	-0.01	-0.01	-0.01
Observations	2,922	2,901	3,427	2,922	2,901	3,427
Kleibergen-Paap	32.607	10.616	10.020	32.607	10.616	10.020

Notes: Two-stage least squares estimates at the MSA-year level. The dependent variables are changes in the average number of rooms of housing per recent migrant household, building size index and building age index. The building size is a categorical variable that captures the number of units in a building from 1-10, with the number 10 representing buildings of 50 or more units. The age index is a number from 1-28 that captures the age of the building, with lower numbers representing older buildings. All indices are taken from the ACS house dwelling characteristics variables. Endogenous regressors are immigrant share growth and immigrant composition changes (youth share, high-skill share, and headship rate), instrumented with shift-share flow and composition instruments. All specifications include year fixed effects with standard errors clustered by MSA. Reported first-stage diagnostics include Kleibergen-Paap statistics. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.18 shows regressions of the composition and migrant share instruments on lagged values of endogenous regressors and house price changes. The instruments are unrelated to lagged values of house price growth. I find that immigrant shares are somewhat correlated over time, with instruments for migrant inflows and headship predicting lagged migrant share growth.

Figure 3.7 and Table 3.17 report event-study pretrend checks around large composition-shift episodes. I identify composition shocks for the combination of the individual instruments for age and skill as top decile increases in the average of the two instruments. This top decile captures sharp increases in youth share, and high-skilled migrants. Event time is defined as the first MSA-year entering the top decile of the absolute composition-change index. Coefficients at leads (event times -5 to -2) are small and statistically imprecise for both headship and house price growth, consistent with limited evidence of differential pretrends in the current sample. I find that housing price changes in response to strong compositional shifts occur only after the event, with house prices increasing in response to large moves in age and skill instruments.

3.5 Conclusion


This paper evaluates the impact of immigrant composition on house price growth. I propose a method to decompose the impact of immigrants in terms of changes to the migrant share or inflows, and demographic variables that plausibly alter housing demand through their effects on the headship rate. The shift-share strategy exploits variation in origin-level shifts in demographic composition, as well as demographic-specific settlement shares in a pre period. I show that such instruments produce independent variation relative to standard shift-share instruments, and can be effectively used for decomposing population effects into different channels. I find that increasing migrant share consistently lowers house price growth, while increasing migrant inflows consistently raises house price growth. The fact that the coefficient on migrant inflows does not fluctuate when including compositional regressors suggests that the inflow coefficient when instrumented is capturing an inflow-specific term, rather than a composite of population and compositional features. These effects are consistent with a story of increased demand pressure from migrants to the extent that they increase local population. The fact that migrant shares predict lower house price growth suggests that relative to local workers, migrants may apply less pressure on housing prices than native workers, either through differences in headship or housing quality preferences.

I also find that younger migrants decrease house prices and increase house price volatility, suggesting that higher youth shares reduce pressure on aggregate housing demand, but also exacerbate price fluctuations. This may be because younger renters and buyers change housing arrangements more often, and are more mobile. The effects of youth share on volatility are maintained when controlling for migrant inflows, suggesting that changing migrant composition has independent effects on housing market fluctuations. Last, increasing migrant high skilled share raises the average number of rooms demanded by recent migrants, and increases the average building size resided in. This suggests that high-skilled migrant groups demand larger houses and multi-family buildings, which changes the composition of units demanded when the type of migrants changes.

3.6 Acknowledgements


Chapter 3, “The Effect of Migrant Composition on Housing Prices,” is currently being prepared for submission for publication of the material therein. The dissertation author, Andre Daniel Gray, is the sole author of this chapter.

3.7 Appendix 1: Additional Figures and Tables



Notes: The stacked bar decomposes total variation in immigrant-group headship rates into (i) variation explained by observable demographics (mean age, mean education, and changes in mean skill), (ii) additional variation explained by birthplace and year differences, and (iii) unexplained residual variation. Estimates use the MSA $\times$ birthplace $\times$ year sample and are weighted by immigrant-group population.

Figure 3.4: Variance decomposition of headship rates



(a) HPI Annual Percent Changes

(b) HPI Volatility

Notes: This figure shows pairwise correlations between different HPI indices. The indices include metro level annual measures from FHFA, Zillow, Redfin and (Bogin et al., 2019). The first panel compares indices in terms of percent changes at the metro-year level, while the second panel looks at volatility defined as a 3-year rolling standard deviation of log changes in the index.

Figure 3.5: Pairwise Correlations Between HPI Indices

Table 3.9: OLS House Price Volatility and Immigrant Composition

	Δ HPI	Δ HPI	Δ HPI	Δ HPI
Δ Immigrant Share	-0.001 [0.001]*	-0.001 [0.001]	-0.001 [0.001]	-0.001 [0.001]
Δ Immigrant Youth Share		-0.001 [0.000]	-0.001 [0.000]	-0.001 [0.000]
Δ Immigrant High Skill Share			-0.000 [0.000]	-0.000 [0.000]
Δ Immigrant Headship Rate				-0.000 [0.000]**
Mean Dep.	0.03	0.03	0.03	0.03
Observations	4,322	4,316	4,310	4,310

Notes: OLS estimates at the MSA-year level. The dependent variable is house price volatility, measured as the rolling 3-year standard deviation of annual house price growth. Regressors include immigrant share growth and changes in immigrant composition variables (youth share, high-skill share, and headship rate). Regressions include year fixed effects, and MSA fixed effects. Standard errors are clustered by MSA. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.10: OLS Single Family House Price and Immigrant Composition

	Δ HPI	Δ HPI	Δ HPI	Δ HPI
Δ Immigrant Share	0.056 [0.145]	0.026 [0.126]	0.052 [0.144]	0.071 [0.143]
Δ Immigrant Youth Share		0.016 [0.013]	0.020 [0.016]	0.020 [0.016]
Δ Immigrant High Skill Share			0.022 [0.025]	0.021 [0.025]
Δ Immigrant Headship Rate				0.015 [0.014]
Mean Dep.	0.07	0.07	0.07	0.07
Observations	1,828	1,981	1,981	1,980

Notes: OLS estimates at the MSA-year level. The dependent variable is the growth of the single-family house price index calculated by Redfin. Regressors include immigrant share growth and changes in immigrant composition variables (youth share, high-skill share, and headship rate). Regressions include year fixed effects, and MSA fixed effects. Standard errors are clustered by MSA. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.11: OLS Multi Family House Price and Immigrant Composition

	Δ HPI	Δ HPI	Δ HPI	Δ HPI
Δ Immigrant Share	-0.677 [0.591]	-1.131 [0.570]**	-1.116 [0.566]*	-0.923 [0.543]*
Δ Immigrant Youth Share		-0.058 [0.031]*	-0.056 [0.029]*	-0.051 [0.029]*
Δ Immigrant High Skill Share			0.014 [0.035]	0.005 [0.036]
Δ Immigrant Headship Rate				0.156 [0.075]**
Mean Dep.	0.10	0.10	0.10	0.10
Observations	1,764	1,912	1,912	1,911

Notes: OLS estimates at the MSA-year level. The dependent variable is the growth of the multi-family house price index calculated by Redfin. Regressors include immigrant share growth and changes in immigrant composition variables (youth share, high-skill share, and headship rate). Regressions include year fixed effects, and MSA fixed effects. Standard errors are clustered by MSA. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.12: OLS Housing Characteristics and Immigrant Composition

	Rooms	Bldg Size	Bldg Age
Δ Immigrant Share	2.727 [1.681]	3.360 [2.217]	4.128 [3.253]
Δ Immigrant Youth Share	-1.149 [0.184]***	0.196 [0.149]	-0.945 [0.297]***
Δ Immigrant High Skill Share	-0.034 [0.172]	1.211 [0.197]**	1.388 [0.282]***
Δ Immigrant Headship Rate	-0.421 [0.179]**	0.471 [0.191]**	0.381 [0.289]
Mean Dep.	-0.01	-0.01	0.12
Observations	4,133	4,133	4,133

Notes: OLS estimates at the MSA-year level. The dependent variables are changes in the average number of rooms of housing per recent migrant household, building size index and building age index. The building size is a categorical variable that captures the number of units in a building from 1-10, with the number 10 representing buildings of 50 or more units. The age index is a number from 1-28 that captures the age of the building, with lower numbers representing older buildings. All indices are taken from the ACS house dwelling characteristics variables. Regressors include immigrant share growth and changes in immigrant composition variables (youth share, high-skill share, and headship rate). Regressions include year fixed effects. Standard errors are clustered by MSA. * p<0.10, ** p<0.05, *** p<0.01.

Table 3.13: Decade Level House Price Growth And Immigration Composition

	Δ HPI	Δ HPI	Δ HPI
Δ Immigrant Share	-0.074 [0.026]***	-0.064 [0.022]***	-0.079 [0.041]*
Δ Immigrant Youth Share		-0.414 [0.618]	
Δ Immigrant High Skill Share			1.430 [5.135]
Mean Dep.	0.45	0.45	0.45
Observations	416	416	415
Kleibergen-Paap Fstat	98.069	1.008	0.084
Sanderson-Windmeijer Flows	98.069	12.476	3.079
Sanderson-Windmeijer Composition		2.090	0.174

Notes: Two-stage least squares estimates at the MSA-year level. The dependent variable is annual house price growth. Endogenous regressors are immigrant share growth and immigrant composition changes (youth share, high-skill share, and headship rate), instrumented with shift-share flow and composition instruments. All specifications include MSA and year fixed effects with standard errors clustered by MSA. Reported first-stage diagnostics include Sanderson-Windmeijer and Kleibergen-Paap statistics. * p<0.10, ** p<0.05, *** p<0.01.

Table 3.14: 2SLS Estimates of House Price Growth and Individual Migrant Characteristics

	Δ HPI	Δ HPI	Δ HPI	Δ HPI
Δ Immigrant Share	-0.045 [0.020]**			
Δ Immigrant Youth Share		-0.045 [0.021]**		
Δ Immigrant High Skill Share			-0.019 [0.034]	
Δ Immigrant Headship				0.020 [0.220]
Mean Dep.	0.03	0.04	0.04	0.04
Observations	2,938	2,671	3,377	3,615
Kleibergen-Paap	9.104	11.915	13.076	1.270

Notes: Two-stage least squares estimates at the MSA-year level. The dependent variable is annual house price growth. Endogenous regressors are immigrant share growth and immigrant composition changes (youth share, high-skill share, and headship rate), instrumented with shift-share flow and composition instruments. All specifications include MSA and year fixed effects with standard errors clustered by MSA. Reported first-stage diagnostics include Kleibergen-Paap statistics. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.15: 2SLS Estimates of HPI Volatility and Migrant Characteristics


	HPI Vol	HPI Vol	HPI Vol	HPI Vol
Δ Immigrant Share	-2.707 [0.965]***	-6.932 [4.757]	-3.499 [1.472]**	-8.263 [8.221]
Δ Immigrant Youth Share		-5.928 [5.220]		
Δ Immigrant High Skill Share			0.366 [0.902]	
Δ Immigrant Headship				9.491 [17.700]
Mean Dep.	0.03	0.06	0.04	0.03
Observations	3,465	3,243	3,377	3,454
Kleibergen-Paap Fstat	9.649	0.964	4.849	0.194
Sanderson-Windmeijer Flows	9.649	3.062	6.035	0.564
Sanderson-Windmeijer Composition		6.779	13.519	0.390

Notes: Two-stage least squares estimates at the MSA-year level. The dependent variable is annual house price growth. Endogenous regressors are immigrant share growth and immigrant composition changes (youth share, high-skill share, and headship rate), instrumented with shift-share flow and composition instruments. All specifications include MSA and year fixed effects with standard errors clustered by MSA. Reported first-stage diagnostics include Sanderson-Windmeijer and Kleibergen-Paap statistics. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.16: 2SLS Estimates of House Price Volatility and Migrant Flows and Composition


	HPI Vol	HPI Vol	HPI Vol	HPI Vol
Δ Immigrant Inflows	40.226 [6.306]***	54.934 [16.863]***	47.747 [6.657]***	98.692 [268.739]
Δ Immigrant Youth Share		1.722 [0.707]**		
Δ Immigrant High Skill Share			-0.119 [0.277]	
Δ Immigrant Headship				16.460 [91.459]
Mean Dep.	0.05	0.09	0.06	0.05
Observations	3,085	2,671	3,243	3,422
Kleibergen-Paap Fstat	51.856	9.863	8.819	0.015
Sanderson-Windmeijer Flows	51.856	16.339	40.517	0.043
Sanderson-Windmeijer Composition		21.557	22.579	0.036

Notes: Two-stage least squares estimates at the MSA-year level. The dependent variable is the rolling standard deviation of house price growth. Endogenous regressors are immigrant inflows and immigrant composition changes (youth share, high-skill share, and headship rate), instrumented with shift-share flow and composition instruments. All specifications include year fixed effects with standard errors clustered by MSA. Reported first-stage diagnostics include Sanderson-Windmeijer statistics. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.



Notes: Coefficients and 95% confidence intervals for composition-instrument first stages, estimated with and without immigrant share growth controls.

Figure 3.6: Tests of Instrument Independence in First Stage



Notes: Event-study coefficients (with 95% confidence intervals) for headship rates and FHFA house price growth around the first large composition-shift event in each treated MSA. Reference period is event time -1.

Figure 3.7: Event Study Tests of Pre-Trends

Table 3.17: Event-Study Lead Coefficients (Pretrend Check)

Outcome	Event time	Coef.	Std. Err.	p-value
Headship rate	-2	-0.0037	0.0035	0.2950
Headship rate	-3	-0.0024	0.0042	0.5595
Headship rate	-4	-0.0042	0.0038	0.2732
Headship rate	-5	-0.0042	0.0032	0.1940
FHFA house price growth	-2	-0.0030	0.0087	0.7264
FHFA house price growth	-3	0.0037	0.0066	0.5741
FHFA house price growth	-4	-0.0025	0.0069	0.7124
FHFA house price growth	-5	-0.0040	0.0069	0.5689

Event year is the first MSA-year in the top decile of absolute composition-IV index.


Reference period is event time -1; models include MSA and year FE with MSA-clustered SE.

Notes: Lead coefficients only (event times -5 to -2) from event-study specifications for headship and house price growth. Models include MSA and year fixed effects and cluster standard errors by MSA.


Table 3.18: IV Correlations with Lagged MSA Trends

Instrument	Lagged trend	Coef.	Std. Err.	p-value	N
Inflow IV	Lagged house-price growth	0.0629	0.0791	0.4278	3069
Inflow IV	Lagged immigrant-share growth	-1.1338	0.4640	0.0154	3069
Inflow IV	Lagged headship change	-0.3258	0.6114	0.5947	3069
Age comp IV	Lagged house-price growth	0.0687	0.1132	0.5449	2904
Age comp IV	Lagged immigrant-share growth	0.3139	0.3083	0.3100	2904
Age comp IV	Lagged headship change	1.4112	1.1640	0.2269	2904
Skill comp IV	Lagged house-price growth	-0.0474	0.1171	0.6861	2766
Skill comp IV	Lagged immigrant-share growth	0.0450	0.3469	0.8969	2766
Skill comp IV	Lagged headship change	-0.0948	0.7179	0.8951	2766
Headship comp IV	Lagged house-price growth	0.0193	0.0737	0.7935	3069
Headship comp IV	Lagged immigrant-share growth	-1.0826	0.4538	0.0180	3069
Headship comp IV	Lagged headship change	-0.3079	0.5988	0.6076	3069


Notes: This table regresses instruments on lagged values of characteristics of interest. Each row regresses one of the instruments on interest on first lags of endogenous variables or house price growth.




(a) Inflows vs. Youth Share Changes



(b) Immigrant Share vs. Youth Share Changes




(c) Inflows vs. Skilled Share Changes




(d) Immigrant Share vs. Skilled Share Changes

Notes: This figure shows pairwise correlations between endogenous variables of interest. Each variable is winsorized, and includes a line of best fit. Youth share is defined as the share of recent migrants that are 18-33. Skilled share is defined as the share of recent migrants that have a bachelor's degree or higher. Variables are in terms of annual % changes.


Figure 3.8: Pairwise Correlations Between Migrant Inflows and Composition




(a) Inflows vs. Youth Share Changes



(b) Inflows vs. Skilled Share Changes



(c) Inflows vs. Headship Changes



(d) Youth Share vs. Skilled Share Changes

Notes: This figure shows pairwise correlations between constructed instruments. Each variable is winsorized, and includes a line of best fit. Variables are in terms of annual % changes.

Figure 3.9: Pairwise Correlations Between Shift-Share Instruments

3.8 Appendix 2: Simulations of Shift-Share Identification with Inflow and Composition Effects

This appendix describes two Monte Carlo exercises used to evaluate shift-share instrumental-variable (IV) designs when migration inflows and migrant composition are jointly determined. The first simulation considers a case of instrumenting for inflows and composition simultaneously. The second simulation considers the consequences of ignoring compositional effects in a standard migration shift-share exercise.

3.8.1 Common Simulation Environment

In both exercises, I simulate a cross-section of counties indexed by $i = 1, \dots, N$ and origin groups indexed by $g = 1, \dots, G$. Let s_{ig} denote county i 's baseline share of origin g , with $\sum_g s_{ig} = 1$. I generate shift-share components using origin-level shocks:

$$z_i^M = \sum_{g=1}^G s_{ig} u_g, \quad z_i^C = \sum_{g=1}^G s_{ig} v_g,$$

where u_g is an inflow shock and v_g is a composition shock. The correlation between u_g and v_g is controlled by

$$v_g = \rho u_g + \sqrt{1 - \rho^2} \varepsilon_g, \quad \varepsilon_g \sim \mathcal{N}(0, 1),$$

so $\rho \in [0, 1)$ governs how strongly inflow and composition shocks comove.

County-level inflow and composition regressors are generated as

$$M_i = a_M z_i^M + b_M z_i^C + \lambda_M \eta_i + v_i^M, \quad C_i = a_C z_i^C + b_C z_i^M + \lambda_C \eta_i + v_i^C,$$

where η_i is a latent county factor and v_i^M, v_i^C are idiosyncratic shocks. The common factor η_i induces endogeneity by affecting both regressors and outcomes.

3.8.2 Exercise 1: Two Endogenous Regressors (Inflow and Composition)

The structural outcome equation is linear:

$$Y_i = \beta_M M_i + \beta_C C_i + \gamma \eta_i + \epsilon_i.$$

Both M_i and C_i are endogenous through η_i .

I compare:

1. OLS: Y_i on (M_i, C_i) .
2. 2SLS: instrument both endogenous regressors using shift-share instruments:

$$M_i \leftarrow z_i^M, \quad C_i \leftarrow z_i^C.$$

I vary ρ to change the correlation between inflow and composition shocks and record, for each ρ :

- Bias of $\hat{\beta}_M$ and $\hat{\beta}_C$ (OLS and IV),
- RMSE of $\hat{\beta}_M$ and $\hat{\beta}_C$,
- First-stage strength (F-statistics),
- Correlation diagnostics among regressors and instruments.

This exercise asks whether a two-instrument shift-share design can separately recover inflow and composition effects when both regressors are endogenous and increasingly collinear as ρ rises.

3.8.3 Exercise 2: Inflow-Only Estimation with Omitted Composition

Suppose now the researcher estimates only:

$$Y_i = \theta M_i + u_i,$$

using OLS or IV with z_i^M as the instrument for inflow.

The true model still includes composition:

$$Y_i = \beta_M M_i + \beta_C C_i + \gamma \eta_i + \epsilon_i$$

in the linear case, and

$$Y_i = \beta_M M_i + \beta_C C_i + \beta_{C^2} C_i^2 + \gamma \eta_i + \epsilon_i$$

in the nonlinear case.

Because composition is omitted, the inflow-only coefficient can absorb both direct inflow effects and composition- related channels. I therefore report:

1. Bias relative to the direct structural inflow effect β_M ,
2. Bias relative to the *total inflow-related projection effect*,


$$\theta^{\text{total}} = \beta_M + \beta_C \frac{\text{Cov}(C_i, M_i)}{\text{Var}(M_i)} \quad (\text{plus } \beta_{C^2} \frac{\text{Cov}(C_i^2, M_i)}{\text{Var}(M_i)} \text{ in the nonlinear case}).$$

I again vary ρ and report how OLS and IV inflow-only estimates move with inflow–composition comovement.


3.8.4 Simulation Protocol

Each (ρ, model) cell is replicated many times (e.g., 1,000 draws). For each cell we summarize mean estimates, bias, RMSE, first-stage strength, and relevant correlations. Figures plot how these quantities evolve as ρ increases.

These simulations clarify when shift-share designs isolate structural parameters versus when estimates reflect reduced-form mixtures due to comovement and omitted channels.




(a) Inflow Coefficient



(b) Composition Coefficient

Notes: This figure shows the results of a simulation with two endogenous regressors. The figures show the relative bias of an OLS strategy vs. a double instrumented 2SLS model where inflows and composition are predicted via shift-share. The x-axis shows varying levels of correlation between shocks to inflows and shocks to composition.

Figure 3.10: Bias with Two Endogenous Regressors by Shock Correlation



Notes: This figure shows the results of a simulation with two endogenous regressors, when only one regressor is included in the estimation. The figures show the relative bias of an OLS strategy vs. a single instrumented 2SLS model where inflows are predicted via shift-share. The x-axis shows varying levels of correlation between shocks to inflows and shocks to the missing composition variable.

Figure 3.11: Bias and Total Effect in Inflow-Only Shift-Share Simulation

Bibliography

- Abramitzky, R., Ager, P., Boustan, L., Cohen, E., and Hansen, C. W. (2023). The Effect of Immigration Restrictions on Local Labor Markets: Lessons from the 1920s Border Closure. *American Economic Journal: Applied Economics*, 15(1):164–191.
- Abramitzky, R. and Boustan, L. (2017). Immigration in american economic history. *Journal of economic literature*, 55(4):1311–1345.
- Abramitzky, R., Boustan, L., and Eriksson, K. (2020a). Do immigrants assimilate more slowly today than in the past? *American Economic Review: Insights*, 2(1):125–141.
- Abramitzky, R., Boustan, L., Eriksson, K., and Hao, S. (2020b). Discrimination and the returns to cultural assimilation in the age of mass migration. In *AEA Papers and Proceedings*, volume 110, pages 340–346. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- Abramitzky, R., Boustan, L. P., and Connor, D. (2020c). Leaving the enclave: Historical evidence on immigrant mobility from the industrial removal office. Technical report, National Bureau of Economic Research.
- Abramitzky, R., Boustan, L. P., and Giuntella, O. (2025). Enclaves and assimilation in the age of mass migration: Evidence from ethnic catholic churches. Technical report, National Bureau of Economic Research.
- Ager, P., Feigenbaum, J. J., Hansen, C. W., and Tan, H. R. (2023). How the Other Half Died: Immigration and Mortality in US Cities*. *The Review of Economic Studies*, page rdad035.
- Albert, C., Glitz, A., and Llull, J. (2021). Labor market competition and the assimilation of immigrants.
- Alesina, A., Devleeschauwer, A., Easterly, W., Kurlat, S., and Wacziarg, R. (2003). Fractionalization. *Journal of Economic growth*, 8:155–194.
- Alesina, A. and Ferrara, E. L. (2005). Ethnic Diversity and Economic Performance. *Journal of Economic Literature*, 43(3):762–800.

- Alesina, A., Harnoss, J., and Rapoport, H. (2016). Birthplace diversity and economic prosperity. *Journal of Economic Growth*, 21:101–138.
- Alesina, A. and Zhuravskaya, E. (2011). Segregation and the quality of government in a cross section of countries. *American Economic Review*, 101(5):1872–1911.
- Alsan, M. (2015). The effect of the tsetse fly on african development. *American Economic Review*, 105(1):382–410.
- Amodio, F. and Chiovelli, G. (2018). Ethnicity and violence during democratic transitions: evidence from south africa. *Journal of the European Economic Association*, 16(4):1234–1280.
- Ang, A., Kwon, E., and Zheng, S. (2026). Asian immigrants, school quality, and the us housing market. *Journal of Urban Economics*, 153:103860.
- Arbatli, C. E., Ashraf, Q. H., Galor, O., and Klemp, M. (2020). Diversity and Conflict. *Econometrica*, 88(2):727–797.
- Arslanalp, M. S., Koepke, M. R., and Verschuur, J. (2021). *Tracking trade from space: an application to Pacific island countries*. International Monetary Fund.
- Ashraf, Q. and Galor, O. (2013). The “out of africa” hypothesis, human genetic diversity, and comparative economic development. *American Economic Review*, 103(1):1–46.
- Bagde, S., Epple, D., and Taylor, L. (2022). The emergence of private high schools in india: The impact of public-private competition on public school students. *Journal of Public Economics*, 215:104749.
- Bazzi, S. (2017). Wealth heterogeneity and the income elasticity of migration. *American Economic Journal: Applied Economics*, 9(2):219–255.
- Bazzi, S. and Blattman, C. (2014). Economic shocks and conflict: Evidence from commodity prices. *American Economic Journal: Macroeconomics*, 6(4):1–38.
- Bazzi, S., Ferrara, A., Fiszbein, M., Pearson, T., and Testa, P. A. (2023). The other great migration: Southern whites and the new right. *The Quarterly Journal of Economics*, 138(3):1577–1647.
- Bazzi, S., Gaduh, A., Rothenberg, A. D., and Wong, M. (2016). Skill transferability, migration, and development: Evidence from population resettlement in indonesia. *American Economic Review*, 106(9):2658–2698.
- BenYishay, A., Rotberg, R., Wells, J., Lv, Z., Goodman, S., Kovacevic, L., and Runfolo, D. (2017). Geocoding afrobarometer rounds 1-6: Methodology & data quality. *AidData*. Available online at <http://geo.aiddata.org>.

- Berman, N. and Couttenier, M. (2015). External shocks, internal shots: the geography of civil conflicts. *Review of Economics and Statistics*, 97(4):758–776.
- Betts, J. R. and Fairlie, R. W. (2003a). Does immigration induce “native flight” from public schools into private schools? *Journal of Public Economics*, 87(5):987–1012.
- Betts, J. R. and Fairlie, R. W. (2003b). Does immigration induce ‘native flight’ from public schools into private schools? *Journal of Public Economics*, 87(5):987–1012.
- Bleakley, H. and Lin, J. (2012). Portage and Path Dependence *. *The Quarterly Journal of Economics*, 127(2):587–644.
- Boberg-Fazlić, N. and Sharp, P. (2024). Immigrant communities and knowledge spillovers: Danish americans and the development of the dairy industry in the united states. *American Economic Journal: Macroeconomics*, 16(1):102–146.
- Bogin, A., Doerner, W., and Larson, W. (2019). Local house price dynamics: New indices and stylized facts. *Real Estate Economics*, 47(2):365–398.
- Borjas, G. J. (2003). The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market. *The quarterly journal of economics*, 118(4):1335–1374.
- Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *The Review of economic studies*, 89(1):181–213.
- Borusyak, K., Hull, P., and Jaravel, X. (2025). A practical guide to shift-share instruments. *Journal of Economic Perspectives*, 39(1):181–204.
- Boustan, L. P., Cai, C., and Tseng, T. (2023). White flight from asian immigration: Evidence from california public schools. Technical report, National Bureau of Economic Research.
- Boustan, L. P., Fishback, P. V., and Kantor, S. (2010). The effect of internal migration on local labor markets: American cities during the great depression. *Journal of Labor Economics*, 28(4):719–746.
- Boyle, E. H., King, M., and Sobek, M. (2024). Ipums demographic and health surveys: version 11. (No Title).
- Brückner, M. (2012). Economic growth, size of the agricultural sector, and urbanization in africa. *Journal of Urban Economics*, 71(1):26–36.
- Buckles, K., Haws, A., Price, J., and Wilbert, H. E. (2025). Breakthroughs in historical record linking using genealogy data: The census tree project. *Explorations in Economic History*, page 101717.

- Burchardi, K. B., Chaney, T., and Hassan, T. A. (2019). Migrants, ancestors, and foreign investments. *The Review of Economic Studies*, 86(4):1448–1486.
- Caiumi, A. and Peri, G. (2024). Immigration’s effect on us wages and employment redux. Technical report, National Bureau of Economic Research.
- Callaway, B. and Sant’Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.
- Castells-Quintana, D. (2017). Malthus living in a slum: Urban concentration, infrastructure and economic growth. *Journal of Urban Economics*, 98:158–173.
- Chiovelli, G., Michalopoulos, S., Papaioannou, E., and Regan, T. (2023). Illuminating africa? *Institute for International Economic Policy Working Paper Series, IIEP-WP-2023-11*, George Washington University, 3(5).
- Choi, J., Hyun, J., and Park, Z. (2024). Bound by ancestors: Immigration, credit frictions, and global supply chain formation. *Journal of International Economics*, 147:103855.
- Christiaensen, L., Lozano-Gracia, N., et al. (2023). Migrants, markets, and mayors [migrants, marchés et maires]. *World Bank Publications-Books*.
- Christiaensen, L., Lozano-Gracia, N., et al. (2025). Africa’s urbanisation dynamics 2025: Planning for urban expansion. *OECD*.
- Cohen-Zada, D. (2009). An alternative instrument for private school competition. *Economics of Education Review*, 28(1):29–37.
- Combes, P.-P., Duranton, G., Gobillon, L., and Roux, S. (2010). Estimating agglomeration economies with history, geology, and worker effects. In *Agglomeration economics*, pages 15–66. University of Chicago Press.
- Desmet, K. and Rossi-Hansberg, E. (2024). Climate change economics over time and space. *Annual Review of Economics*, 16.
- Diamond, R. (2016a). The determinants and welfare implications of us workers’ diverging location choices by skill: 1980–2000. *American Economic Review*, 106(3):479–524.
- Diamond, R. (2016b). The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980-2000. *American Economic Review*, 106(3):479–524.
- Farnham, M., Schmidt, L., and Sevak, P. (2011). House prices and marital stability. *American Economic Review*, 101(3):615–619.
- Gagliarducci, S. and Tabellini, M. (2022). Faith and assimilation: Italian immigrants in the us.

- Technical report, National Bureau of Economic Research.
- García, D. and Paciorek, A. (2022). The remarkable recent rebound in household formation and the prospects for future housing demand.
- Glaeser, E. and Gyourko, J. (2025). America's housing supply problem: The closing of the suburban frontier? *Brookings Papers on Economic Activity*, 2025(1):375–425.
- Glaeser, E. L. and Gottlieb, J. D. (2009). The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States. *Journal of Economic Literature*, 47(4):983–1028.
- Goldin, C. (1999). A brief history of education in the united states.
- Goldin, C. and Katz, L. F. (1999). Human Capital and Social Capital: The Rise of Secondary Schooling in America, 1910-1940. *The Journal of Interdisciplinary History*, 29(4):683–723. Publisher: The MIT Press.
- Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.
- Gollin, D., Jedwab, R., and Vollrath, D. (2016). Urbanization with and without industrialization. *Journal of Economic Growth*, 21:35–70.
- Gollin, D., Kirchberger, M., and Lagakos, D. (2021). Do urban wage premia reflect lower amenities? evidence from africa. *Journal of Urban Economics*, 121:103301.
- Gross, R. N. (2017). *Public vs. private: The early history of school choice in America*. Oxford University Press.
- Guedes, R., Iachan, F. S., and Sant'Anna, M. (2023). Housing supply in the presence of informality. *Regional Science and Urban Economics*, 99:103875.
- Gyourko, J., Saiz, A., and Summers, A. (2008). A new measure of the local regulatory environment for housing markets: The wharton residential land use regulatory index. *Urban studies*, 45(3):693–729.
- Hamory, J., Kleemans, M., Li, N. Y., and Miguel, E. (2021). Reevaluating agricultural productivity gaps with longitudinal microdata. *Journal of the European Economic Association*, 19(3):1522–1555.
- Helfer, F., Grossmann, V., and Osikominu, A. (2023). How does immigration affect housing costs in switzerland? *Swiss Journal of Economics and Statistics*, 159(1):5.
- Henderson, J. V., Nigmatulina, D., and Kriticos, S. (2021). Measuring urban economic density. *Journal of Urban Economics*, 125:103188.

- Henderson, J. V., Storeygard, A., and Deichmann, U. (2017). Has climate change driven urbanization in africa? *Journal of development economics*, 124:60–82.
- Hjort, J. (2014). Ethnic divisions and production in firms. *The Quarterly Journal of Economics*, 129(4):1899–1946.
- Hoxby, C. M. (2003). School choice and school productivity. *The economics of school choice*, pages 287–342.
- Huber, S. J. and Schmidt, T. (2019). *Cross-country differences in homeownership: a cultural phenomenon?* Number 40/2019. Deutsche Bundesbank Discussion Paper.
- Imbert, C., Seror, M., Zhang, Y., and Zylberberg, Y. (2022). Migrants and firms: Evidence from china. *American Economic Review*, 112(6):1885–1914.
- IMF (2025). Imf primary commodities database. *IMF*.
- Jaeger, D. A., Ruist, J., and Stuhler, J. (2018). Shift-share instruments and the impact of immigration. Technical report, National Bureau of Economic Research.
- Jedwab, R., Ianchovichina, E., and Haslop, F. (2025). The employment profile of cities around the world: Consumption vs. production cities and economic development. *World Development*, 188:106883.
- Jedwab, R., Kerby, E., and Moradi, A. (2017). History, Path Dependence and Development: Evidence from Colonial Railways, Settlers and Cities in Kenya. *The Economic Journal*, 127(603):1467–1494.
- Jedwab, R. and Moradi, A. (2016). The Permanent Effects of Transportation Revolutions in Poor Countries: Evidence from Africa. *The Review of Economics and Statistics*, 98(2):268–284.
- Kamuiken, C. and Naito, H. (2024). The effect of climate change on internal migration: Evidence from micro census data of 16 sub-saharan african countries.
- Kaplan, G. (2012). Moving back home: Insurance against labor market risk. *Journal of Political Economy*, 120(3):446–512.
- Kelly, T. D., Matos, G. R., Buckingham, D. A., DiFrancesco, C. A., Porter, K. E., Berry, C., Crane, M., Goonan, T., and Sznoppek, J. (2010). Historical statistics for mineral and material commodities in the united states. *US Geological Survey data series*, 140:01–006.
- Kiszewski, A., Mellinger, A., Spielman, A., Malaney, P., Sachs, S. E., and Sachs, J. (2004). A global index representing the stability of malaria transmission. *The American journal of tropical medicine and hygiene*, 70(5):486–498.

- Kramon, E., Hamory, J., Baird, S., and Miguel, E. (2022). Deepening or diminishing ethnic divides? the impact of urban migration in kenya. *American Journal of Political Science*, 66(2):365–384.
- Lagakos, D., Marshall, S., Mobarak, A. M., Vernot, C., and Waugh, M. E. (2020). Migration costs and observational returns to migration in the developing world. *Journal of Monetary Economics*, 113:138–154.
- Lehner, B. and Grill, G. (2013). Global river hydrography and network routing: baseline data and new approaches to study the world’s large river systems. *Hydrological Processes*, 27(15):2171–2186.
- Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics*, 126(2):1029–1069.
- Li, X., Zhou, Y., Zhao, M., and Zhao, X. (2020). A harmonized global nighttime light dataset 1992–2018. *Scientific data*, 7(1):168.
- Lochmann, A., Rao, N., and Rossi, M. A. (2023). The long-run effects of south africa’s forced resettlements on employment outcomes. Technical report, Harvard’s Growth Lab.
- Logan, J. R., Jindrich, J., Shin, H., and Zhang, W. (2011). Mapping america in 1880: The urban transition historical gis project. *Historical Methods*, 44(1):49–60.
- Louie, S., Mondragon, J., Najjar, R., and Wieland, J. F. (2026). Housing affordability and housing demand. *FRBSF Economic Letter*, 2026(03):1–5.
- Marchiori, L., Maystadt, J.-F., and Schumacher, I. (2012). The impact of weather anomalies on migration in sub-saharan africa. *Journal of Environmental Economics and Management*, 63(3):355–374.
- McCarthy, J. (1973). History of Black Catholic Education in the Chicago, 1871-1971. *Dissertations*.
- McCord, G. C. and Rodriguez-Heredia, M. (2022). Nightlights and subnational economic activity: Estimating departmental gdp in paraguay. *Remote Sensing*, 14(5):1150.
- McGuirk, E. and Burke, M. (2020). The economic origins of conflict in africa. *Journal of Political Economy*, 128(10):3940–3997.
- McGuirk, E. F. and Nunn, N. (2024). Transhumant pastoralism, climate change and conflict in africa. *Review of Economic Studies*, page rdae027.
- Michalopoulos, S. (2012). The origins of ethnolinguistic diversity. *American Economic Review*, 102(4):1508–1539.
- Michalopoulos, S. and Papaioannou, E. (2013). Pre-colonial ethnic institutions and contemporary african development. *Econometrica*, 81(1):113–152.

- Montalvo, J. G. and Reynal-Querol, M. (2021). Ethnic Diversity and Growth: Revisiting the Evidence. *The Review of Economics and Statistics*, 103(3):521–532.
- Müller-Crepon, C., Pengl, Y., and Bormann, N.-C. (2022). Linking ethnic data from africa (leda). *Journal of Peace Research*, 59(3):425–435.
- Mullins, J. T. and Bharadwaj, P. (2021). Weather, climate, and migration in the united states. Technical report, National Bureau of Economic Research.
- Murray, T. J. (2016). Public or private? The influence of immigration on native schooling choices in the United States. *Economics of Education Review*, 53:268–283.
- Neal, D. (1997). The Effects of Catholic Secondary Schooling on Educational Achievement. *Journal of Labor Economics*, 15(1):98–123. Publisher: [University of Chicago Press, Society of Labor Economists, NORC at the University of Chicago].
- Notowidigdo, M. J. (2020). The incidence of local labor demand shocks. *Journal of Labor Economics*, 38(3):687–725.
- Nunn, N. and Puga, D. (2012). Ruggedness: The blessing of bad geography in africa. *Review of Economics and Statistics*, 94(1):20–36.
- Paciorek*, A. (2016). The long and the short of household formation. *Real Estate Economics*, 44(1):7–40.
- Porteous, O. (2019). High trade costs and their consequences: An estimated dynamic model of african agricultural storage and trade. *American Economic Journal: Applied Economics*, 11(4):327–366.
- Posner, D. N. (2004). The political salience of cultural difference: Why chewas and tumbukas are allies in zambia and adversaries in malawi. *American political science review*, 98(4):529–545.
- Ramankutty, N., Foley, J. A., Norman, J., and McSweeney, K. (2002). The global distribution of cultivable lands: current patterns and sensitivity to possible climate change. *Global Ecology and biogeography*, 11(5):377–392.
- Robinson, A. L. (2020). Ethnic Diversity, Segregation and Ethnocentric Trust in Africa. *British Journal of Political Science*, 50(1):217–239. Publisher: Cambridge University Press.
- Saiz, A. (2007). Immigration and housing rents in american cities. *Journal of urban Economics*, 61(2):345–371.
- Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125(3):1253–1296.

- Saiz, A. and Wachter, S. (2011a). Immigration and the neighborhood. *American Economic Journal: Economic Policy*, 3(2):169–188.
- Saiz, A. and Wachter, S. (2011b). Immigration and the Neighborhood. *American Economic Journal: Economic Policy*, 3(2):169–188.
- Sanchis-Guarner, R. (2023). Decomposing the impact of immigration on house prices. *Regional Science and Urban Economics*, 100:103893.
- Sanders, J. W. (1977). The education of an urban minority: Catholics in Chicago, 1833-1965. (*No Title*).
- Schubert, G., Stansbury, A., and Taska, B. (2024). Employer concentration and outside options. *Available at SSRN 3599454*.
- Sequeira, S., Nunn, N., and Qian, N. (2017). Migrants and the making of America: The short-and long-run effects of immigration during the age of mass migration. Technical report, National Bureau of Economic Research.
- Sharpe, J. (2019). Re-evaluating the impact of immigration on the US rental housing market. *Journal of Urban Economics*, 111:14–34.
- Shertzer, A. (2016a). Immigrant group size and political mobilization: Evidence from European migration to the United States. *Journal of Public Economics*, 139:1–12.
- Shertzer, A. (2016b). Immigrant group size and political mobilization: Evidence from European migration to the United States. *Journal of Public Economics*, 139:1–12.
- Shertzer, A., Walsh, R. P., and Logan, J. R. (2016). Segregation and neighborhood change in northern cities: New historical GIS data from 1900–1930. *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 49(4):187–197.
- Shrestha, M. (2017). Push and pull: A study of international migration from Nepal. *World Bank Policy Research Working Paper*, (7965).
- Unal, U., Hayo, B., and Erol, I. (2026). The effect of immigration on housing prices: Evidence from 382 German districts: The effect of immigration on housing prices: Evidence from 382 German districts. *The Journal of Real Estate Finance and Economics*, 72(2):386–424.
- Urquiola, M. (2016). Competition among schools: Traditional public and private schools. In *Handbook of the Economics of Education*, volume 5, pages 209–237. Elsevier.
- Vicente-Serrano, S. M., Beguería, S., and López-Moreno, J. I. (2010). A multiscalar drought index sensitive to global warming: the standardized precipitation evapotranspiration index. *Journal of climate*, 23(7):1696–1718.

Wang, Y. (2024). Linguistic distance, internal migration and welfare: Evidence from indonesia. *Internal Migration and Welfare: Evidence from Indonesia**(February 01, 2024).

WB (2025). World bank pink sheet data. *World Bank Group*.

Young, A. (2013). Inequality, the urban-rural gap, and migration. *The Quarterly Journal of Economics*, 128(4):1727–1785.