Essays in Urban Economics

HARVARD UNIVERSITY
Kenneth C. Griffin
Graduate School of Arts and Sciences

DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the

Department of Public Policy

have examined a dissertation entitled

"Essays in Urban Economics"

presented by Robert French

candidate for the degree of Doctor of Philosophy and hereby certify that it is worthy of acceptance.

Signature ______________________

Typed name: Professor Gordon Hanson, Chair

Signature ______________________

Typed name: Professor Edward Glaeser

Signature ______________________

Typed name: Professor Mark Shepard

Signature ______________________

Typed name: Professor Stefanie Stantcheva

Signature ______________________

Typed name: Professor Matthew Weinzierl

Date: May 1, 2024

Essays in Urban Economics

A dissertation presented

by

Robert French

to

The Department of Public Policy

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Public Policy

Harvard University

Cambridge, Massachusetts

May 2024

Dissertation Advisors: Professor Gordon Hanson, Professor Edward Glaeser

Robert French

Essays in Urban Economics

Abstract

The papers in this dissertation use large-scale administrative US Census Bureau data to understand the consequences of urban sorting for spatial inequality, housing affordability, and the welfare of low-income households. The first chapter (co-authored with Ashvin Gandhi and Valentine Gilbert) presents evidence suggesting gentrification primarily affects low-income renters by changing the characteristics of other neighborhoods in their choice sets. This finding has important policy implications: it suggests that policies promoting citywide affordability (e.g., zoning deregulation) may increase welfare for low-income renters more than policies like rent control which prioritize incumbent renters' welfare in the short-run. The second chapter (co-authored with Valentine Gilbert) compares the characteristics of residential vacancy chains created by new suburban single-family homes to those created by new urban multifamily housing. Our current results imply that new suburban housing supply has little effect on urban housing affordability or on the welfare of low-income urban households. The third chapter develops and estimates a quantitative spatial equilibrium model to demonstrate how optimal urban economic policy varies based on three sets of parameters that are subject to contentious debate, both popular and academic: parents' preferences, the technology governing children's upward economic mobility, and social preferences over children and their parents.

Contents

List of Tables

List of Figures

Acknowledgments

I thank my advisors, Edward Glaeser, Gordon Hanson, Mark Shepard, Stefanie Stantcheva, and Matthew Weinzierl, for their generous guidance and mentoring. I am grateful to many other mentors, coauthors, and colleagues who have supported me throughout my academic studies. I especially thank Shari Eli, Ashvin Gandhi, Valentine Gilbert, and Kory Kroft. I would also like to thank my family, friends, and all those who have helped me become a better person. I am particularly grateful for Joon's patience and support throughout my Ph.D.

To Joon and Papi.

Introduction

As cities grow and evolve, spatial sorting–the movement of individuals into different neighborhoods based on economic, social, and personal factors–fundamentally alters the spatial distribution of amenities, economic opportunity, and the set of neighborhoods affordable for low-income households. As certain areas become more desirable due to better consumption amenities, schools, or safer environments, rising demand can increase housing costs. Low-income households may benefit from these changes if they value the changing amenities more than the amount housing costs increase. Conversely, they may be harmed if the changing amenities are insufficient to compensate for the rising housing costs. Low-income households with strong attachments to their home neighborhoods are especially vulnerable to rising housing costs as they are less willing to relocate given any cost increase. In these and other ways, spatial sorting markedly shapes the welfare of low-income households living in urban environments.

By combining economic theory, structural estimation, and large-scale administrative US Census Bureau data, this dissertation provides a nuanced understanding of the consequences of spatial sorting and policies designed to alleviate its negative consequences. In the first chapter of my dissertation, my co-authors and I estimate a dynamic model of residential and workplace choice for low-income renter households. In our model, heterogeneous agents choose their neighborhood and workplace locations for each period to maximize their expected lifetime utility. We incorporate insights from the recent quasi-experimental shift-share literature to credibly identify structural parameters in households' utility functions. We then use these parameters to quantify the welfare effects of gentrification.

Our analysis in this chapter is retrospective. We attempt to understand how gentrification affected the welfare of low-income households in initially poor inner-city US neighborhoods over the past two decades. We find that low-income renter households are highly mobile within cities and have, on average, small attachments to their home neighborhoods. As a result, simulated neighborhood-level measures of changes in expected welfare are uncorrelated with neighborhood-level measures of gentrification throughout our analysis period. If one's home neighborhood changes in a way incumbent low-income renters do not value, our model estimates imply that low-income renters can easily move to an alternative neighborhood with similar characteristics if it exists. Therefore, we argue that gentrification primarily affects incumbent renters by changing the characteristics of other neighborhoods in their choice sets; where low-income renters lived within US metro areas mattered comparatively less than the US metro areas they lived in over the past two decades.

In the second chapter of my dissertation, Valentine Gilbert and I examine how *residential vacancy chains* link suburban and urban housing markets. We define *residential vacancy chains* as the series of moves across housing units initiated by the construction of new housing. Using administrative data on the residential histories of the U.S. population, we compare the characteristics of vacancies created by new suburban single-family homes to those created by new urban multifamily housing. We find that vacancy chains are short, with 90% ending within three rounds of moves. These short vacancy chains imply that each new suburban home leads to only .015 moves in low-income urban neighborhoods. We conduct a simulation exercise to understand what the observed patterns of vacancy chains imply about the welfare and price effects of new housing supply. We show that the geographic distribution of moves created by vacancy chains is correlated with the geographic distribution of welfare and price effects, and that the number of vacancies created in a neighborhood is as strong a predictor of price effects as are model-derived cross-neighborhood substitution effects. Our current results imply that new suburban housing supply has little effect on urban housing affordability or the welfare of low-income urban households.

In the third chapter of my dissertation, I move from conducting retrospective analyses in Chapters 1 and 2 to conducting a prospective analysis describing optimal urban economic policy to promote children's upward economic mobility. I consider when policy makers should prioritize providing opportunities for families to leave disadvantaged neighborhoods in Greater Boston and when they should prioritize investing resources in these disadvantaged neighborhoods and their schools. The third chapter develops a quantitative spatial general equilibrium model amenable to optimal policy analysis for a broad class of spillover functions and agents with heterogeneous preferences. I discuss how to quantify the model using administrative US Census data on the residential histories of most ACS respondents in Greater Boston. I then argue that as college-educated households' preferences for endogenous district-wide amenities increases vis-à-vis non-college-educated households' preferences, the social planner prioritizes investing in disadvantaged school districts. The intuition behind this result is that as college-educated households' preference for endogenous amenities increases, the more they are willing to pay to live in a school district with a higher share of other college-educated households. The social planner can then raise metro-wide taxes on college-educated households who self-segregate and spend the resulting funds on targeted housing assistance and school financing in the poorer district. Importantly, though, these results rely on school spending and peer composition being substitutes in the production of children's upward economic mobility. When school spending and peer composition are instead complementary in production, the additional funds raised from taxing college-educated households cannot be effectively spent in the poorer school district.

Chapter 1

Quantifying the Welfare Effects of Gentrification on Incumbent Low-Income Renters[1,2]

1.1 Introduction

Gentrification is associated with simultaneous increases in housing costs and changes in both public neighborhood amenities (e.g., schools and public safety) and private ones (e.g., restaurants and bars) (Couture and Handbury, 2020; Su, 2022). Residents may benefit from

[1]Chapter 1 is coauthored with Ashvin Gandhi and Valentine Gilbert. I thank Ed Glaeser, Gordon Hanson, Mark Shepard, Stefanie Stantcheva, and Matthew Weinzierl for their patience, support, and guidance. We also thank Nathaniel Baum-Snow, Tridevi Chakma, Rebecca Diamond, Myrto Kalouptsidi, Mengwei Lin, Erica Moszkowski, Mikko Silliman, and Ron Yang for helpful discussions and feedback. James Davis, Shahin Davoudpour, Shital Sharma, and Adam Shum helped with the Census data. Oleg Kolev, Kevin Wang, and Li Zhijian provided excellent research assistance. All errors are our own.

[2]Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2358 (CBDRB-FY24-P2358-R10957, CBDRB-FY24-P2358-R10936). This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by National Science Foundation grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation. This research was supported by The Horowitz Foundation for Social Policy and a Stone Research Grant from Harvard Kennedy School's James M. and Cathleen D. Stone Program in Wealth Distribution, Inequality, and Social Policy.

gentrification if they value the change in amenities more than the amount that housing costs increase (Vigdor, 2010). Conversely, residents may be harmed if the changes in amenities are insufficient to compensate for the rising housing costs. Residents with strong attachments to their home neighborhoods are especially vulnerable to rising housing costs as they are less willing to relocate given any cost increase. This paper leverages extensive US Census Bureau data to quantify these trade-offs for incumbent renters in gentrifying neighborhoods from 2000 to 2019.[3]

Beginning in the 1990s and intensifying after the year 2000, gentrification transformed the socioeconomic composition of vast areas within American inner cities (Couture and Handbury, 2023).[4] Between 2000 and 2017, the share of residents with a college degree in census tracts near their metro areas' central business districts (CBDs) increased by an average of 15 percentage points from a baseline of 24 percent. This contrasts with a 7-percentage-point average increase in suburban neighborhoods' shares of college graduates, reflecting a secular increase in educational attainment.[5] How this transformation of American inner cities affected the incumbent residents of gentrifying neighborhoods remains an open question.

The gentrification of inner-city America reversed the postwar urban decline, in which middle- and upper-class households left inner cities in favor of suburban life (Jackson, 1987; Boustan, 2010; Mieszkowski and Mills, 1993). This postwar suburbanization is considered a

[3]We focus on low-income renter households because of their vulnerability to the financial costs of gentrification and that, in 2010, 64 percent of urban housing units occupied by households with incomes below $50,000 were rented (Manson *et al.*, 2022). Recent research shows the potential for property taxes to harm homeowners in gentrifying neighborhoods (Ding and Hwang, 2020; Berry, 2021; Fu, 2022). Evaluating the welfare effects of gentrification on incumbent homeowners is an interesting question for future research.

[4]The economic forces causing demand for inner-city living to rise among college graduates were multifaceted. Rising top incomes increased demand for local service amenities concentrated in downtown neighborhoods (Couture *et al.*, 2023) and raised the time-cost of commuting (Edlund *et al.*, 2022; Su, 2022). Declining urban crime (Ellen and O'Regan, 2010; Ellen *et al.*, 2019), shifting preferences for urban amenities (Brueckner *et al.*, 1999; Glaeser *et al.*, 2001; Baum-Snow and Hartley, 2020; Couture and Handbury, 2020), evolving transportation infrastructure (LeRoy and Sonstelie, 1983; Glaeser *et al.*, 2008), and delayed childbearing (Moreno-Maldonado and Santamaria, 2022) all likely contributed to demand for downtown living. Increases in the valuation of downtown amenities *caused* by the growing presence of college graduates compounded these forces (Berkes and Gaetani, 2023; Diamond, 2016; Guerrieri *et al.*, 2013).

[5]See Appendix A.2 for more details on neighborhood change since the turn of the century.

significant contributor to concentrated inner-city disadvantage (Wilson, 1987). As primarily white middle- and upper-class households left for the suburbs, inner-city violent crime increased tenfold (Cullen and Levitt, 1999; Curci and Masera, 2023), city public finances declined (Derenoncourt, 2022), and many employers relocated to the suburbs (Kain, 1968; Glaeser and Kahn, 2001; Miller, 2021). If suburbanization was a significant factor behind these changes, its reversal over the past few decades might have caused environmental shifts favored by incumbent inner-city residents. Residents may have benefited from more local job opportunities, improved public amenities funded by rising land values, and greater access to private consumption amenities such as grocery stores and restaurants (Vigdor, 2002).

We conduct our analysis on deidentified person-level data from the US Census Bureau's Master Address File (MAF), which records the near universe of US adults' residential migration histories from 2000 onward. We link these data to persons' earnings and workplace locations from the employer–employee linked Longitudinal Employer–Household Dynamics (LEHD) database, which records the near universe of private sector, state, and local government workers' employment histories between 2000 and 2019. We finally link these data to person-level sociodemographic information from all American Community Survey (ACS) respondents (2005–2021) and property-level data from the Census Bureau's Master Address File Extract (MAF-X) and CoreLogic's residential property databases (2006–2019). Together, these data allow us to observe the residential locations, earnings, and workplaces of over 1 million low-income urban renter households across 50 large metro areas for up to 20 years during the most intense period of recent gentrification. Our panel data allow us to study how gentrification affects residents' welfare living in different neighborhoods within the same metro area—a hitherto underexplored question.

To quantify the welfare effects of gentrification on incumbent residents, we estimate a dynamic model of residential and workplace choice. In our model, heterogeneous agents choose their neighborhood and workplace locations for each period to maximize their expected lifetime utility. Agents are subject to a rich set of moving costs, can accumulate neighborhood-specific capital, and are forward-looking. They obtain flow utilities comprised of expected housing and nonhousing consumption, neighborhood amenities, and their accumulated neighborhood capital. Conditional on neighborhood rents, agents' expected consumption varies across neighborhoods because of differences in commute time–discounted proximity to jobs. This feature of our model captures that neighborhoods farther from employment centers are less desirable because of the increased financial cost of commuting (Le Barbanchon *et al.*, 2021).

We model neighborhood amenities as a function of neighborhoods' shares of college graduates. This choice is motivated by literature documenting a robust positive relationship between the provision of local public and private amenities and the local share of college graduates (Glaeser *et al.*, 2001; Diamond, 2016; Autor *et al.*, 2017; Su, 2022; Almagro and Domínguez-Iino, 2022; Hoelzlein, 2023). We further allow neighborhood amenities to vary over time with unobserved factors not caused by changes in the local share of college graduates. While the evolution of these unobserved exogenous neighborhood amenities is not caused by shifts in the local college graduate share, they may nonetheless be correlated with households' residential location choices, presenting a challenge to identification.

We identify our model parameters by combining establishment-level employment data from the Census Bureau's Longitudinal Business Database (LBD) with advances in the quasi-experimental shift-share literature. Specifically, we construct two sets of instrumental variables (IVs) to disentangle preferences for observed neighborhood characteristics and unobserved exogenous neighborhood amenities. The first set of IVs aggregates skill-specific shocks to potential commuting destinations. The intuition behind these instruments is that as neighborhoods' access to high-skill employment opportunities improves, they become more desirable for college graduates and, thus, are more likely to gentrify. These sets of IVs build on recent work in Baum-Snow *et al.* (2019) and Baum-Snow and Han (2023) that microfound measures of employment access with commuting data in a workplace choice model à la Tsivanidis (2022). Relative to this prior work, our establishment-level business data make identification from national industry shocks plausible, permitting the correlation between business establishment locations and unobserved neighborhood characteristics (Borusyak *et al.*, 2022).

Our second set of IVs is motivated by the observation that gentrification tends to occur near neighborhoods with already high shares of college graduates (Guerrieri *et al.*, 2013). For each census tract in our data, we construct distance-weighted measures of proximity to other neighborhoods' shares of college graduates. We then exploit that our analysis spans 50 metro areas and interact our neighborhood-level proximity measures with core-based statistical area– (CBSA–) wide Bartik labor demand shocks. Identification then proceeds analogously to that under a triple difference-in-difference estimator: we compare differences in gentrification between neighborhoods near and far from already gentrified tracts in CBSAs experiencing large labor demand shocks to differences in gentrification between neighborhoods near and far from already gentrified tracts in CBSAs not experiencing large labor demand shocks (Brummet and Reed, 2021).

We use our structural model to approximate the impact of changing neighborhood rents and shares of college graduates throughout 2000–2019 on the welfare of low-income incumbent renters. To do so, we use our parameter estimates to compute expected welfare separately for low-income renter households living in each low-income urban census tract in the year 2000. These calculations are based on the observed changes in the distribution of neighborhood rents and shares of college graduates throughout 2000–2019 but hold unobserved neighborhood amenities fixed at their 2000 levels. We then compare these measures to the expected welfare the same households would have obtained if the economy were instead in steady-state in the year 2000. This exercise yields census tract-level measures of how the changing distribution of neighborhood rents and shares of college graduates throughout 2000-2019 affected incumbent renters, as viewed from the standpoint of the year 2000. By finally comparing these measures across gentrified census tracts and census tracts that stayed poor, we uncover how living in gentrifying neighborhoods affected the welfare of incumbent low-income renters.

Our welfare analysis implies that, *on average*, incumbent renters initially living in census tracts that gentrified after the year 2000 were not made significantly worse off relative to incumbent renters initially living in census tracts that remained poor. As discussed above, incumbent renters will experience welfare losses from gentrification if their willingness to pay for accompanying amenity changes is less than the rise in real rental costs. These welfare losses can become large if households have high moving costs, a strong degree of neighborhood attachment, or there are few desirable alternative neighborhoods in households' choice sets. Our estimates suggest that low-income renter households living in poor neighborhoods both valued living in neighborhoods with a marginally greater share of college graduates and had only moderate moving costs.[6] These findings imply that low-income renters initially living in poor but gentrifying neighborhoods did not experience significant welfare losses relative to their counterparts initially living in neighborhoods that stayed poor from 2000 to 2019.

Our modest estimated moving costs underlie a core insight of our welfare analysis. Namely, where low-income renters lived within US metro areas mattered comparatively less than which US metro area they lived in from the standpoint of the year 2000. Because the average low-income renter household faced only modest costs from relocating to other low-income tracts within their metro area each period, changes to one's home census tract mattered less than changes to their home metro area overall. Policies directed at keeping metro areas broadly affordable for low-income renters may improve welfare for this population more than policies designed to ensure incumbent renters can remain in their home neighborhoods over extensive time horizons.

Relation to Literature Our paper contributes to three strands of literature. First, we contribute to the literature on the welfare implications of spatial sorting. Grounded in the

[6]We find that low-income Black (non-Black) renters are willing to pay $1,224 ($312) in annual rents to live in a neighborhood with a 10% higher share of college graduates (we calculate all dollar-denominated welfare costs in year-2010 dollars). We additionally find the fixed cost of moving within one's own CBSA is $3,578 ($1,692) for Black (non-Black) households and increases by $612 ($237) for households who have lived in the same census tract for at least five years. Households' baseline residential mobility informs these moderately sized structural moving cost estimates. We find that low-income renter households are highly mobile, with just 50.4 (51.7) percent of Black (non-Black) households remaining in their home census tracts for at least five years at a time.

canonical spatial equilibrium models stemming from Rosen (1974) and Roback (1982), an empirical literature has sought to quantify the implications of urban spatial sorting for households differentiated by their educational attainment. Moretti (2013) studies the implication of cross-metro sorting for real income inequality, while Diamond (2016) incorporates the endogenous supply of citywide amenities to study the implications of cross-metro sorting for welfare inequality. Similarly, Su (2022) and Couture *et al.* (2023) examine the welfare implications of spatial sorting but focus on within-metro sorting.[7] Other closely related papers focus on understanding the emergence of endogenously provided local amenities (Couture and Handbury, 2020; Almagro and Domínguez-Iino, 2022; Hoelzlein, 2023; Glaeser *et al.*, 2023). With few exceptions, the existing literature quantifies the effect of spatial sorting on the expected welfare of *prospective* city residents.[8] Our paper instead exploits rich panel data to examine the impact of gentrification on the welfare of *incumbent* renters.[9] To the best of our knowledge, we are the first to rigorously examine how the welfare effects of gentrification varied across neighborhoods for incumbent residents within US cities.

Second, we contribute to the empirical residential choice literature that estimates households' willingness to pay for housing and neighborhood characteristics. Earlier static models of residential choice (Brock and Durlauf, 2002; Bayer *et al.*, 2007; Vigdor, 2010) have given way to dynamic models that account for moving frictions and forward-looking behavior (Kennan and Walker, 2011; Bishop, 2012; Bayer *et al.*, 2016). Researchers have used these dynamic neighborhood choice models to estimate preferences over the racial composition of

---

[7]See Kuminoff *et al.* (2013) and Diamond and Gaubert (2022) for a comprehensive review of these and other papers examining the implications of spatial sorting on inequality. This research, in turn, contributes more broadly to the quantitative spatial equilibrium literature summarized in Redding and Rossi-Hansberg (2017).

[8]Balboni *et al.* (2020) is a notable exception, using a repeated static commuting model coupled with the exact-hat algebra of Dekle *et al.* (2008) to estimate the welfare impacts of transit infrastructure investments and the resulting sorting of households on the welfare of incumbent residents in Dar es Salaam, Tanzania. Couture *et al.* (2023) also use a static model to quantify the welfare effects of gentrification, comparing the expected welfare from living in different types of neighborhoods (downtown vs. suburbs) over time. The paper's static model necessarily abstracts from heterogeneity in residents' initial conditions conditional on income.

[9]Urban housing policies such as rent control and eviction protections prioritize incumbent renters' welfare over that of landlords and residents unprotected by such policies (Glaeser and Luttmer, 2003; Diamond *et al.*, 2018; Collinson *et al.*, 2023; Abramson, 2023). Understanding the difference in welfare effects between prospective and incumbent residents is thus critical to discerning the appropriate set of policy responses to gentrification.

neighborhoods (Davis *et al.*, 2023), the insurance value of rent control (Diamond *et al.*, 2018), the willingness to pay to avoid violent crime and air pollution (Bishop and Murphy, 2019), and horizontally differentiated consumption amenities (Almagro and Domínguez-Iino, 2022), among others. We contribute to this literature by providing estimates on low-income renters' preferences over welfare-relevant neighborhood characteristics, levels of neighborhood attachment, and a rich set of moving costs by combining our detailed census data with advances in the quasi-experimental shift-share literature.[10] We show how Black households appear to place less weight than non-Black households on access to employment opportunities but more weight on neighborhoods with higher shares of college graduates. We also show how the cost of moving between neighborhoods varies with the physical distance between neighborhoods and the *social distance* between neighborhoods, where we define *social distance* as the absolute difference between neighborhoods' shares of college graduates.

Third, our exploratory analyses in Section 1.3 contribute to empirical research documenting the effects of gentrification on observable outcomes for low-income residents. Much of this research has focused on gentrification's impact on the propensity of incumbent residents to leave their home neighborhoods (Freeman and Braconi, 2004; Freeman, 2005; Ellen and O'Regan, 2011; Ding *et al.*, 2016; Dragan *et al.*, 2020; Pennington, 2021). This research has recently broadened to consider a broader range of outcomes. Baum-Snow *et al.* (2019) examine the impact of neighborhood change on children's long-run outcomes. Brummet and Reed (2021) and Beauregard (2024) consider effects on employment alongside effects on residential mobility. Ferreira *et al.* (2023) explore minority households' local migration networks. Lester and Hartley (2014) and Meltzer and Ghorbani (2017) focus on the employment impacts of neighborhood change on incumbent residents. This literature broadly finds economically insignificant *average* effects on incumbents' household-level outcomes.[11] Statistical power,

[10] Our instrumental variable construction builds on Baum-Snow *et al.* (2019) and Baum-Snow and Han (2023), who are the first to construct and microfound shocks to employment access in a model of workplace choice à la Tsivanidis (2022). Brummet and Reed (2021) and Glaeser *et al.* (2023) use proximity to already gentrified tracts as an instrument for gentrification.

[11] Baum-Snow *et al.* (2019) find meaningful impacts on incumbent children's future credit outcomes, consistent

however, limits these studies' ability to examine heterogeneity across critical dimensions such as the environment in households' origin neighborhoods. We advance this literature by employing our extensive panel data to show that these average results are robust in our accounting for the environment in incumbent households' neighborhoods of origin. This is a surprising finding, as one might expect the impacts of gentrification to differ markedly based on local housing supply elasticities and the baseline level of neighborhood amenities, for example.

Roadmap This paper is structured as follows. Section 1.2 introduces our data and our sample of low-income households. Section 1.3 defines our measure of gentrification and presents descriptive evidence on the effects of gentrification on low-income incumbent renters. Sections 1.4 and 1.5 detail our dynamic model of neighborhood and workplace choice. Section 1.6 outlines our identification strategy and reports our parameter estimates. Finally, Section 1.7 presents our welfare analysis, and Section 1.8 concludes.

1.2 Data and Sample Construction

We conduct our analyses using person- and establishment-level administrative microdata from the US Census Bureau spanning 2000 to 2019. Table 1.1 provides an overview of these data sources. We postpone to Appendix A.1 the detailed discussion of the raw data and how we use it to construct our analysis samples, presenting only a cursory discussion here.

We conduct our analyses on an annual panel that records the earnings, workplaces, and residential addresses of over 1 million low-income urban renter households from 2000 to 2019. To construct this panel, we first form an annual panel of persons' residential histories using the MAF-ARF. We then merge to these residential histories persons' annual earnings and workplace locations from the LEHD and additional sociodemographic characteristics from the ACS. These merges are facilitated by a unique person identifier called a protected

---

with the potential for neighborhood environments to affect children's outcomes in adulthood (Chyn and Katz, 2021).

Table 1.1. *Data Overview*

Source	Coverage	Description
A. Household Panel		
Master Address File – Auxiliary Reference File (MAF-ARF)	2000–2019	Annual address-level residential locations
Longitudinal Employer–Household Dynamics (LEHD) database	2000–2019	Annual earnings, workplace locations, basic sociodemographics
American Community Survey (ACS)	2005–19	Detailed sociodemographics
B. Housing Characteristics		
CoreLogic	2006–2017	Address-level housing transactions and multiple listing service entries
Master Address File Extract (MAF-X)	2019	Address-level unit characteristics
C. Business Data		
Longitudinal Business Database (LBD)	2000–2019	Establishment-level employment and payroll aggregates

Notes: The MAF-X is a continuously updated inventory of all known living quarters in the US maintained by the US Census Bureau. The 2019 MAF-X is a snapshot of the MAF-X in 2019. Addresses verified in the past, but that are no longer known living quarters, remain in the MAF-X except in rare circumstances. The MAF-X 2019 therefore contains an inventory of all known addresses spanning our entire sample period.

identity key (PIK), assigned to individuals across data sets by the Census Bureau via probabilistic linking (Wagner and Layne, 2014). We aggregate earnings by housing unit and designate the highest earner of each unit as the household head for that year.[12] Our household panel is then restricted to persons identified as household heads who occupy a rental housing unit.[13]

We restrict our sample to household heads that are between 25 and 65 years of age. To focus our analysis on low-income households, we further restrict our sample to household heads earning in the bottom tercile of their respective CBSA and decadal age band in the

[12] We define housing units by their addresses in the MAF-X. Persons must have positive earnings in the given year to be considered a household head.

[13] We describe how we impute rental unit status in Appendix A.1. We also detail in Appendix A.1 how we smooth residential histories and handle changes to household formations, out-of-sample migrations, and missing observations.

year that they were first assigned household-head status.[14] We finally limit our study to household heads living in the urban cores of the 50 largest CBSAs, located within the 28 states for which data are accessible in the LEHD. We define urban cores using a method similar to the ones in Hwang and Lin (2016) and Couture and Handbury (2020). We denote as urban cores the set of census tracts associated with each CBSA that contain 50 percent of the CBSA's population closest to its central business district (CBD).[15] Table 1.2 presents descriptive statistics from a 2010 cross-section of our panel.[16]

1.3 Exploratory Analyses

This section reports findings from descriptive regressions of gentrification on an array of household-level outcomes. Our findings motivate a structural analysis and inform key features of our welfare analysis in Section 1.7. Before presenting our findings, however, we define our measure of gentrification and discuss our analysis period.

Gentrification We follow the existing literature and define gentrification at the census tract level (e.g., Ding and Hwang (2020), Dragan *et al.* (2020), Brummet and Reed (2021)). Specifically, we define gentrification from period t_0 to t as the increase in the number of college graduates in a census tract, n, during those years. We then normalize this measure by tracts' total adult population in time t_0:

$$\text{Gent}_{n,t_0,t} \equiv \frac{\text{College}_{n,t} - \text{College}_{n,t_0}}{\text{Adult Population}_{n,t_0}} \tag{1.1}$$

[14]CBSAs consist of counties associated with an urban core of at least 10,000 persons and adjacent counties deemed integrated through commuting ties. The Office of Management and Budget (OMB) defines CBSAs.

[15]Our CBD definitions come from Fee and Hartley (2013). These CBD definitions are with respect to 2008 CBSA delineations. To maintain consistency with these CBD definitions, we therefore use the 2008 delineations of CBSAs throughout our analysis. Moreover, while each CBSA is associated with a primary urban center, some CBSAs additionally contain secondary urban centers called metropolitan divisions with their own CBD. Although we construct our low-income cutoff from the earnings distribution of the entire CBSA, our urban core cutoffs are specific to each urban center's population, including metropolitan divisions. We believe these choices best capture our target population of low-income urban residents.

[16]The 2010 cross-section captures characteristics of the full sample that we use to estimate our descriptive regressions presented in Section 1.3. We are postponing the release of sample statistics for the complete panel to help streamline the US Census Bureau disclosure review processes.

Table 1.2. *Sample Characteristics: Household Heads in 2010*

	Black Households	Non-Black Households
Panel A: Household Head Characteristics		
Household Income	21,250	21,660
	(14,220)	(13,570)
Commute Time	27.64	25.6
	(12.43)	(12.82)
College Degree	0.11	0.165
	(0.313)	(0.371)
Immigrant	0.212	0.45
	(0.409)	(0.498)
Age	42.38	43.08
	(10.64)	(10.8)
Female	0.611	0.519
	(0.488)	(0.5)
Household Size	2.544	2.505
	(1.549)	(1.578)
Parent	0.275	0.266
	(0.447)	(0.442)
Panel B: Household Heads' Tract Characteristics		
Median Rents	834.4	934.7
	(232.4)	(286.2)
Median Property Value	247,000	289,200
	(205,100)	(214,600)
Share White	0.435	0.692
	(0.264)	(0.208)
Share College Educated	0.219	0.288
	(0.147)	(0.183)
Share College Educated and White	0.124	0.213
	(0.137)	(0.169)
Distance to CBD	9.552	9.56
	(5.172)	(5.33)
Unique Households	314,000	688,000

Notes: Table 1.2 reports mean characteristics for household heads with standard errors in parentheses. The sample comprising Table 1.2 consists of all household heads in the 2010 cross-section of the panel. Panel A reports household-level characteristics during 2010. Panel B reports characteristics of household heads' census tracts, also in 2010. Dollars are deflated to 2010 levels, and census tracts are delineated by 2010 boundaries. Commute time is measured in minutes. The Parent variable is calculated only for household heads present in the ACS 2005–2021 and is inferred by the reported age of the child in the year in which the household head is a survey respondent. The college degree variable is computed only for PIKs for whose education variables are not imputed in the LEHD or for whom we ascertain educational attainment through our ACS surveys. Details on the construction of tract aggregates are in Appendix A.1. *Sources*: ACS (2005–2021), LEHD (2010), CoreLogic (2006–2017), MAF-X (2019), and MAF-ARF (2010). Sample characteristics were disclosed by the US Census Bureau's Disclosure Review Board. Project Number 2358. Disclosure Clearance Number CBDRB-FY24-P2358-R10936.

We follow the economics literature by normalizing the change in the number of college graduates by the total adult population in t_0 (Brummet and Reed, 2021; Card *et al.*, 2008; Böhlmark and Willén, 2020). One alternative definition is the change in the *share* of college graduates between periods t and t_0. We nonetheless follow the literature's convention since this measure minimizes the mechanical relationship between gentrification and a primary outcome of interest: the out-migration rates of incumbent renters. This is because the out-migration rates of low-income incumbent renters have little influence on our measure of gentrification; only 11 percent and 16.5 percent of Black and non-Black households in our sample possess a college degree, and our sample comprises a small fraction of total households in each neighborhood.[17] Our choice to define gentrification on the basis of educational attainment also follows convention in the economics literature (Diamond, 2016; Brummet and Reed, 2021; Su, 2022).[18]

Analysis Period We choose to focus on the years 2010–2019 to establish our descriptive findings. We do so for two primary reasons. First, this choice mitigates the influence of potentially confounding factors resulting from the Great Recession, especially since we can control for changing neighborhood-level characteristics before 2010. Second, restricting our panel to 2010–2019 allows us to control for household characteristics throughout 2000–2009. Household characteristics such as length of prior residential tenure are likely correlated with households' residential location choices.

[17]All the results that we present below are quantitatively similar to those from specifications that exclude college-educated adults from our sample of low-income renters.

[18]Some research defines gentrification based on changes in income (e.g., Dragan *et al.* (2020) and Ding and Hwang (2020)). Other research considers gentrification through the lens of changing racial compositions (Baum-Snow and Hartley, 2020). We nonetheless believe that focusing on educational composition offers the clearest connection to the existing literature. We have experimented with alternative definitions based on racial composition and income and found quantitatively similar results. We will disclose these results in due course.

1.3.1 Is Gentrification Associated with Higher Neighborhood Out-Migration Rates?

Gentrification is not associated with higher neighborhood out-migration rates among low-income incumbent renter households. We show that this finding holds across a range of baseline neighborhood environments, including census tracts that are already partially gentrified and highly developed.

We estimate a set of Cox proportional hazard models to understand the relationship between gentrification and incumbent renter households' out-migration rates. These models estimate gentrification's impact on the probability that an incumbent household leaves its origin neighborhood in any year (i.e., on incumbents' hazard rates).[19] Our Cox models take the following form:

$$\log(h(t|i)) = \alpha^{Cox} + \beta_{\text{NC}}^{Cox}\text{Gent}_{n(i),10,19} + \gamma^{Cox}X_i + \delta^{Cox}X_{n(i)} + \alpha_{CBSA}^{Cox} + \varepsilon_i^{Cox} \tag{1.2}$$

where $h(t|i)$ is the hazard in period t for household i. $n(i)$ denotes household i's origin neighborhood, X_i is a vector of household-level controls, $X_{n(i)}$ is a vector of controls characterizing the origin neighborhood of household i, and α_{CBSA}^{Cox} is a CBSA-level fixed effect. We detail and motivate our choice of controls in Appendix A.3. To mitigate concerns over sample selection, we often restrict our sample to longtime renters, defined as renter households that have lived in their origin tract for at least five years before 2010. We also postpone the discussion of this choice and identification more broadly to Appendix A.3. Our descriptive analyses cluster standard errors at the census tract level, our treatment unit (Abadie *et al.*, 2023).

Estimates of β_{NC}^{Cox} from equation 1.2 are reported in Panel A of Table 1.3. Columns (1) and (2) of Panel A in Table 1.3 report estimates for our full sample of renter householders

[19]Our choice to estimate Cox models is motivated by incumbent renters' short unconditional neighborhood tenures. Only 50.4 (51.7) percent of Black (non-Black) incumbent renter households remained in their 2010 origin census tract until at least 2015; these unconditional survival probabilities fall to 31.3 and 32.8 percent for 2019, respectively. Existing research that relies on the intermittent sampling of residents misses potential identifying variation from incumbent residents with short unconditional residential tenures. The Cox proportional hazard model allows us to efficiently utilize our annual residential histories panel to identify gentrification's effects on incumbent renters' neighborhood tenures.

separately for Black- and non-Black-headed households. Columns (3) and (4) report the same analyses but are restricted to longtime renters. Finally, columns (5) and (6) further restrict to census tracts with an initial share of college graduates below the sample-weighted median among all tracts in our sample. We document an economically insignificant relationship between gentrification and incumbent renters' hazard rates for all subsets of our data.

Consider the effect of neighborhood change on non-Black longtime incumbent renters' hazard rates (column (4) in Panel A of Table 1.3). A 10-percentage-point increase in gentrification corresponds to a 1.87 percent increase in the probability that these renters leave their origin neighborhood in any given year between 2010 and 2019. Since the unconditional probability of leaving one's origin neighborhood in any year (i.e., baseline hazard) reaches at most 20 percent, these effect sizes are negligible.[20] These results are consistent with the extant literature attributing neighborhood change to changes in in-migration patterns as opposed to increased out-migration among incumbent residents (e.g., Ding *et al.* (2016), Dragan *et al.* (2020), and Brummet and Reed (2021)).

We advance our understanding of the impacts of gentrification on out-migration by using our expansive panel data to show that these null results do not mask meaningful underlying heterogeneity. Column (5) shows that, if anything, out-migration rates *decrease* among Black low-income renters in neighborhoods with an initially low share of college graduates. A 10-percentage-point increase in gentrification corresponds to a 7.7 percent decrease in the probability that a longtime incumbent Black renter leaves her origin neighborhood if it has an initially low share of college graduates. Again, however, with a baseline hazard rate of at most 20 percent, these are economically small effect sizes. Table A.2 in Appendix A.3 explores additional sources of potential heterogeneity, including the density of urban development. We continue to find economically insignificant effect sizes. Finally, we test the robustness of our Cox proportional hazards models by estimating linear probability models where the dependent variable is an indicator equal to one if an incumbent renter remains in

[20] A ten-percentage-point increase in gentrification is associated with at most a $(20/100) \times 1.87 = 0.374$ percentage point increase in the probability that an incumbent renter household leaves its origin neighborhood in any year.

her origin census tract until at least 2019.[21] We obtain similar results in this framework.

1.3.2 Is Gentrification Associated with Changing Economic Outcomes for Incumbent Renters?

Gentrification is not meaningfully associated with changing economic outcomes for incumbent renters This is true across most baseline neighborhood environments. We do, however, find that gentrification is associated with lower average annual earnings for longtime Black renters living in tracts with an initially low share of college graduates.

To understand the relationship between gentrification and incumbent renters' economic outcomes, we estimate linear regression models regressing changes in incumbent residents' annual earnings and commute distances on gentrification,

$$\Delta y_i = \alpha^{LP} + \beta^{LP}_{\text{NC}}\text{Gent}_{n(i),10,19} + \gamma^{LP} X_i + \delta^{LP} X_{n(i)} + \alpha^{LP}_{CBSA} + \varepsilon^{LP}_i \tag{1.3}$$

where Δy_i denotes the change in either annual earnings or commute distances between 2010 and 2019. X_i, $X_{n(i)}$, and α^{LP}_{CBSA} are the same set of control variables and fixed effects used in equation 1.2. Estimates of β^{LP}_{NC} are reported in Panel B of Table 1.3. To interpret the magnitude of the coefficients in Panel B of Table 1.3, consider the effect of gentrification on Black incumbent renters' annual earnings and commute times. A 10-percentage-point increase in our measure of gentrification is associated with \$287 higher annual earnings and .2 minute quicker commutes, economically small effect sizes. We do find some evidence, however, that gentrification is associated with lower annual earnings for longtime Black incumbent renters in tracts with a low initial share of college graduates. Here, a 10-percentage-point increase in gentrification is associated with \$914 lower average annual earnings after ten years.

[21] Cox proportional hazards models assume that treatment effects are constant over time (i.e., proportional hazards). Our linear probability model specification mimics the existing literature that relies on intermittent sampling of residents (e.g., Brummet and Reed (2021)).

Table 1.3. *Effect of Gentrification on Incumbent Renters*

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Cox Proportional Hazards Model Outcomes						
Hazard Rate	-0.0688	-0.107*	-0.0793	0.187*	-0.777***	-0.264
	(0.0542)	(0.0439)	(0.132)	(0.0781)	(0.286)	(0.325)
Panel B: Linear Regression Model Outcomes						
Leave Tract	-0.0176	-0.0263	-0.016	0.07*	-0.305***	-0.0888
	(0.0223)	(0.0174)	(0.0509)	(0.0278)	(0.106)	(0.0894)
Annual Earnings	2,872**	2,268**	-932	-86.05	-9,138*	-551.8
	(948.8)	(736.4)	(2,170)	(1,375)	(3,854)	(3,427)
Commute Time	-0.2280	-1.905***	0.637	0.5900	-1.854	3.146
	(0.776)	(0.531)	(1.609)	(1.002)	(3.132)	(2.775)
Rent	0.271***	0.0960***	0.389***	0.222***	0.863***	0.607***
	(0.0472)	(0.0288)	(0.0620)	(0.0329)	(0.145)	(0.096)
College Share	0.726***	0.299***	1.079***	0.567***	3.726***	4.108***
	(0.204)	(0.0598)	(0.131)	(0.0788)	(0.250)	(0.236)
Controls						
Full Controls	✓	✓	✓	✓	✓	✓
CBSA Fixed Effects	✓	✓	✓	✓	✓	✓
Sample Restrictions						
Race	Black	Non-Black	Black	Non-Black	Black	Non-Black
Longtime Renters			✓	✓	✓	✓
Low Initial College Share					✓	✓
N (1,000s)	314	688	56	156	35	67

Notes: We discuss how to interpret all coefficients in the main text. Leave Tract is an indicator equal to 1 if the household leaves its origin tract before 2019. We measure annual earnings in 2010 dollars. We measure rent and college shares in percent changes. We measure commute times in minutes. Every specification includes the full set of controls listed and detailed in Appendix A.3. Standard errors clustered at the census tract level are in parentheses. Longtime renters are renters who have lived in their origin census tract since at least 2005. Tables A.3 and A.2 in Appendix A.3 report results for a wider range of baseline neighborhood environments. *Sources*: ACS (2005–2021), LEHD (2010), CoreLogic (2006–2017), MAF-X (2019), and MAF-ARF (2010). Estimates were disclosed by the US Census Bureau's Disclosure Review Board. Project Number 2358. Disclosure Clearance Number CBDRB-FY24-P2358-R10936.

1.3.3 Is Gentrification Associated with Changing Neighborhood Conditions for Incumbent Renters?

Gentrification is associated with meaningful changes in the neighborhood conditions incumbent renters experience. That is, incumbent renters who initially live in gentrifying census tracts experience greater changes in their neighborhood conditions than similar residents in census tracts not gentrifying. This is not a mechanical result, as incumbent renters can move across tracts throughout our analysis period.

Panel A in Table 1.3 reports the relationship between gentrification and changes in incumbent renters' neighborhood conditions. For example, a ten-percentage-point increase in gentrification is associated with Black incumbent renters living in tracts during 2019 that had, on average, 2.7 percent higher rents and 7.3 percent higher shares of college graduates. We see in columns (5) and (6) that gentrification was strongly associated with the neighborhood conditions of incumbent renters initially residing in tracts with a low share of college graduates.

Our estimates on experienced neighborhood characteristics suggest moving frictions for at least some low-income renter households in our sample. Without moving frictions, renters would reoptimize their location choice each period to ensure proximity to their ideal bundle of neighborhood characteristics, yielding economically insignificant estimates on incumbent renters' neighborhood conditions. The presence of moving frictions, in turn, suggests the potential for differences in welfare effects from gentrification across tracts *within* CBSAs, motivating our paper's focus on incumbent renters. Moving frictions make incumbent renters averse to leaving their home census tract irrespective of changes in its characteristics. Whether incumbent renters then benefit from gentrification depends on their relative valuations of rents, job market access, and amenities vis-à-vis the actual change in these neighborhood characteristics. Our structural analysis in Section 1.4 examines whether the implied moving frictions translate into large average moving costs and quantifies their importance for households' welfare.

It is worth noting that there is limited correspondence between household residential mobility and welfare. Indeed, our estimates documenting no meaningful increase in neighborhood exit rates in response to gentrification are consistent with either positive or negative welfare effects from gentrification, depending on the aforementioned trade-off between rents, job market access, and amenities. Similarly, if we instead observed economically significant increases in neighborhood exit rates in response to gentrification, we would need to understand whether these estimates reflect insignificant moving frictions and dense choice sets or declines in incumbents' relative valuations of their origin neighborhoods (Vigdor, 2002). We now turn to estimating a dynamic model of residential and workplace choice to estimate the welfare effects from gentrification on incumbent low-income renter households.

1.4 A Dynamic Model of Neighborhood and Workplace Choice

To quantify the welfare impact of neighborhood change on low-income incumbent renters, we estimate a single-agent dynamic discrete neighborhood and workplace choice model (Bayer *et al.*, 2016; Diamond *et al.*, 2018; Davis *et al.*, 2021; Almagro and Domínguez-Iino, 2022; Davis *et al.*, 2023). Our single-agent framework considers the neighborhood and workplace choice problem that low-income renter households face each period, treating neighborhood and workplace characteristics as exogenous.[22] We use this framework to obtain parameter estimates of low-income renter households' preferences over neighborhood characteristics pertinent to understanding the welfare effects of gentrification.

Among papers in the dynamic discrete choice literature, our setup is most similar to the neighborhood demand model of Almagro and Domínguez-Iino (2022), which analyzes the endogenous formation of horizontally differentiated private consumption amenities in Amsterdam's 2010–2019 tourism boom. One important departure from their demand model is that we incorporate differences in within-CBSA access to employment opportunities

[22]It is common to treat neighborhood characteristics as exogenous in estimating households' preferences even when these neighborhood characteristics are functions of neighborhoods' socioeconomic composition (Bayer *et al.*, 2016; Davis *et al.*, 2021, 2023). We discuss the restrictions that this assumption imposes on our welfare calculations below.

through a first-step workplace choice problem. We combine our LEHD and LBD data to compute microfounded, time-varying, and neighborhood-level measures of job access for our sample population. This addition models an important determinant of households' location choices (e.g., Su (2022), Gu *et al.* (2021)). Moreover, coupled with our job market access instrument described in Section 1.6, these measures help facilitate the credible identification of households' preferences.

We use our estimated model parameters to quantify the welfare effects of gentrification for incumbent renters in Section 1.7. To do so, we use our parameter estimates to compute expected welfare separately for renter households living in each low-income urban census tract in the year 2000. We do this using the observed distribution of neighborhood rents and shares of college graduates from 2000 to 2019. We then compare these measures to the expected welfare that the same renter households would have obtained if the economy were instead in steady steady in the year 2000. This exercise yields census tract–level estimates of the welfare effects of gentrification for incumbent renter households. We finally conduct counterfactual experiments to unpack these welfare effects

1.4.1 Households and Timing of Choices

In each period, t, household heads, i, must decide which neighborhood in the city, c, they should live in.[23] In addition to choosing their residential neighborhood, $n_{i,t}$, household heads must decide which neighborhood to work in, $m_{i,t}$. Their workplace choice maximizes commute time–discounted period income and is decided after neighborhood residence is known. Longer commute times reduce the time that household heads spend working, effectively discounting the wage offered in workplace m. Finally, conditional upon deciding where to live and work, households must decide how much to spend on housing given the neighborhood-wide and period-specific rental rate.

Households differ ex–ante by the household head's race, which we denote by $k \in \{\text{Black, Non-Black}\}$. Households can further differ in their previous neighborhood resi-

[23] Households take their city as given.

dences, which inform both their current neighborhood residence, $n_{i,t-1}$, and how long they have lived there as of period $t-1$, $\tau_{i,t-1}$. We collect these observable household-level state variables in $x_{i,t} \equiv (n_{i,t-1}, \tau_{i,t-1})$.

Just before households make their workplace and neighborhood choices in each period, they receive idiosyncratic productivity and preference shocks, respectively. These preference shocks are unobserved by the econometrician but rationalize observed variation in households' choices within types conditional on $n_{i,t-1}$ and $\tau_{i,t-1}$. The remainder of this section presents the household head's problem, starting with her workplace choice problem.

1.4.2 Workplace Choice

Upon making their residential neighborhood and housing consumption choices in period t, households receive two independent productivity shocks. The first productivity shock is denoted by b_t^{kc} and is common across all type-k households in city c. The second productivity shock is household- and workplace tract–specific. We denote this second productivity shock by $z_{m,t}^i$, where m denotes the workplace tract. Conditional on living in neighborhood n, households choose their work location to maximize their commute time–discounted income:

$$I_{n,t}^k \equiv \; b_t^{kc} \; \cdot \; \max_m \frac{z_{m,t}^i}{d_{n,m}} w_{m,t},$$

where $w_{m,t}$ is the wage offered in workplace tract m and period t measured in efficiency units. $d_{n,m} > 0$ is the time that it takes to commute between neighborhood n and m. We assume that households spend a fixed amount of time each day working or commuting, so $d_{n,m}$ effectively discounts the total wage offered in m: $z_{m,t}^i \cdot w_{m,t}$. We assume that $z_{m,t}^i$ is drawn independently from a Frechet distribution with shape parameter $\epsilon^c \; \forall \, i \in c$. These shape parameters are specific to each city, which we make explicit with the superscript c. We further assume the Frechet shocks are independent across years, implying no cost to switching jobs. The expected income for a type-k household living in tract n at time t is therefore

$$\bar{I}_{n,t}^k = \Gamma\left(1 - \frac{1}{\epsilon^c}\right) \cdot b_t^{kc} \; \cdot \; RMA_n^{1/\epsilon^c}$$

where $\Gamma(\cdot)$ is the gamma function and $RMA_n \equiv \sum_{m \in \mathcal{N}^c} \left(\frac{w_m}{d_{n,m}}\right)^{\epsilon^c}$ is a summary measure of access to employment, which we follow the literature in terming residential market access. We derive these equations and describe how we construct their empirical analogs in Appendix A.4.3.

1.4.3 Neighborhood Choice

Households' Neighborhood Choice Problem Households choose their residential locations to maximize the sum of their expected discounted utilities,

$$\max_{\{n \in \mathcal{N}^c\}_t^\infty} \quad \mathbb{E}\left[\sum_{t'=t}^{\infty} \delta^{t'-t} \cdot u_n^k \left(s_{i,t'}^c\right) \,\middle|\, \mathcal{I}_{i,t}\right] \tag{1.4}$$

where δ is a known discount factor and $s_{i,t'}^c$ is a vector of state variables that determine household i's flow utility u_n^k from choosing neighborhood n. $s_{i,t'}^c$ includes the measures of expected income, $\bar{I}_{n,t}^k$, derived in Section 1.4.2. $\mathbb{E}\left[\cdot|\mathcal{I}_{i,t}\right]$ denotes the expectation operator conditioned on household i's information set at time t. $\mathcal{N}^c \equiv \{OO^c, 1^c, \ldots, N^c\}$ is the city-specific choice set, where OO^c denotes the outside option of leaving the city entirely. In each period, households observe the state variables s_{it}^c before choosing their residential location. Flow utilities are then realized, and states evolve. Each household's information set $\mathcal{I}_{it}$ therefore includes all current and past state variables that households may use to form expectations over their future evolution. We specify households' belief formation in Section 1.5.1.

State Variables Households' flow utilities depend on the vector of state variables $s_{it}^c \equiv \left(x_{it}, \varepsilon_{int}, \omega_t^c, \xi_t^{kc}\right)$, where $(x_{it}, \varepsilon_{int})$ are household-level observable and unobservable state variables, respectively. By contrast, $\left(\omega_t^c, \xi_t^{kc}\right)$ are city-specific observable and unobservable state variables, respectively. The household-level observable state variables, x_{it}, are comprised of households' residential tenure and neighborhood choice in the previous period, $x_{it} = (n_{it-1}, \tau_{it-1})$. Household i's residential choices determine the evolution of these observable household-level state variables. ε_{int} is the household's unobservable state, which we assume is i.i.d. across households, neighborhoods, and time. We conceptualize ε_{int} as an unobserved-to-the-econometrician time-varying household and neighborhood-specific preference shock. As is common, we assume that ε_{int} is distributed according to a type I extreme value distribution.

ω_t^c denotes observable city-specific state variables. The collection of city-specific state variables includes vectors for each neighborhood's housing costs, r_{nt}^k, the share of college graduates in the neighborhood, $\frac{Coll_{nt}}{Pop_{nt}}$, each neighborhood's commute time–discounted expected income, $\bar{I}_{nt}^k$, and an index for the period t[24]:

$$\omega_t^c = \left(\{r_{nt}\}_{n \in \mathcal{N}^c}, \left\{ \frac{Coll_{nt}}{Pop_{nt}} \right\}_{n \in \mathcal{N}^c}, \left\{ \bar{I}_{nt}^k \right\}_{n \in \mathcal{N}^c}, t \right)$$

Finally, ξ_t^{kc} is a city-specific vector of unobservable time-varying neighborhood-level amenity valuations among type k households. For example, ξ_t^{kc} could include time-varying valuations among type-k residents for suburban life, independent of gentrification. To facilitate exposition, we define $\bar{\omega}_t^{kc} \equiv \left(\omega_t^c, \xi_t^{kc}\right)$ as the vector containing both observable and unobservable city-specific state variables.

Flow Utility Preferences over neighborhood characteristics net of moving costs for a type-k household can be represented by

$$A_{n,t}^k \, Q_{n,t}^k \, \tau_{it}^{\beta_\tau^k} \, \exp(\varepsilon_{int})$$

where $A_{n,t}^k$ is a type-k's valuation of amenities in neighborhood n and $Q_{n,t}^k$ is a consumption composite that is Cobb–Douglas over nonhousing consumption, $C_{n,t}^k$, and housing consumption, $H_{n,t}^k$:

$$Q_{n,t}^k \equiv \left(C_{n,t}^k\right)^{\beta_C^k} \left(H_{n,t}^k\right)^{1-\beta_C^k}$$

Households' expected period- and neighborhood-specific budget constraint is given by

$$C_{n,t}^k \geq \bar{I}_{n,t}^k - r_{n,t} \cdot H_{n,t}^k$$

[24]We include a time index in the set of observable neighborhood-level state variables to explicitly incorporate nonstationarity so that the remaining observed state variables' evolutions can depend on time.

Type-specific neighborhood amenities are

$$A_{n,t}^{k} \equiv \left(\frac{Coll_{nt}}{Pop_{nt}}\right)^{\beta_A^k} \exp\left(\xi_{nt}^{k}\right)$$

We can decompose unobserved neighborhood- and period-specific amenities, ξ_{nt}^{k}, into time-invariant neighborhood-specific components, time-varying city-level components, and neighborhood-specific time-varying components:

$$\xi_{nt}^{k} \equiv \alpha_n^k + \alpha_t^{kc} + \tilde{\xi}_{nt}^{k}$$

Taking logs, incorporating moving costs ($MC_t^k(n_t, n_{it-1})$ defined below), solving for expected optimal housing consumption, and substituting in the amenity specification yields the following expected flow utility specification for a type-k household choosing neighborhood n with state s_{it}^{c} that is consistent with households' preferences over neighborhood characteristics[25]:

$$\begin{aligned} u_n^k\left(s_{it}^c\right) = \alpha_n^k + \alpha_t^{kc} + \beta_w^k \ln\left(\bar{I}_{n,t}\right) - \beta_r^k \log\left(r_{n,t}\right) + \beta_A^k \ln\left(\frac{Coll_{nt}}{Pop_{nt}}\right) \\ + \beta_\tau^k \ln(\tau_{it}) - MC_t^k(n_t, n_{it-1}) + \tilde{\xi}_{nt}^{k} + \varepsilon_{int} \end{aligned}$$

Moving Costs If a type-k household decides to leave its current neighborhood for another neighborhood in the same city, it incurs a nonmonetary moving cost, $MC_t^k(n_{it}, n_{it-1})$. This nonmonetary moving cost comprises a fixed disutility from moving, the physical straight-line distance between the household's origin and destination neighborhoods, and the *social distance* (defined below) between these two neighborhoods. Conversely, if a type-k household decides to leave its city entirely, it incurs a single, city-specific fixed cost. Specifically,

$$MC_t^k(n_t, n_{t-1}) = \begin{cases} 0 & \text{if } n_t = n_{t-1} \\ MC^k + \boldsymbol{\beta}_d^{k\prime} \boldsymbol{d}(n_t, n_{t-1}) + \boldsymbol{\beta}_s^{k\prime} \boldsymbol{s}(n_t, n_{t-1}) & \text{if } n_t \neq n_{t-1} \text{ and } n_t, n_{t-1} \neq OO^c \\ MC^{kc} & \text{if } n_{it} \neq n_{t-1} \text{ and } n_t \text{ or } n_{t-1} = OO^c \end{cases}$$

[25]Flow utility is expected in that it is the value that households expect to obtain before they realize their period-specific workplace-tract productivity shocks. It is this flow utility specification that is relevant for households' neighborhood choice and thus for our structural estimation.

where MC^k and MC^{kc} are the fixed intensive- and extensive-margin moving costs. $\boldsymbol{d}(n_t, n_{t-1})$ is a vector describing the physical distance between n_t, n_{t-1} and $\boldsymbol{s}(n_t, n_{t-1})$ is a vector describing the social distance between n_t, n_{t-1} in period t:

$$\boldsymbol{d}(n_t, n_{t-1}) \equiv \begin{bmatrix} |Dist(n_t, n_{t-1})| \\ |Dist(n_t, n_{t-1})|^2 \end{bmatrix} \quad \boldsymbol{s}_t(n_t, n_{t-1}) \equiv \begin{bmatrix} |(S(n_t) - S(n_{t-1}))/(S(n_t) + S(n_{t-1})| \\ |(S(n_t) - S(n_{t-1}))/((S(n_t) + S(n_{t-1}))|^2 \end{bmatrix}$$

where $Dist(n_t, n_{t-1})$ is the straight-line distance between the centroids of neighborhood n_t and n_{t-1} and $S(n_t)$ is the share of college graduates in neighborhood n at time t.[26] Our measure of social distance captures that, while low-income renter households may value residing in neighborhoods with a high share of college graduates, it may be costly to assimilate to neighborhood environments different from one's own (Gans, 1982; Jargowsky, 2009). Indeed, recent experimental research suggests that low-income households' moving costs are poorly approximated by the physical distance of residents' potential moves but strongly predicted by differences in the sociodemographic composition of households' origin and potential destination neighborhoods (Bergman *et al.*, 2023).

Value Functions, Choice Probabilities, and Expectational Errors We denote $V^k(s^c_{it})$ as the value function of the dynamic programming problem associated with equation 1.4. By Bellman's principle of optimality,[27]

$$V^k\left(s^c_{it}\right) = \max_{n \in \{OO^c, 1^c, \ldots, N^c\}} \left\{ \mathbb{E}_{x'|x,n}\left[u^k_n\left(s^c_{it}\right)\right] + \delta \mathbb{E}_t\left[V^k\left(s^c_{it+1}\right)\middle| n, s^c_{it}\right]\right\}$$

We define household i's ex-ante continuation value function as the expectation of the value function with respect to ε_{int}:

$$\bar{V}^k\left(x_{it}, \bar{\omega}^{kc}_t\right) \equiv \int V^k\left(s^c_{it}\right) dF^{\varepsilon}(\varepsilon_{int}) \tag{1.5}$$

[26] We normalize our social distance measure by the sum of shares so that the percentage changes are invariant to the direction of the residential move.

[27] The expectation operator $\mathbb{E}_{x'|x,n}[\cdot]$ is with respect to the future value of households' observed household-level state variables, x', conditional on households' current state and their neighborhood choice. While the current deterministic setup renders this operator redundant, we include it here to be consistent with our empirical application that models the evolution of households' residential tenure stochastically conditional on their neighborhood choice.

and define household i's conditional value function as

$$
\begin{aligned}
v_n^k\left(x_{it},\bar{\omega}_t^{kc}\right) &\equiv \mathbb{E}_{x'|x,n}\left[u_n^k\left(s_{it}^c\right)\right] - \varepsilon_{int} + \delta\mathbb{E}_t\left[\bar{V}^k\left(x_{it+1},\bar{\omega}_{t+1}^{kc}\right)|n,x_{it},\bar{\omega}_t^{kc}\right] \\
&\equiv \bar{u}_n^k\left(x_{it},\bar{\omega}_t^{kc}\right) + \delta\mathbb{E}_t\left[\bar{V}^k\left(x_{it+1},\bar{\omega}_{t+1}^{kc}\right)|n,x_{it},\bar{\omega}_t^{kc}\right]
\end{aligned} \tag{1.6}
$$

Then, given our assumption that ε_{int} are distributed i.i.d type I extreme value, the probability that a type-k household with state variables $\left(x_{it},\bar{\omega}_t^{kc}\right)$ chooses neighborhood n in period t is given by

$$
p_n^k\left(x_{it},\bar{\omega}_t^{kc}\right) = \frac{\exp\left(v_n^k\left(x_{it},\bar{\omega}_t^{kc}\right)\right)}{\sum\limits_{n'\in\mathcal{N}^c}\exp\left(v_{n'}^k\left(x_{it},\bar{\omega}_t^{kc}\right)\right)}, \tag{1.7}
$$

and the ex ante value function in 1.5 has the value

$$
\bar{V}^k\left(x_{it},\bar{\omega}_t^{kc}\right) = \ln\left(\sum_{n\in\mathcal{N}^c}\exp\left(v_n^k\left(x_{it},\bar{\omega}_t^{kc}\right)\right)\right) + \gamma
$$

where γ is Euler's constant. Combining these two expressions yields the following well-known result, which is critical to deriving our estimating equations (Hotz and Miller, 1993):

$$
\bar{V}^k\left(x_{it},\bar{\omega}_t^{kc}\right) = v_n^k\left(x_{it},\bar{\omega}_t^{kc}\right) - \ln\left(p_n^k\left(x_{it},\bar{\omega}_t^{kc}\right)\right) + \gamma \tag{1.8}
$$

Another expression critical for deriving our estimating equations is the difference between households' expected ex-ante continuation values and their realized counterparts:

$$
e^{\bar{V}}(x',\bar{\omega}_t^{kc},\bar{\omega}_{t+1}^{kc}) \equiv \underbrace{\bar{V}(x',\bar{\omega}_{t+1}^{kc})}_{\text{realized}} - \mathbb{E}_{\bar{\omega}'|\bar{\omega}_t^{kc}}\left[\bar{V}(x',\bar{\omega}')|\bar{\omega}_t^{kc}\right] \tag{1.9}
$$

We follow Kalouptsidi *et al.* (2021) and term the differences *expectational errors*. These expectational errors allow us to discard households' actual expectations in estimation. Solving for households' expectations would be prohibitively costly given the high-dimensional nature of a household's state space (some urban cores have over a thousand 2010-delineated census tracts).

Now that the function dependencies are clear, going forward, we suppress their arguments and remove the city superscripts unless we require them for explicative purposes: $\bar{V}_{xnt}^k \equiv \bar{V}^k\left(x_{it},\bar{\omega}_t^{kc}\right)$, $V_{int}^k \equiv V^k(s_{it}^c)$, $\bar{u}_{xnt}^k \equiv \bar{u}_n^k\left(x_{it},\bar{\omega}_t^{kc}\right)$, $v_{xnt}^k \equiv v_n^k\left(x_{it},\bar{\omega}_t^{kc}\right)$, and

$p^k_{xnt} \equiv p^k_n \left(x_{it}, \bar{\omega}^{kc}_t\right)$

1.5 Structural Estimation

We estimate our neighborhood demand parameters with what is termed in the dynamic discrete choice literature an Euler equations in conditional choice probabilities (ECCP) estimator.[28] Similarly to other conditional choice probability estimators, ECCP estimation involves two steps. In the first step, we estimate households' conditional choice probabilities, the soon-to-be-introduced transition distributions for the household-level state variables, and households' variable moving costs, $\boldsymbol{\beta}^{k\prime}_d$ and $\boldsymbol{\beta}^{k\prime}_s$. We then estimate the remaining model parameters in a second step, conditional on our first-step estimates. We estimate the second step using moment restrictions implied by the dynamic optimization of households, which we derive and explain in Section 1.5.2. With flow utilities linear in the model's parameters, we can evaluate these moment restrictions in a standard 2SLS framework.[29]

The ECCP estimator has many advantages in our setting. First, the ECCP estimator is computationally light. Since our analysis covers the residential history of low-income households for ten years in 50 large US metropolitan areas at the census tract level, traditional dynamic discrete choice estimation procedures that explicitly solve for households' value functions are infeasible (e.g., Rust (1987)). Second, our focus on gentrification implies an inherently nonstationary environment, making modeling the evolution of neighborhood change conceptually and computationally challenging. As we demonstrate when deriving our moment restrictions, ECCP estimation requires neither the complete specification of households' information sets nor a description of how the city-specific state variables evolve. Third, by yielding moment conditions we can evaluate in a 2SLS framework, we

[28]ECCP estimators are so named given their likeness to Euler equations in models with continuous choice variables (Aguirregabiria and Magesan, 2013; Kalouptsidi *et al.*, 2021).

[29]ECCP estimation has been used in a variety of applied settings, from choices over agricultural land use (Scott, 2013) and new technology adoption (Groote and Verboven, 2019) to occupational choice (Traiberman, 2019; Gendron-Carrier, 2023) and, most relevant to our setting, residential neighborhoods (Diamond *et al.*, 2018; Davis *et al.*, 2021; Almagro and Domínguez-Iino, 2022). See Kalouptsidi *et al.* (2021) for a comprehensive econometric treatment of linear regression techniques with ECCP estimators.

can relate our instrumental variables (detailed in Section 1.6) to the recent literature on quasi-experimental shift-share instruments (Goldsmith-Pinkham *et al.*, 2020; Borusyak *et al.*, 2022).

1.5.1 Estimation Assumptions

To identify our neighborhood demand parameters, we must make the following set of assumptions:

1. *State Transitions*: The state variables s^c_{it} evolve according to a controlled first-order Markov process with a transition distribution that factors as[30]

$$f\left(s^c_{it+1}|n_{it}, s^c_{it}\right) = f^x(x_{it+1}|n_{it}, x_{it}) \cdot f^{\bar{\omega}}\left(\bar{\omega}^{kc}_{t+1}|\bar{\omega}^{kc}_t\right) \cdot f^{\varepsilon}(\varepsilon_{int+1})$$

2. *Utility Normalization*: We normalize the utility offered by the outside option in every city to α^{ck} for each period:

$$\alpha^k_{OO^c} + \beta^k_w \ln(\bar{I}_{OO^c,t}) - \beta^k_r \ln(r^k_{OO^c,t}) + \beta^k_A \ln\left(\frac{Coll_{nt}}{Pop_{nt}}\right) + \tilde{\xi}^k_{OO^c t} = \alpha^{ck} \quad \forall\, t$$

3. *Rational Expectations*: Households' expectations over the evolution of the CBSA-level state variables conditional on their information set $\mathcal{I}_{it}$ correspond to the conditional expectations of the true data generating process given $\mathcal{I}_{it}$:

$$\mathbb{E}\left[e^{\bar{V}}(x', \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t+1})|\mathcal{I}_{it}\right] = 0$$

   where $e^{\bar{V}}(x', \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t+1})$ are the expectational errors defined in equation 1.9.

An important implication of assumption 1 is that the market-level state variables $\bar{\omega}^{kc}_t$ are *perceived* as exogenous by individual households; a household cannot expect to individually

[30]The evolution of the individual-level state variables, x_{it}, are "controlled" in that their evolution is influenced by the household's choices. While the current setup ensures that the household's neighborhood choice fully determines the evolution of x_{it}, our empirical implementation assumes that τ_{it} evolves stochastically conditional on the household's choice to reduce the dimensionality of our problem. Whether τ_{it} evolves stochastically or not, our discussion surrounding households' choice sets and moving costs in Section 1.4.3 should make clear that x_{it+1} evolves independently of $\bar{\omega}^{kc}_t$ conditional on n_{it} and x_{it}.

affect the evolution of $\bar{\omega}_t^{kc}$ with its own neighborhood choice (cf. Assumption (1) in Kalouptsidi *et al.* (2021)). Given that the typical 2010-delineated US census tract contains around 4,000 residents, we believe that this assumption is plausible.[31] Note that Assumption 1 does not require the observed and unobserved city-specific state variables to evolve independently. We highlight this to foreshadow the econometric challenge we face when attempting to identify preferences over functions of ω_t^c.

Assumption 2 says that residents who choose to reside outside of their respective CBSA's urban core obtain a time-invariant and CBSA-specific mean utility.[32] This assumption normalizes each CBSA's neighborhood mean utilities to a constant and time-invariant level, which is necessary to compare welfare across households within CBSAs given that logit models identify only differences in mean utilities. The assumption facilitates exposition and, because each α^{ck} is unobserved, highlights the incommensurability of expected welfare across different k-types and CBSAs.

Last, Assumption 3 says that, on average, households correctly anticipate the evolution of $\bar{\omega}_t^{kc}$. An important corollary of Assumption 3 is that the contents of households' information sets in time t are mean independent of their expectational errors at time t as well (cf. Lemma 1 in Kalouptsidi *et al.* (2021)). The importance of this corollary will become clear in Section 1.5.3 when we discuss our choosing among our set of IVs to estimate households' preferences.

1.5.2 Deriving Our Estimating Equations

Our goal now is to take our model setup and show how one can derive estimating equations that are linear in households' demand parameters. With such estimating equations, we can estimate our model with standard linear regression techniques. To derive these equations,

[31]While households' dynamic optimization implies our estimating equations, they are not informed by any equilibrium conditions, allowing us to remain agnostic over how households' individual choices influence the evolution of the CBSA-wide state variables.

[32]Recall that *all* residents of a given CBSA (i.e., including those outside the urban core) additionally receive a time-varying but neighborhood-invariant utility shock, α_t^k. The value of the outside option can, therefore, shift over time, albeit always in proportion to the mean utilities in the urban cores.

however, we must first introduce the concept of renewal actions.

Renewal Actions To derive our estimating equations, we use what are termed renewal actions in the dynamic discrete choice literature (Hotz and Miller, 1993; Arcidiacono and Miller, 2011). Renewal actions are actions that, when taken in period t, lead to the same distribution of states at the beginning of period $t+1$, regardless of the household's state in period t. In our setting, simply moving to a new neighborhood is a renewal action; moving to a new neighborhood resets a household's residential tenure to 0 regardless of the household's origin neighborhood or its current residential tenure. Moreover, because the city-specific state variables, $\bar{\omega}_{t+1}^{kc}$, and unobserved idiosyncratic preference shocks, ε_{int}, are independent of the household's state in period t, *all* the remaining state variables reset to a common value upon moving to a new neighborhood.[33]

We exploit such renewal actions when deriving our estimating equation. To see how, consider the residential choices of two hypothetical type-k households between periods $t-1$ and $t+1$. Assume that, in period $t-1$, these households reside in the same neighborhood, n_{t-1}, but not in period t (i.e., at least one household chooses a new neighborhood in period t). Further, assume that in period $t+1$, both households move to the same neighborhood, n_{t+1}. Our estimation procedure involves relating the difference in the expected discounted utilities of the two households' neighborhood choices to the difference in the probability that households make these neighborhood choices. Critically, because moving to neighborhood n_{t+1} in period $t+1$ constitutes the same renewal action for both households, their state variables are reset to a common value, equalizing their continuation values. Differences in the expected discounted utilities associated with the two sets of neighborhood choices are thus a function only of households' flow utilities. Relating such differences in expected discounted utilities to households' choice probabilities in this way helps disentangle the observed variation in households' flow utilities from households' unobserved continuation

[33]Note that renewal actions depend on our construction of households' neighborhood tenure in Section 1.4.3. This construction assumes that the length of households' prior residential tenures does not impact the value of future residential tenure. While this is a strong assumption, it is necessary to keep the dimension of the household-level state space manageable.

values.

Moving forward, we term these consecutive residential location choices *residential paths* to ease exposition.

Our Estimating Equation Consider the set of residential paths that we just described for type-k households but with an additional requirement that one of the households chooses the outside option in period t:

1. In period $t-1$, both households reside in the same neighborhood, n_{t-1}.

2. In period t, one household chooses neighborhood n, while the other household chooses n'. n' also happens to be the outside option. While it must be the case that $n' \neq n$, it may be that $n_{t-1} = n$ or $n_{t-1} = n'$.

3. In period $t+1$, both households convene at $\tilde{n}$, where $\tilde{n} \neq n'$ and $\tilde{n} \neq n$.

Given these residential paths, we can derive equation 1.10, which is linear in households' preference parameters. It is this equation we bring to the data to identify households' preferences. We derive it by equating the difference in the expected discounted utilities associated with the two residential paths to the difference in the probability that households take these paths:

$$Y^k_{xnn'\tilde{n}t} = \tilde{\alpha}^k_n + \alpha^k_t + \beta^k_w \ln(\bar{I}_{OO^c,t}) - \beta^k_r \ln(r^k_{OO^c,t}) + \beta^k_A \ln\left(\frac{Coll_{nt}}{Pop_{nt}}\right) + \beta^k_\tau \tilde{\tau}_x - \widetilde{MC}^k_t + \nu^k_{xnn'\tilde{n}t} \tag{1.10}$$

where

$$\begin{aligned}
Y^k_{xnn'\tilde{n}t} &\equiv \ln\left(\frac{p^k_{xnt}}{p^k_{xn't}}\right) - \delta\left(\sum_{x'} \ln\left(p^k_{x\tilde{n}t}\right) f^x\left(x'|n, x_t, \bar{\omega}^{kc}_t\right) - \sum_{x'} \ln\left(p^k_{x\tilde{n}t}\right) f^x\left(x'|n', x_t, \bar{\omega}^{kc}_t\right)\right) \\
\tilde{\alpha}^k_n &\equiv \alpha^k_n - \alpha^{ck} \\
\tilde{\tau}_x &\equiv \sum_{x'} \ln(\tau_{xt}(x'))f^x\left(x'|n, x_t, \bar{\omega}^{kc}_t\right) - \sum_{x'} \ln(\tau_{xt}(x'))f^x\left(x'|n', x_t, \bar{\omega}^{kc}_t\right) \\
\widetilde{MC}^k_t &\equiv MC^k_t(n, n_{xt-1}) - MC^k_t(n', n_{xt-1}) - \delta\left(MC^k_t(\tilde{n}, n) - MC^k_t(\tilde{n}, n')\right) \\
\nu^k_{xnn'\tilde{n}t} &\equiv \tilde{e}(x_t, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t+1}) + \tilde{\xi}^k_{nt}
\end{aligned}$$

Since the full derivation of this estimating equation is becoming well known, we relegate it to Appendix A.4.1. With estimates of households' conditional choice probabilities, $\hat{p}^k$, and estimates of the household-level transition distributions, $\hat{f}^x$, we can estimate equation 1.10 using linear regression techniques. The following subsection details our two-step estimation procedure.

1.5.3 Two-Step Estimation Procedure

In the first step of the estimation procedure, we estimate i) transition probabilities for the household-level state variables, ii) households' conditional choice probabilities, and iii) households' variable moving cost parameters. With these estimates, we can estimate equation 1.10 using 2SLS.

Household Transition Distributions To keep the dimension of the household state space manageable, we follow the literature stemming from Rust (1987) and discretize our household-level residential tenure measure into two buckets.[34] Specifically, we aggregate tenure into two buckets:

$$\bar{\tau} = \begin{cases} 1 & \text{if } \tau \leq 4 \\ 2 & \text{otherwise} \end{cases}$$

We assume that this aggregated location tenure variable evolves stochastically according to the following distribution function[35]:

$$f^{\bar{\tau}}(\bar{\tau}_t = 1 | n_t, x(n_{t-1}, \bar{\tau}_{t-1}), \bar{\omega}^{kc}) = 1 \text{ if } n_t \neq n_{t-1}$$

$$f^{\bar{\tau}}(\bar{\tau}_t = 2 | n_t, x(n_{t-1}, \bar{\tau}_{t-1}), \bar{\omega}^{kc}) = \begin{cases} 1, & \text{if } n_t = n_{t-1} \text{ and } \bar{\tau}_{t-1} = 2 \\ g^k_{n_{t-1}}, & \text{if } n_t = n_{t-1},\ \bar{\tau}_{t-1} = 1, \text{ and } i \in k \end{cases}$$

[34]The remaining household-level state variable is the household's residential location in the previous year. This state variable evolves deterministically depending on the residential path under consideration. We, therefore, do not need to specify any transition probability for this component of x_t.

[35]Our notation here corresponds to the marginal transition distribution of τ_{it}, with n_{it} taken as given. This approach follows Almagro and Domínguez-Iino (2022).

where we estimate $g^k_{n_{t-1}}$ directly from the data:

$$\hat{g}^k_{n_{t-1}} = \frac{\sum\limits_{i \in n_t, k} \mathbb{1}\{\tau_{xt} = 5\}}{\sum\limits_{i \in n_{t-1}, k} \mathbb{1}\{\tau_{xt-1} \leq 4\}}$$

Given that our analysis is at the census tract level, estimated in this way, $\hat{g}^k_{n_{t-1}}$ is a sparse measure of $g^k_{n_{t-1}}$. In practice, we therefore take a weighted average of $\hat{g}^k_{n_{t-1}}$ across census tracts in each county and each year.

Conditional Choice Probabilities Researchers typically face a trade-off between sparsity and flexibility when estimating first-step conditional choice probabilities, $\hat{p}^k_{xnt}$. Our setting is no different. On the one hand, we may estimate $\hat{p}^k_{xnt}$ directly from the data by calculating the probability that a type-k household with state x_{it} moves to each neighborhood $n_t \in \mathcal{N}^c$.[36] While this approach does not impose any restrictions on the implied data-generating process, it leads to very sparse estimates of $\hat{p}^k_{xnt}$ in our setting given the number of census tracts in our largest CBSAs. On the other hand, we may impose some structure on the implied data-generating process to smooth $\hat{p}^k_{xnt}$. Given that our setting yields particularly sparse empirical choice probabilities, we choose this latter option.

We model the count of type-k households in neighborhood $n \in \mathcal{N}^c$ choosing neighborhood $n' \in \mathcal{N}^c$ between periods t and $t+1$ as being derived from a Poisson distribution.[37] We parameterize the mean of the Poisson distribution in a way that imposes minimal additional restrictions on the data-generating process implied by our dynamic model.[38] Specifically, we estimate the following flexible Poisson regression separately for each type-k household:

$$\log\left(\mathbb{E}\left[\left|n^{k\bar{\tau}}_{t-1} \rightarrow n'^{k\bar{\tau}}_t\right|\right]\right) = \gamma^k_{n't} + \mu_{\bar{\tau}} \cdot \mathbb{1}\{n' = n_{t-1})\} + \gamma^k_{n't} \cdot \lambda_{\bar{\tau}} \cdot \mathbb{1}\{n' = n_{t-1})\} - MC^k_t(n', n_{t-1}) \tag{1.11}$$

[36]We could similarly employ any nonparametric method to compute $\hat{p}^k_{xnt}$ directly from the data.

[37]We choose to model the data as a Poisson distribution because of its ability to account for sparse data and its computational efficiency (Correia *et al.*, 2020).

[38]Note that the independence of households' neighborhood moves in any one period implied by the Poisson distribution is embedded in Assumption 1; f^x for household i is independent of all other households' actions. Appendix A.4.2 shows how our Poisson regression specification imposes minimal additional restrictions on the neighborhood choice problem of households in our dynamic model.

where $\left|n_{t-1}^{k\bar{\tau}} \to n_t'^{k\bar{\tau}}\right|$ is the count of type-k households with aggregated tenure status $\bar{\tau}$ in neighborhood n that choose neighborhood n' between periods $t-1$ and t. $\gamma_{n't}^k$ is a fixed effect that captures the neighborhood- and period-specific component of utility associated with choosing neighborhood n' in period t, $\mu_{\bar{\tau}}$ is a fixed effect that captures the additional utility that residents obtain from staying in their origin neighborhood given their tenure status, $\bar{\tau}$, and $\lambda_{\bar{\tau}}$ is a parameter capturing how the additional value one obtains from staying in her origin neighborhood varies with neighborhood mean utilities. $\mathbb{1}\{n' = n_{t-1})\}$ is an indicator variable that equals 1 if the household stays in its origin tract, and $MC_t^k(n', n_{t-1})$ are the same moving costs described in Section 1.4.3. We use our estimates from the Poisson model to predict the probability a type-k household with aggregated tenure status $\bar{\tau}$ living in neighborhood n chooses neighborhood n' in each year: $\hat{p}_{xn't}^k$.[39] As expected, the predicted probabilities are strongly correlated with their empirical counterparts; the correlation coefficients are 0.951 and 0.984 for Black and non-Black households, respectively.

Equation 1.11 additionally identifies our variable moving cost parameters, $\boldsymbol{\beta}_d^{k\prime}$ and $\boldsymbol{\beta}_s^{k\prime}$. Since the cost of moving to neighborhood n differs for each type-k household depending on its origin neighborhood, we can separately identify the parameters governing variable moving costs from γ_{nt}^k, $\mu_{\bar{\tau}}$, and $\lambda_{\bar{\tau}}$. How the cost of moving varies with the physical distance of the move, $\boldsymbol{\beta}_d^{k\prime}$, is identified with variation in the distance that households move within their urban core, conditional on moving. Similar variation for social distance identifies the cost of moving to neighborhoods socially different to one's origin neighborhood, $\boldsymbol{\beta}_s^{k\prime}$.[40] We report our variable moving cost estimates with the rest of our parameter estimates in Table 1.4.

---

[39]Many census tracts lack type-k households or in-migrants in a given year. Consequently, we cannot compute their conditional choice probabilities and exclude them from type-k households' choice sets. These tracts tend to be suburban or affluent, and their removal is unlikely to affect our estimates.

[40]Note that, in this repeated cross-sectional framework, we are unable to separately identify the fixed cost of moving from the value of residential tenure. We must instead estimate households' fixed moving costs in a second step using equation 1.10. Note also that our assumption of a Poisson distribution in equation 1.11 does not impact the variable moving cost estimates we obtain here. This is because of the isomorphism between the score of the Poisson distribution and of the conditional logit (represented in equation 1.7) for these continuous variables, yielding identical maximum likelihood estimation (MLE) estimates (Guimarães *et al.*, 2003).

Step Two Given our estimated variable moving cost parameters, our estimated conditional choice probabilities, and our estimated transition probabilities from step 1, we may now construct the empirical analogue of equation 1.10:

$$\hat{Y}^k_{xnn'\tilde{n}t} = \tilde{\alpha}^k_n + \alpha^k_t + \beta^k_w \ln(\bar{I}_{n,t}) - \beta^k_r \ln(r_{n,t}) + \beta^k_A \ln\left(\frac{Coll_{nt}}{Pop_{nt}}\right) + \beta^k_\tau \hat{\tilde{\tau}}_{xt} - \widetilde{\widehat{MC^F}}^k_t + \nu^k_{xnn'\tilde{n}t} \quad (1.12)$$

where

$$\begin{aligned}
\hat{Y}^k_{xnn'\tilde{n}t} &\equiv \ln\left(\frac{\hat{p}^k_{xnt}}{\hat{p}^k_{xn't}}\right) - \delta\Bigg(\sum_{x'} \ln\left(\hat{p}^k_{x\tilde{n}t}\right) \hat{f}^x\left(x'|n, x_t, \bar{\omega}^{kc}_t\right) \\
&\qquad - \sum_{x'} \ln\left(\hat{p}^k_{x\tilde{n}t}\right) \hat{f}^x\left(x'|n', x_t, \bar{\omega}^{kc}_t\right)\Bigg) + \widetilde{\widehat{MC^V}}^k_t \\
\hat{\tilde{\tau}}_{xt} &\equiv \sum_{x'} \ln(\bar{\tau}_{xt}(x'))\hat{f}^x\left(x'|n, x_t, \bar{\omega}^{kc}_t\right) - \sum_{x'} \ln(\bar{\tau}_{xt}(x'))\hat{f}^x\left(x'|n', x_t, \bar{\omega}^{kc}_t\right) \\
\widetilde{\widehat{MC^v}}^k_t &\equiv \widehat{MC^v}^k_t(n, n_{t-1}) - \widehat{MC^v}^k_t(n', n_{t-1}) - \delta\left(\widehat{MC^v}^k_t(\tilde{n}, n) - \widehat{MC^v}^k_t(\tilde{n}, n')\right) \\
\nu^k_{xnn'\tilde{n}t} &\equiv \tilde{e}(x_t, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t+1}) + \tilde{\xi}^k_{nt}
\end{aligned}$$

and where $\hat{\cdot}$ represents estimates from the first step. $\widetilde{\widehat{MC^v}}^k_t$ is the difference in either the fixed, $v = F$, or variable, $v = V$, portion of moving costs. We detail the difference in expectational errors, $\tilde{e}(x_t, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t+1})$, in Appendix A.4.1.

To be precise about identification, it is worth unpacking the error term, $\nu^k_{xnn'\tilde{n}t}$, in equation 1.12. $\nu^k_{xnn'\tilde{n}t}$ is comprised of both unobserved neighborhood-specific amenities, $\tilde{\xi}^{kc}_{nt}$, and expectational errors, $\tilde{e}(x_t, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t+1})$. We consider each of these in turn starting with the unobserved neighborhood-specific amenities, $\tilde{\xi}^{kc}_{nt}$. Since we do not restrict the relationship between $\tilde{\xi}^{kc}_{nt}$ and the remaining time-varying observable neighborhood characteristics, ordinary least squares (OLS) estimates of 1.12 would be biased. Expected income, neighborhood-level housing costs, and the share of college graduates are invariably correlated with many unobserved neighborhood-level factors, such as proximity to natural amenities that we do not observe as econometricians.

To distinguish between preferences for observed versus unobserved neighborhood amenities, we start by differencing equation 1.12 using residential paths starting in 2017 (i.e., $t-1=2017$) and residential paths starting in 2010 (i.e., $t-1=2010$), obtaining,

$$\Delta \hat{Y}^k_{xnn'\tilde{n}} = \Delta \alpha^k_t + \beta^k_w \Delta \ln(\bar{I}_{n,t}) - \beta^k_r \Delta \ln(r_{n,t}) + \beta^k_A \Delta \ln\left(\frac{Coll_n}{Pop_n}\right) + \Delta \nu^k_{xnn'\tilde{n}} \tag{1.13}$$

where the Δs correspond to the difference in the associated variables between $t=2011$ and $t=2018$. Differencing equation 1.12 removes the time-invariant component of exogenous neighborhood amenities, $\tilde{\alpha}^k_n$, the measures of residential tenure, $\beta^k_\tau \tilde{\tau}_x$, and the time-invariant components of the moving costs variables, MC^k, MC^{kc}, and $\boldsymbol{\beta}^{k\prime}_d \boldsymbol{d}(n_t, n_{t-1})$.[41] Our main concern now is that *changes* in the observed components of households' flow utilities are correlated with *changes* in unobserved neighborhood amenities and household expectational errors. We must, therefore, construct neighborhood-level instruments, z_n, for our endogenous regressors that are orthogonal to both of these components:

$$\begin{aligned} 0 &= \mathbb{E}\left[z_n \, \Delta \nu^k_{xnn'\tilde{n}}\right] \\ &= \mathbb{E}\left[z_n \, \left(\Delta \xi^k_n + \Delta \, \tilde{e}(x, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t'})\right)\right] \\ &= \mathbb{E}\left[z_n \, \left(\Delta \xi^k_n + \Delta \left(e^{\bar{V}}(n, x, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t'}) - e^{\bar{V}}(n', x, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t'})\right)\right)\right] \end{aligned} \tag{1.14}$$

where $e^{\bar{V}}(n, x, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t'})$ is the difference between the realized type-k ex-ante continuation value and type-k households' expectations of these continuation values, integrated over the potential realizations of the household-level states:

$$e^{\bar{V}}(n, x, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t'}) \equiv \sum_{x'} \left(\bar{V}(x', \bar{\omega}^{kc}_{t'}) - \mathbb{E}_{\bar{\omega}'|\bar{\omega}^{kc}_t}\left[\bar{V}(x', \bar{\omega}')\middle|\bar{\omega}^{kc}_t\right]\right) f^x\left(x'|n, x_n \bar{\omega}^{kc}_t\right).$$

In addition to being orthogonal to changes in unobserved neighborhood amenities, equation 1.14 shows that our instruments must also be mean independent of changes in households' expectational errors—the second component of $\nu^k_{xnn'\tilde{n}t}$. Recall Assumption 3 states that households have rational expectations over the evolution of the model's state variables. A

[41] We estimate the remaining time-invariant parameters in a final stage. Specifically, we estimate equation 1.12 conditional on the estimates from our differenced regressions and our first-step multinomial choice model. We assume that the effect of residential tenure and moving costs on the likelihood of different residential paths is uncorrelated with unobserved neighborhood amenities.

corollary of this assumption is that the contents of households' information sets at time t are mean independent of their expectational errors (Kalouptsidi *et al.*, 2021). Conversely, elements of households' future information sets that cannot be predicted from their period-t information sets will be correlated with their expectational errors. For this reason, our instruments must not predict future values of $\bar{\omega}_t^{kc}$ in a way that cannot simultaneously be predicted from the information in households' period-t information sets. To see why, consider an instrument that shocks the neighborhood-n elements of $\omega_{t'}^{kc}$ for any $t \leq 2017$.[42] Assume that this shock is uncorrelated with changes in unobserved neighborhood amenities but cannot be predicted with the information in households' contemporaneous information sets, $\mathcal{I}_{it}$. If it is relevant, the instrument will be mechanically correlated with the realized values of households' time-t ex-ante continuation values, $\bar{V}(x', \bar{\omega}_{t'}^{kc})$, but uncorrelated with households' time-t expectations, $\mathbb{E}_{\bar{\omega}'|\bar{\omega}_t^{kc}}\left[\bar{V}(x', \bar{\omega}') \middle| \bar{\omega}_t^{kc}\right]$, violating the exclusion restriction embodied in 1.14. For this reason, the instruments we detail immediately below are designed to predict changes in the CBSA-level state variables through 2011–2018 using only variation that can be predicted from households' 2010 information sets.

1.6 Identification Strategy

We now present the IVs we use to identify low-income renters' preferences for our three endogenous variables, β_w^k, β_r^k, and β_A^k. Our first set of instruments predicts changes in neighborhood-level job market access for our target population of low-income renters, helping to identify their preference for job market access, β_w^k. Our second set of instruments is similar but predicts changes in neighborhood-level job market access for college graduates. By predicting neighborhood demand among college graduates, these instruments help identify preferences for endogenous amenities, β_A^k. Our third set of instruments interacts predicted changes in neighborhood-level job market access for college graduates with the

[42] If the instrument shocks elements of $\omega_{t'}^{kc}$ for $t' > 2017$, we must also consider how the instrument affects the difference in expectational errors over time. The current example shows that we must construct z_n using variation that can be predicted from households' 2010 information sets.

intensity of neighborhoods' ex-ante urban development.[43] Neighborhoods with an ex-ante high share of urban development tend to fall on the inelastic segment of local housing supply curves (Baum-Snow and Han, 2023). This positioning makes it more probable that demand shocks will increase rents. These interaction terms are thus helpful in identifying households' distaste for paying rent, β_r^k.

We find that, conditional on our controls, the job market access instruments predict rent changes and job market access for low-income workers well. However, they are less predictive of changes in neighborhoods' college shares. To increase first-stage power, we therefore include a fourth set of instruments. The fourth set of instruments uses the proximity to other neighborhoods' shares of college graduates to help predict neighborhood demand among college graduates. These instruments are particularly helpful in identifying households' preferences for endogenous amenities, β_A^k. We now discuss each instrument in turn.

1.6.1 Job Market Access IV

Neighborhoods differ in their access to employment opportunities. Neighborhoods located near establishments with a high demand for skilled labor will be attractive to college graduates because of shorter expected commute times among this group, all else equal. A similar argument holds for low-income households and neighborhoods near establishments employing these workers. Our job market access instruments predict changes in neighborhoods' desirability based on changes in their expected commute times to employment opportunities. We construct these instruments in two steps. First, we define industry- and neighborhood-specific measures of job market access in a baseline year separately for college graduates and low-income households. Second, we interact these baseline measures with national industry employment trends to predict changes in job market access that are plausibly uncorrelated with underlying trends in neighborhoods' exogenous amenities.

[43] Specifically, we interact changes in job market access for college graduates with the share of land in the census tract that is covered by urban development in 2011. We obtain these measures from Baum-Snow and Han (2023), who, in turn, construct them using data from the National Land Cover Database.

We base our job market access instruments on the instruments constructed in Baum-Snow *et al.* (2019) and Baum-Snow and Han (2023). We build on these instruments by using our employer–employee linked (LEHD) and business establishment (LBD) data to construct more precise and granular measures of job market access.[44] Moreover, because our business data are disaggregated at the 6-digit NAICS level, we identify our model parameters by appealing to the conditional exogeneity of national industry shocks (Borusyak *et al.*, 2022).[45] We discuss how we construct these instruments, their identifying assumptions, and threats to identification below.

Instrument Construction We formally define our neighborhood-level measures of job market access in terms of each neighborhood n's access to employment in industry d at time t as,

$$JMA_{ndt} = \sum_{m \in \mathcal{N}^c \setminus n} e^{-\eta^c \tau_{nm}} l_{mdt} \tag{1.15}$$

where l_{mdt} is the number of jobs in workplace tract m and 6-digit NAICS industry d at time t, $\tau_{n,m}$ is the travel time between tracts n and m, and η^c is a spatial decay parameter governing the importance of faraway jobs relative to closer jobs in determining a tract's employment access. We derive equation 1.15 from households' workplace choice problem in Appendix A.4.3.

We obtain measures of $\tau_{n,m}$ for college graduates using reported commute destinations and times from our ACS data. However, since the number of neighborhood–pairs in each CBSA is large relative to the number of college graduates surveyed in the ACS, we follow Baum-Snow *et al.* (2019) and estimate a simple forecasting model to predict $\tau_{n,m}$ for all neighborhood pairs in each CBSA. We obtain CBSA-specific measures of η^c by using our employer–employee linked data to estimate gravity equations derived from a workplace

[44] See Chow *et al.* (2021) and Giroud and Rauh (2019) for details on the LBD and Abowd *et al.* (2009) for details on the LEHD.

[45] This identification assumption allows business establishments to select into tracts based on their exposure to changes in unobserved neighborhood characteristics, $\Delta\xi^k_{n,t}$ (Borusyak *et al.*, 2022). Out identification assumption instead requires establishments in industries *experiencing positive nationwide employment shocks* did not concentrate near neighborhoods experiencing positive or negative changes in their unobserved characteristics.

choice model à la Tsivanidis (2022). We detail our forecasting model and gravity equations in Appendix A.4.3.[46]

We use these industry-specific measures of employment access to construct our job market access IVs:

$$\Delta\,\widetilde{JMA}_{n,t_0,t} = \sum_{d\in\mathcal{T}} \underbrace{\frac{JMA_{n,d,t_0}\theta^c_d}{\sum_{d'} JMA_{n,d',t_0}\theta^c_{d'}}}_{Share} \underbrace{\frac{L^c_{d,t} - L^c_{d,t_0}}{L^c_{d,t_0}}}_{Shift} \tag{1.16}$$

where $L^c_{d,t}$ is national employment in industry d less employment in industry d located in neighborhood n's CBSA. $\mathcal{T}$ denotes the set of tradable industries that we use to predict changes in job market access.[47] As we do not observe the educational level or race of workers in our establishment-level data, we scale our industry-level measures of job market access by the share of workers employed in each industry and in each state who have a college degree, θ^c_d. Remember that we construct analogous IVs for low-income renter households.

We select $t_0 = 2002$ and $t = 2007$ for all our job market access instruments. As our discussion around households' expectational errors in Section 1.5.3 highlights, our instruments may not use information outside households' information sets to shift the endogenous variables. If the instruments predict changes in the endogenous variables that households do not—on average—expect, then the instruments will be correlated with their expectational errors. By setting $t_0 = 2002$ and $t = 2007$, we ensure our instruments rely on information in households' information sets throughout our analysis period (2010-2019). Serial correlation in the endogenous variables ensures that our instruments remain relevant. Our choice of $t_0 = 2002$ is motivated by the US Census Bureau's Economic Census occurring on years ending in 2 and 7 (Chow *et al.*, 2021). The allocation of firm employment data across establishments is most accurate in these years, increasing the precision of our baseline

[46] The travel times and spatial decay parameters are defined separately for college graduates and low-income renters.

[47] We define our set of tradable industries using trade costs for 6-digit NAICS manufacturing and service industries, as estimated in Gervais and Jensen (2019). We label an industry as tradable if its estimated trade costs are in the bottom three quartiles of the manufacturing and service industries analyzed by Gervais and Jensen (2019). This threshold ensures sufficient first-stage power while excluding industries whose establishment locations are likely endogenous to the spatial sorting within cities, such as local retailers. Below, we discuss our choice to focus on tradable industries.

shares.[48] We maximize first-stage power for our job market access instrument by choosing $t = 2007$; local labor demand shocks induced by the great recession do not appear to influence households' within-CBSA location choices throughout 2011–2018.

Identifying Assumptions Recent studies on shift-share instruments show how the exclusion restriction (e.g., Equation 1.14) can hold with either conditionally exogenous shares (Goldsmith-Pinkham *et al.*, 2020) or with conditionally exogenous shifts (Borusyak *et al.*, 2022). The shares in equation 1.16 correspond to neighborhoods' baseline commute time–discounted exposure to employment in each industry. The shifts correspond to the national employment growth rate in each industry. As it is less plausible that establishments' baseline neighborhood locations are unrelated to underlying trends in nearby neighborhoods' exogenous amenities, we argue that identification comes from the conditional exogeneity of national industry employment shifts.

Borusyak *et al.* (2022) show that three conditions are together sufficient to ensure that our employment "shifts" are conditionally exogenous. First, establishments in industries with nationwide employment shocks (positive or negative) must not be concentrated near neighborhoods experiencing trends (positive or negative) in their unobserved exogenous amenities. Second, no small subset of industries may comprise a large portion of the baseline shares. Third, industries' national employment shifts must be mutually uncorrelated given trends in unobserved amenities and baseline shares.[49]

We consider several threats to identification. First, researchers have argued that, throughout our analysis period, there was a general trend toward suburbanization among low-income households irrespective of gentrification (Bartik and Mast, 2023). Suppose establishments concentrated near suburban neighborhoods were overrepresented in industries

[48] We considered using 1997 as our baseline year, but this would have required manually geocoding establishments' addresses since the LBD had not started reporting establishments' census tracts then.

[49] The first condition can be represented formally as $\mathbb{E}[g_d|\{\bar{\Delta\nu}_d\}_d, \{s_d\}_d] = \mu$, the second condition as $\mathbb{E}\left[\sum_d s_d^2\right] \to 0$, and the third condition as $\text{Cov}\left(g_d, g_{d'}|\{\bar{\Delta\nu}_d\}_d, \{s_d\}_d\right) = 0 \,\forall\, (d, d')$, where $s_d = \sum_n s_{n,d} = \sum_n \frac{JMA_{n,d,t_0}\theta_d^c}{\sum_{d'} JMA_{n,d',t_0}\theta_{d'}^c}$, $g_d = \frac{L_{d,t}^c - L_{d,t_0}^c}{L_{d,t_0}^c}$, and $\bar{\Delta\nu}_d = \frac{\sum_{nn'\tilde{n}} s_{n,d}\Delta\nu_{nn'\tilde{n}}}{s_d}$.

experiencing negative nationwide employment shocks. In that case, we may mistake a secular migration trend toward the suburbs as a distaste for market access. To account for this possibility, we residualize $\Delta\,\widetilde{JMA}_{n,t_0,t}$ on measures of proximity to the metro division's CBD. These measures include a quadratic in the physical distance between the CBD and neighborhood n's centroid, a quadratic in the population-weighted distance between the CBD and neighborhood n's centroid, and fixed effects for five equally sized concentric rings centered on the CBD. We partition the concentric rings using population-weighted distance. Together, these measures ensure that our job market access instruments induce variation in the endogenous variables among neighborhoods equidistant from each metro division's CBD.

A second threat to identification is that changing consumer preferences may jointly influence households' location choices *and* industry employment trends. For example, changing preferences for different types of nontradable services may simultaneously influence employment in those industries and households' within-CBSA residential location decisions. To address this concern, we construct $\Delta\,\widetilde{JMA}_{n,t_0,t}$ using only employment shifts in *tradable* 6-digit NAICS industries.[50] This ensures we are not using variation in job market access caused by households' residential location choices. We also construct our national employment shifters excluding employment in the CBSA for which we predict changes in job market access.

A third threat to identification relates to the secular decline in manufacturing (NAICS 31–33) employment throughout 2002–2007 (Autor *et al.*, 2013). Correlated employment shocks to 6-digit NAICS industries within the manufacturing sector threaten the consistency of our estimates as manufacturing establishments tend to be spatially concentrated (e.g., in suburban neighborhoods). To account for correlated employment shocks within the manufacturing sector, we residualize $\Delta\,\widetilde{JMA}_{n,t_0,t}$ on neighborhoods' baseline exposure to

[50] By restricting to tradable industries, the effective "shares" in $\Delta\,\widetilde{JMA}_{n,t_0,t}$ do not sum to 1 across industries and within neighborhoods. To control for the possibility that neighborhoods near high concentrations of tradable industry establishments are systematically exposed to increasing/decreasing exogenous neighborhood amenities, we also control for neighborhoods' exposure to the total share of tradable employment (Borusyak *et al.*, 2022).

manufacturing employment. Identification then requires that industry employment shifts *within* the manufacturing sector be mutually uncorrelated conditional on baseline shares and unobserved amenities, a significantly less stringent requirement (Borusyak *et al.*, 2022).

We now move on to detailing our final set of instruments. These final instruments are particularly helpful in identifying households' preferences for observed neighborhood amenities, β_A^k.

1.6.2 Proximity IV

Guerrieri *et al.* (2013) show that during a positive city-wide employment demand shock, among low-income neighborhoods, those closest to other high-income neighborhoods experience the greatest home price appreciation—a proxy for gentrification. Motivated by these findings, we construct a neighborhood-level and distance-weighted measure of proximity to other neighborhoods' share of college graduates. We then interact this measure with CBSA-wide Bartik labor demand shocks constructed from the initial CBSA-wide shares of college graduates in tradable 6-digit NAICS industries. Again, though our instrument construction is novel, we are not the first to use proximity to other high-income neighborhoods as an instrument for gentrification (Brummet and Reed, 2021; Glaeser *et al.*, 2023).

Instrument Construction We first form neighborhood-level distance-weighted measures of exposure to other neighborhoods' shares of college graduates in the same CBSA. We then interact these neighborhood-level measures with shocks to CBSA-wide tradable industry employment:

$$\Delta \widetilde{\text{Prox}}_{n,t_0,t',t} = \sum_{d \in \mathcal{T}} \underbrace{\sum_{m \in \mathcal{N}^c \setminus n} e^{-\rho \tau_{n,m}} \frac{Coll_{m,t_0}}{Pop_{m,t_0}}}_{Proximity} \cdot \underbrace{\frac{l_{d,t'} \cdot \theta_d^c}{\sum_{d'} l_{d',t'} \cdot \theta_{d'}^c}}_{Share} \underbrace{\frac{L_{d,t}^c - L_{d,t'}^c}{L_{d,t'}^c}}_{Shift}$$

where $\frac{Coll_{m,t_0}}{Pop_{m,t_0}}$ is the share of residents in neighborhood m who are college graduates and where $l_{d,t'} = \sum_n l_{d,n,t'}$. The variables $l_{d,n,t'}$, $L_{d,t}^c$ and the parameters $\theta_{d'}^c$ are defined as before.

ρ is a spatial decay parameter governing the importance of faraway neighborhood college graduate shares relative to closer neighborhood shares. As $\rho \to \infty$, only the neighborhoods immediately adjacent to neighborhood n matter for determining the value of $\Delta\,\widetilde{\text{Prox}}_{n,t_0,t}$. Conversely, as $\rho \to 0$, the instrument loses all relevance, as every neighborhood matters equally in determining $\Delta\,\widetilde{\text{Prox}}_{n,t_0,t}$, ensuring that $\Delta\,\widetilde{\text{Prox}}_{n,t_0,t}$ is constant within each CBSA. Since we do not have a good prior for ρ, we calibrate its value using k-fold cross-validation in a set of first-stage regressions, regressing $\text{Gent}_{n,t_0,t}$ on $\Delta\,\widetilde{\text{Prox}}_{n,t_0,t}$ and selecting the value of ρ that best predicts the changes in neighborhood share of college graduates.

In contrast to local labor demand, industry-wide employment shocks throughout the Great Recession predict future *CBSA-wide* labor demand. For this reason, we use employment shifts between $t' = 2007$ and $t = 2010$ to construct our Bartik shift share, but construct the proximity shares in $t_0 = 2002$ to ensure consistency with our job market access instrument.

Identification Identification here proceeds differently from that under our job market access instruments. Because our final estimating equations include CBSA-wide fixed effects, the identifying variation must come from within CBSAs, and so identification from the CBSA-wide shift shares alone is ruled out. Identification instead proceeds from the interaction between the "proximity" term and the Bartik shift share.

Neighborhoods' proximity to other neighborhoods with a high share of college graduates may be correlated with underlying trends in exogenous amenities. To account for this possibility, we take two steps. First, we use lagged shares of neighborhoods' proximity to other neighborhoods with a high share of college graduates. Second, we residualize $\Delta\,\widetilde{\text{Prox}}_{n,t_0,t',t}$ on the proximity terms themselves, so that identification comes solely from the *interaction* between the shift shares and the proximity terms. Identification thus proceeds analogously to that under a triple difference-in-difference estimator: we compare differences in gentrification between neighborhoods near and far from already gentrified tracts in CBSAs experiencing large labor demand shocks to differences in gentrification between neighborhoods near and far from already gentrified tracts in CBSAs not experiencing large labor demand shocks (Brummet and Reed, 2021). We, therefore, assume that differences in changes in unobserved amenities between neighborhoods near and far already gentrified tracts are uncorrelated with the size of CBSAs' labor demand shocks.

1.6.3 Moment Conditions

The final collection of instruments is:

$$z_{1,n} = \Delta\, \widetilde{JMA}^{Coll}_{n,02,07} \tag{1.17}$$

$$z_{2,n} = \Delta\, \widetilde{JMA}^{Low-Income}_{n,02,07} \tag{1.18}$$

$$z_{3,n} = \Delta\, \widetilde{JMA}^{Coll}_{n,02,07} \cdot \varphi_n, \tag{1.19}$$

$$z_{4,n} = \Delta\, \widetilde{\text{Prox}}_{n,02,07,10} \tag{1.20}$$

where φ_n are neighborhood-level measures of the share of land in the census tract covered by urban development in 2011. The superscripts *Coll* and *Low-Income* refer to the group for whom the instrument is constructed.[51] We estimate the moment restriction in equation 1.14 via 2SLS separately for each type-k household:

$$\mathbb{E}\left[z_n\, \Delta\nu^k_{xnn'\tilde{n}}\right] = 0$$

A corresponding observation in this moment restriction is a set of residential paths for a type-k household with individual state x. Each type-k household of city c has $2 \cdot (N^c + 1)$ initial states in period $t-1$, N^c possible neighborhood choices in that same period, and $N^c - 1$ possible neighborhood choices in period t, implying $\sum_c 2 \cdot (N^c + 1) \cdot N^c \cdot (N^c - 1)$ observations for for both Black and non-Black low-income renter households. Since the urban cores of our largest CBSAs contain approximately two thousand census tracts, the number of potential residential paths reaches the tens of billions. We restrict the set of

[51]We estimate the commute elasticities and the employment share parameters for Black and non-Black low-income households jointly when constructing $\Delta\, \widetilde{JMA}^{Low-Income}_{n,02,07}$. This ensures that the same variation identifies both type-k households' preferences.

residential paths we analyze to ease the computational burden. Specifically, we restrict the set of residential paths to those that start from neighborhoods with the highest shares of type-k low-income renters in their CBSA. The cutoff for "highest shares" varies across CBSAs. For CBSAs with over 500 census tracts in their urban core, we select the top ten percent of tracts in terms of their share of type-k low-income renters. For CBSAs with under 500 census tracts in their urban core, we select the 50 highest tracts in terms of their share of type-k low-income renters. This choice ensures that a minimum number of census tracts from each CBSA informs our estimates. We report our estimates in Table 1.4.

1.6.4 Parameter Estimates

The leftmost columns of Panel A in Table 1.4 report parameter estimates for households' valuations of neighborhood characteristics, while the same columns in Panel B report estimates on households' moving frictions. Since logit models identify only relative changes in welfare, the rightmost columns in Table 1.4 translate the estimates into households' willingness to pay measured in annual rents. The annual rent values for college share (β_A) and market access (β_w) reflect the extra annual rent payments that households would incur to reside in a neighborhood with a 10 percent higher share of college graduates or market access, respectively.[52]

The interpretation of the annual rent valuations differs across the moving friction parameter estimates. The annual rent values for physical move distance (β_d and β_d^{Sq}) represent the amount in annual rents that households would be willing to pay to move to a census tract one mile closer to their origin neighborhood, conditional on moving. The annual rent values for social distance (β_s and β_s^{Sq}) represent the amount in annual rents that households would be willing to pay to move to a census tract 10 percentage points more similar to their current census tract in terms of neighborhood college shares, again

[52] We use the average neighborhood rent levels for Black and non-Black households in Table 1.2 as the basis of the percentage change in rents.

conditional on moving.[53] The fixed moving cost annual rent valuations represent how much additional annual rent households would be willing to pay to avoid moving in any one year. The high residential tenure annual rent valuations represent how much the fixed cost of moving increases for households who have lived in the same census tract for at least five years. Because of institutional differences in CBSAs' housing regulations, the cost of moving neighborhoods and leaving one's home CBSA likely differs across CBSAs.[54] To account for the impact of these institutional differences on households' residential choices in a parsimonious way, we estimate CBSA-specific measures of neighborhood attachment, $\beta^k_\tau \hat{\bar{\tau}}$, and extensive-margin fixed moving costs, MC^c. In adherence with US Census Bureau disclosure guidelines, Table 1.4 reports the pseudo-median of these estimates.[55]

Our parameter estimates all have the expected sign, though Black households have a surprisingly low valuation for market access. One explanation for Black households' low valuation for market access is that the variation in market access induced by our job market access instruments is concentrated in industries primarily employing non-Black workers; such variation would not induce migration responses from households not employed in these industries, suggesting no preference for market access. Future iterations of the paper will test the robustness of our parameter estimates to skill- *and* race-specific measures of job market access.

Our parameter estimates reveal that it is moderately costly to move across neighborhoods. Low-income Black renter households' within-CBSA migration decisions suggest that the fixed cost of moving neighborhoods is over $3,000 and that this cost increases by approximately $600 for residents with a high amount of accumulated neighborhood

---

[53]We calculate the value of physical distance based on an initial move of 1 mile. We calculate the value of social distance assuming the average baseline share of college graduates for both Black and non-Black households, as reported in Table 1.2.

[54]In CBSAs with strong rent control/stabilization laws, the financial premium from staying in one's home residence increases with residential tenure. We would thus expect households in these CBSAs to display different levels of neighborhood attachment. Indeed, we see in our data that households' residential tenures are highest in CBSAs with such policies, like New York-Newark-Jersey City and San Francisco-Oakland-Fremont.

[55]The pseudo-median of these parameter estimates is the average of the median and the four CBSA-level estimates closest to this median value (two on either side of the median).

capital ($\bar{\tau} = 2$). Surprisingly, the social distance between neighborhoods plays little role in households' residential migration decisions.[56] One explanation for this result is that low-income households initially residing in gentrifying tracts move to tracts with a lower share of college graduates, conditional on moving. Our model would interpret these moves as revealing a preference for social distance. A different model specification may incorporate asymmetric moving costs, capturing the distinct cost of moving into a neighborhood with a *higher* share of college graduates than one's origin tract.

1.7 The Welfare Effects of Gentrification

In this section, we examine what our parameter estimates imply about the welfare effects of gentrification for incumbent renters in the year 2000. We compute our measures of expected welfare as follows. Consider a representative low-income renter in a gentrifying urban-core neighborhood, n, in 2000. To this renter, the distribution of neighborhood-level rents and shares of college graduates across her CBSA are exogenous; her migration decisions alone do not affect their distribution or evolution (cf. assumption 1). Given the exogenous distribution of these market-level state variables, $\bar{\omega}_t^k$, and her expectations over their evolution, we compute her expected welfare from the perspective of the year 2000.[57] Then, to approximate the welfare impact of gentrification, we compare this measure to the expected welfare that the same renter in neighborhood n would obtain in the year 2000 if the economy were in a steady state. If her expected welfare is lower in the steady-state equilibrium, the changing distribution of neighborhood characteristics across her

[56]While Black households appear to value some social distance conditional on moving, they have a distaste for moves that involve significant differences in neighborhood educational composition.

[57]For all these welfare calculations, we keep the level of exogenous amenities and the observed distribution of job market access at their 2000 levels. We hold the exogenous amenities fixed at their 2000 levels to help isolate the causal impact of gentrification on households' welfare. We hold the observed distribution of job market access at their 2000 levels because of our reduced-form findings in Section 1.3 showing that gentrification has negligible impacts on the employment outcomes of incumbent renter households. While job market access factors into households' location choices, gentrification does not appear to affect market access in a way that meaningfully impacts incumbent low-income renter households. This could result, for example, from gentrification skewing the composition of nontradable employment opportunities toward higher-skilled workers even within 6-digit NAICS industries. These possibilities suggest interesting questions for further research.

Table 1.4. *Parameter Estimates*

	Estimates		Value in Annual Rent	
	Black	Non-Black	Black	Non-Black
Panel A: Neighborhood Characteristics				
Rents (β_r)	-10.31***	-26.40***	–	–
	(1.494)	(3.341)		
College Share (β_A)	12.61***	7.343***	$1,224	$312
	(0.390)	(1.437)		
Market Access (β_w)	0.122	20.04***	$12	$851.4
	(1.635)	(3.487)		
Panel B: Moving Frictions				
Distance (β_d)	-0.2782***	-.3091***		
	(0.0004)	(0.0003)	-$270	-$130
Distance Squared (β_d^{Sq})	0.0039***	0.0044***		
	(0.0000)	(0.0000)		
Social Distance (β_s)	0.1335***	-0.1751***		
	(0.0108)	(0.0084)	$5	-$15
Social Distance Squared (β_s^{Sq})	-0.5598***	-0.4096***		
	(0.0074)	(0.0060)		
Within-Urban Core Fixed Moving Cost (MC)	-4.558***	-4.318***	-$3,578	-$1,692
	(0.000)	(0.0000)		
Outside-Urban Core Fixed Moving Cost (MC^c)	-4.458***	-4.621***	-$3,515	-$1,801
	(0.000)	(0.0000)		
High Residential Tenure (β_τ^k)	0.612***	0.551***	$612	$237
	(0.0002)	(0.0001)		
Controls				
Origin-Tract Fixed Effects	✓	✓		
Exposure to Tradable Industry JMA	✓	✓		
Exposure to Manufacturing JMA	✓	✓		
Proximity To CBD	✓	✓		
Fraction Tract Developed	✓	✓		
First Stage F-Stat	63.35	27.92		
N (1,000s)	2,573,000	10,630,000		

Notes: An observation is a set of residential paths described in Section A.4.1. We include fixed effects for each residential path's origin tract (i.e., n_{t-1}). Details on the remaining controls are in the main text. Standard errors are in parentheses. The number of residential paths in our analysis differs across Black and non-Black low-income renter households because we predict fewer conditional choice probabilities for Black households in the first step of the analysis, as discussed in footnote 39. *Sources*: ACS (2005–2021), LEHD (2010–2012, 2017–2019), CoreLogic (2006–2017), MAF-X (2019), and MAF-ARF (2010–2012, 2017–2019). Estimates were disclosed by the US Census Bureau's Disclosure Review Board. Project Number 2358. Disclosure Clearance Number CBDRB-FY24-P2358-R10957.

CBSA throughout 2000–2019 increased her welfare. The opposite is true if her expected welfare is higher in the steady-state equilibrium. We make these calculations separately for representative low-income renters in every low-income neighborhood in the year 2000.

In the steady-state equilibria used to construct our welfare estimates, the distribution of market-level state variables, $\bar{\omega}_t^k$, are fixed at their 2000 levels for every period t in the future. The value of exogenous neighborhood amenities, $\{\xi_n\}_n$, are further fixed at levels that induce a stationary distribution among low-income renter households.[58] We compare households' expected welfare in the year 2000 to the expected welfare they would have obtained if the economy was instead in steady-state in order to facilitate comparisons of welfare effects across neighborhoods in the same CBSA. Without comparing households' expected welfare to baseline steady-state measures, the initial *levels* of neighborhood characteristics would drive differences in our neighborhood-level measures of expected welfare within CBSAs.

We now describe our welfare calculations formally. For every low-income neighborhood across each of our 50 large CBSAs, we evaluate the following expression for each type-k household with tenure $\bar{\tau}$:

$$\Delta W^k(n,\bar{\tau}) = \frac{\bar{V}^k(x=(n,\bar{\tau}),\bar{\omega}_{2000}^{kc}) - V_{ss}^k(n,\bar{\tau})}{\beta_r^k} \quad \forall\, k,\, n \tag{1.21}$$

where $\bar{V}^k(x=(n,\bar{\tau}),\bar{\omega}_{2000}^{kc})$ is the expected welfare of a type-k incumbent renter living in neighborhood n with tenure status $\bar{\tau}$ in the year 2000. $V_{ss}^k(n,\bar{\tau})$ is the same calculation but computed under the assumption that the economy is in steady state. To compare type-k households, we normalize expected utility by households' marginal utility of log rent, β_r^k. $\Delta W^k(n,\bar{\tau})$ thus captures the impact of a CBSA's changing neighborhood-level shares of college graduates and rents starting in 2000 on incumbent renters initially living in neighborhood n, measured in log-rent units. We provide further details on our welfare calculations in Appendix A.4.4.

$\Delta W^k(n,\bar{\tau})$ differs across neighborhoods to the extent that moving frictions mediate

[58]A stationary distribution implies that the share of each type-k household of tenure level $\bar{\tau}$ in each neighborhood remains constant over time. We formally define a steady-state equilibrium and a stationary distribution in Appendix A.4.4.

the welfare impacts of gentrification. If it were costless to move across neighborhoods and households did not accumulate neighborhood-level capital, households' individual states, $x(n_{i,t-1}, \tau_{i,t-1})$, would become irrelevant, and differences in welfare effects across neighborhoods within CBSAs disappear. This happens because, in each period, households would choose to live in the neighborhood that offers the most ideal bundle of neighborhood characteristics regardless of their current location. By contrast, with high moving costs and neighborhood capital accumulation, incumbent renters are averse to leaving their current neighborhood irrespective of changes in its characteristics. In this sense, incumbent renters bear the incidence of changing neighborhood environments if they have substantial moving frictions.

Before we present our results, it is worth discussing two nuances surrounding our welfare analysis. First, when computing $\bar{V}^k(x = (n, \bar{\tau}), \bar{\omega}_t^{kc})$, we hold the distribution of exogenous neighborhood amenities fixed at their 2000 levels. As a result, the distribution of low-income renters across neighborhoods predicted by our model does not match its observed counterpart after 2000. While fixing exogenous neighborhood amenities at their 2000 levels helps isolate the welfare effects of gentrification, the discrepancy between the model's predicted choice probabilities and their observed counterparts highlights a limitation of our single-agent framework—namely, that we cannot use our framework to assess counterfactual *policies* such as rent control, as these policies will affect the distribution of $\bar{\omega}_t^{kc}$ in ways that we cannot predict. Therefore, all counterfactual analyses we run directly manipulate households' preferences or elements in $\bar{\omega}_t^{kc}$.

Second, we do not have a good measure of *absolute* changes in amenities during our analysis period. Recall that we model observed neighborhood amenities as a function of the share of college graduates in each neighborhood. The share of college graduates increased nationwide from 26.8 percent to 36.5 percent between 2000 and 2017.[59] While changes in the distribution of college graduates across neighborhoods within a city likely reflect changes in the relative distribution of neighborhood amenities (Su, 2022), it is unlikely that the absolute

[59]These are our calculations based on publicly available IPUMS census data (Manson *et al.*, 2022).

level of amenities increased proportionately to nationwide educational attainment. For this reason, our welfare analysis focuses on differences in welfare effects across renters living in different neighborhoods in the year 2000.

Minimal Variation in Welfare Effects Across Neighborhoods Table 1.5 reports results from regressions of $\Delta W^k(n, \bar{\tau})$ on our measure of gentrification defined in Section 1.3.[60] The purpose of these regressions is to understand the welfare effects of gentrification for incumbent renters living in different neighborhoods within the same CBSA. The unit of analysis in these regressions is a 2010-delineated census tract. The sample includes the poorest 50 percent of all census tracts among the 50 largest urban cores in the US. These regressions include measures of neighborhoods' proximity to their CBSA's CBD. Neighborhoods closer to the periphery of the urban core have an artificially lower level of access to other neighborhoods than tracts closer to the CBD.[61] Tracts closer to the periphery of the urban core also experience less gentrification. We design the CBD proximity controls to account for this correlation. The regressions also contain CBSA fixed effects and baseline neighborhood controls, including baseline rents and college shares. These controls ensure that our regression results compare the impact of gentrification across neighborhoods within the same CBSA and across neighborhoods with the same baseline levels of rents and college shares. Our results are not sensitive to excluding any combination of these controls.

The coefficients on gentrification in Table 1.5 reflect how gentrification impacts incumbent renters' welfare. The economically insignificant coefficients in Table 1.5 suggest that the economic impact of gentrification in one's current neighborhood is negligible. A 10-percentage-point increase in gentrification is associated with a 0.02638 (-0.00123) increase in log consumption units for Black (non-Black) renters with high neighborhood tenure (only

[60] Here, we define gentrification with the end period, t, set to 2006. We found that gentrification defined over 2000–2006 explained the most variation in outcomes across various exploratory analyses.

[61] We define urban cores in Section 1.2. There are neighborhoods outside each urban core (and, hence, inside each CBSA's outside option) that residents of neighborhoods on the periphery of their urban core could, in practice, easily move to. Since, in our model, it is equally costly to move to the outside option from anywhere in the urban core, our setup artificially lowers the level of access to neighboring tracts for incumbent residents in the urban periphery. Our CBD proximity controls attempt to account for these facts.

\$267.65 (-\$13.79) in total lifetime rent payments). Results are similar across households with high and low levels of accumulated neighborhood capital.

At least three explanations exist for the null results reported in Table 1.5. First, it may be that the benefits of improving amenities exactly offset the costs of higher rents as incumbent renters stay in their home neighborhoods. Under this scenario, gentrification is welfare neutral regardless of one's home census tract. Second, gentrification may sometimes benefit incumbent renters and sometimes harm them, depending on the characteristics of their baseline neighborhood environment. In this scenario, the results in Table 1.5 mask underlying heterogeneity. Third, incumbent renters may be sufficiently mobile to render changes in their current neighborhood relatively unimportant for expected utility. Here, gentrification is not welfare-neutral, but *where* incumbent renters initially live within their city is unimportant. We show that this third scenario drives the results in Table 1.5.

Table 1.5. *Welfare Effects of Gentrification on Incumbent Renters*

$\Delta W^k(n, \bar{\tau})$	(1)	(2)	(3)	(4)
Gentrification	0.157***	-0.0217**	0.2638***	-0.0123
	(0.008)	(0.0087)	(0.0100)	(0.0091)
	[1.521]	[1.065]	[1.490]	[1.032]
Controls				
CBD Proximity Controls	✓	✓	✓	✓
Baseline Neighborhood Controls	✓	✓	✓	✓
CBSA Fixed Effects	✓	✓	✓	✓
Sample Restrictions				
Race	Black	Non-Black	Black	Non-Black
Tenure Status	$\bar{\tau} = 1$	$\bar{\tau} = 1$	$\bar{\tau} = 2$	$\bar{\tau} = 2$
Number of Tracts	6,422	6,468	6,422	6,468

Notes: Table 1.5 reports results from regressions of $\Delta W^k(n, \bar{\tau})$ on $\text{Gent}_{n,2000,2006}$. Standard errors clustered at the CBSA level are in parentheses. We report the average change in log consumption across all neighborhoods in square brackets. We compute $\Delta W^k(n, \bar{\tau})$ for the poorest 50 percent of all 2010-delineated census tracts among the largest 50 CBSAs aside from Los Angeles–Anaheim–Riverside and New York-Newark-Jersey City, which we exclude for computational efficiency reasons. Future iterations of this paper will include these CBSAs in our welfare analysis. We do not expect our results to change qualitatively. *Sources*: Publicly available ACS (2005–2021) and Decennial Census (2000) (Manson *et al.*, 2022).

Changes in Neighborhood Choice Sets Govern Welfare Impacts We show here that the null results in Table 1.5 are not because gentrification is welfare neutral or because they mask underlying heterogeneity. Rather, we show that gentrification's effects operate primarily through changing the characteristics of *other* neighborhoods in renters' choice sets. We show this by conducting the following welfare decomposition. Consider a representative incumbent low-income renter residing in neighborhood n in 2000. Assume that neighborhood n's observable characteristics evolve as observed in the data before. However, also assume that all other neighborhoods' characteristics remain fixed at their 2000 levels. We then ask: Do the changing characteristics of neighborhood n benefit or harm the incumbent renter in neighborhood n? By holding the characteristics of other neighborhoods fixed, we decompose the role of changing neighborhood characteristics in one's home neighborhood and the changing characteristics in other neighborhoods across the CBSA.

We perform this welfare decomposition by constructing a modified measure of $\Delta W^k(n, \bar{\tau})$:

$$\Delta \widetilde{W}^k(n, \bar{\tau}) = \frac{\widetilde{V}^k(x = (n, \bar{\tau}), \bar{\omega}^{kc}_{2000}) - V^k_{ss}(n, \bar{\tau})}{\beta^k_r} \quad \forall\, k\, n \tag{1.22}$$

where $\widetilde{V}^k(x = (n, \bar{\tau}), \bar{\omega}^{kc}_{2000})$ is constructed precisely like $\bar{V}^k(x = (n, \bar{\tau}), \bar{\omega}^{kc}_{2000})$ except that the market-level state variables for $n' \neq n$ are fixed at their 2000 values for every year from 2000 onward.

Table 1.6 reports results from regressions of $\widetilde{V}^k(x = (n, \bar{\tau}), \bar{\omega}^{kc}_{2000})$ on our measure of gentrification. We observe economically and statistically significant positive effects of gentrification on incumbent renters' welfare. This is true for both Black and non-Black households (though the results are stronger for Black households), which have meaningfully different valuations of neighborhood characteristics, and for households with high and low levels of accumulated neighborhood capital. For example, a 10-percentage-point increase in gentrification increases Black (non-Black) renters' welfare by 1.0671 (.07910) log consumption units ($19,094.35 ($923.25) in total lifetime rent payments) when these renters have a high level of accumulated neighborhood capital (Table 1.6, columns (3) and (4)).

Table 1.6 shows that gentrification meaningfully affects incumbent renters' welfare when we hold other neighborhoods' characteristics constant. Conversely, as we let other neighborhoods' characteristics vary, where an incumbent renter initially lives has little impact on her expected welfare. Tables 1.5 and 1.6 together suggest that how CBSAs as a whole changed throughout 2000–2019 was more important for incumbent renters than how their neighborhood alone changed throughout this time period. The quality of low-income renter households' choice sets is far more important for this group's welfare than how their home neighborhoods changed throughout our analysis period.

Table 1.6. *Mediation of Welfare Impacts by Neighborhood Choice Sets*

$\Delta\widetilde{W}^k(n,\bar{\tau})$	(1)	(2)	(3)	(4)
Gentrification	11.9438***	0.7376***	10.6712***	.7910***
	(0.2425)	(0.1353)	(0.2059)	(0.1033)
	[0.9274]	[0.0316]	[0.7958]	[-0.0462]
Controls				
CBD Proximity Controls	✓	✓	✓	✓
Baseline Neighborhood Controls	✓	✓	✓	✓
CBSA Fixed Effects	✓	✓	✓	✓
Sample Restrictions				
Race	Black	Non-Black	Black	Non-Black
Tenure Status	$\bar{\tau}=1$	$\bar{\tau}=1$	$\bar{\tau}=2$	$\bar{\tau}=2$
Number of Tracts	6,422	6,468	6,422	6,468

Notes: Table 1.6 reports results from regressions of $\Delta\widetilde{W}^k(n,\bar{\tau})$ on $\text{Gent}_{n,2000,2006}$. The table is otherwise identical to Table 1.5.

Did Gentrification Benefit Incumbent Residents? The results reported in Table 1.5 suggest that gentrification benefited incumbent renters on average from the perspective of the year 2000. Despite negligible variation in welfare across tracts within CBSAs, the average tract-level increase in presented-discounted expected welfare over 2000–2019 for an incumbent Black (non-Black) resident with a low level of accumulated capital was \$35,814 (\$21,321) in total lifetime rent payments.[62] There are several reasons to interpret these

[62] These estimates are obtained based on the average increase in log consumption units reported in the square brackets in Table 1.5 and the average annual rents in Table 1.2.

average effect sizes with caution. First, as mentioned above, we do not have a good measure of *absolute* changes in amenities during our analysis period. It is unlikely that the secular increase in nationwide educational attainment led to proportional increases in absolute levels of neighborhood amenities. Second, our model assumes that households can continuously adjust their housing consumption in response to rising rents. A model that assumes unit housing demand and, therefore, income effects may yield different average welfare effects in response to gentrification. Third, our analysis assumes constant preferences within each set of type-k households. Our results show that with comparatively small moving costs, households' changing neighborhood choice sets govern the welfare effects of gentrification. However, for households with higher-than-average moving costs or a strong degree of neighborhood attachment, gentrification may affect incumbent residents through changes in their home neighborhood's characteristics. Future iterations of this paper will test the robustness of our results to these caveats.

1.8 Discussion

Beginning in the 1990s and intensifying after 2000, gentrification transformed the socioeconomic composition of vast areas within American inner cities. This paper shows how this transformation affected incumbent renters in these cities. We first show that gentrification did not meaningfully affect the employment outcomes of renters living in gentrifying neighborhoods nor the length of time that incumbent renters stayed in their home neighborhoods. We show that these results are robust to heterogeneity in neighborhoods' baseline environments, an a priori surprising finding. Gentrification did, however, impact the characteristics of the neighborhoods that incumbent renters lived in throughout our analysis period. This last finding suggests the possibility that gentrification meaningfully affected incumbent renters' welfare. We estimate a dynamic model of residential and workplace choice to test whether and how gentrification affected incumbent renters' welfare. We use our parameter estimates from this model to approximate the welfare effects of gentrification for low-income renters initially living in each low-income urban neighborhood in the US. We show using our framework that gentrification affected low-income incumbent renters primarily by changing the characteristics of *other* neighborhoods in renters' choice sets. We finally conclude with some caveats on interpreting the average level of our welfare estimates.

Chapter 2

Suburban Housing and Urban Affordability: Evidence from Residential Vacancy Chains[1,2]

2.1 Introduction

Over the past thirty years, real home prices and rents in dense urban centers in the U.S. grew more rapidly than housing costs in suburban metropolitan neighborhoods (Howard and Liebersohn, 2021).[3] At the same time, the vast majority of housing supply growth has been

[1]Chapter 2 is coauthored with Valentine Gilbert. Valentine thanks his advisors, Ed Glaeser, Stefanie Stantcheva, and Adrien Bilal, for their valuable feedback, patience, and support. All errors are our own. We thank participants of the Economics and Social Policy Workshop for helpful comments. We also thank Jim Davis, Sari Kerr, Shawn Ratcliffe, and Shital Sharma for help with the Census data. Oleg Kolev provided excellent research assistance. Valentine gratefully acknowledge's receiving financial support from the John R. Meyer Dissertation Fellowship of the Harvard Joint Center for Housing Studies while working on this research.

[2]Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2733. (CBDRB-FY24-P2733-R10951).

[3]Between 1990 and 2018, real rents in the 10% of tracts closest to the city center increasing by an average of 45%, substantially more than the average increase in real metropolitan rents of 30%. We make these and subsequent aggregate-level calculations using National Historical Geographic Information System (NHGIS) data (Manson *et al.*, 2023).

in low-density suburbs, which accounted for less than half of the metropolitan area housing stock in the U.S. in 1990 but more than 80% of the increase in housing supply between 1990 and 2018 (Baum-Snow, 2023).[4] This pattern of rising prices in dense urban centers and expanding supply on the urban periphery reflects both a secular increase in the demand for urban amenities and the difficulty of building in already developed central neighborhoods (Couture *et al.*, 2021; Couture and Handbury, 2020; Glaeser and Gyourko, 2003; Baum-Snow and Han, 2023; Gyourko *et al.*, 2008). But the pattern also raises an important question about how U.S. cities can best address what is widely regarded as a housing-affordability crisis. Namely, can continued suburban expansion alleviate rising housing costs in the urban center, or will cities have to grow denser to become more affordable?

In this paper, we seek to answer this question by examining how different housing submarkets are connected by residential mobility. If residential mobility is high, then an increase in the supply of housing in one submarket can decrease housing costs in other disparate submarkets.[5] If, on the other hand, there is little mobility between housing submarkets, the benefits of expanding the supply of one type of housing will be concentrated among the residents of that type of housing. These two possibilities are represented in two sides of an ongoing debate about urban housing policy, with one side advocating for dramatically expanding the overall supply of housing and the other arguing that increasing the supply of market-rate housing will do little to make cities more affordable for low- and middle-income households that are unable to afford these new market-rate housing units (Been *et al.*, 2019; Gray, 2021; Demsas, 2022; Dougherty, 2020; Friedrich *et al.*, 2023).[6]

Importantly, expanding the supply of one type of housing can make other types of housing more affordable even if there is little residential mobility directly between the two

[4]By contrast, the supply of housing in the 10% of tracts closest to city centers increased by less than one million units, accounting for less than 3% of the growth in metro area housing supply.

[5]In the extreme case of perfect mobility, then spatial equilibrium implies that all households are equally well off and thus that expanding the supply of housing of one type will benefit all residents of the city equally.

[6]The pro-supply and supply-skeptical sides of this debate are respectively summarized by the contrasting headlines, "Build Build Build Build Build Build Build Build Build Build Build Build Build Build" (Dougherty, 2020) and "More Building Won't Make Housing Affordable" (Friedrich *et al.*, 2023).

submarkets, so long as the submarkets are connected by a chain of residential moves that pass through other submarkets. This paper directly studies these *residential vacancy chains* – the series of moves initiated by the construction of a new housing unit.[7] The first migration round, or "link", in a vacancy chain consists of moves into a newly constructed housing unit, which potentially leave the movers' origin units vacant. The second link consists of moves into these vacated units – moves that generate their own set of vacancies. The chain continues on in this way until it ends, either because the origin unit is not vacated,[8] because the vacated origin unit remains vacant or is demolished, or mover's origin unit lies outside the market under consideration, as in the case of a move originating outside the U.S.

In the first part of this paper, we motivate our focus on vacancy chains with a simple model and then present new descriptive facts about the vacancy chains initiated by different kinds of new housing. We show that vacancy chains are relatively short and that new suburban housing supply generates few moves in urban neighborhoods. In the second part of the paper, we conduct a simulation exercise to understand the economic consequences of the descriptive facts we document. We find that the number of vacancies created in a neighborhood is strongly correlated with the price and welfare effects of new housing. These simulation results, when applied to our descriptive findings, imply that the geographic and distributional incidence of the benefits of new housing supply depend importantly on where and what kind of new housing is built.

To fix ideas, we begin by presenting a simple stylized model of a differentiated housing market and derive a general expression for the effect of additional supply in one submarket on prices in other submarkets. We show that even in this simple model, this price effect can be decomposed into a direct effect and an indirect effect, with the indirect effects depending on a chain of cross-submarket residential substitution terms.[9] The expression for

[7]Vacancy chains may be initiated by other events, such as a death that creates a new vacancy or the consolidation of two households. Because we are interested in understanding the effects of new housing construction, we focus only on vacancy chains initiated by the creation of a new housing unit.

[8]The origin unit of a move is not vacated when a new household is formed, as in the case when a roommate moves out to live on her own.

[9]More precisely, the indirect effects depend on the product of residential diversion ratios. The residential

the indirect price effect in our model illustrates the potential importance of vacancy chains as a mechanism connecting different housing submarkets.

Our descriptive characterization of residential vacancy chains is one of the main contributions of our paper. We use newly available administrative data from the Census Bureau on the residential histories of the entire U.S. population from 2000 to 2021 and on the inventory of U.S. residential housing units in 2022. We use these data to identify 1.5 million new single family suburban and multifamily urban housing units built between 2009 and 2018 in the 17 most populous metropolitan areas in the U.S. We then construct the vacancy chains initiated by these units, tracing their paths through different kinds of neighborhoods. We document that while vacancy chains that grow long enough do connect disparate housing submarkets, vacancy chains are generally quite short, with 90% ending within three migration rounds.[10] We also show that the majority of vacancies created by a new unit are created within one year of the initial move into that unit, implying that the vacancy chains we construct would not grow substantially longer if followed over a longer period of time.

We find that each unit of new high-income urban multifamily and new low-density suburban single family housing creates an average of .9 vacancies that subsequently become occupied within four years.[11] Because new housing units are typically more expensive, the number of vacancies created in below-median income tracts is much lower: New high-income urban multifamily housing generates about .15 such vacancies and low-density suburban single family homes generate .25 such vacancies. The number of vacancies created in low-income high-density tracts is even lower still, with high-income urban multifamily and low-density suburban single family housing respectively creating only .03 and .015 vacancies in below-median income tracts in the top decile of population density.

---

diversion ratio between neighborhoods m and n captures the share of households that leave m in response to rising housing costs that substitute towards n.

[10] Existing work on residential vacancy chains by Mast (2021) and Bratu *et al.* (2023) similarly documents the connections between disparate submarkets created by residential mobility. The fact that vacancy chains end quickly is a new fact that we are able to establish because of the comprehensive data we use.

[11] To avoid counting demolished units and units that are unavailable for occupancy because of renovation, we only count vacated units that subsequently become occupied.

While the connectivity between the submarkets for new suburban single family housing and for housing in low-income high-density tracts is especially low, the large increase in the supply of suburban homes means that new suburban construction has generated more vacancies in low-income high-density tracts than has new high-income urban multifamily construction. The 1.2 million new low-density suburban single family units identified in our data collectively created about 17 thousand vacancies in low-income high density tracts, compared to the 11 thousand such vacancies created by the 356 thousand new high-income urban multifamily units identified in our data.

These descriptive results are suggestive but cannot tell us about the effects of new suburban and urban housing construction on prices and the welfare of residents without additional structure. In the second part of this paper, we conduct a simulation exercise that connects the observed characteristics of vacancy chains to the unobserved price and welfare consequences of new housing construction. We conduct these simulations using a model and preference parameters taken from.Bayer *et al.* (2007) and data drawn from the IPUMS 1990 5% sample[12]

The simulation exercise is conceptually simple: We first simulate an initial equilibrium set of prices and matches between households and housing units; then, iterating many times, we add a small number of new housing units to a randomly chosen neighborhood and simulate the new equilibrium prices and matches. The difference between the initial equilibrium and the new equilibrium implies a set of vacancy chains, price effects, and welfare effects, which we analyze to understand what vacancy chains can tell us about the price and welfare effects of new housing. This exercise is the second main contribution of our paper.

We show that, despite using a relatively small sample of data for our simulation, our initial equilibrium prices and assignment of households to housing units replicate the key stylized patterns found in the underlying data. We then simulate the effect of a 5% increase

[12]We use the IPUMS 1990 5% sample rather than the Census Bureau microdata used for our descriptive analysis because it more closely corresponds to the data used to estimate the preference parameters in Bayer *et al.* (2007).

in housing supply in one neighborhood at a time, running 1,000 simulations in total. The simulated welfare and price effects of new housing are economically meaningful – the average elasticity of the returns to living in the city with respect to an increase in supply is 1, and the average elasticity of the urban rent premium with respect to supply is -0.3.

Underlying these average effects is considerable variation in the impact of new housing supply. We show that the number of vacancies created in a neighborhood by vacancy chains initiated by new housing is strongly correlated with this variation. We then compare the predictive power of these vacancies with the direct and indirect cross-neighborhood substitution effects implied by the model underlying our simulation. A key result is that the observed number of vacancies is as predictive of variation in the price effects of new housing as are the model-derived substitution effects. The fact that vacancy chains are relatively easy to observe, whereas the model-derived substitution effects require the estimation of a large number of own- and cross-price demand parameters, makes them especially useful for predicting the non-local price effects of new housing.

Related Literature This paper builds on and contributes to several strands of the extensive literature on housing supply and affordability. Our work is most closely related to a small literature on vacancy chains that goes back to White (1971). The data required to observe vacancy chains means that earlier work was mostly theoretical or relied on statistical modeling of vacancy chains (Marullo, 1985; Turner, 2008). The availability of more detailed data on residential histories has only recently made it possible to construct vacancy chains, as in recent work by Mast (2021) and Bratu *et al.* (2023) who study residential vacancy chains in the U.S. and Finland, respectively. Our data allow us to contribute to this literature by documenting new facts about vacancy chains. In addition, we contribute to this literature by providing insights into the economic implications of these descriptive patterns.

Previous work has examined other ways in which housing submarkets interconnect, either through filtering (Rosenthal, 2014; Liu *et al.*, 2022), search (Piazzesi *et al.*, 2019; Landvoigt *et al.*, 2015), aggregate substitution between housing submarkets (Nathanson, 2020), or the hyper-local effects of new housing construction (Asquith *et al.*, 2021; Damiano and Frenier, 2020; Diamond and McQuade, 2019; Pennington, 2021; Li, 2019). Vacancy chains represent an important micro-level mechanism underlying these higher-level mechanisms.

We also contribute to the literature on the city-wide effects of increases in housing supply (e.g., Anenberg and Kung (2020); Molloy *et al.* (2022)). Our approach allows us to contribute to this literature by going beyond city-wide averages to better understand the geographic and sociodemographic incidence of the benefits of new housing.

Finally, this paper relates to extensive literatures on supply constraints (Song, 2021; Gyourko *et al.*, 2008; Saiz, 2010; Baum-Snow and Han, 2023; Kulka *et al.*, 2022) and urban housing affordability (Couture *et al.*, 2021; Couture and Handbury, 2020; Su, 2022; Handbury, 2021). The consequences of urban supply constraints on urban housing costs depend crucially on residential mobility across submarkets. This, in turn, has significant implications for housing policy. If households are able to move easily between submarkets, then policies that increase supply in areas with few constraints may be effective. In contrast, limited mobility between submarkets would recommend policies that relax existing constraints. We contribute to these literatures by documenting the extent to which increases in the supply of housing in one submarket ripple across other submarkets.

The rest of this paper proceeds as follows: Section 2.2 presents a simple model of differentiated housing demand to fix ideas. Section 2.3 describes our data and the construction of vacancy chains. Section 2.4 presents new descriptive facts about vacancy chains initiated by new housing construction in the U.S. Section 2.5 describes our simulation exercise and results. Section 2.6 concludes.

2.2 A Simple Model

In this section, we develop a simple static model of residential demand for housing in different submarkets. Despite the model's simplicity, we are able to express the effect of an increase in the supply of housing in one submarket on housing prices in other submarkets as a function of direct and indirect effects, with the indirect effects consisting of a chain of residential substitution effects mediated by intermediary submarkets. This chain of substitution effects is naturally interpreted as a residential vacancy chain, thus motivating the descriptive analysis of residential vacancy chains in Section 2.4.

Setup We model the housing market of a single city as being made up of housing submarkets $n \in \mathcal{N} = \{1, ..., N\}$. Housing supply in each submarket is perfectly inelastic, with supply in submarket n denoted S_n. We model demand for housing in submarket n with the reduced form $D_n(\mathbf{p})$, where $\mathbf{p} = (p_1, ..., p_N)$ is the vector of prices for housing in all submarkets. We assume that $\frac{\partial D_n(\mathbf{p})}{\partial p_n} < 0$ and $\frac{\partial D_n(\mathbf{p})}{\partial p_m} \geq 0$ for $m \neq n$. We denote the own-price effect on demand for submarket n by $\epsilon_n \equiv -\frac{\partial D_n}{\partial p_n}$; and the cross-price effect on demand for submarket n with respect to prices in $m \neq n$ by $\gamma_{nm} \equiv \frac{\partial D_n}{\partial p_m}$. Prices in equilibrium adjust to clear each submarket.

In addition, we denote the "residential diversion ratio" of m to n by $\lambda_{nm} \equiv \frac{\gamma_{nm}}{\epsilon_m}$. This measure captures the share of the change in demand for housing in submarket m that comes from substitution to or from submarket n when price p_m changes. In the event of a *decrease* in p_m, λ_{nm} represents the share of *increased* demand for housing in m due to substitution *away* from n. When aggregate residential demand is microfounded on household preferences that feature unit housing demand, this interpretation of λ_{nm} is further refined to the share of households that move to submarket m that come from submarket n.

Cross-market effects We are interested in how an increase in supply in one submarket affects prices in other submarkets. Without loss of generality, we consider the impact of an exogenous increase in supply in submarket 1 on prices in submarket 2.

For clarity of exposition, first consider the cross-market price effects when $N = 3$. Differentiating the market clearing conditions and solving the model yields the following:

$$\frac{dp_2}{dS_1} \propto -\epsilon_2^{-1} \left(\underbrace{\lambda_{21}}_{\textit{direct effect}} + \underbrace{\lambda_{23}\lambda_{31}}_{\textit{indirect effect}} \right).$$

This expression shows that the cross-market price effect depends on a *direct effect*, which is proportional to the residential diversion ratio of 2 to 1, and an *indirect effect*, which is proportional to the product of the residential diversion ratios of 3 to 1 and of 2 to 3. The interpretation of this indirect effect is as follows: The increased supply in 1 leads households to substitute housing demand away from 3 towards 1; this, in turn, leads households to shift housing demand away from 2 towards 3. The natural interpretation of this chain of substitution effects is a residential vacancy chain.

The fact that this cross-market price effect depends on a direct effect and on indirect effects extends to the general case with N neighborhoods. We demonstrate this by first deriving a general expression for the cross-price effects that takes a recursive form.

Proposition 1 *Given a set of neighborhoods $\mathcal{N} = \{1, ..., N\}$, the price-effects of an increase in S_1 take the form:*

$$\frac{dp_i}{dS_1} = -\epsilon_i^{-1} \frac{\Phi_i(\mathcal{N})}{\sum_{j \in \mathcal{N}} \lambda_{1j} \Phi_j(\mathcal{N})}, \tag{2.1}$$

where $\Phi_i(\mathcal{N})$ is defined recursively:

$$\Phi_{i \neq 1}(\mathcal{N}) = \sum_{j \in \mathcal{N} \setminus i} \lambda_{ij} \Phi_j(\mathcal{N} \setminus i) \tag{2.2}$$

$$\Phi_1(\mathcal{N}) = -\sum_{j \in \mathcal{N} \setminus 1} \lambda_{Nj} \Phi_j(\mathcal{N} \setminus 1). \tag{2.3}$$

The recursive form of this expression shows that the cross-market price effects of an increase in S_1 on submarket i depend on the cross-market effects that would exist in submarkets $j \neq i$ if submarket i were removed from $\mathcal{N}$. These cross-market effects are then transmitted to i based on the residential diversion ratio of i for j.

Given that we are able to express the cross-market effects in the case when $N = 3$ as a function of a direct and indirect effect, and given the general expression in the proposition above, it follows from inspection that the price effects consist of a direct effect and a set of indirect effects for any $N \geq 3$. Further, the indirect effects themselves consist of direct and indirect effects, resulting in indirect effects running between submarkets 1 and 2 that are mediated by up to $N - 2$ other neighborhoods. While the simple setup of this model doesn't allow us to make explicit predictions about vacancy chains, it illustrates the intuition for how an increase in the supply of housing in one submarket impacts prices in another submarket through multiple chains of residential substitution between neighborhoods. This motivates our descriptive analysis of vacancy chains in Section 2.4.

2.3 Data

2.3.1 Data sources

Administrative data from the Master Address File The primary data sources we use to construct vacancy chains are derived from the Census Bureau's Master Address File (MAF). The MAF database is an inventory of all known living quarters in the United States and was created for the 2000 Census. It is updated semi-annually from the US Postal Services delivery sequence file. Additional updates occur through partnerships with local and state governments, address canvassing activities, and other sources. The MAF defines the base sampling frame for the American Community Survey, decennial censuses, and other Census Bureau data products.

We use the 2022 MAF Extract (MAF-X), a snapshot of the MAF database in which we observe the inventory of US housing units in 2022. Housing units in the MAF-X are assigned a MAFID, a unique identifier that can be used to link records across data sources. In addition, we observe housing units' addresses and geographic locations.

In addition, we use the 2000-2021 MAF Auxiliary Reference Files (MAF-ARF) to construct residential histories of the US population at an annual frequency (Sullivan and Genadek, 2024). The MAF-ARF is derived from several underlying data sources, including individual tax returns, Selective Service registration data, and Medicare enrollment. Each year of the MAF-ARF is at the individual-level and consists of an individual-level unique identifier and an associated MAFID.

Two key housing unit characteristics for our analysis are the age of the unit and the number of units in the unit's building. While we do not directly observe these characteristics, we are able to impute them from the MAF-X and MAF-ARF.

To impute the year in which a unit first becomes occupied, we use the first year in which a MAFID is associated with a PIK in the MAF-ARF as that unit's first year of occupancy. We construct vacancy chains initiated by new housing in each year from 2009 to 2018. Because our residential histories extend back to 2000, units identified as being new had no associated PIKs for at least nine years before first appearing in the MAF-ARF.

To impute the number of units in a unit's building, we simply take the number of valid MAFIDs in the same census tract with the same street address. In doing so, we exclude units that are indicated to be trailers or mobile homes, or have an address indicating that units are located on separate lots.

Additional data sources We combine these administrative data sources with data from the American Community Survey (ACS) covering 2005-2018. We use the ACS to measure the tract-level characteristics with which we characterize the vacancies created by new housing units and the types of new housing constructed.

2.3.2 Moves across submarkets

Connections between housing submarkets Figure 2.1 illustrates connections between housing submarkets within Core-Based Statistical Areas (CBSAs) with a 2010 population of three million or greater between 2010 and 2017.

Figure 2.1A shows the share of individuals aged 25 and older moving from a tract at a given within-CBSA decile of household income that increase or decrease their tract's decile by a given amount. Individuals who leave the CBSA account for the difference between the total share of movers who change their tract's decile and 1. The figure shows that, while there is some stickiness in the kinds of tracts that individuals move to, tracts at different deciles of household income are still connected by a substantial number of moves. For example, while 41% of all moves out of top-decile tracts are also to top-decile tracts, about 27% of moves are to lower-decile tracts within the same CBSA. In addition, the figure illustrates how vacancy chains can connect submarkets even when there are few direct moves between them. For example, while there are very few direct moves between

Figure 2.1. *Moves between tracts within CBSAs*

A. *Median Household Income*

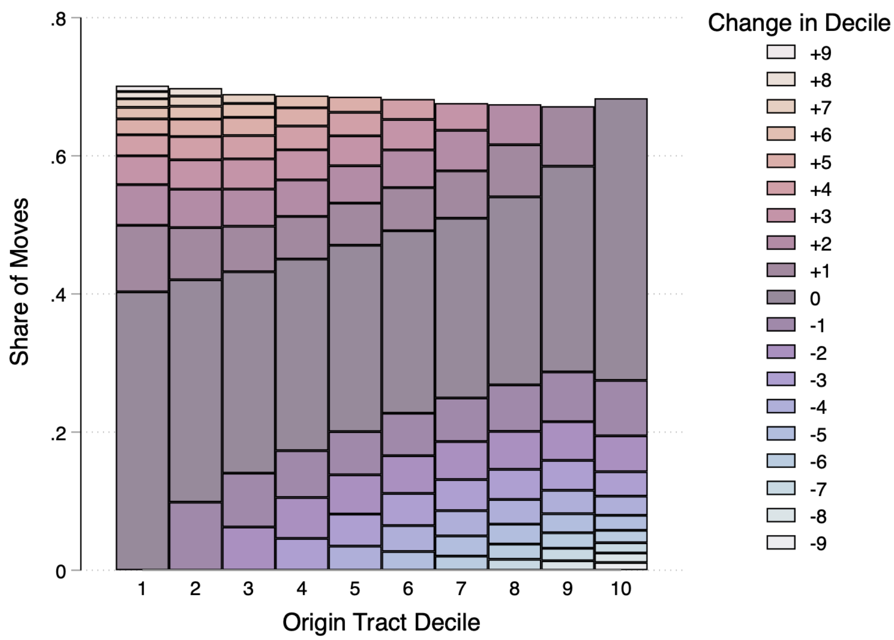

B. *Median Two-Bedroom Rent*

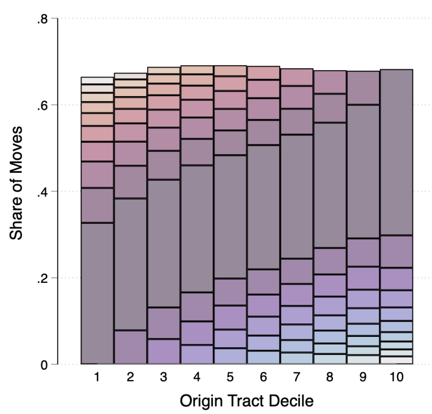

C. *Percent College-Educated*

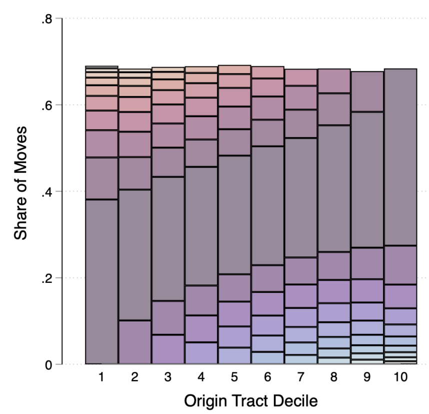

Notes: This figure shows the distribution of changes in tract characteristics conditional on origin tract characteristics for individuals aged twenty-five and older who moved between 2010 and 2017 and originated from a tract in a CBSA with a 2010 population of three million or more. In each panel, each column of stacked bars represents movers out of an origin tract at the given within-CBSA decile of the indicated characteristic. The size of each bar indicates the share of moves out of the given origin-tract decile that result in the change of tract decile indicated in the legend of panel A. The stacked bars sum to less than one because of moves out of the CBSA. We calculate tract characteristics using the 2013-2017 ACS.

the bottom- and top-decile tracts, these two submarkets are connected indirectly through chains of moves from bottom-decile tracts that pass through tracts at other deciles and end at top-decile tracts. An example of one such path is the series of moves from bottom-decile tracts to top-decile tracts that increase the mover's tract decile by one.

Panels B and C of Figure 2.1 illustrate similar patterns of migratory flows across submarkets defined by median rents and college share.

2.3.3 Constructing vacancy chains

New housing gives rise to vacancy chains by initiating a series of moves in which households move into newly available units and vacate their origin units. While the idea is simple, some complications arise when considering how to construct a vacancy chain and connect moves over time. First, not all moves leave a unit vacant, although they may leave a room within the unit vacant. For example, if a roommate moves out, this may initiate a series of moves even though the unit was not vacated. For our analysis, we simplify things by only considering moves that vacate the origin unit. If a unit is not vacated, the chain ends.

Second, units may sit vacant for some time before becoming occupied. When constructing and describing vacancy chains, we would like to describe not only the composition of neighborhoods and movers that are part of the chain but also how the chain evolves over time. To do so, we construct vacancy chains over different time horizons. Figure 2.2 illustrates how we construct a hypothetical vacancy chain initiated by a new housing unit that was first occupied in 2010. To construct the first link of the chain, we identify the occupants of the new housing unit and trace them back to their 2009 origin unit or units. For origin units that were vacated in the first link of the chain, we search for new occupants in 2010 and trace them back to their 2009 origin unit or units, constructing the second link of the chain.

We continue in this fashion until the chain ends, either because no origin unit is found, an origin unit isn't left vacant, or a vacated unit doesn't become occupied within the chosen time horizon. The vacancy chain illustrated in Figure 2.2 would end after two rounds of

Figure 2.2. *Vacancy Chain Construction*

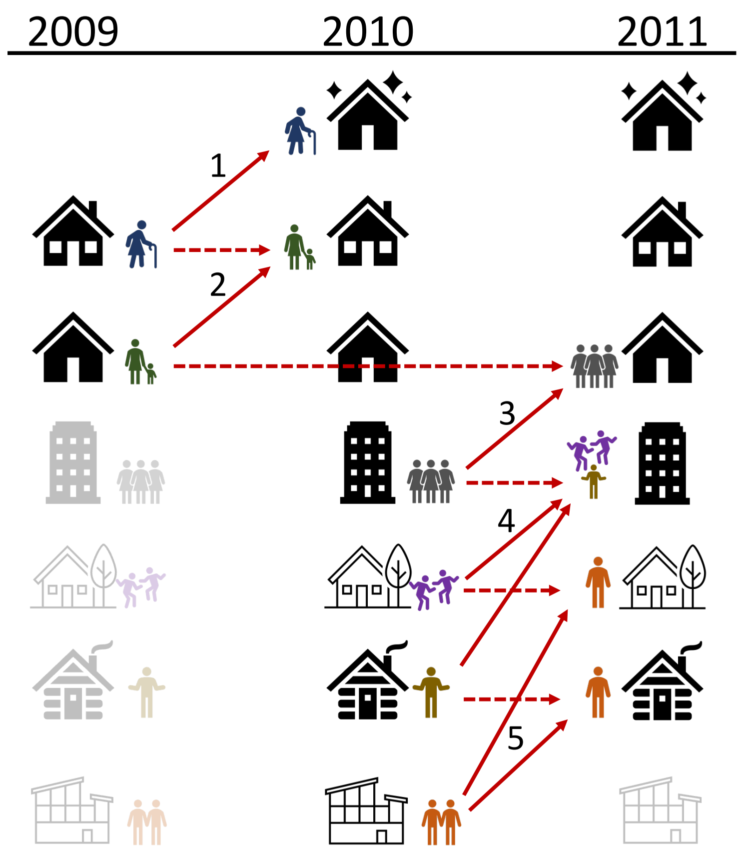

moves if we were to construct the chain over a one-year horizon. Over a two-year horizon, we are able to extend the chain by searching for a new occupant of the vacated housing unit in 2011. We then continue building the chain over a one-year horizon, with 2011 as the reference year. This generalizes for longer horizons straightforwardly.

Describing vacancy chains To characterize vacancy chains, we consider the number of effective vacancies created in different kinds of neighborhoods. Effective vacancies are calculated as a weighted sum, where the weights are inversely proportional to the number of distinct chains connected to a given vacancy. For example, if a unit is vacated by a move in which some household members move to a unit that is part of one vacancy chain and the remaining household members move to a unit that is part of a separate chain, we attribute half of the resulting vacancy to each chain. In addition, we assign a weight of zero to units that are vacated but are not observed to be filled within the time horizon under consideration. This is to avoid counting vacated units that are demolished or are unavailable for occupancy due to renovation or redevelopment.

We focus on vacancy chains initiated by two types of new housing in CBSAs with populations greater than three million: low-density suburban single family homes and high-income urban multifamily buildings. The first type consists of single family homes located in below-median density tracts outside the metropolitan area's principal city. New construction in these tracts accounts for 80% of the increase in housing stock in the US since 1980 (Baum-Snow, 2023). The second type of new construction we consider consists of units in 20+ unit buildings located in above-median income tracts within five miles of the metropolitan area's central business district, which corresponds to the type of new housing studied in the existing literature on vacancy chains (Mast (2021); Bratu *et al.* (2023)). Our primary analysis sample consists of 1,159,000 initiated by low-density suburban single family homes and 356,000 vacancy chains initiated by high-income urban multifamily units.

Figure 2.3. *Origin Tract Incomes by Migration Round*

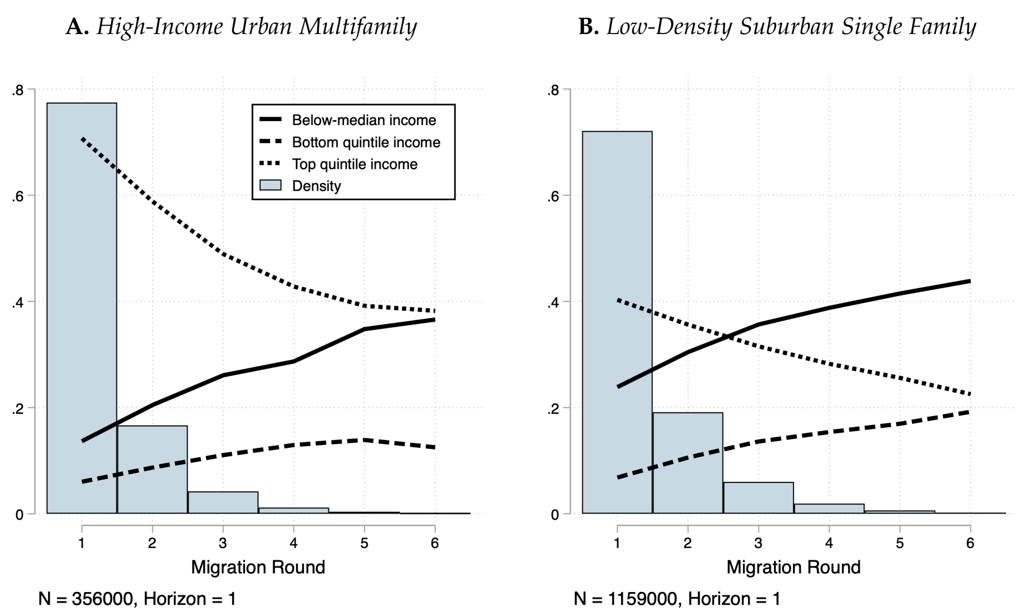

Notes: This figure shows the share of vacated units in each migration round located in a tract with a given characteristic. In each panel, migration round 1 represents vacancies created by the set of moves into the new units that initiated the chain; migration round 2 represents the vacancies created by the set of moves into the vacancies created in round 1; and so on. Chains are constructed over a one-year horizon. Tract characteristics for vacancy chains initiated in year t are taken from the 5-year ACS covering years $t-4$ to t and quantiles correspond to the national distribution. Income is median household income per capita.

2.4 Descriptive Facts

We now turn to our descriptive characterization of vacancy chains. We begin by considering how the composition of vacated units changes as the vacancy chain grows longer. Figure 2.3 shows the share of effective vacancies created in tracts with a given characteristic, conditional on the migration round and over a one-year horizon. Panel A shows these shares for vacancy chains initiated by high-income urban multifamily housing. The first round of moves into these new high-income units create vacancies in predominantly high-income tracts, with 71% of vacancies created in top quintile income tracts, 14% in below-median tracts, and only 8% in bottom quintile tracts. This is unsurprising, given that these kinds of new units are typically very expensive.

As the chain progresses, however, the share of vacancies created in high-income tracts in each round declines and the share created in low-income tracts rises. By the sixth round of moves, 38% of vacancies are created in top quintile income tracts and 37% of vacancies are created in below-median income tracts.

Panel B of Figure 2.3 displays similar trends as the vacancy chains initiated by new single family homes in low-density suburban tracts grow longer. The share of vacancies in top quintile tracts declines from 40% in the first round of moves to 22% in the sixth round while the share of vacancies created in below-median income tracts increases from 24% in the first round to 44% in the sixth. A notable difference from the chains initiated by high-income urban multifamily housing is that the vacancies created by the initial set of moves is concentrated in tracts with lower median incomes. Again, this is unsurprising, since housing costs for new single family homes in low-density suburbs are typically lower than those for high-income urban multifamily units.

Overall, Figure 2.3 illustrates how different housing submarkets are connected by residential mobility. The fact that new housing units create vacancies in lower income tracts suggests that building new housing – even in high-income neighborhoods – can loosen demand for housing in lower segments of the market and lower housing costs.

Taken in isolation, this fact might suggest that a viable strategy for lowering housing costs for low-income renters is to build more housing of any kind. However, another salient fact illustrated in Figure 2.3 is that vacancy chains are relatively short. In both panels, the density of the total number of vacancies is shown by the blue histogram. Panel A shows that 77% of all vacancies created by high-income urban multifamily housing are created in the first round of moves, and each round of moves creates about 25% as many vacancies as the previous round. In panel B, a qualitatively similar pattern holds. An important implication of this pattern is that the location and type of new housing construction matters. Even though increases in the supply of new high-end housing can loosen demand in lower-end segments if the resulting vacancy chains go on long enough, our current results imply vacancy chains are very unlikely to continue for many rounds and these supply increases are unlikely to have meaningful effects on costs.

Given the short length of vacancy chains, we turn our focus to the cumulative number of effective vacancies created by vacancy chains. Figure 2.4 shows the cumulative number of effective vacancies created in each round of moves. Panels A and B show vacancies created over a one-year time horizon, with panel A showing vacancies created by high-income urban multifamily housing and panel B showing vacancies created by low-density suburban single family homes. In both panels, the number of vacancies quickly plateaus as the chain grows longer due to the relatively low probability of a chain advancing from one round to the next. This suggests that there would be very few additional vacancies created in migration rounds beyond the sixth round, and the number of vacancies created by the sixth round is a close approximation of the total number of vacancies created by each type of new construction.

Panels C and D show cumulative vacancies created over a four-year horizon, with panel C showing vacancies created by high-income urban multifamily housing and panel D showing vacancies created by low-density suburban single family homes. While chains constructed over this longer time horizon tend to be longer, the distribution of vacancies is still strongly skewed towards earlier migration rounds such that the cumulative number of vacancies levels off by the sixth round.

Figure 2.5 shows how the cumulative number of effective vacancies created by both kinds of new construction changes over time. Panels A and B demonstrate that both types of housing produce more vacancies in high-income neighborhoods than in low-income neighborhoods, though high-income urban multifamily housing produces substantially more high-income vacancies than does low-density suburban single family housing. In particular, panel A shows that high-income urban multifamily housing produces .44 and .58 vacancies in top quintile income tracts over a one- and four-year horizon, respectively. This represents about two thirds of the total number of vacancies created by this type of housing. In contrast, high-income urban multifamily housing produces only .1 and .15 vacancies in below-median income tracts over a one- and four-year horizon, respectively, which represents about 15% of the total number of vacancies created.

Panel B shows that low-density suburban housing produces .28 and .34 vacancies in

Figure 2.4. *Cumulative Vacancies by Migration Round*

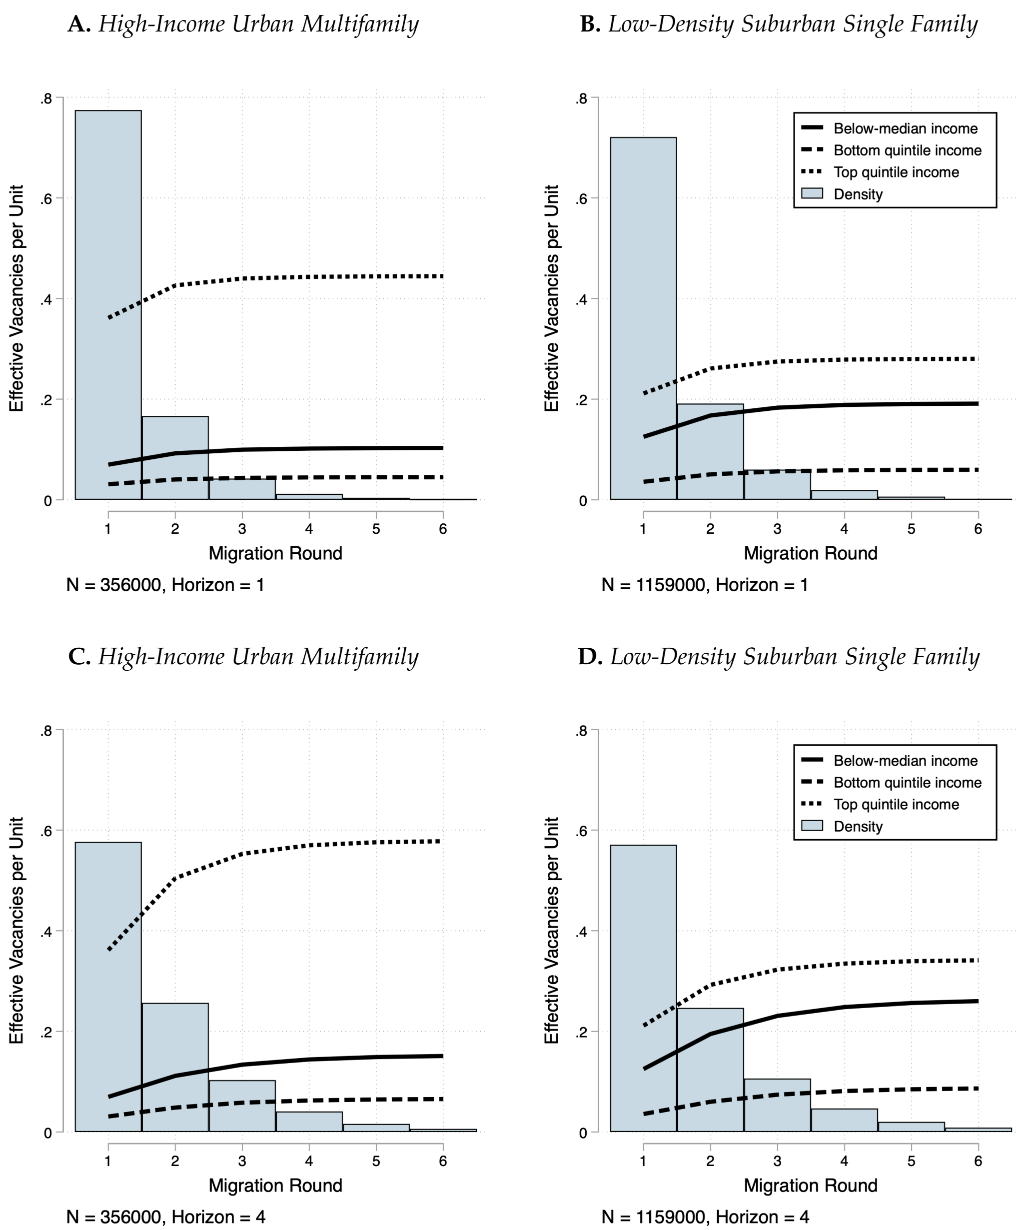

Notes: This figure shows the cumulative number of effective vacancies created in each migration round located in a tract with a given characteristic. In each panel, migration round 1 represents vacancies created by the set of moves into the new units that initiated the chain; migration round 2 represents the vacancies created by the set of moves into the vacancies created in round 1; and so on. Chains are constructed over a one-year horizon. Tract characteristics for vacancy chains initiated in year t are taken from the 5-year ACS covering years $t-4$ to t and quantiles correspond to the national distribution. Income is median household income per capita.

Figure 2.5. *Cumulative Vacancies per Unit over Time*

A. *High-Income Urban Multifamily* **B.** *Low-Density Suburban Single Family*

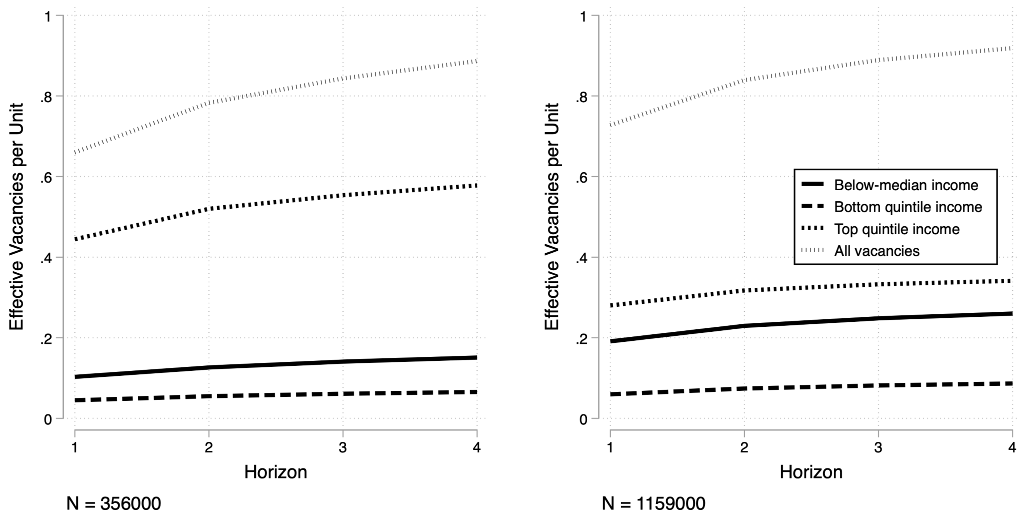

C. *High-Income Urban Multifamily* **D.** *Low-Density Suburban Single Family*

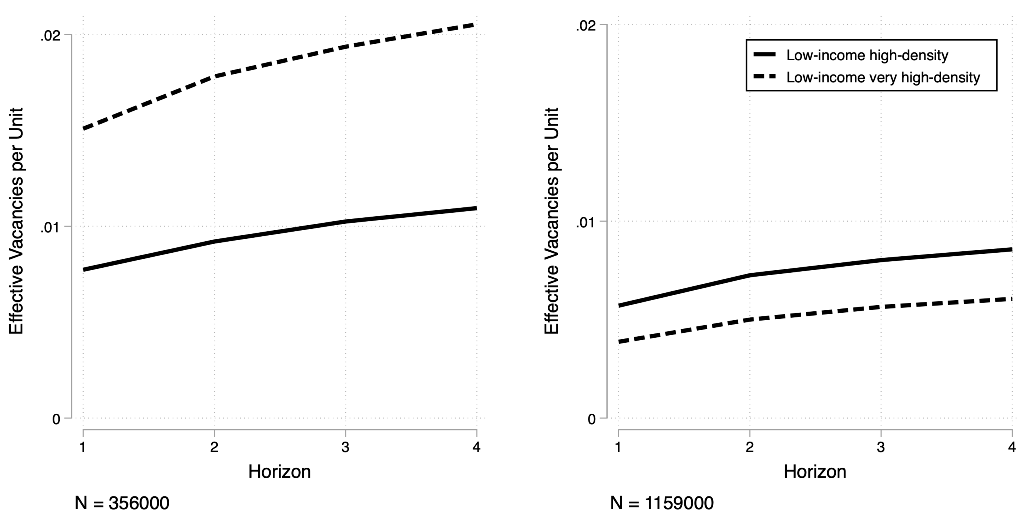

Notes: This figure shows the number of effective vacancies created over time in tracts with a given characteristic per unit of new housing. In each panel, migration round 1 represents vacancies created by the set of moves into the new units that initiated the chain; migration round 2 represents the vacancies created by the set of moves into the vacancies created in round 1; and so on. Each point represents the number of effective vacancies created over six rounds of moves over the indicated time horizon. Tract characteristics for vacancy chains initiated in year t are taken from the 5-year ACS covering years $t-4$ to t and quantiles correspond to the national distribution. Income is median household income per capita. Low-income tracts are those in the bottom quintile of the income distribution, high-density tracts are in the 19th vintile of the distribution of population density, and very high-density tracts are those in the top vingtile of the distribution.

top quintile income tracts over a one- and four-year horizon, respectively, which represents about 35% of the total number of vacancies created. New low-density suburban housing creates a comparable number of vacancies in below-median income tracts: .19 and .26 vacancies over one- and four-year horizons, representing about 25% of the total number of vacancies created.

Both panels A and B show that the majority of vacancies created by new housing construction of either type are created within a one-year horizon. In addition, the number of new vacancies created in each year is diminishing. This pattern suggests that few additional vacancies are created over time horizons beyond four years and we are therefore capturing the majority of vacancies created by new housing construction.

Panels C and D of Figure 2.5 focus on how both types of new housing construction connect to the low-income high-density submarkets where the households most exposed to rising housing costs are most likely to live. Both panels show the number of vacancies created over time in low-income tracts that are either high-density or very high-density. We define low-income tracts as those in the bottom quintile of the national distribution of income; high-density tracts as those in the 19th ventile of the national distribution of population density; and very high-density tracts as those in the top ventile of the national distribution of population density.

Panel C shows the number of effective vacancies created in low-income and high-density tracts by high-income urban multifamily housing. In general, the number of vacancies created is very low – over a four-year horizon, it takes 50 new high-income urban multifamily units to generate one vacancy in a low-income very high-density tract and 100 new units to generate a vacancy in a low-income high-density tract. Panel D shows that the number of effective vacancies created in these submarkets by new low-income suburban single family housing is even lower, requiring more than 100 new units to generate a vacancy in either high- or very high-density low income tracts over a four-year horizon.

While each unit of high-income urban multifamily housing generates more vacancies in low-income urban neighborhoods than a new unit of low-density suburban single family

Figure 2.6. *Total Vacancies over Time*

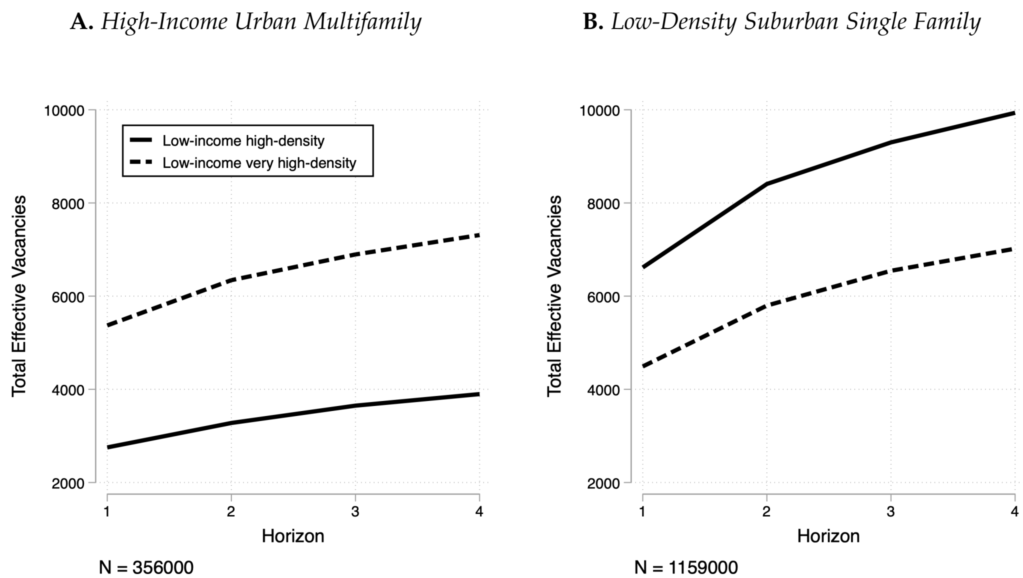

Notes: This figure shows the total number of effective vacancies created over time in tracts with a given characteristic. In each panel, migration round 1 represents vacancies created by the set of moves into the new units that initiated the chain; migration round 2 represents the vacancies created by the set of moves into the vacancies created in round 1; and so on. Each point represents the number of effective vacancies created over six rounds of moves over the indicated time horizon. Tract characteristics for vacancy chains initiated in year t are taken from the 5-year ACS covering years $t-4$ to t and quantiles correspond to the national distribution. Income is median household income per capita. Low-income tracts are those in the bottom quintile of the income distribution, high-density tracts are in the 19th vingtile of the distribution of population density, and very high-density tracts are those in the top vingtile of the distribution.

housing, Figure 2.6 shows that low-density suburban single family homes have created more total vacancies in low-income high-density tracts and a comparable number of vacancies in low-income very high-density tracts. This is due to the much larger number of new suburban housing units constructed between 2008 and 2018.

2.5 The Economic Content of Vacancy Chains

We have presented a detailed picture of how vacancy chains vary according to the type of new housing that is built and the location of construction. In this section, we conduct a simulation exercise that connects the observed characteristics of vacancy chains to unobserved price and welfare effects.

The simulation exercise is conceptually simple: We first simulate an initial equilibrium matching of households to housing units and a vector of prices that sustains that matching; then, iterating many times, we add a small number of new housing units to a randomly chosen neighborhood and simulate the new equilibrium prices and matching. The difference between the initial equilibrium and the new equilibrium implies a set of vacancy chains, price effects, and welfare effects, which we analyze to understand what vacancy chains can tell us about the price and welfare effects of new housing.

We use a modified version of the model estimated by Bayer *et al.* (2007) and calibrate it using their parameter estimates. We sample neighborhoods, housing units, and households from the 1990 IPUMS 5% sample. Given these preferences and the sampled neighborhoods, housing units, and households, we apply a tatonnement algorithm to find an initial equilibrium consisting of a market-clearing set of prices and the corresponding matching of households to housing units. We then repeatedly draw new units and add them to the housing stock in the simulated data, recomputing the equilibrium prices and matching of households to units in each iteration. We show how the resulting simulated vacancy chain characteristics correlate with the characteristics of the new unit types and locations, as well as the price and welfare effects.

2.5.1 Model

The following is a modified version of the residential choice model presented in Bayer *et al.* (2007), modified to make it easily replicated using the IPUMs data. There is a finite number of households indexed by i, housing units indexed by h, and neighborhoods indexed by n. The set of available housing units consists not only of units in the city, but also units in an outside option neighborhood oo.

Households choose a unit to live in to solve

$$\max_h V^i_{hn} = \underbrace{\alpha^i_X \mathbf{X}_h + \alpha^i_Z \mathbf{Z}_n - \alpha^i_p p_h + \xi^i_h}_{\equiv v^i_{hn}} + \epsilon^i_{nh},$$

where $\mathbf{X}_h$ is a vector of non-price housing unit characteristics that includes housing unit age bin indicators and the number of rooms; $\mathbf{Z}_n$ is a vector of neighborhood characteristics that includes racial and ethnic composition, college-educated share, and average income; p_h is the price of unit h, ξ_n is unobserved neighborhood quality; and ϵ_h^i is household i's idiosyncratic preference for unit h, where ϵ_h^i is drawn from a type-1 extreme value distribution. Units in the outside option are normalized such that $v_{h,oo}^i = 0$. In addition, households are indifferent between units in the outside option, so $\epsilon_h^i = \epsilon_{oo}^i$ for all units h in the outside option.

Household preferences are permitted to vary with the following observable household characteristics: The presence of children under 18; capital and non-capital income; capital income; race and ethnicity; educational attainment; employment status; and age.

We impute unobserved neighborhood quality ξ_n^i by estimating a hedonic regression and taking the mean residual variation in rents across across housing units within a neighborhood. We assume that this residual variation reflects a willingness to pay to live in n that is common across all households. Because the marginal utility of consumption is permitted to vary across households, this assumption implies that ξ_n^i varies across households.

Parameter estimates are computed from the tables in Bayer *et al.* (2004). This model is particularly well suited to our application because it models residential choices at the housing unit level and the parameters are estimated to maximize the likelihood of observing each household matched with the unit in which it resides. This is in contrast to many residential discrete choice models in which a continuum of households choose over neighborhoods and the parameters are estimated to match the neighborhood choice shares observed in the data.

2.5.2 Equilibrium and Iteration

We now describe the tatonnement algorithm we use to compute equilibrium prices for a given set of households and housing units. This is an implementation of the Hungarian algorithm (Demange, Gale, and Sotomayor, 1986; Easley and Klineberg, 2010).

We begin by setting all prices equal to zero. In each iteration of the algorithm, we find each household's utility-maximizing set of housing units given the current vector of prices, which we refer to as their preferred units. If there is a perfect matching of households to housing units in which each household matches with one of its preferred units, we have found an equilibrium. If there is no such perfect matching, then there must exist a *constricted set* of units $\mathcal{S}$ – a set of units such that: (a) the households that prefer units in $\mathcal{S}$ prefer no units outside of $\mathcal{S}$; and (b) there are more households that prefer units in $\mathcal{S}$ than there are units in $\mathcal{S}$. We identify the constricted set and raise prices for all units in $\mathcal{S}$ by one price increment. We then begin the next iteration of the algorithm and continue until a perfect matching is found.

Because the algorithm requires that valuations and prices have discrete support, we normalize the data in several ways. First, we convert all preferences into a willingness to pay by rescaling each household's preference parameters such that the marginal utility of consumption is unity. This implies that the scale of idiosyncratic preferences varies across households. Second, we rescale the utility achieved with each choice to be integer-valued. We do so by dividing by the price increment used in the algorithm and rounding to the nearest integer.[13]

Finally, we normalize the value of the outside option to be equal to the minimum utility achieved by a household choosing an inside option when prices are equal to 0. We then shift all utilities by a constant such that the utility achieved by a household choosing the outside option is 0. After these normalizations, each household's utility when matched to housing unit h reflects their willingness to pay (in units implied by the price increment) to live in h rather than in the outside option.

Once we have computed an initial equilibrium, we repeatedly simulate the effects of new housing construction. In each iteration, we begin with the initial equilibrium and randomly sample a small number of new housing units and add them to the set of existing housing units in a single neighborhood. We then find the new equilibrium, construct the resulting vacancy chains, and calculate the price and welfare effects of the increase in supply.

[13]Rounding to the nearest integer naturally introduces some error into the algorithm. Using smaller price increments leads to lower approximation error but at the expense of computation time.

One limitation of this exercise is that it does not allow for vacancy chains to end as a result of new household formation or because a unit remains vacant, both of which are important reasons that vacancy chains end in the observed data. For this exercise, vacancy chains can only end because they reach the outside option. Despite this limitation, the simulation exercise provides insights that help us interpret the descriptive patterns described in the previous section.

2.5.3 Data

Our simulation exercise uses microdata from the 1990 IPUMS 5% sample.[14] While we define neighborhoods at the tract level in our descriptive analysis of vacancy chains, the most granular geographic units we observe in our simulation data are PUMAs. We, therefore, define neighborhoods at the PUMA level for this exercise. While PUMAs are much more populous than tracts, containing at least 100,000 individuals, they are geographically compact in dense metro areas, making them a reasonable proxy for neighborhoods.

The tatonnement algorithm we use to find equilibrium matchings and prices is computationally intensive and fails to converge when using a realistic number of households, housing units, and neighborhoods. Because of this, we use a reduced sample to simulate the effects of increased housing supply. For our primary specification, we construct a bootstrapped sample by randomly sampling 10 PUMAs from the Chicago CBSA with replacement. We then sample 100 housing units and 133 households from each of the sampled PUMAs. In addition, we take all housing units in our sampled PUMAs that are less than one year old as the pool from which we draw new housing units in our simulation.

Because there are more households than housing units in our sample, we augment the sample by adding additional units to represent the outside option. The value of these outside option units is normalized such that households have zero willingness to pay to

[14]We use the 1990 IPUMS data to facilitate the use of the parameter estimates from Bayer *et al.* (2007). One concern with using these data rather than more recent data is that the demographic composition and amenity value of cities have changed substantially since 1990. This might mean that our simulation results do not generalize to the time period used for our descriptive analysis of vacancy chains.

live in them and are indifferent between all outside option units. When computing an equilibrium, we thus have a perfect matching when every housing unit in the CBSA is matched to a household and the remaining households are matched to units in the outside option.

Table 2.1 reports mean characteristics of the PUMAs used in the simulation exercise. The main point worth noting is that there is substantial variation in neighborhood characteristics, with PUMAs ranging from very low-income and low college share to high-income high college share. Figure 2.7 shows the locations of these neighborhoods in the Chicago Metropolitan Area. The PUMAs used in our simulation exercise are also geographically varied, with some located in high-density areas near the city center and others in more distant suburbs.

2.5.4 Results

Initial Equilibrium We begin by describing the initial equilibrium of our simulation exercise. Reassuringly, we observe that the patterns in this initial equilibrium are similar to those in the underlying data. Figure 2.8 shows that the simulated housing unit prices in our initial equilibrium are highly correlated with observed housing prices, with an increase in a unit's observed price being accompanied on average by a one-for-one increase in that unit's simulated price.

Figure 2.9 shows that the matching of households to units preserves the sorting patterns observed in the underlying data. Each panel shows the mean of a given characteristic of sampled households conditional on the mean of the PUMA they reside in. The red squares show this relationship for the underlying data, while the blue circles show the relationship between the mean characteristics of sampled households conditional on the mean of the PUMA they match with in our initial equilibrium. The pattern of sorting that results from our simulation is remarkably similar to the pattern observed in the underlying data, with high-income households sorting to high-income PUMAs, college-educated households

Table 2.1. *Mean Characteristics of Sampled PUMAs*

PUMA	ξ_n	Black	Hispanic	College	Household Income	Owner-occupied	Rooms	Age	Employed	Child Present	Built in 80s
1	-12.4	0.11	0.03	0.25	48,628	0.78	6.2	44.9	0.81	0.50	0.20
2	-82.7	0.22	0.04	0.24	45,301	0.77	6.0	48.8	0.71	0.43	0.10
3	-45.3	0.81	0.01	0.25	23,573	0.16	4.2	48.9	0.44	0.34	0.05
4	-17.1	0.00	0.02	0.27	52,834	0.84	6.2	47.8	0.77	0.41	0.36
5	-0.9	0.03	0.04	0.22	47,093	0.75	5.6	49.6	0.72	0.35	0.09
6	34.7	0.03	0.02	0.50	59,717	0.72	6.2	43.0	0.86	0.42	0.30
7	46.6	0.21	0.03	0.17	42,211	0.78	5.7	47.7	0.73	0.43	0.14
8	-50.4	0.49	0.19	0.10	31,548	0.64	5.3	50.7	0.58	0.44	0.02
9	-93.6	0.00	0.02	0.20	41,389	0.75	5.8	47.3	0.74	0.42	0.16
10	-8.9	0.79	0.16	0.04	20,989	0.29	5.2	47.1	0.44	0.55	0.08

Notes: This table reports mean characteristics of the PUMAs sampled for our simulation exercise, as reported in the 1990 Census via IPUMS. The column ξ_n reports the mean residual from a hedonic regression of rent on PUMA- and unit-level characteristics. The mean individual-level characteristics Black, Hispanic, College, Age, and Employed are calculated based on the characteristics of the household head. Child Present indicates the mean number of households with a member under 18 years of age. All means are calculated using household weights.

Figure 2.7. *Chicago Metropolitan Area*

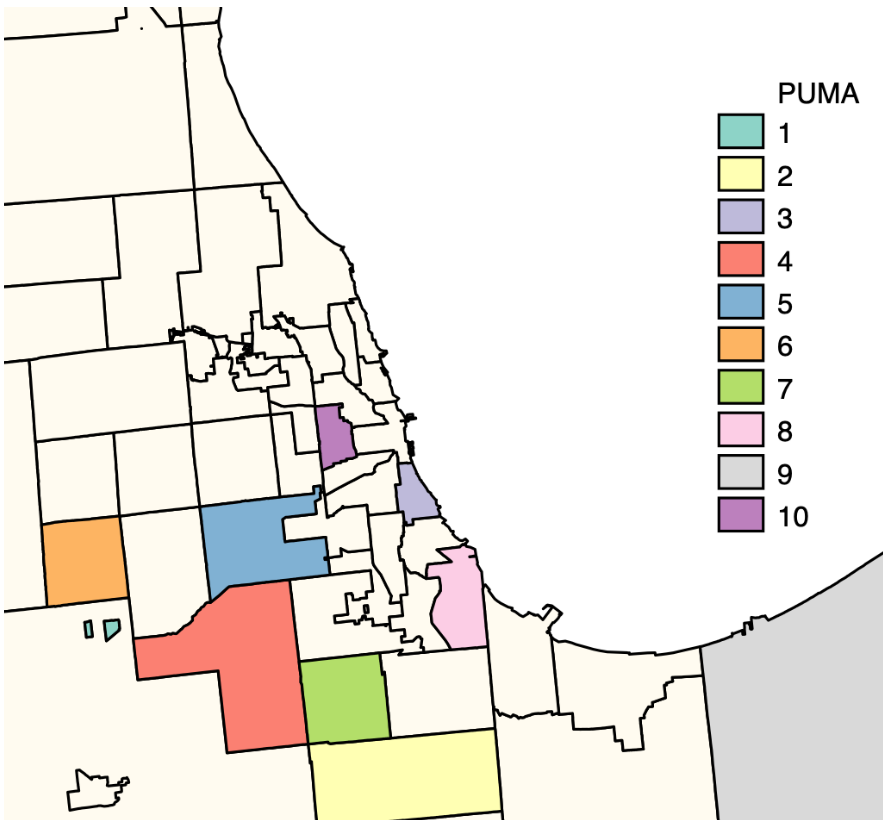

Notes: This figure shows the location of the ten 1990 PUMAs used in our simulation exercise. PUMA numbers correspond to those in Table 2.1.

Figure 2.8. *Simulated Unit Prices and Observed Rent*

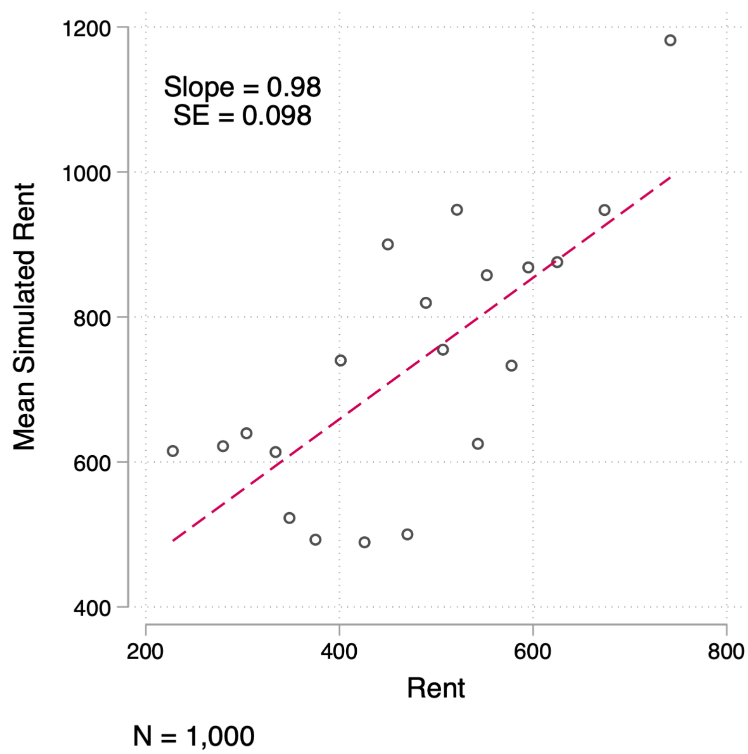

Notes: This figure shows the mean simulated price of housing units conditional on their observed rent in the IPUMS 1990 5% sample. We estimate a hedonic regression to impute the rent of owner-occupied housing units in the IPUMS sample.

sorting to PUMAs with a high college share, Black households sorting PUMAs with a higher share of Black households, and Hispanic households sorting to PUMAs with a higher share of Hispanic households.

While the preference parameters we use in our simulation are estimated to match similar data, there are several reasons why it is not ex-ante obvious that our simulation exercise would be able to replicate these features of the underlying data so closely. First, and most importantly, the preference parameters estimated by Bayer *et al.* (2007) are estimated via maximum-likelihood, taking prices as given and with no structure placed on how households match with housing units. By contrast, we simulate equilibrium prices via a tatonnement algorithm to find a one-to-one matching of households to housing units such

Figure 2.9. *Simulated and Observed Residential Sorting*

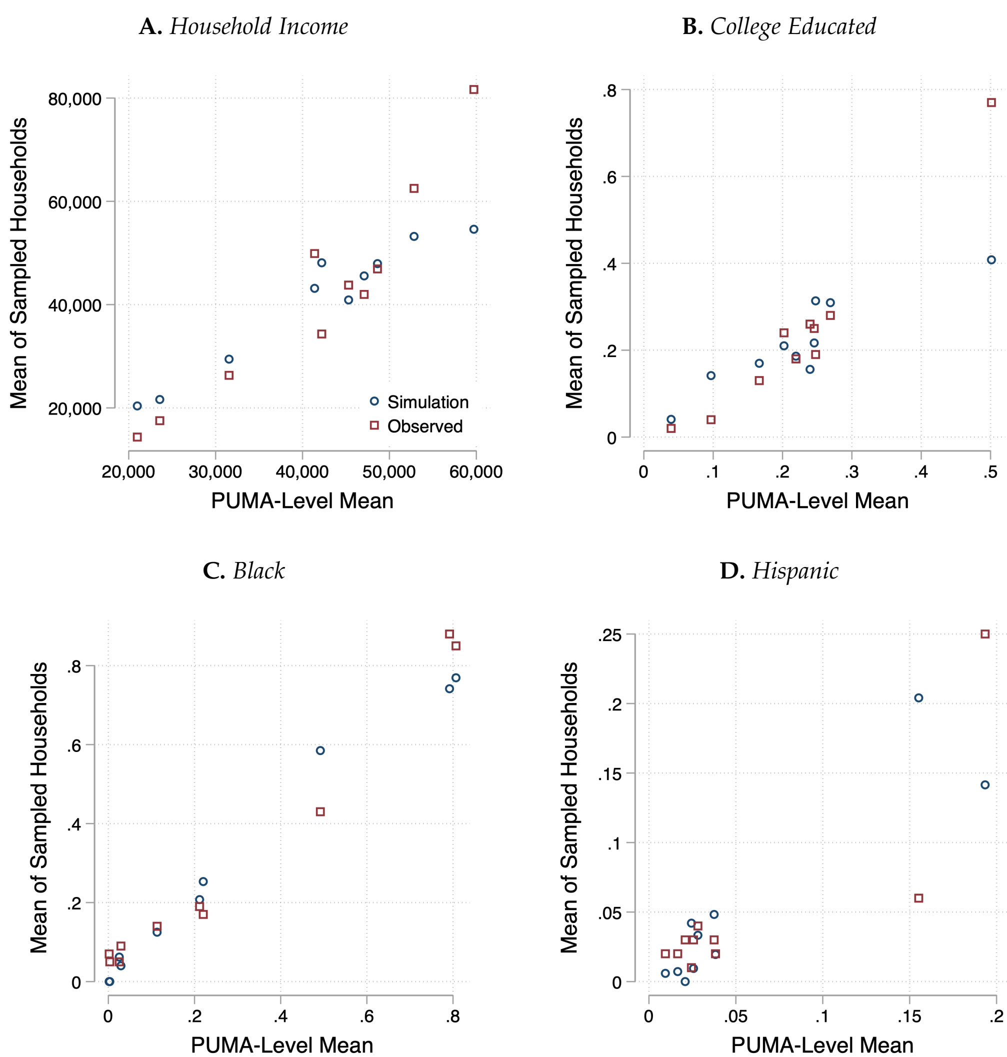

Notes: This figure shows the mean characteristics of households used in our simulation exercise conditional on the PUMA-level mean of those characteristics in the PUMAs they are assigned to in the initial simulated equilibrium. PUMA-level means are estimated using the IPUMS 1990 5% sample. Each panel represents the 1,000 households in our sample that are matched with a sampled housing unit in the initial equilibrium.

that no household wants to switch units.

Second, the preference parameters are estimated using data on households and housing units in San Francisco, while our simulation exercise uses data from Chicago. If there were unobserved heterogeneity in preferences across cities, applying the preferences of San Franciscans to the residential choices facing Chicagoans might have resulted in a simulated equilibrium that failed to match the observed patterns of residential choices.

Finally, we conduct our simulation exercise using only a small subset of the data. The fact that households in our sample have a limited choice set might have resulted in a different pattern of sorting than in the observed data. Overall, the similarity between our simulation results and the observed data gives us more confidence when applying our simulation results to interpret the descriptive facts on vacancy chains presented in the first part of this paper.

Price and Welfare Effects We now turn to our main objects of interest for this exercise – the simulated price and welfare effects of new housing supply. Figure 2.10 shows the distribution of these effects generated by 1000 simulations. Panel A shows the distribution of welfare effects as a percent of the initial level of aggregate welfare. The dashed line indicates the mean welfare effect of 0.5%. Given that we normalize the utility of the outside option to be zero and that, in each simulation, we add five housing units to the existing sample of 1,000, this implies an elasticity of the returns to living in the city of 1. Panel B shows the distribution of price effects, which approximately mirrors the distribution of welfare effects. The dashed line corresponds with the mean price effect of -.14%, which implies that the elasticity of the urban rent premium with respect to supply is 0.3.

Table 2.2 shows the incidence of these effects on different types of households. Columns two and four, respectively, show aggregate utility and mean prices for each group in our simulation's initial equilibrium, while columns three and five show the change in aggregate utility and average prices for each group. The welfare effects reported in columns one and two show that households that move and local households (i.e., those residing in the PUMA that receives new housing supply) experience the largest percent increases in welfare – 2.1%

Figure 2.10. *Simulated Welfare and Price Effects of New Housing*

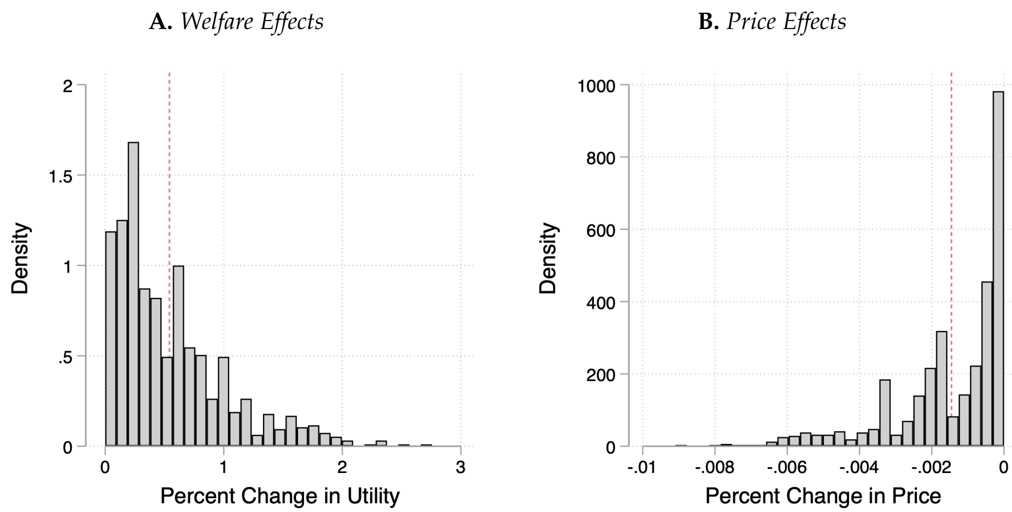

Notes: This figure shows the distribution of welfare and price effects from new housing calculated over 1000 simulations. The dashed red line indicates the mean effect size.

and 1.7% respectively. The fact that 43% of the aggregate welfare gains accrue to movers is attributable mostly to better matches. While local households experience relatively large percent increases in welfare, 83% of the aggregate welfare effect accrues non-locally, which suggests that the non-local price effects of new housing supply are economically important.

How does the variation in welfare and price effects documented in Figure 2.10 and Table 2.2 correlate with the vacancy chains that result from the addition of new housing? Figure 2.11 shows the mean simulated price and welfare effects of new housing conditional on the number of vacancies created in a neighborhood. The mean welfare effects are calculated for households that lived in a PUMA in which a given number of vacancies was created, while mean price effects are calculated for housing units in a PUMA with a given number of vacancies. Both price and welfare effects are strongly correlated with the number of vacancies. Households residing in PUMAs in which no vacancies were created experienced an increase in welfare of only .26%, while those in PUMAs with five or more vacancies experienced an average increase in welfare of 1.1%. Similarly, units in PUMAs with no vacancies saw a fall in prices of less than .001% while those with five or more vacancies saw

Table 2.2. *Price and Welfare Effects*

	# of Households	Utility	ΔUtility	Price	ΔPrice
All	1,330	2,612	14.1	547.5	-0.008
			(12.3)		(0.009)
Low-Income Households	665	1,119	4.9	528.1	-0.005
			(5.0)		(0.007)
Low-Income PUMAs	500	993	4.6	655.7	-0.007
			(4.6)		(0.008)
Movers	240	281	6.0	599.8	-0.010
	(10)	(18)	(3.7)	(14.8)	(0.012)
Stayers	1,095	2,331	8.2	536.3	-0.007
	(10)	(18)	(9.4)	(3.4)	(0.009)
Local	100	211	2.4	728.2	-0.018
		(64)	(3.3)	(433.7)	(0.027)
# of Simulations: 1,000					

Notes: This table reports means of initial utility and prices and mean welfare and price effects for different PUMAs and households over 1000 simulations. Standard errors are reported in parentheses.

a more than three-fold greater fall in prices of .0024%.

Demand Substitution and Vacancy Chains We also examine how our simulated price effects compare to those predicted by the underlying residential choice model. To do so, we compute the individual own- and cross-price partial derivatives of demand implied by the model. Following the notation introduced in the earlier section, we denote the own-price partial derivative of demand for neighborhood i by $\epsilon_i \equiv -\frac{\partial D_i}{\partial p_i}$ and the cross-price partial derivative of demand for neighborhood i with respect to the average price of units in j by $\gamma_{ij} \equiv \frac{\partial D_i}{\partial p_j}$. We denote the residential diversion ratio of j for i by $\lambda_{ji} \equiv \frac{\gamma_{ji}}{\epsilon_i}$.

To better understand what vacancies reveal about price and welfare effects, we estimate a series of regressions in which we regress the simulated change in prices on these substitution terms and the number of simulated vacancies. To make the estimates easier to interpret, we normalize all variables to be mean zero with unit variance. Table 2.3 reports these regression estimates. Column 1 reports estimates from a regression of the simulated change in prices in PUMA j on the direct and indirect substitution chains between PUMAs i and j,

Figure 2.11. *Simulated Price and Welfare Effects Conditional on Number of Vacancies*

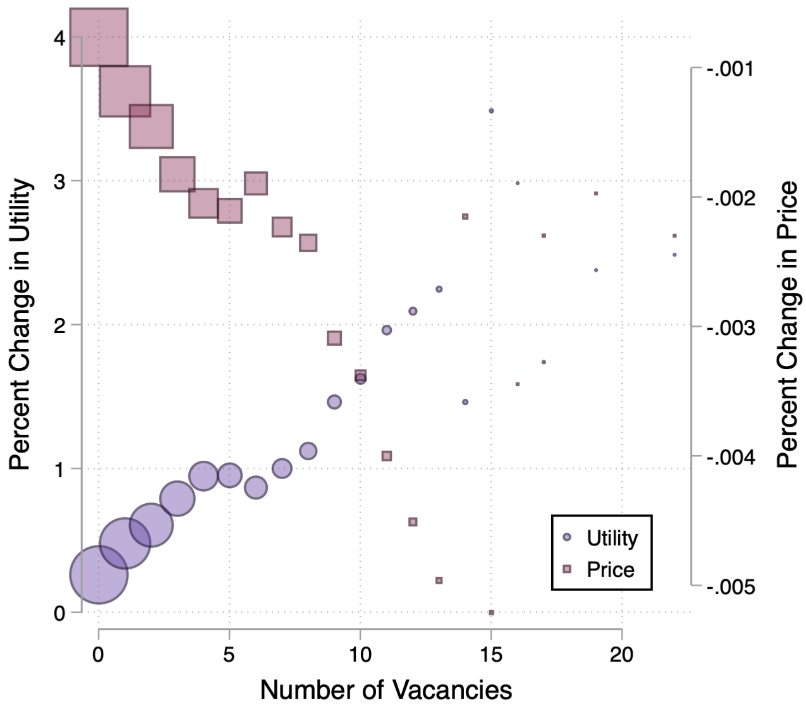

Notes: Welfare effects are calculated at the household level and are conditional on the number of vacancies in the household's PUMA of residence in the initial equilibrium. Price effects are calculated at the housing unit level. Point sizes are proportional to the number of observations.

where i is the PUMA in which new housing was added. We include indirect substitution effects that pass through up to two different PUMAs. We find that both direct and indirect substitution effects are highly significant predictors of the price effects of new housing supply – a one standard deviation increase in direct substitution between i and j leads to a .79 standard deviation increase in the magnitude of the price effect while a one standard deviation increase in indirect substitution mediated by one and two other neighborhoods leads, respectively, to .33 and .5 standard deviation increases in the magnitude of the price effect.

Column 2 adds the number of vacancies created in PUMA j as a regressor. We find that the number of vacancies strongly predicts variation in the simulated price effects, with a one standard deviation increase in vacancies (i.e. an increase of 2.6) predicting a .27 standard

Table 2.3. *Regression Estimates*

	(1) $\Delta Price_j$	(2) $\Delta Price_j$	(3) $\Delta Price_j$	(4) $\Delta Price_j$	(5) $\Delta Price_j$
$\epsilon_j^{-1}\lambda_{ji}$	-0.785***	-0.513***	-0.384***		
	(0.033)	(0.035)	(0.037)		
$\epsilon_j^{-1}\sum_k \lambda_{jk}\lambda_{ki}$	-0.332***	-0.321***	-0.298***		
	(0.019)	(0.019)	(0.019)		
$\epsilon_j^{-1}\sum_{k,\ell} \lambda_{jk}\lambda_{k\ell}\lambda_{\ell i}$	-0.497***	-0.407***	-0.293***		
	(0.036)	(0.035)	(0.037)		
$Vacancies_j$		-0.272***	0.228***	-0.282***	0.424***
		(0.014)	(0.055)	(0.010)	(0.054)
$Vacancies_j \times \epsilon_j^{-1}$			-0.532***		-0.718***
			(0.057)		(0.054)
Constant	0.107***	0.087***	0.060***	-0.056***	-0.056***
	(0.013)	(0.013)	(0.013)	(0.010)	(0.010)
R^2	0.077	0.110	0.118	0.075	0.091
Observations	10,000	10,000	10,000	10,000	10,000

Notes: All variables are transformed to be mean zero with unit variance. $\epsilon_i \equiv -\frac{\partial D_i}{\partial p_i}$ is the own-price partial derivative of demand for neighborhood i; $\gamma_{ij} \equiv \frac{\partial D_i}{\partial p_j}$ is the cross-price partial derivative of demand for neighborhood i with respect to the average price of units in j. $\lambda_{ji} \equiv \frac{\gamma_{ji}}{\epsilon_i}$ is the residential diversion ratio of j for i.

deviation increase in the magnitude of the price effect. Column 3 adds the number of vacancies interacted with the inverse own-price effect for PUMA j. This, too, is a highly significant predictor of variation in price effects and is associated with the largest variation in effect sizes of any regressor.

Columns 4 and 5 consider the predictive power of the number of vacancies alone. Notably, variation in the number of vacancies alone is just as predictive of variation in price effects as the direct and indirect substitution terms in column 1, explaining 7% of the variation in price effects. In column 5, adding the interaction between the number of vacancies and the inverse own-price elasticity of demand for PUMA j explains three-quarters of the variation explained by the full set of regressors in column 3.

Overall, the results reported in Table 2.3 show that the number of vacancies created by vacancy chains strongly predicts the incidence of price effects generated by new housing supply. While the direct and indirect substitution effects have independent predictive power, these terms are much harder to observe. With just ten PUMAs, forty-five distinct pairwise cross-price terms and ten distinct own-price terms are required to calculate these effects. In a realistic setting with many more PUMAs, the number of parameters to estimate quickly becomes infeasibly large. In contrast, the number of vacancies created in vacancy chains is easily observed, regardless of the number of neighborhoods, and is as predictive of variation in price effects as the substitution parameters.

2.6 Conclusion

The effect of new housing supply in one submarket on housing costs in other submarkets depends crucially on how residential mobility connects these submarkets. The impact that the large increase in suburban housing supply over the past four decades has had on the costs facing low-income renters in the urban core thus depends on whether the chains of moves it initiated reached low-income high-density neighborhoods. Our results show that they did not – instead, residential vacancy chains initiated by new low-density suburban single family housing end quickly, before they can reach the urban neighborhoods in which residents are most exposed to rising housing costs. This descriptive feature of vacancy chains, when viewed in light of our simulation results, suggests that the non-local price effects of new housing supply are concentrated in nearby submarkets and that the incidence of the benefits of additional housing therefore depends crucially on what kind of housing is built and where.

These results have important implications for housing policy that seeks to increase housing affordability. While many supply advocates argue that increasing supply of any kind will be effective at decreasing housing costs for all, this paper suggests that a more targeted approach is required if policymakers want to reduce costs in the least affordable neighborhoods or for the most rent-burdened households. Our results also suggest that a more targeted approach can be effectively guided by the distribution of vacancies created by new supply in a given submarket. Our simulation results show that the distribution of these vacancies is as predictive of variation in price effects as the cross-neighborhood substitution effects derived from individual demand elasticities. While housing policy will have to be guided by the costs of construction in different neighborhoods, as well as potential effects on local amenities, the observed number of vacancies connected to different kinds of new housing can help policymakers evaluate the non-local price effects of a given policy.

Chapter 3

Schools or Suitcases? Optimal Urban Policy and Economic Opportunity in Greater Boston[1]

3.1 Introduction

"Would [you] like to move to a community in Greater Boston with good schools and low crime?". So reads a 2018 letter from the Boston Housing Authority (BHA) to thousands of low-income Bostonians (Graham, 2018). The letter was advertising a new pilot program–Expanding Choice in Housing Opportunities (ECHO)–designed to help low-income families with young children move to neighborhoods outside Boston City.[2] The unintentionally

[1] I thank Dionissi Aliprantis, Chris Avery, Nathaniel Baum-Snow, Peter Blair, Pat Bayer, Raj Chetty, Ingrid Gould Ellen, Ashvin Gandhi, Mike Gilraine, Ed Glaeser, Gordon Hanson, Nathaniel Hendren, Gabriel Kreindler, Robin Lee, Katherine O'Regan, Bruce Sacerdote, Mark Shepard, Stefanie Stantcheva, Winnie van Dijk, Daniel Waldinger, and especially Matthew Weinzierl for helpful discussions and feedback. Valentine Gilbert provided helpful contributions to an early version of the project. I thank Stephanie Bailey, Sari Kerr, Shawn Ratcliff, and Shital Sharma for help with the Census data. Oleg Kolev has provided outstanding programming assistance. This research was supported by a Stone Research Grant from Harvard Kennedy School's James M. and Cathleen D. Stone Program in Wealth Distribution, Inequality, and Social Policy.

[2] ECHO is among a growing number of programs that help low-income families relocate to neighborhoods with less poverty and greater economic opportunities. Other residential mobility programs include the Housing and Urban Development's Moving to Opportunity demonstration program and Small Area Fair Market programs, as well as Seattle's Creating Moves to Opportunity demonstration program.

offensive advertisement insinuated to some residents that living in Boston City meant one's children must grow up in a high-crime environment with failing schools, and it eventually led the BHA Administrator to call its own advertisement "an insult to the residents of Boston" (Szaniszlo, 2018).

A core tension underlying the response to the BHA letter is how policymakers should address concentrated urban poverty. Should policy makers focus on investing in disadvantaged urban neighborhoods and their schools, or should they prioritize helping low-income families leave these neighborhoods? Policies targeted at disadvantaged neighborhoods and their schools may improve outcomes for incumbent residents as well as new residents who are attracted by these investments. The effectiveness of such policies may be muted, however, if there are strong complementarities between neighborhood prosperity and investment levels. By contrast, policies facilitating moves out of disadvantaged neighborhoods offer large reductions in neighborhood poverty for movers but offer few benefits for residents with strong attachments to their home neighborhood. General equilibrium migration responses among households across the income distribution may reinforce or undo the intended effects of both targeted investment policies and residential mobility programs.

In this chapter, I revisit the debate between investing in place and investing in people in the context of Greater Boston's 142 school districts. Specifically, I show how optimal urban economic policy varies with three sets of parameters for which there is contentious popular or academic debate: parents' preferences, the technology governing children's upward economic mobility, and social preferences over children and their parents. I develop and estimate a quantitative spatial equilibrium model with district-wide endogenous amenities, school quality, housing prices, and metro-wide endogenous skill-specific wages. Parents and non-parents in the model differ in their level of educational attainment, which shapes the provision of the district- and metro-wide endogenous objects just noted. I embed the quantitative model in a social planning problem, wherein the social planner chooses among district-wide levels of school spending and housing assistance, as well as metro-wide income transfers. I quantify the model using administrative US Census Bureau data on the residential histories of nearly all ACS respondents living in Greater Boston between 2000 and 2019.

With redistributive preferences, the social planner trades off between direct and indirect redistribution. Direct redistribution involves taxing college-educated households who self-segregate in suburban districts to spend on schools in disadvantaged inner-city districts. In contrast, indirect redistribution integrates college- and non-college-educated households across school districts, potentially raising the endogenous provision of district-wide economic mobility and amenities for non-college-educated households. The tension between direct and indirect redistribution results from households' threat to leave Greater Boston entirely. By lowering the average level of district-wide amenities and school quality for college-educated households, residential integration reduces the taxes the planner can levy on college-educated households before marginal residents leave Greater Boston entirely.

I conduct my analysis in three steps. In the first step, I obtain a baseline set of technological parameters governing children's upward economic mobility by calibrating a nested CES human capital production function to district-level estimates of children's economic outcomes in adulthood aggregated from Chetty *et al.* (2018). To construct model-implied levels of children's upward economic mobility, I use historical school spending data from the National Center for Education Statistics and historical population data from the 1990 Decennial Census. I then use the calibrated technological parameters governing children's upward economic mobility to predict contemporary district- and parental-education-specific levels of children's upward economic mobility in Greater Boston.

In the second step, I estimate households' preferences over predicted district-wide measures of upward economic mobility, consumption, and endogenous and exogenous district amenities. Because unobserved district-wide characteristics are invariably correlated with predicted and observed district-wide characteristics, I instrument for price using BLP-style instruments. I additionally exploit my administrative Census data to obtain variation in parents' and non-parents' district choices that identifies preferences over children's expected upward economic mobility and endogenous district-wide amenities. This second step is ongoing.

In the third and final step, I invert the school district choice model to characterize optimal urban economic policy in Greater Boston. I characterize optimal policy by efficiently searching over feasible district-level population distributions and selecting the equilibrium that maximizes social welfare. By conducting this analysis for different combinations of the social planner's preferences and technological parameters governing children's upward economic mobility I show how optimal policy is shaped by social preferences as well as *how* place impacts children's upward economic mobility.

While a complete analysis of the social planner's problem is in progress, in this chapter I report preliminary and informal findings from a simple two-district version of the model. In particular, I present optimal policy simulations varying non-college-educated households' preferences over endogenous district-wide amenities. I show that as college-educated households' preferences for endogenous district-wide amenities increases vis-à-vis non-college-educated households' preferences, the social planner prioritizes direct redistribution over indirect redistribution. The intuition behind this result is that as college-educated households' preference for endogenous amenities increases, the more they are willing to pay to live in a school district with a higher share of other college-educated households. The social planner can then raise metro-wide taxes on college-educated households who self-segregate and spend the resulting funds on targeted housing assistance and school financing in the poorer district. Importantly, though, these results rely on school spending and peer composition being substitutes in the production of children's upward economic mobility. When school spending and peer composition are instead complementary in production, the additional funds raised from taxing college-educated households cannot be effectively spent in the poorer school district. Future research based on this chapter will fully explore how parental and social preferences as well as technological parameters mediate the optimal design of urban economic policy to promote children's economic mobility.

Relation to Literature Debates over whether to prioritize investing in disadvantaged neighborhoods or whether to prioritize helping residents leave these neighborhoods erupted in the mid-1960's amid rising US inner-city poverty and racially charged riots (Harrison, 1974). Kain and Persky (1969) arguing for what they called "ghetto dispersal" as the most sustainable solution to urban poverty and racial segregation captures one side of these early debates. Labrie (1970) championing "rebuilding the central city ghetto" in the inaugural volume of The Review of Black Political Economy captures another side. Sixty years later, the conversation continues (Ellen and Steil, 2019; Imbroscio, 2023; Ellen, 2023), though economic analyses have broadened to consider class and income segregation in addition to racial segregation,[3] and are primarily empirical (Cutler and Glaeser, 1997; Lundberg and Startz, 1998; Oreopoulos, 2012; Fryer Jr and Katz, 2013; Chyn and Daruich, 2022). Moreover, most contemporary empirical studies of neighborhood- or residential mobility-based interventions focus on understanding how policies applied in isolation impact their intended recipients.[4] Relating well-identified treatment effects to interventions' costs in a unified framework then shows where the government's marginal dollar most effectively raises social welfare (Hendren and Sprung-Keyser, 2020). My quantitative spatial general equilibrium framework complements this important line of research.

My general equilibrium framework incorporates insights from seminal theoretical neighborhood sorting models studying inequality and residential segregation (De Bartolome, 1990; Benabou, 1996; Durlauf, 1996; Fernandez and Rogerson, 1996, 1998; Nechyba, 2000), and empirical residential choice models estimating households' preferences over home and neighborhood characteristics (Brock and Durlauf, 2001; Bayer *et al.*, 2007).[5] It also closely relates to a growing quantitative macroeconomics literature studying parental location deci-

[3]See Anderson (2010) and Shelby (2014) for a contemporary debate over the philosophical and empirical merits of policies targeting residential racial integration.

[4]A sample of the voluminous literature examining neighborhood- or residential mobility-based interventions includes Kling *et al.* (2007), Chetty *et al.* (2016), Dobbie and Fryer Jr (2011), Abdulkadiroğlu *et al.* (2011), Bergman *et al.* (2023). Jackson *et al.* (2015), Lafortune *et al.* (2018), Biasi (2023), Jackson and Mackevicius (2024), and Handel and Hanushek (2024) study the impact of school financing on children's outcomes, and Oreopoulos (2003), Durlauf (2004), Chyn (2018), Chetty and Hendren (2018), Laliberté (2021), and Chyn and Katz (2021) study the causal impact of residential environments on children's upward economic mobility.

[5]Additional related and recent residential choice models include Galiani *et al.* (2015), Caetano (2019), Davis *et al.* (2021), Couture and Handbury (2020), Almagro *et al.* (2023), and Campos and Kearns (2024). These papers estimate preferences over neighborhood characteristics, but do not consider endogenous children's upward economic mobility nor optimal urban policy.

sions, children's human capital development or school quality, and residential segregation (Ferreyra, 2007; Chyn and Daruich, 2022; Zheng and Graham, 2022; Eckert and Kleineberg, 2024; Gilraine *et al.*, 2023; Agostinelli *et al.*, 2024).[6] While the quantitative macroeconomics literature considers counterfactual welfare gains from select place-based policies, I embed my quantitative model in a social planning problem and search over spatial equilibria to understand how optimal policy varies with social preferences and the technology governing children's endogenous upward economic mobility. My approach facilitates nuanced comparisons among a variety of potentially coexisting place-based policies.

This chapter naturally relates to a literature on optimal place-based policies and taxation (Glaeser and Gottlieb, 2008; Fajgelbaum and Gaubert, 2020; Gaubert *et al.*, 2021; Rossi-Hansberg *et al.*, 2019).[7] Fajgelbaum and Gaubert (2020) and Rossi-Hansberg *et al.* (2019) characterize first-best optimal policy in closed-economy spatial equilibrium models with local spillovers across worker types. I instead characterize optimal redistributive policy with local spillovers in an open-city framework, where households' threats to leave Greater Boston constrain the planner's redistributive goals.[8] Gaubert *et al.* (2021) do characterize optimal place-based redistribution but use a perturbation argument and focus on correcting for fiscal externalities instead of local spillovers. I characterize optimal redistributive policy in the presence of local spillovers using an equilibrium search algorithm. Equilibrium selection can yield large welfare gains relative to local price distortions when local spillovers

[6]Other related quantitative spatial equilibrium papers that consider children's upward mobility include Fogli and Guerrieri (2019) who examine how residential segregation has contributed to income inequality since the 1980s. Agostinelli *et al.* (2020) examine how parenting styles interact with spatial determinants of upward economic mobility. Aliprantis and Carroll (2018) and Gregory *et al.* (2023) focus on racial segregation and neighborhood spillovers. Anstreicher (2024) takes a partial equilibrium approach to understanding how cross-state migration influences economic mobility.

[7]Other recent research characterizing optimal policy in quantitative settings include Blouri and Ehrlich (2020) and Henkel *et al.* (2021). Blouri and Ehrlich (2020) examine EU regional transfers while Henkel *et al.* (2021) examine fiscal transfers in Germany. Both papers take a direct approach to ascertaining optimal policy, searching over policy instruments themselves. Neither paper jointly optimizes over the considered policy instruments, however. See also Kline and Moretti (2014) for a review of place-based policies.

[8]Fajgelbaum and Gaubert (2020) and Rossi-Hansberg *et al.* (2019)'s analytical results characterize optimal policy when the social planner's objective is globally concave. I place no restriction on the shape of planner's objective function and focus on equilibrium selection in the presence of potentially strong spillovers, which makes a globally concave planning problem unlikely.

are large (Kline, 2010).

3.2 Model

In this section, I develop a quantitative spatial general equilibrium model that captures the relevant equilibrium forces that matter for determining optimal urban policy when children's upward economic mobility is endogenous. Despite the myriad general equilibrium interactions governing the spatial distribution of workers and economic opportunity, the model remains highly tractable.[9] This tractability permits optimal policy analysis and a clear exposition of the core trade-offs the social planner faces when choosing across policy instruments.

In a static framework, I model a single city embedded within a larger economy. The city consists of N distinct school districts, indexed by n, and heterogeneous workers, indexed by i. School districts differ ex-ante with respect to their natural amenities, their housing supply elasticities, and the number of existing housing units. School districts further endogenously supply amenities and "district effects" that influence children's expected human capital accumulation. Workers differ with respect to their educational level, their parental status, and other socio-demographic characteristics that together determine their preferences for district amenities and consumption, and—if parents—the technology governing their child's expected human capital accumulation. Workers who choose to live outside the city obtain a type-specific fixed mean utility.

Production occurs at the city level, where identical firms imperfectly combine college- and non-college-educated workers to produce a homogeneous good that is traded at constant prices with the larger economy. Firms are perfectly competitive and pay workers their marginal product. A worker's education level therefore determines her pre-tax income.[10]

I embed this quantitative spatial equilibrium model into a social planning problem. In the

[9]I use the term workers to refer to the adult in each household.

[10]Given that a worker's education level determines her income, I refer to workers by their income level and educational attainment interchangeably.

social planning problem, a city-wide social planner jointly chooses school district-level policy instruments—namely housing assistance and school spending—and a place-blind income tax schedule to maximize social welfare, defined as the weighted sum of the prospective residents' expected utility. Redistribution by the social planner is possible despite the open city setup. Workers have idiosyncratic preferences for living in each neighborhood within the city as well as for exercising their outside option of leaving the city. Inframarginal city residents thus obtain a rent from living in the city that the social planner can attempt to extract for redistributive purposes.[11]

3.2.1 Workers

Workers in the economy are indexed by i. Workers differ with respect to their highest level of educational attainment (college = ω^H, non-college = ω^L), their parental status (parent = a^p, non-parent = a^{np}), and other socio-demographic characteristics that determine their preferences for neighborhood amenities, consumption, and—if parents—their child's expected human capital accumulation, z. I define θ as the triple defining a workers' type, $\theta \equiv (\omega, a, z)$. Parents have one child and value their child's expected human capital, which, as described below, is a function of both household consumption and the school district the child grows up in. The set of θ-type workers in the economy is denoted by I^θ, with $\int I^\theta dz = I^{\omega,a}$ and $\int I^{\omega,a} da = I^\omega$. Measures of these sets are in turn denoted $|I^\theta|$, $|I^{\omega,\alpha}|$, and $|I^\omega|$.

3.2.2 Production

I follow Diamond (2016) and assume the city is populated with many homogeneous firms, indexed by d. These firms produce a single tradable good by employing non-college-educated workers (L_d) and college-educated workers (H_d) according to the production

[11]This setup contrasts with canonical open city models where workers of the same type have homogeneous preferences for neighborhood and city amenities (Brueckner, 1987). With homogeneous preferences, spatial equilibrium ensures utility from residing in the city equals the value of the workers' outside option. City size adjusts to ensure this equality maintains.

technology,[12]

$$Y_d = \left[(A_L L_d)^{\varrho} + (A_H H_d)^{\varrho}\right]^{1/\varrho}.$$

Thus, each firm combines workers of different education groups as imperfect substitutes into production with a constant elasticity of substitution, given by $1/(1-\varrho)$. The productivity of college- and non-college-educated workers is exogenous and measured by A_H and A_L, respectively. The labor market is perfectly competitive yielding the following education-level wage profile:

$$w_H = A_H^{\varrho}\left[A_H^{\varrho} + A_L^{\rho}(L_d/H_d)^{\varrho}\right]^{(1-\varrho)/\varrho},$$
$$w_L = A_L^{\varrho}\left[A_L^{\varrho} + A_H^{\rho}(H_d/L_d)^{\varrho}\right]^{(1-\varrho)/\varrho}.$$

Firms' labor demand corresponds to city-level aggregate labor demand as firms have identical and constant returns to scale production technologies.

3.2.3 Housing

There are N school districts in the city indexed by n. Each school district has an existing housing stock $\mathcal{A}_n$. Housing is produced using land and construction materials. Equilibrium in the housing market yields local prices, r_n, and housing unit quality is constant within school districts.

The cost of land is a function of the aggregate district-wide demand for housing units. Households each demand exactly one unit of housing, so aggregate district demand equals the mass of workers residing in the district.[13] The district-level housing unit supply equations are,

$$r_n = \kappa_n \left(HD_n \cdot \mathcal{A}_n^{-1}\right)^{\epsilon_n},$$

[12]All workers of the same education level are endowed with the same effective productivity and are perfectly substitutable in production, irrespective of their parental status.

[13]I follow Couture *et al.* (2023) and model housing as a necessity. This contrasts with much of the economic geography literature that models housing consumption assuming Cobb-Douglas preferences across housing and non-housing consumption. By modeling housing as a necessity, I capture the fact that many households are "priced-out" of expensive school districts.

$$HD_n = H_n + L_n,$$

where $HD_n \cdot \mathcal{A}_n^{-1}$ is the aggregate housing demand in district n scaled by the district's existing housing stock. H_n and L_n are the mass of college- and non-college-educated workers residing in district n, respectively. The district-specific elasticity of rent with respect to housing demand is given by ϵ_n. κ_n shifts the relative cost of building new housing in each district. ϵ_n and κ_n capture both topographical and institutional features of each district's housing environment, such as minimum lot size requirements.

Workers residing in the city own a fraction of all land and sell this land to perfectly competitive developers. These ownership shares can differ by workers' educational status, and are determined outside of the model. I denote the effective share of total rents in the city that each $\omega-$type worker receives as Π^{ω}, with $\Pi^{\omega_L}\sum_n L_n + \Pi^{\omega_H}\sum_n H_n \leq 1$ and $\Pi^{\omega} \geq 0\,\forall\omega$. The case where $\Pi^{\omega} = 0\;\forall\,\omega$ corresponds to the absentee-landlord assumption often invoked in the urban economics literature to close open-city models.

3.2.4 Neighborhood Amenity Supply

School districts differ in the amenities A_n^i they offer each worker, i. These amenities differ across districts due to exogenous factors that are determined outside of this model, X_n^i (e.g., access to employment opportunities, proximity to CBD, natural amenities like hills or coastlines, and worker-level factors like proximity to family), as well as due to endogenous factors, $\bar{Q}_n\left(\{H_n, L_n\}_n\right)$ that are determined by the educational mix of workers across districts in the city:

$$A_n^i = \bar{Q}_n\left(\{H_n, L_n\}_n\right) \cdot \exp\left(X_n^i\right).$$

$\bar{Q}_n$ is an aggregator of the distribution of workers across the city from the perspective of district n. As such, a workers' valuation for a district's amenities may depend not only on the composition of workers in their own district, but also on the composition of workers in nearby districts.

3.2.5 Children's Human Capital Accumulation

Children's human capital accumulation is a Cobb-Douglas function of household consumption and a "district effect" term:[14]

$$HC_n^i = \left(C_n^{\omega,p}\right)^{\alpha_c^i} \cdot \mathcal{N}_n\left(S_n, \bar{q}_n^{\omega}, \Xi_n^{\omega}; \omega\right)^{\left(1-\alpha_c^i\right)}.$$

The importance of consumption, $C_n^{\omega,p}$, vis-à-vis district effects, $\mathcal{N}_n\left(S_n, \bar{q}_n^{\omega}; \Xi_n^{\omega}, \omega\right)$, in determining a child's upward mobility is governed by the Cobb-Douglas share parameter α_c^i. This share parameter can vary across parents, hence the i superscript.

School districts affect children's long-term outcomes through per-student spending on districts' schools, S_n, the composition and pre-tax wages of workers in the district, $\bar{q}_n^{\omega}$, as well as other environmental factors that are determined outside of the model Ξ_n^{ω}.[15] I refer to $\bar{q}_n^{\omega}$ as peer effects. These components of the district effect term are constant across parents of the same education level. Districts' schools are financed both locally and by the social planner. I denote s_n as the component of local school spending financed by the the social planner, so that

$$S_n = s_n + \frac{H_n + L_n}{I^{\omega_H,p} \cdot \tilde{\pi}_n^{\omega_H} + I^{\omega_L,p} \cdot \tilde{\pi}_n^{\omega_L}} \tau_n^s r_n,$$

where $\tilde{\pi}_n^{\omega_H}$ and $\tilde{\pi}_n^{\omega_L}$ are the shares of college- and non-college-educated parents residing in district n and τ_n^s is the exogenously set property tax rate of district n.[16]

District effects are determined by combining spending on districts' schools and the education mix of workers in the district as imperfect substitutes into production with a constant elasticity of substitution. The exogenous environmental factors contribute to district effects multiplicatively. The parameters governing the shape of children's human capital accumulation can vary according to the educational attainment of the child's parent, and

[14] My use of the term "district effect" mimics the term "neighborhood effect" which comes from a long line of empirical research on the impacts of neighborhood environment on children's later-life outcomes (Chyn and Katz, 2021).

[15] Going forward, I suppress the arguments of $\mathcal{N}_n\left(S_n, \bar{q}_n, \Xi_n^{\omega}; \omega\right)$ when their values are clear.

[16] I will show in the resulting paper that the core tradeoffs in the model are maintained when the property tax rate is politically determined.

thus have the superscript ω, but are constant within education groups:[17]

$$\mathcal{N}_n^{\omega} = (\gamma_s^{\omega} S_n^{\rho^{\omega}} + (1-\gamma_s^{\omega})(\bar{q}_n^{\omega})^{\rho^{\omega}})^{1/\rho^{\omega}} \cdot \exp(\Xi_n^{\omega}).$$

ρ^{ω} governs the degree of complementarity between peer effects and school spending in the production of children's human capital. As $\rho^{\omega} \to 1$, peer effects and school spending are perfect substitutes in the production of human capital, and as $\rho^{\omega} \to -\infty$, peer effects and school spending become perfect compliments in the production of human capital. γ_s governs the relative efficiency of school spending and peer effects in determining district effects.

The education mix of workers in a district enters into the district effect term through $\bar{q}_n$. This function captures the non-pecuniary effects residential educational composition has on children's human capital production at the school-district level. I follow Benabou (1996) and require $\bar{q}_n$ to be increasing in the local distribution of pre-tax income.[18] In particular, I specify $\bar{q}_n^{\omega}$ to combine pre-tax income of college- and non-college-educated workers with constant elasticity of substitution:

$$\bar{q}_n^{\omega} = \left(\frac{H_n}{H_n+L_n} w_H^{\eta^{\omega}} + \frac{L_n}{H_n+L_n} w_L^{\eta^{\omega}}\right)^{\frac{1}{\eta^{\omega}}}.$$

This specification is flexible enough to capture different responses to mean-preserving spreads in the distribution of pre-tax wages. When η is equal to one, $\bar{q}_n$ simply equals the mean pre-tax wage level in the district, whereas when $\eta^{\omega} \to -\infty$, $\bar{q}_n^{\omega}$ converges to the pre-tax wage level of non-college-educated workers, and when $\eta^{\omega} \to \infty$, $\bar{q}_n^{\omega}$ converges to the pre-tax wage level of college-educated workers.

[17]For example, it may be that school spending levels are less important for high-income households, as parents of these households can more easily compensate for low school spending in resource-poor districts through tutoring or attending private schools.

[18]I define $\bar{q}_n$ to be a function of *pre-tax* income to focus on the non-fiscal channels of peer effects. Pre-tax income in this context can proxy for labor market attachment and associated social determinants of upward mobility, such as the presence of role models, the presence of violent crime, and social network effects.

3.2.6 Parents' Problem

Each parent observes their idiosyncratic utility, ε_n^i, from residing in district n and from living outside of the city ($n = oo$). The ε_n^i are drawn i.i.d. from a Type I Extreme Value Distribution, with a common variance parameter σ normalized to 1. Upon observing these idiosyncratic utilities, parents simultaneously choose whether to live in the city, and if they do decide to live in the city, which district to reside in. If a parent decides to live outside the city, they obtain a mean reservation utility that is normalized to zero as well as their idiosyncratic utility component from living outside of the city.[19] Parents have Cobb-Douglas preferences over the tradable consumption good, $C_n^{\omega,a}$, district amenities, A_n^i, and the expected human capital accumulation of their child, HC_n^i.

Parents maximize their utility subject to their budget constraint by choosing which school district to live in (or to live outside of the city),[20]

$$\max_n \quad \beta_c^i \log(C_n^{\omega,a}) + \beta_H^i \log\left(HC_n^i\right) + \beta_A^i \log\left(A_n^i\right) + \varepsilon_n^i,$$

subject to,

$$C_n^{\omega,a} = w_\omega - T(w_\omega) - r_n(1 + \tau_n^s - \tau_n 1\{\omega = \omega_L\}) + \Pi^\omega R \geq 0.$$

Parents' budget constraints are a function of the government's policy instruments: τ_n is housing assistance offered to non-college educated parents living in district n;[21] τ_n^s is the local property tax rate in district n, which, recall, is fixed exogenously; and $T(w_\omega)$ is a non-linear place-blind income tax schedule that taxes workers according to their income. Recall also that Π^ω is the effective share of total city rents (R) an ω-type worker receives. The positive consumption constraint requires that parents locate in a school district they can

[19]Since one cannot identify the relative levels of workers' utility, there is no loss in information by assuming all workers' outside option mean utilities are zero.

[20]I assume that $\beta_c^i \log(C_{oo}^{\omega,a}) + \beta_H^i \log\left(HC_{oo}^i\right) + \beta_A^i \log\left(A_{oo}^i\right) = 0 \, \forall \, i$, where oo refers to the outside option of residing outside of the city.

[21]The precise form of housing assistance is determined ex-post, depending on the desired distribution of households across the city.

afford. As such, district choice sets may differ across parents of different incomes.

Substituting the budget constraint into the parents' utility function and rearranging slightly yields the following conditional indirect utility for worker i living in school district n:

$$
\begin{aligned}
V_n^i = {} & \bar{\beta}_c^i \log\left(w_\omega - T(w_\omega) - r_n(1 + \tau_n^s - \tau_n 1_{\{\omega=\omega_L\}}) + \Pi^\omega R\right) \\
& + \bar{\beta}_{\mathcal{N}}^i \log(\mathcal{N}_n^\omega) + \beta_A^i \log(A_i^i) + \varepsilon_n^i, \\
\equiv {} & v_n^i + \varepsilon_n^i,
\end{aligned}
$$

where $\bar{\beta}_c^i \equiv \beta_c^i + \alpha_c^i \cdot \beta_H^i$ and $\bar{\beta}_{\mathcal{N}}^i \equiv \beta_H^i \cdot \alpha_{\mathcal{N}}^i$.

This setup is conditional logit. The probability that a parent of education level ω lives in neighobrhood n is therefore,

$$
\tilde{\pi}_n^\omega \equiv \frac{1}{|I^{\omega,p}|} \int_{i \in I^{\omega,p}} \frac{\exp\left(v_n^i\right)}{1 + \sum\limits_{m \in N^{\omega,p}} \exp\left(v_m^i\right)},
$$

where $N^{\omega,p}$ is the neighborhood choice set for an ω-type parent.

3.2.7 Non-Parents' Problem

The problem for non-parents' is identical to that of parents, aside from the fact that parents are ineligible for housing assistance, $\tau_n = 0 \,\forall\, \omega$, and have no child they are altruistic toward, which I model by assuming $\beta_H^i = 0 \,\forall\, i \in I^{\omega,np}$.

3.2.8 Social Planner's Problem

Policy Equilibrium To discuss the social planner's problem, I must first present an equilibrium concept. A *policy equilibrium* $\mathcal{P} \equiv \{\{H_n, L_n\}_n, \mathcal{T}, \boldsymbol{r}, \boldsymbol{w}\}$ consists of an allocation of college- and non-college-educated workers to neighborhoods, $\{H_n, L_n\}_n$, a tax system $\mathcal{T} \equiv \{T, \tau_n, s_n\}$, a vector of housing prices $\boldsymbol{r} = \{r_n\}_n$, and a couple of wages $\boldsymbol{w} = \{w_L, w_H\}$ such that the following conditions jointly hold:

1. No worker wants to switch neighborhoods or leave/enter the city;

2. Consumption is weakly greater for college-educated workers in each neighborhood: $C_n^{\omega^H,a} \geq C_n^{\omega^L,a}, \forall\, n, a;$[22]

3. The (i) production, (ii) housing, (iii) neighborhood amenity, and (iv) neighborhood upward mobility markets clear;

4. District-level school spending, consumption, and housing assistance are weakly positive: $s_n, \tau_n, C_n^{\omega,a} \geq 0 \,\forall\, n\,,\, \forall\, \omega, a;$

5. The tax system $\mathcal{T}$ satisfies the government budget constraint:

$$H \cdot T(w_H) + L \cdot T(w_L) \geq \sum_{n \in N} \left(I^{\omega_L,p} \cdot \tilde{\pi}_n^{\omega_L} \cdot r_n \tau_n + \left(I^{\omega_L,p} \cdot \tilde{\pi}_i^{\omega_L} + I^{\omega_H,p} \cdot \tilde{\pi}_i^{\omega_H} \right) \cdot s_n \right)$$

Social Planner's Problem The social planner's problem is to choose the policy equilibrium, $\mathcal{P} \equiv \{\{H_n, L_n\}_n, \mathcal{T}, \boldsymbol{r}, \boldsymbol{w}\}$, that maximizes a weighted sum of workers' expected utility from residing in the city:

$$\max_{\mathcal{P}} \int_i \lambda^i \frac{\mathcal{V}^i(\mathcal{P}) - \mathcal{V}^i(\mathcal{P}^0)}{\bar{\beta}_c^i},$$

where $\mathcal{V}^i(\mathcal{P}) \equiv \mathbb{E}\left[\max_{n \in \mathcal{N}_{\omega(i),a(i)}} \left\{V_n^i(\mathcal{P})\right\}\right]$ represents expected utility for worker i under the policy equilibrium environment, $\mathcal{P}$, with the expectation taken over workers' idiosyncratic school district utility shocks.[23]

In anticipation of comparisons in optimal policy across different technological and preference parameter values, the measure of social welfare shifts workers' expected utilities

[22] I incorporate this constraint into the definition of a policy equilibrium to ensure that college-educated workers are not incentivized to obtain the same income as non-college-educated workers. Since, conditional on living in a school district, the only difference in the arguments entering workers' indirect utility functions is the level of consumption, constraining college-educated workers to have higher income than non-college-educated workers ensures these workers are not incentivized to masquerade as non-college-educated workers. I further implicitly assume that the utility cost of non-college-educated workers mimicking college-educated workers is arbitrarily high. In sum, these assumptions imply that I am constraining the equilibrium concept to 'separating equilibria', wherein neither college- nor non-college-educated workers have an incentive to masquerade as a worker of the other education level.

[23] I assume that workers' idiosyncratic school district preferences are redrawn for each policy equilibrium. Since workers' idiosyncratic preferences are residual of an arbitrarily detailed set of worker-specific district-wide amenities, this assumption is innocuous.

by the expected utility they would each obtain in a given status-quo policy equilibrium: $\mathcal{V}^i(\mathcal{P}^0)$. Differences in workers' expected utilities from this status-quo equilibrium are additionally normalized by their marginal utility of log consumption so that changes in welfare are measured in log consumption units. The social planner aggregates these expected utility differences using worker-specific Pareto welfare weights, λ^i.

3.3 Equilibrium Analysis

In this section, I discuss how I solve for a unique policy equilibrium and, in doing so, demonstrate my framework's tractability.

3.3.1 Solving for Optimal Redistributive Policy

Without strong assumptions on the shape of the social planner's objective function, solving for optimal government policy in quantitative spatial general equilibrium frameworks is ordinarily prohibitively computationally costly.[24] In the absence of strong assumptions over the shape of the social planner's objective function, researchers have taken direct and numerical approaches to ascertaining optimal policy. However, in multi-district models with social spillovers, obtaining an equilibrium price vector and equilibrium allocation of workers across districts can be computationally intensive for a single choice of policy instruments. Searching over the entire space of these policy instruments quickly becomes infeasible.[25]

[24] Under conditions of global concavity, it is possible to derive closed-form expressions for policies that implement the optimal allocation of workers across space (Fajgelbaum and Gaubert, 2020; Rossi-Hansberg *et al.*, 2019). Supposing global concavity is often overly restrictive, however. For instance, global concavity in Fajgelbaum and Gaubert (2020)'s planning problem requires that the weakest congestion forces across all types of workers dominate the strongest agglomeration forces across these types. In the current framework, this would imply that adding a college-educated worker to a disadvantaged school district always lowers social welfare, as the resulting rent increase must harm social welfare more than any increase of district amenities or district effects would raise it.

[25] Some recent attempts at ascertaining optimal policy in quantitative settings take a direct approach–searching over the policy instruments themselves—include Blouri and Ehrlich (2020) and Henkel *et al.* (2021). Both papers partition their policy instruments and search over these partitions in turn, presumably for computational tractability. In contrast, my framework allows me to determine optimal policy when the social planner is jointly deploying many sets of policy instruments. Jointly deploying many policy instruments is critical in my setting, as only under specific *combinations* of policy instruments may certain configurations of households across a city be sustained in equilibrium.

Additionally, given any choice of policy instruments, with endogenous spillovers there will almost surely be many equilibria consistent with the chosen policy instruments. The current framework provides a tractable alternative to directly searching over the planner's policy instruments and sidesteps issues of multiplicity. Critical to the model's tractability is the following proposition.

Proposition If consumption and district effects are normal goods for every household, then, given

1. the model's parameters and worker-specific levels of exogenous district amenities, $\boldsymbol{\theta} \equiv \{\varrho, A_L, A_H, \kappa_n, \epsilon_n, X_n^{\omega,a,i}, \beta_c^i, \beta_H^i, \beta_A^i, \alpha_c^i, \alpha_{\mathcal{N}}^i, \gamma_s^\omega, \rho^\omega, \Xi_n^\omega, \Pi^\omega\}$,
2. an allocation of college- and non-college-educated workers across school districts, $\{H_n, L_n\}_n$, and
3. a non-linear income tax schedule, $T(w)$,

any Policy Equilibrium is characterized by a unique vector of prices $\{\boldsymbol{r}, \boldsymbol{w}\}$ and a unique vector of place-based policy instruments, $\{\boldsymbol{\tau}, \boldsymbol{s}\}$.

Searching over the $2 \times (N+1)$-dimensional vector of college- and non-college-educated worker allocations and place-blind nonlinear income tax schedule is remarkably efficient. For (i), all equilibrium prices obtain close-form solutions; and (ii), equilibrium objects containing the place-specific policy instruments are uniquely determined through a pair of N-dimensional contraction mappings. These equilibrium objects are subsequently inverted to obtain closed-form expressions for the corresponding unique place-specific policy instruments associated with the candidate policy equilibrium vector.

Obtaining the equilibrium objects in (ii) resembles model inversion procedures in canonical quantitative spatial equilibrium models initially developed in economic geography and international trade frameworks (Redding and Rossi-Hansberg, 2017). The current procedure differs from the canonical approach in one important regard, however. Unlike in the canonical framework, households of the same type have heterogeneous preferences over the resulting equilibrium objects. This departure from the canonical framework ensures households of the same educational background and parental status can differentially value the planner's policy interventions. This is especially important if, for example, one wants to examine how ostensibly race-blind policy interventions impact outcomes like racial segregation—a question I plan to examine carefully in future iterations of the resulting paper.

3.4 Quantifying the Model

In this section, I quantify the model's structural parameters. I do this in two stages. In the first stage, I quantify the technological parameters governing children's upward mobility, district-level housing supply, and tradable good production using publicly available data. In the second stage, I estimate workers' preference parameters using administrative Census microdata. In this second stage I condition on the district effects implied by the technological parameters obtained in the first stage. In this chapter, I only outline how I will estimate the second-stage parameters; the resulting paper will detail the full procedure. Table 3.1 lists the technological parameters I estimate in the first stage alongside the values I select for the baseline calibration. I discuss each of the calibrations briefly below.

Children's Upward Mobility Chetty et al. (2020) construct a publicly available data set reporting children's outcomes in adulthood by each 2010 Census tract for children growing up in the 1980s and early 1990s. Included in these reports are children's earnings distributions in adulthood, computed separately by the national income rank of the children's parents. I use these reported earnings distributions for children growing up in Greater Boston to calibrate the human capital production function. For each school district in Greater Boston, I obtain the median tract-level estimate of expected earnings in adulthood for children growing up with parents who have incomes at the 25th and 75th percentile of the national income distribution. I then calibrate the children's human capital production function to these statistics using a method-of-moments estimator and publicly available Census and CCD data from 1990.

Since the mobility estimates are constructed for children growing up in the 1980s and early 1990s, I simulate the children's human capital production function using data from the 1990 Decennial Census and the 1989-1990 Local Education Agency (School District) Finance and Universe Surveys. I use the 1990 Decennial Census to obtain district-level counts of households headed by college- and non-college-educated adults, district-wide median annual earnings for college- and non-college-educated households, nationwide annual earnings for households at the 25th and 75th percentiles of the nation's earnings distribution, as well as median annual housing expenditures by school district.

Table 3.1. *Technological Parameters*

	Parameter	Value	Calibration Source
A. Children's Upward Mobility			
Consumption share:	α_c	0.04	CFHJP (2020)
Curvature of $\mathcal{N}$:	ρ	-1.79	CFHJP (2020)
Spending efficiency:	γ_s	.04	CFHJP (2020)
Curvature of $\bar{q}$:	η	.38	CFHJP (2020)
Exogenous factors:	Ξ_n	District residuals	CFHJP (2020)
B. Housing Supply			
Supply elasticities:	ϵ_n	District-specific	Baum-Snow and Han (2022)
Supply shifters:	κ_n	District-specific	ACS
Property tax rates:	τ_n^s	District-specific	Internal calibration
C. Production			
Non-college efficiency:	A_L	$3,290$	IPUMS
College efficiency:	A_H	$46,100$	IPUMS
Substitution elasticity:	ρ	.39	Diamond (2016)

Notes: See the main text for details on each of the parameters. As outlined in the main text, I estimate common technological parameters across parents with different levels of college attainment.

I combine the district-level counts of college- and non-college-educated households with the district-wide annual earnings to compute district-level peer effects conditional on technological parameters, $\bar{q}_n^{\omega}$.[26] I additionally combine the nation-wide annual earnings measures with the district-level housing expenditures (inclusive of property taxation and computed using homeowners with a mortgage) and average effective tax rates in 1990 to obtain model-consistent post-tax measures of non-housing consumption across school districts for households at the 25th and 75th percentiles of the nationwide income distribu-

[26] I use district-wide as opposed to metro-wide earnings so that I may identify the technological parameters governing $\bar{q}_n^{\omega}$ in a single cross-section of the data.

tion.[27,28] I finally combine the 1989-1990 Local Education Agency Finance survey with the 1989-1990 Local Education Agency Universe survey to obtain total district-level measures of per-student spending. The Finance survey provides total spending by school district, while the Universe survey provides counts of total students under each district's care.[29]

I use these data with a method-of-moments estimator to calibrate i) the share of children's upward mobility governed by consumption vis-à-vis neighborhood effects, α_c^i; ii) the complementarity of per-student school spending and peer effects, ρ^ω; iii) the efficiency of per-student school spending, γ^ω; iv) the complementarity between college-educated and non-college-educated workers' incomes in producing district-wide peer effects, η^ω; and v) district-level exogenous neighborhood effects, Ξ_n^ω. In order to identify these parameters, however, I must make the following simplifying assumptions. First, I must assume that these technological parameters governing children's human capital production are constant across parents' levels of educational attainment: $\alpha_c^{\omega,z} = \alpha_c^z$, $\rho^\omega = \rho$, $\gamma^\rho = \gamma$, $\eta^\omega = \omega$, and $\Xi_n^\omega = \Xi_n \;\; \forall\, \omega, n$. I thus assume that a child with parents of incomes at the 25th percentile of the national income distribution would have had the same expected adult earnings as a child with parents of incomes at the 75th percentile of the national income distribution if they grew up in the same school district and had access to the same level of material resources. This assumption is necessary to identify the consumption share of children's human capital production α_c^z. Second, due to incomplete measures of upward mobility by sociodemographic group across Greater Boston, I must assume that the share of consumption in the production of children's human capital is independent of the child's race, $\alpha_c^z = \alpha_c \,\forall\, z$.

Identification of α_c comes from differences in upward mobility within school districts

---

[27]Historical income tax rates are obtained from the Tax Foundation (https://taxfoundation.org/historical-income-tax-rates-brackets/).

[28]I use percentiles from the district-wide income distribution to compute the production of peer effects, but percentiles from the nationwide income distribution to compute district-level post-tax non-housing consumption. While the peer effects are produced locally, the measures of upward mobility are calculated for parents at different percentiles of national income distribution.

[29]All statistics are crosswalked using household count weights from 1990 boundaries to their 2010 equivalents. I thank Peter Rich for providing a crosswalk from 1990 school districts to 2010 Census tract boundaries.

but across parents at the 25th and 75th percentiles of the national income distribution, for the only difference between children of these parents is their level of consumption. Identification of ρ and γ comes from differences in upward mobility across districts that vary in per-student spending and the average pre-tax income of district residents. Holding average pre-tax income of district residents constant, the degree to which higher per-student school spending is associated with higher upward mobility identifies γ. The extent to which a given increase in per-student school spending is associated with higher upward mobility across districts with higher or lower average pre-tax incomes then identifies ρ.

η is identified by the sensitivity of upward mobility to mean preserving spreads in the distribution of pre-tax wages, conditional on a level of per-student school spending. To understand this, consider two school districts with the same level of per-student school spending and the same average pre-tax income. Assume, however, that wages in the first district are more dispersed than in the second. If upward mobility is lower in the first district than in the second, η must lie below 1, and upward mobility is shaped more by the presence of the poorest households in the district than the presence of the richest.[30] Given these parameter estimates, the exogenous components of the neighborhood effect function, Ξ_n, are finally calculated as the residuals from the moments equations.

Housing Supply Baum-Snow and Han (2023) construct a publicly available data set containing tract-level estimates of housing supply elasticities covering the majority of urban America. I aggregate these tract-level estimates for all school districts in the Greater Boston Metro Area under the assumption that district-wide housing demand shocks are propagated through identical tract-level housing demand shocks. This assumption yields the following aggregation formula:

$$1/\epsilon_n = \sum_{t \in n} \frac{\frac{HS_t}{1+1/\epsilon_t}}{\sum_t \frac{HS_n}{1+1/\epsilon_t}},$$

[30]While the model presented above assumes constant wages across school districts, given imperfect substitution between college- and non-college-educated workers, the sensitivity of peer effects to mean preserving spreads in the pre-tax income distribution remains an important determinant of upward mobility in the optimal policy simulations.

where HS_t is housing supply in tract t and ϵ_t is the corresponding tract-level inverse housing supply elasticity. With these district-wide housing supply estimates in hand, I calibrate the district-level housing supply shifters, κ_n, to match ACS estimates of district-wide median annual housing costs of homeowners less property taxes.

Production I obtain a value of 0.39 for ϱ from Diamond (2016). This corresponds to a labor substitution elasticity between college- and non-college-educated workers of 1.6, a value in line with the broader labor inequality literature (Katz and Autor, 1999). Conditional on ϱ, I calibrate the education-specific productivity shifters A_l and A_H to match the mean household income among non-college and college-headed households in the Greater Boston Metro area as reported in the 2015-2019 5-year ACS samples.

3.4.1 Parameter Estimates

I will include in the resulting paper estimates of households' preferences over consumption, $\bar{\beta}_c$, endogenous district-wide amenities, β_A, and if parents, the district-effect component of children's upward economic mobility, $\bar{\beta}_{\mathcal{N}}$. I will estimate these parameters separately by household type and sociodemographic group, θ. I will estimate these parameters using administrative Census micro data tracking the residential location choices of nearly all ACS respondents in Greater Boston. I will instrument for price and identify preferences over $\mathcal{N}_n$ using differences in parents' and non-parents' district choices.

3.5 Optimal Policy Analysis

In this section, I demonstrate the applicability of my framework for characterizing optimal urban economic policy with endogenous economic mobility for children. I show for different values of the preference parameters—and for a hypothetical set of technological parameters—how optimal policy and household allocations vary in a two-district version of the model outlined in this chapter. The exercise demonstrates a key force underlying the economic framework I have constructed.[31] In particular, I present optimal policy simulations varying college-educated households' preferences over endogenous district-wide amenities, $\beta_A^{\omega_H}$. In Figure 3.1 I ask, for a given level of residential integration across educational types (i.e., neighborhood dissimilarity), what is the maximal amount of social welfare the planner can achieve, and what are the corresponding policy instruments that implement the corresponding equilibrium? I report further details in the figure notes.

Figure 3.1 shows that as college-educated households' preferences for endogenous district-wide amenities increases vis-à-vis non-college-educated households' preferences, equilibria with high spending in disadvantaged neighborhoods and less residential integration become relatively more desirable (greater direct redistribution). The intuition behind this result is that as college-educated households' preference for endogenous amenities increases, the more they are willing to pay to live in a school district with other college-educated households. The social planner can then tax these households more as they self segregate, and spend the resulting funds on targeted housing assistance and school financing. Importantly, though, these simulations rely on school spending and peer composition being substitutes in the production of children's upward economic mobility ($\rho = .8$). When school spending and peer composition are instead complementary in production, the additional funds raised from taxing college-educated households in their segregated school district cannot be effectively spent in the poorer school district.

3.6 Concluding Remarks

In this chapter, I outline a tractable framework for characterizing optimal urban policy with endogenous upward economic mobility. I quantify technological parameters governing children's upward economic mobility, district-wide housing supply, and metro-wide tradable good production. I discuss plans for estimating households' preference parameters and

[31]The resulting paper will focus on values and deviations from the technological parameter estimates reported above. For the sake of this chapter, I consider fixed hypothetical values of the technological parameters to demonstrate the forces underlying optimal policy characterizations.

Figure 3.1. *Optimal Policy and Relative Preferences for Endogenous Amenities*

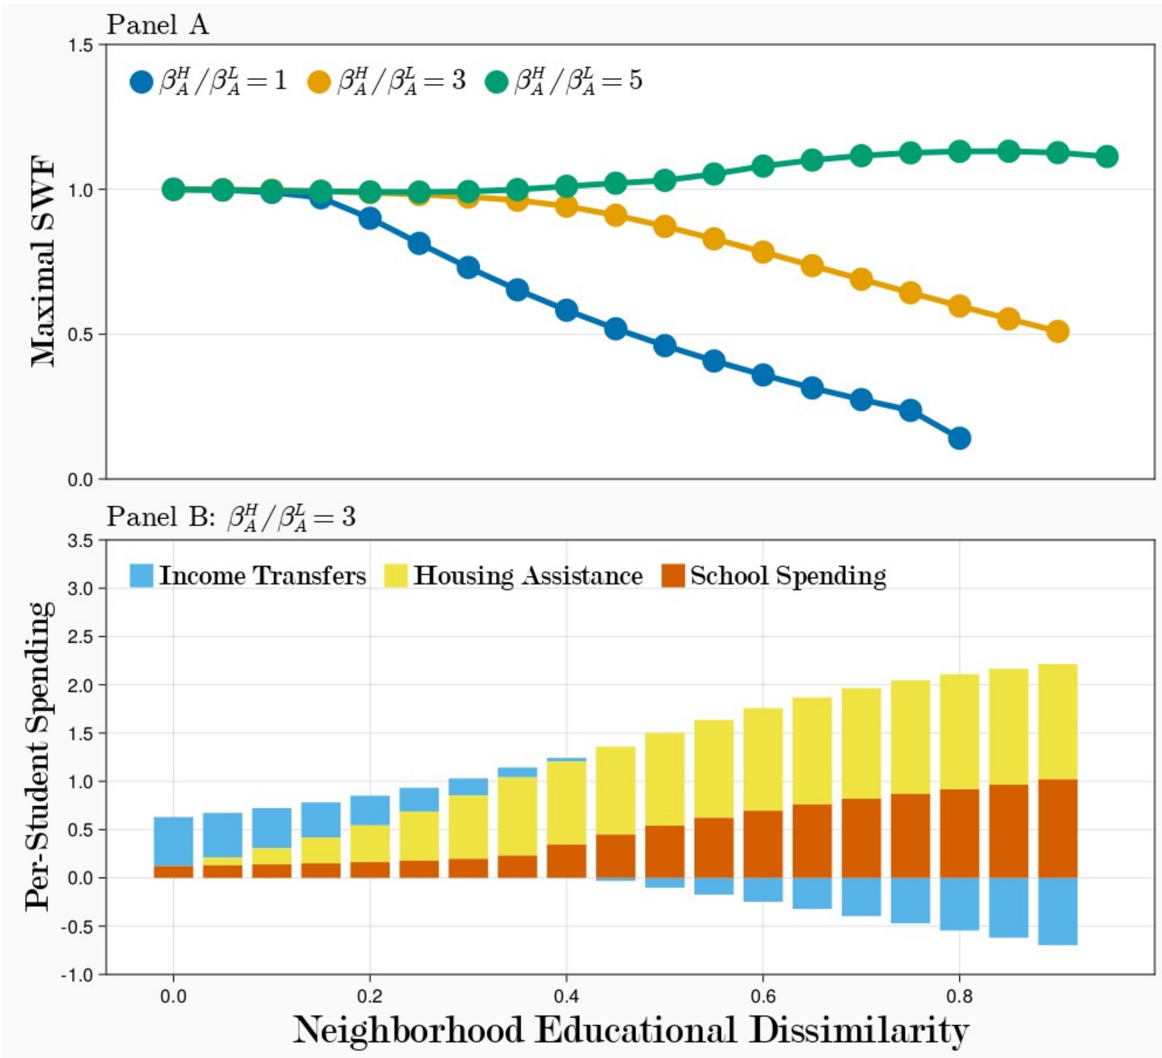

Notes: Panel A shows the maximal amount of social welfare obtainable conditional on a level of neighborhood (read, district) residential integration across educational groups. The three lines correspond to different levels of relative preferences between college- and non-college-educated households for endogenous district-wide amenities. Panel B shows, for the case where $\beta_A^H/\beta_A^L = 3$, the corresponding per-household levels of income transfers to non-college-educated households, housing assistance to non-college-educated households, and per-student spending on non-college-educated households. These simulations take households as homogeneous within educational group (aside from their idiosyncratic district preference). The values of the technological parameters and remaining preference parameters used in these simulations are constant across households and are as follows: $A_L = 4$, $A_H = 6$, $\varrho = .8$, $\kappa_n = .1 \,\forall\, n$, $\epsilon_n = 1.5 \,\forall\, n$, $\tau_n^s = 1/1,000$, $\Xi_n = 0 \,\forall\, n$, $\eta = 1$, $\gamma_s = .5$, $\rho = .8$, $\alpha_c = 1$, $\alpha_{\mathcal{N}} = 1$, $\xi_n = 0 \,\forall\, n$, $\beta_c = \beta_{\mathcal{N}} = 1$. Only interior solutions are considered in these simulations.

present preliminary optimal policy simulation results. The resulting paper will involve a comprehensive analysis of optimal urban economic policy when children's upward economic mobility is endogenous, and show how optimal policy varies with the technological parameters and the planner's social preferences.

References

Abadie, A., Athey, S., Imbens, G. and Wooldridge, J. (2023). When Should You Adjust Standard Errors for Clustering? *The Quarterly Journal of Economics*, **138** (1), 1–35.

Abdulkadiroğlu, A., Angrist, J. D., Dynarski, S. M., Kane, T. J. and Pathak, P. A. (2011). Accountability and Flexibility in Public Schools: Evidence from Boston's Charters and Pilots. *The Quarterly Journal of Economics*, **126** (2), 699–748.

Abowd, J., Stephens, B., Vilhuber, L., Andersson, F., McKinney, K., Roemer, M. and Woodcock, S. (2009). The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators. In T. Dunne, B. Jensen and M. Roberts (eds.), *Producer Dynamics: New Evidence from Micro Data*, *5*, University of Chicago Press, pp. 149–230.

Abramson, B. (2023). The Equilibrium Effects of Eviction and Homelessness Policies. *Working Paper*.

Agostinelli, F., Doepke, M., Sorrenti, G. and Zilibotti, F. (2020). *It Takes a Village: The Economics of Parenting with Neighborhood and Peer Effects*. Tech. rep., National Bureau of Economic Research.

—, Luflade, M. and Martellini, P. (2024). *On the Spatial Determinants of Educational Access*. Tech. rep., National Bureau of Economic Research.

Aguirregabiria, V. and Magesan, A. (2013). Euler Equations for the Dynamic Estimation of Discrete Choice Structural Models. *Structural Econometric Methods: Advances in Econometrics*, **31**, 3–44.

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M. and Wolf, N. (2015). The Economics of Density: Evidence From the Berlin Wall. *Econometrica*, **83** (6), 2127–2189.

Aliprantis, D. and Carroll, D. R. (2018). Neighborhood Dynamics and the Distribution of Opportunity. *Quantitative Economics*, **9** (1), 247–303.

Almagro, M., Chyn, E. and Stuart, B. A. (2023). *Urban Renewal and Inequality: Evidence from Chicago's Public Housing Demolitions*. Tech. rep., National Bureau of Economic Research.

— and Domínguez-Iino, T. (2022). Location Sorting and Endogenous Amenities: Evidence from Amsterdam. *Working Paper*.

Anderson, E. (2010). *The Imperative of Integration*. Princeton University Press.

Anenberg, E. and Kung, E. (2020). Can more housing supply solve the affordability crisis? Evidence from a neighborhood choice model. *Regional Science and Urban Economics*, **80**, 103363.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton university press.

Anstreicher, G. (2024). Spatial Influences in Upward Mobility. *Journal of Political Economy*, **132** (1).

Arcidiacono, P. and Miller, R. (2011). Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity. *Econometrica*, **79** (6), 1823–1867.

Asquith, B. J., Mast, E. and Reed, D. (2021). Local Effects of Large New Apartment Buildings in Low-Income Areas. *The Review of Economics and Statistics*, pp. 1–46.

Autor, D. H., Dorn, D. and Hanson, G. H. (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American economic review*, **103** (6), 2121–2168.

—, Palmer, C. J. and Pathak, P. A. (2017). Gentrification and the Amenity Value of Crime Reductions: Evidence from Rent Deregulation. *Working Paper 23914, National Bureau of Economic Research*.

Balboni, C., Bryan, G., Morten, M. and Siddiqi, B. (2020). Transportation, Gentrification, and Urban Mobility: The Inequality Effects of Place-Based Policies. *Working Paper*.

Bartik, A. and Mast, E. (2023). Black Suburbanization: Causes and Consequences of a Transformation of American Cities. *Working Paper*.

Baum-Snow, N. (2023). Constraints on City and Neighborhood Growth: The Central Role of Housing Supply. *Journal of Economic Perspectives*, **37** (2), 53–74.

— and Han, L. (2023). The Microgeography of Housing Supply. *Working Paper*.

— and Hartley, D. (2020). Accounting for central neighborhood change, 1980-2010. *Journal of Urban Economics*, **117**, 103228.

—, — and Lee, K. O. (2019). The Long-Run Effects of Neighborhood Change on Incumbent Families. *Working Paper*.

Bayer, P., Ferreira, F. and McMillan, R. (2007). A Unified Framework for Measuring Preferences for Schools and Neighborhoods. *Journal of Political Economy*, **115** (4), 588–638.

—, McMillan, R., Murphy, A. and Timmins, C. (2016). A Dynamic Model of Demand for Houses and Neighborhoods. *Econometrica*, **84** (3), 893–942.

—, — and Rueben, K. (2004). An Equilibrium Model of Sorting in an Urban Housing Market. *Working Paper 10865, National Bureau of Economic Research*.

Beauregard, P.-L. (2024). Gentrification, Displacement, and Income Trajectory of Incumbents. *Working Paper*.

Been, V., Ellen, I. G. and O'Regan, K. (2019). Supply Skepticism: Housing Supply and Affordability. *Housing Policy Debate*, **29** (1), 25–40.

Benabou, R. (1996). Equity and Efficiency in Human Capital Investment: The Local Connection. *The Review of Economic Studies*, **63** (2), 237–264.

Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L. and Palmer, C. (2023). Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice. *Working Paper*.

Berkes, E. and Gaetani, R. (2023). Income Segregation and the Rise of the Knowledge Economy. *American Economic Journal: Applied Economics*, **15** (2), 69–102.

Berry, C. (2021). Reassessing the Property Tax. *Working Paper*.

Biasi, B. (2023). School Finance Equalization Increases Intergenerational Mobility. *Journal of Labor Economics*, **41** (1), 1–38.

Bishop, K. C. (2012). A Dynamic Model of Location Choice and Hedonic Valuation. *Mimeo, Washington University in St. Louis*, pp. 1–38.

— and Murphy, A. D. (2019). Valuing Time-Varying Attributes Using the Hedonic Model: When is a Dynamic Approach Necessary? *Review of Economics and Statistics*, **101** (1), 134–145.

Blouri, Y. and Ehrlich, M. V. (2020). On the optimal design of place-based policies: A structural evaluation of EU regional transfers. *Journal of International Economics*, **125**, 103319.

Böhlmark, A. and Willén, A. (2020). Tipping and the Effects of Segregation. *American Economic Journal: Applied Economics*, **12** (1), 318–47.

Borusyak, K., Hull, P. and Jaravel, X. (2022). Quasi-Experimental Shift-Share Research Designs. *Review of Economic Studies*, **89** (1), 181–213.

Boustan, L. P. (2010). Was Postwar Suburbanization "White Flight"? Evidence from the Black Migration. *The Quarterly Journal of Economics*, **125** (1), 417–443.

Bratu, C., Harjunen, O. and Saarimaa, T. (2023). JUE Insight: City-wide effects of new housing supply: Evidence from moving chains. *Journal of Urban Economics*, **133**, 103528.

Brock, W. A. and Durlauf, S. N. (2001). Discrete Choice with Social Interactions. *The Review of Economic Studies*, **68** (2), 235–260.

— and — (2002). A Multinomial-Choice Model of Neighborhood Effects. *American Economic Review, P&P*, **92** (2), 298–303.

Brueckner, J. (1987). The structure of urban equilibria: A unified treatment of the muth-mills model. *Handbook of Regional and Urban Economics*, **2**, 821–845.

Brueckner, J. K., Thisse, J.-F. and Zenou, Y. (1999). Why is central paris rich and downtown detroit poor?: An amenity-based theory. *European Economic Review*, **43** (1), 91–107.

Brummet, Q. and Reed, D. (2021). The Effects of Gentrification on Incumbent Residents. *Working Paper*.

Caetano, G. (2019). Neighborhood sorting and the value of public school quality. *Journal of Urban Economics*, **114**, 103193.

Campos, C. and Kearns, C. (2024). The Impact of Public School Choice: Evidence from Los Angeles's Zones of Choice. *The Quarterly Journal of Economics*, **139** (2), 1051–1093.

Card, D., Mas, A. and Rothstein, J. (2008). Tipping and the Dynamics of Segregation. *Quarterly Journal of Economics*, **123** (1), 177–218.

Chetty, R., Friedman, J. N., Hendren, N., Jones, M. R. and Porter, S. R. (2018). The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility.

— and Hendren, N. (2018). The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects. *The Quarterly Journal of Economics*, **133** (3), 1107–1162.

—, — and Katz, L. F. (2016). The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. *American Economic Review*, **106** (4), 855–902.

Chow, M., Fort, T., Goetz, C., Goldschlag, N., Lawrence, J., Perlman, E., Stinson, M. and White, K. (2021). Redesigning the Longitudinal Business Database. *NBER Working Paper*.

Chyn, E. (2018). Moved to Opportunity: The Long-Run Effects of Public Housing Demolition on Children. *American Economic Review*, **108** (10), 3028–3056.

— and Daruich, D. (2022). *An Equilibrium Analysis of the Effects of Neighborhood-Based Interventions on Children*. Tech. rep., National Bureau of Economic Research.

— and Katz, L. F. (2021). Neighborhoods Matter: Assessing the Evidence for Place Effects. *Journal of Economic Perspectives*, **35** (4), 197–222.

Collinson, R., Humphries, J. E., Mader, N., Reed, D., Tannenbaum, D. and van Dijk, W. (2023). Eviction and Poverty in American Cities. *Quarterly Journal of Economics*.

Correia, S., Guimarães, P. and Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal*, **20** (1), 95–115.

Couture, V., Gaubert, C., Handbury, J. and Hurst, E. (2021). Income Growth and the Distributional Effects of Urban Spatial Sorting. p. 87.

—, —, — and — (2023). Income Growth and the Distributional Effects of Urban Spatial Sorting. *The Review of Economic Studies*.

— and Handbury, J. (2020). Urban Revival in America. *Journal of Urban Economics*, **119**, 103267.

— and — (2023). Neighborhood change, gentrification, and the urbanization of college graduates. *Journal of Economic Perspectives*, **37** (2), 29–52.

Cullen, J. B. and Levitt, S. D. (1999). Crime, Urban Flight, and the Consequences for Cities. *The Review of Economics and Statistics*, **81** (2), 159–169.

Curci, F. and Masera, F. (2023). Flight from Urban Blight: Lead Poisoning, Crime, and Suburbanization. *The Review of Economics and Statistics*, pp. 1–45.

— and Yousef, H. (2022). Gentrifying Cities, Amenities and Income Segregation: Evidence from San Francisco. *Working Paper*.

Cutler, D. M. and Glaeser, E. L. (1997). Are Ghettos Good or Bad? *The Quarterly Journal of Economics*, **112** (3), 827–872.

Damiano, A. and Frenier, C. (2020). Build Baby Build?: Housing Submarkets and the Effects of New Construction on Existing Rents. p. 48.

Davis, M., Gregory, J. and Hartley, D. (2023). Preferences over the Racial Composition of Neighborhoods: Estimates and Implications. *Working Paper*.

—, —, — and Tan, K. (2021). Neighborhood effects and housing vouchers. *Quantitative Economics*, **12** (4), 1307–1346.

De Bartolome, C. A. (1990). Equilibrium and Inefficiency in a Community Model with Peer Group Effects. *Journal of Political Economy*, **98** (1), 110–133.

Dekle, R., Eaton, J. and Kortum, S. (2008). Global Rebalancing with Gravity: Measuring the Burden of Adjustment. *IMF Staff Papers*, **55** (3), 511–540.

Demsas, J. (2022). Housing Breaks People's Brains. Section: Ideas.

Derenoncourt, E. (2022). Can You Move to Opportunity? Evidence from the Great Migration. *American Economic Review*, **112** (2), 369–408.

Diamond, R. (2016). The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *American Economic Review*, **106** (3).

— and Gaubert, C. (2022). Spatial Sorting and Inequality. *Annual Review of Economics*, **14**, 795–819.

— and McQuade, T. (2019). Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development. *Journal of Political Economy*, p. 55.

—, — and Qian, F. (2018). The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco. *NBER Working Paper*, (24181).

Ding, L. and Hwang, J. (2020). Effects of Gentrification on Homeowners: Evidence from a Natural Experiment. *Regional Science and Urban Economics*, **83**, 103536.

—, — and Divringi, E. (2016). Gentrification and Residential Mobility in Philadelphia. *Regional Science and Urban Economics*, **61**, 38–51.

Dobbie, W. and Fryer Jr, R. G. (2011). Are High-Quality Schools Enough to Increase Achievement Among the Poor? evidence from the harlem children's zone. *American Economic Journal: Applied Economics*, **3** (3), 158–187.

Donaldson, D. and Hornbeck, R. (2016). Railroads and American Economic Growth: A "Market Access" Approach. *The Quarterly Journal of Economics*, **131** (2), 799–858.

Dougherty, C. (2020). Build Build Build Build Build Build Build Build Build Build Build Build Build Build. *The New York Times*.

Dragan, K., Ellen, I. G. and Glied, S. (2020). Does gentrification displace poor children and their families? new evidence from medicaid data in new york city. *Regional Science and Urban Economics*, **83**, 103481.

Durlauf, S. N. (1996). A Theory of Persistent Income Inequality. *Journal of Economic growth*, **1**, 75–93.

— (2004). Neighborhood effects. *Handbook of Regional and Urban Economics*, **4**, 2173–2242.

Eckert, F. and Kleineberg, T. (2024). The geography of opportunity: Education, work, and intergenerational mobility across us counties. *Working Paper*.

Edlund, L., Machado, C. and Sviatschi, M. (2022). Gentrification and the Rising Returns to Skill. *Economica*, **89** (354), 258–292.

Ellen, I. G. (2023). A Response to David Imbroscio: Neighborhoods Matter, and Efforts to Integrate Them Are Not Futile. *Housing Policy Debate*, **33** (4), 789–792.

—, Horn, K. M. and Reed, D. (2019). Has falling crime invited gentrification? *Journal of Housing Economics*, **46**, 101636.

— and O'Regan, K. (2010). Crime and urban flight revisited: The effect of the 1990s drop in crime on cities. *Journal of Urban Economics*, **68** (3), 247–259.

— and O'Regan, K. M. (2011). How low income neighborhoods change: Entry, exit, and enhancement. *Regional Science and Urban Economics*, **41** (2), 89–97.

— and Steil, J. P. (2019). *The Dream Revisited: Contemporary Debates about Housing, Segregation, and Opportunity in the Twenty-first Century*. Columbia University Press.

Fajgelbaum, P. and Gaubert, C. (2020). Optimal Spatial Policies, Geography, and Sorting. *The Quarterly Journal of Economics*, **135** (2), 959–1026.

Fee, K. and Hartley, D. (2013). The Relationship Between City Center Density and Urban Growth or Decline. In K. Zeuli and S. Wachter (eds.), *Revitalizing American Cities*, *3*, University of Pennsylvania Press, pp. 45–64.

Fernandez, R. and Rogerson, R. (1996). Income Distribution, Communities, and the Quality of Public Education. *The Quarterly Journal of Economics*, **111** (1), 135–164.

— and — (1998). Public education and income distribution: A dynamic quantitative evaluation of education-finance reform. *American Economic Review*, pp. 813–833.

Ferreira, F. V., Kenney, J. H. and Smith, B. (2023). Household Mobility, Networks, and Gentrification of Minority Neighborhoods in the US. *NBER Working Paper*, (31480).

Ferreyra, M. M. (2007). Estimating the Effects of Private School Vouchers in Multidistrict Economies. *American Economic Review*, **97** (3), 789–817.

Fogli, A. and Guerrieri, V. (2019). *The End of the American Dream? Inequality and Segregation in US Cities*. Tech. rep., National Bureau of Economic Research.

Freeman, L. (2005). Displacement or Succession? Residential Mobility in Gentrifying Neighborhoods. *Urban Affairs Review*, **40** (4), 463–491.

— and Braconi, F. (2004). Gentrification and Displacement New York City in the 1990s. *Journal of the American planning association*, **70** (1), 39–52.

Friedrich, M., Rosenthal, T., Rosenthal, T., Sandman, B., Sandman, B., Byrnes, S. T., Byrnes, S. T., Semley, J., Semley, J., Stern, S. W., Stern, S. W., Bradfield, S. and Bradfield, S. (2023). More Building Won't Make Housing Affordable. *The New Republic*.

Fryer Jr, R. G. and Katz, L. F. (2013). Achieving Escape Velocity: Neighborhood and School Interventions to Reduce Persistent Inequality. *American Economic Review*, **103** (3), 232–237.

Fu, E. (2022). The Financial Burdens of Property Taxes: Evidence from Philadelphia. *Working Paper*.

Galiani, S., Murphy, A. and Pantano, J. (2015). Estimating neighborhood choice models: Lessons from a housing assistance experiment. *American Economic Review*, **105** (11), 3385–3415.

Gans, H. J. (1982). *Urban Villagers*. Simon and Schuster.

Gaubert, C., Kline, P. and Yagan, D. (2021). Place-Based Redistribution. *NBER Working Paper*, (28337).

Gendron-Carrier, N. (2023). Prior Work Experience and Entrepreneurship: The Careers of Young Entrepreneurs. *IZA Discussion Paper*, (16145).

Gervais, A. and Jensen, J. B. (2019). The tradability of services: Geographic concentration and trade costs. *Journal of International Economics*, **118**, 331–350.

Gilraine, M., Graham, J. and Zheng, A. (2023). *Public Education and Intergenerational Housing Wealth Effects*. Tech. rep., National Bureau of Economic Research.

Giroud, X. and Rauh, J. (2019). State Taxation and the Reallocation of Business Activity: Evidence from Establishment-Level Data. *Journal of Political Economy*, **127** (3), 1262–1316.

Glaeser, E. and Gottlieb, J. (2008). The Economics of Place-Making Policies. *Brookings Papers on Economic Activity*, **2008**, 155–239.

—, Kahn, M. and Rappaport, J. (2008). Why do the poor live in cities? the role of public transportation. *Journal of Urban Economics*, **63**, 1–24.

—, Kolko, J. and Saiz, A. (2001). Consumer city. *Journal of Economic Geography*, **1**, 27–50.

—, Luca, M. and Moszkowski, E. (2023). Gentrification and retail churn: Theory and evidence. *Regional Science and Urban Economics*, **100**, 103879.

Glaeser, E. L. and Gyourko, J. (2003). The Impact of Building Restrictions on Housing Affordability. *FRBNY Economic Policy Review*, pp. 21–39.

— and Kahn, M. E. (2001). Decentralized Employment and the Transformation of the American City. pp. 1–63.

— and Luttmer, E. F. P. (2003). The Misallocation of Housing Under Rent Control. *American Economic Review*, **93** (4), 1027–1046.

Goldsmith-Pinkham, P., Sorkin, I. and Swift, H. (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review*, **110** (8), 2586–2624.

Graham, J. (2018). Mayor Walsh slams Boston Housing Authority message encouraging residents to leave city. *Boston Herald*.

Gray, M. N. (2021). America Needs More Luxury Housing, Not Less. Section: Ideas.

Gregory, V., Jordan-Wood, S., Kozlowski, J. and Rubinton, H. (2023). *Residential Segregation and the Black-White College Gap*. Tech. rep., Federal Reserve Bank of St. Louis.

Groote, O. D. and Verboven, F. (2019). Subsidies and Time Discounting in New Technology Adoption: Evidence from Solar Photovoltaic Systems. *American Economic Review*, **109** (6), 2137–72.

Gu, Y., Guo, N., Wu, J. and Zou, B. (2021). Home Location Choices and the Gender Commute Gap. *Journal of Human Resources*.

Guerrieri, V., Hartley, D. and Hurst, E. (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics*, **100**, 45–60.

Guimarães, P., Figueirdo, O. and Woodward, D. (2003). A Tractable Approach to the Firm Location Decision Problem. *Review of Economics and Statistics*, **85** (1), 201–204.

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.

Handbury, J. (2021). Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities. *Econometrica*, **89** (6), 2679–2715.

Handel, D. V. and Hanushek, E. A. (2024). Contexts of Convenience: Generalizing from Published Evaluations of School Finance Policies. *Evaluation Review*, **48** (3), 461–494.

Hanson, G. H. (2005). Market potential, increasing returns and geographic concentration. *Journal of International Economics*, **67** (1), 1–24.

Harris, C. D. (1954). The Market as a Factor in the Localization of Industry in the United States. *Annals of the Association of American Geographers*, **44** (4), 315–348.

Harrison, B. (1974). Ghetto Economic Development: A Survey. *Journal of Economic Literature*, **12** (1), 1–37.

Hendren, N. and Sprung-Keyser, B. (2020). A Unified Welfare Analysis of Government Policies. *The Quarterly Journal of Economics*, **135** (3), 1209–1318.

Henkel, M., Seidel, T. and Suedekum, J. (2021). Fiscal Transfers in the Spatial Economy. *American Economic Journal: Economic Policy*, **13** (4), 433–68.

Herzog, I. (2021). National transportation networks, market access, and regional economic growth. *Journal of Urban Economics*, **122**, 103316.

Hoelzlein, M. (2023). Two-Sided Sorting and Spatial Inequality in Cities. *Working Paper*.

Hotz, J. and Miller, R. (1993). Conditional Choice Probabilities and the Estimation of Dynamic Models. *The Review of Economic Studies*, **60** (3), 497–529.

Howard, G. and Liebersohn, J. (2021). Why is the rent so darn high? The role of growing demand to live in housing-supply-inelastic cities. *Journal of Urban Economics*, **124**, 103369.

Hwang, J. and Lin, J. (2016). What Have We Learned About the Causes of Recent Gentrification? *Cityscape*, **18** (3), 9–26.

Imbroscio, D. (2023). Beyond Opportunity Hoarding: Interrogating Its Limits as an Account of Urban Inequalities. *Housing Policy Debate*, **33** (4), 770–788.

Jackson, C. K., Johnson, R. C. and Persico, C. (2015). The Effects of School Spending on Educational and Economic Outcomes: Evidence from School Finance Reforms. *The Quarterly Journal of Economics*, **131** (1), 157–218.

— and Mackevicius, C. L. (2024). What Impacts Can We Expect from School Spending Policy? Evidence from Evaluations in the United States. *American Economic Journal: Applied Economics*, **16** (1), 412–446.

Jackson, K. T. (1987). *Crabgrass Frontier: The Suburbanization of the United States*. Oxford University Press.

Jargowsky, P. A. (2009). Immigrants and Neighbourhoods of Concentrated Poverty: Assimilation or Stagnation? *Journal of Ethnic and Migration Studies*, **35** (7), 1129–1151.

Kain, J. (1962). The journey-to-work as a determinant of residential location. *Papers of the Regional Science Association*, **9** (1), 137–160.

Kain, J. F. (1968). Housing Segregation, Negro Employment, and Metropolitan Decentralization. *The Quarterly Journal of Economics*, **82** (2), 175–197.

KAIN, J. H. and PERSKY, J. J. (1969). Alternatives to the gilded ghetto. *The Public Interest*, **14**, 74.

KALOUPTSIDI, M., SCOTT, P. and SOUZA-RODRIGUES, E. (2021). Linear IV regression estimators for structural dynamic discrete choice models. *Journal of Econometrics*, **222**, 778–804.

KENNAN, J. and WALKER, J. R. (2011). The Effect of Expected Income on Individual Migration Decisions. *Econometrica*, **79** (1), 211–251.

KLEINBAUM, D. G. and KLEIN, M. (1996). *Survival Analysis: A Self-Learning Text*. Springer.

KLINE, P. (2010). Place Based Policies, Heterogeneity, and Agglomeration. *American Economic Review, Papers & Proceedings*, **100** (2), 383–387.

— and MORETTI, E. (2014). People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs. *Annual Review of Economics*, **6**, 629–662.

KLING, J. R., LIEBMAN, J. B. and KATZ, L. F. (2007). Experimental Analysis of Neighborhood Effects. *Econometrica*, **75** (1), 83–119.

KULKA, A., SOOD, A. and CHIUMENTI, N. (2022). How to Increase Housing Affordability? Understanding Local Deterrents to Building Multifamily Housing.

KUMINOFF, N. V., SMITH, V. K. and TIMMINS, C. (2013). The New Economics of Equilibrium Sorting and Policy Evaluation Using Housing Markets. *Journal of Economic Literature*, **51** (4), 1007–1062.

LABRIE, P. (1970). Black Central Cities: Dispersal or Rebuilding (Part I). *The Review of Black Political Economy*, **1** (2), 3–27.

LAFORTUNE, J., ROTHSTEIN, J. and SCHANZENBACH, D. W. (2018). School Finance Reform and the Distribution of Student Achievement. *American Economic Journal: Applied Economics*, **10** (2), 1–26.

LALIBERTÉ, J.-W. (2021). Long-Term Contextual Effects in Education: Schools and Neighborhoods. *American Economic Journal: Economic Policy*, **13** (2), 336–377.

LANDVOIGT, T., PIAZZESI, M. and SCHNEIDER, M. (2015). The Housing Market(s) of San Diego. *American Economic Review*, **105** (4), 1371–1407.

LE BARBANCHON, T., RATHELOT, R. and ROULET, A. (2021). Gender Differences in Job Search: Trading off Commute against Wage. *The Quarterly Journal of Economics*, **136** (1), 381–426.

LEROY, S. and SONSTELIE, J. (1983). Paradise lost and regained: Transportation innovation, income, and residential location. *Journal of Urban Economics*, **13**, 67–89.

LESTER, T. W. and HARTLEY, D. A. (2014). The long term employment impacts of gentrification in the 1990s. *Regional Science and Urban Economics*, **45**, 80–89.

LI, X. (2019). Do New Housing Units in Your Backyard Raise Your Rents? *Working Paper*.

LIU, L., MCMANUS, D. and YANNOPOULOS, E. (2022). Geographic and temporal variation in housing filtering rates. *Regional Science and Urban Economics*, **93**, 103758.

LUNDBERG, S. and STARTZ, R. (1998). On the Persistence of Racial Inequality. *Journal of Labor Economics*, **16** (2), 292–323.

MANSON, S., SCHROEDER, J., RIPER, D. V., KUGLER, T. and RUGGLES, S. (2022). IPUMS National Historical Geographic Information System: Version 17.0 [dataset].

—, —, VAN RIPER, D., KNOWLES, K., KUGLER, T., ROBERTS, F. and RUGGLES, S. (2023). National Historical Geographic Information System: Version 18.0.

MARULLO, S. (1985). Housing Opportunities and Vacancy Chains. *Urban Affairs Quarterly*, **20** (3), 364–388, publisher: SAGE Publications.

MAST, E. (2021). JUE Insight: The effect of new market-rate housing construction on the low-income housing market. *Journal of Urban Economics*, p. 103383.

MELTZER, R. and GHORBANI, P. (2017). Does gentrification increase employment opportunities in low-income neighborhoods? *Regional Science and Urban Economics*, **66**, 52–73.

MIESZKOWSKI, P. and MILLS, E. S. (1993). The Causes of Metropolitan Suburbanization. *Journal of Economic Perspectives*, **7** (3), 135–147.

MILLER, C. (2021). When Work Moves: Job Suburbanization and Black Employment. *Working Paper*.

MOLLOY, R., NATHANSON, C. G. and PACIOREK, A. (2022). Housing supply and affordability: Evidence from rents, housing consumption and household location. *Journal of Urban Economics*, **129**, 103427.

MORENO-MALDONADO, A. and SANTAMARIA, C. (2022). Delayed Childbearing and Urban Revival. *Working Paper*.

MORETTI, E. (2013). Real Wage Inequality. *American Economic Journal: Applied Economics*, **5** (1), 65–103.

NATHANSON, C. G. (2020). Trickle-down housing economics. *Working Paper*.

NECHYBA, T. J. (2000). Mobility, Targeting, and Private-School Vouchers. *American Economic Review*, **91** (1), 130–146.

NOTOWIDIGDO, M. J. (2020). The Incidence of Local Labor Demand Shocks. *Journal of Labor Economics*, **38** (3), 687–725.

OREOPOULOS, P. (2003). The Long-Run Consequences of Living in a Poor Neighborhood. *The Quarterly Journal of Economics*, **118** (4), 1533–1575.

— (2012). Moving Neighborhoods Versus Reforming Schools: A Canadian's Perspective. *Cityscape*, pp. 207–212.

Pennington, K. (2021). Does Building New Housing Cause Displacement?: The Supply and Demand Effects of Construction in San Francisco. *Working Paper*.

Piazzesi, M., Schneider, M. and Stroebel, J. (2019). Segmented Housing Search. *Working Paper*.

Redding, S. J. and Rossi-Hansberg, E. (2017). Quantitative Spatial Economics. *Annual Review of Economics*, **9**, 21–58.

Roback, J. (1982). Wages, Rents, and the Quality of Life. *Journal of Political Economy*, **90** (6), 1257–1278.

Rosen, S. (1974). Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, **82** (1), 34–55.

Rosenthal, S. S. (2014). Are Private Markets and Filtering a Viable Source of Low-Income Housing? Estimates from a "Repeat Income" Model. *American Economic Review*, **104** (2), 687–706.

Rossi-Hansberg, E., Sarte, P.-D. and Schwartzman, F. (2019). Cognitive Hubs and Spatial Redistribution. *NBER Working Paper*, (26267).

Rust, J. (1987). Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, **55** (5), 999–1033.

Saiz, A. (2010). The Geographic Determinants of Housing Supply. *Quarterly Journal of Economics*, **125** (3), 1253–1296.

Scott, P. (2013). Dynamic Discrete Choice Estimation of Agricultural Land Use. *Discussion Paper, New York University*.

Shelby, T. (2014). Integration, Inequality, and Imperatives of Justice: A Review Essay. *Phil. & Pub. Aff.*, **42**, 253.

Silva, J. S. and Tenreyro, S. (2006). The Log of Gravity. *The Review of Economics and statistics*, **88** (4), 641–658.

Song, J. (2021). The Effects of Residential Zoning in U.S. Housing Markets. *Working Paper*.

Su, Y. (2022). The Rising Value of Time and the Origin of Urban Gentrification. *American Economic Journal: Economic Policy*, **14** (1), 402–439.

Sullivan, J. and Genadek, K. (2024). Using the Census Bureau's Master Address File for Migration Research. *ADEP Working Paper Series*.

Szaniszlo, M. (2018). Get out of town! - BHA pitch to public aid. *Boston Herald*.

Traiberman, S. (2019). Occupations and Import Competition: Evidence from Denmark. *American Economic Review*, **109** (12), 4260–4301.

Tsivanidis, N. (2022). Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio. *Working Paper*.

Turner, L. M. (2008). Who Gets What and Why? Vacancy Chains in Stockholm's Housing Market. *European Journal of Housing Policy*, **8** (1), 1–19.

Vigdor, J. L. (2002). Does Gentrification Harm the Poor? *Brookings-Wharton Papers on Urban Affairs*, pp. 133–182.

— (2010). Is urban decay bad? Is urban revitalization bad too? *Journal of Urban Economics*, **68** (3), 277–289.

Wagner, D. and Layne, M. (2014). The Person Identification Validation System (PVS): Applying the Center for Administrative Records Research and Applications' (carra) Record Linkage Software. *CARRA Working Paper Series*.

White, H. C. (1971). Multipliers, Vacancy Chains, and Filtering in Housing. *Journal of the American Institute of Planners*, **37** (2).

Wilson, W. J. (1987). *The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy*. University of Chicago Press.

Zheng, A. and Graham, J. (2022). Public education inequality and intergenerational mobility. *American Economic Journal: Macroeconomics*, **14** (3), 250–282.

Appendix A

Appendix to Chapter 1

A.1 Data Appendix

This appendix provides a brief description of how we construct our panel, impute unit rental status, and construct our tract aggregates. We will include a detailed version of this appendix in future iterations of the resulting paper.

Sample Construction To construct our analysis panel, we first form an annual panel of persons' residential histories using the MAF-ARF. The MAF-ARF links PIKs to MAFIDs annually from 2000. For years before 2012, multiple MAFIDs can be assigned to PIKs. We randomly choose a single MAFID for PIKs assigned multiple MAFIDs before 2012. We merge PIKs across years to construct an annual panel of MAFIDs for each PIK reported in the MAF-ARF. As MAFIDs come from multiple (potentially conflicting) sources, MAFIDs may switch between years for PIKs that did not change residential addresses. To minimize this potential, we smooth over PIKs' MAFID histories. Specifically, we find the first and last time a PIK is reported living at a MAFID and assume that any changes in MAFIDs between these years are artifacts of the data. We also fill in missing MAFIDs between years in which a PIK is assigned a MAFID. If the MAFIDs on either side of the missing years are the same, we simply assign the same MAFID to these missing years. If the MAFIDs are different, we assume the PIK moved between the MAFIDs at the midpoint of the two years MAFIDs are observed for the PIK.

We merge workers' total annual earnings across all jobs from the LEHD to this residential panel via workers' PIKs. Earnings come from the LEHD Employment History File. We assign the workplace establishment for each PIK-year as the workplace establishment that the PIK received the most earnings during each calendar year. We finally assign sociodemographic information to PIKs from the LEHD's Individual Characteristics File and the ACS.

Rental Unit Imputation We use a combination of the ACS, the MAF-X, and the CoreLogic data sets to assign MAFIDs as rental units. To do so, we take the following approach:

1. We assign MAFIDs surveyed in the ACS as rental units if the survey respondents at the respective MAFID reported being renters. For MAFIDs that appear multiple times in the ACS, we assign rental unit status separately for each year they appear in the sample.

2. We assign rental unit status to MAFIDs reported in the CoreLogic Multiple Listing Services (MLS) file as being for rent. For MAFIDs that appear multiple times in the MLS file, we record rental status separately for each year they appear in the data.

3. We assign MAFIDs in the CoreLogic Tax History and CoreLogic transactions data with absentee owners as rental units. Again, we assign rental status separately for each year the same MAFID appears in these files.

4. We take the set of MAFIDs in the MAF-X that had not been assigned as rental units in any year between 2000 and 2019. We then remove MAFIDs from this set that appear in any of the CoreLogic databases as being occupied by their owner. Among this reduced set of MAFIDs, we assign all MAFIDs as rental units recorded as part of a multi-unit building.

These procedures yield a small false-positive rate; i.e., these four steps assign few owner-occupied units as rental units. Future iterations of the resulting paper will release precise false-positive rates and additional data validation statistics. We also expect to assign more

MAFIDs as rental units as we further understand the CoreLogic data and utilize additional real estate data from Black Knight.

Tract Aggregates We construct tract aggregates for rents, property values, and socioeconomic/demographic characteristics from the ACS annually. We use five-year averages to construct rents and property values. We hedonically adjust rents using housing characteristics and geographical variables reported in the ACS. We construct socioeconomic and demographic tract characteristics using the residential panel before subsetting to our analysis sample. We account for varying PIK coverage rates across years in the MAF-ARF (Sullivan and Genadek, 2024). We will release details explaining these calculations in future iterations of the paper.

A.2 Aggregate Neighborhood Change Since 2000

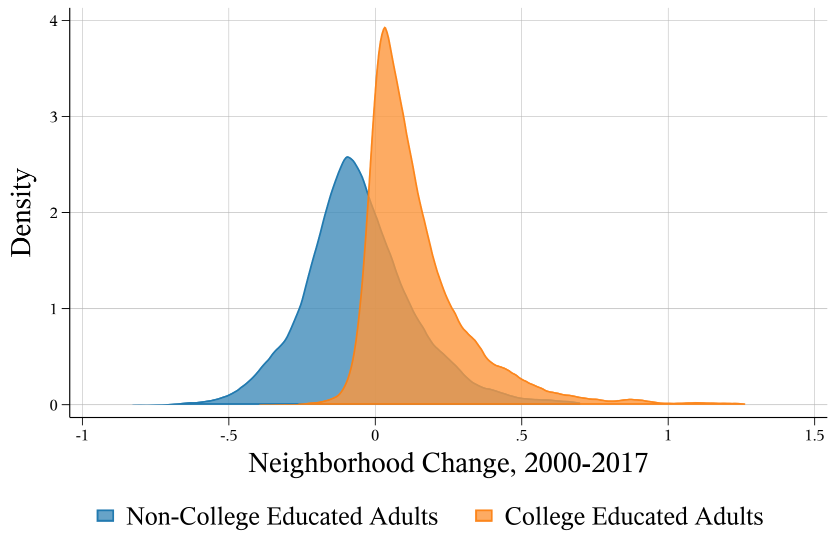

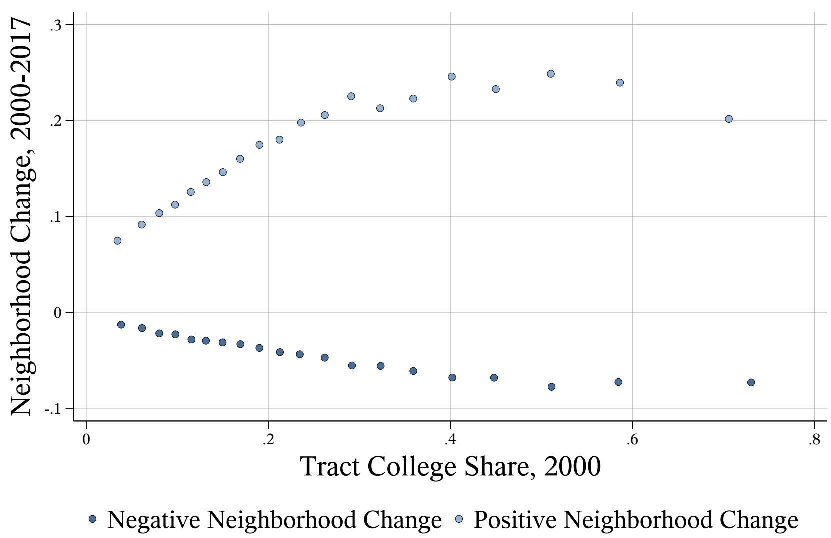

Figure A.1. *Neighborhood Change, 2000 - 2017*

These figures are constructed using 2010 delineated census tract population counts for the 100 largest (ranked by population) US CBSAs from the publicly available 2000 NHGIS Census and the American Community Survey 2015-2019 5-year aggregates (Manson *et al.*, 2022). We excluded tracts with fewer than 1,000 adult residents in 2000. Neighborhood change for non-college-educated and college-educated households is defined analogously to equation 1.1. The top panel plots kernel densities of neighborhood change between 2000-2017 among census tracts that contain 20 percent of the CBSA's population closest to its CBD. The bottom panel plots our tract-level measure of neighborhood change (y-axis) against the tract's initial share of college-educated workers in 2000, separately for tracts that experienced positive and negative changes in their share of college-educated households.

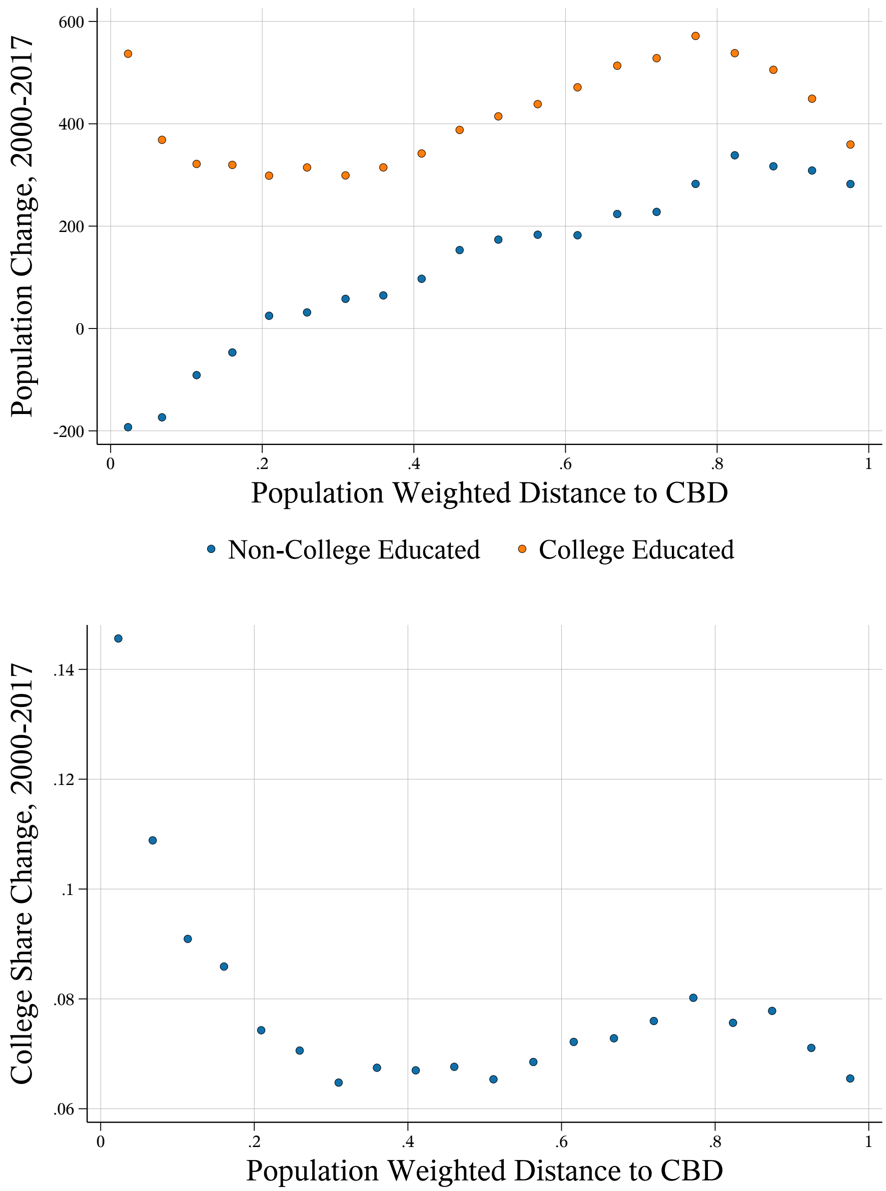

Figure A.2. *Neighborhood Change and Initial College Share*

The figure is constructed using 2010 census tract population counts for the 100 largest (ranked by population size) US metropolitan divisions from the publicly available 2000 NHGIS Census and the American Community Survey 2015-2019 5-year aggregates (Manson *et al.*, 2022). We excluded tracts with fewer than 1,000 adult residents in 2000. We define neighborhood change in equation 1.1. The top panel plots the degree of population change for college- and non-college-educated households between 2000-2017 by the population-weighted distance to each metropolitan division's CBD. The bottom panel plots the change in the share of college-educated households, also by the population-weighted distance to each metropolitan division's respective CBD.

A.3 Reduced-Form Appendix

A.3.1 Control Variables

Household-Level Controls Our vector of household-level controls, X_i, is designed to capture household characteristics that jointly determine one's origin location and our set of outcome variables, Δy_i and $h(t|i)$. These controls include i) average household income, defined as the mean income of all adults residing at the same address during 2010; ii) household size, defined as the number of adults residing at the same address;[1] iii) a second-order polynomial in the household head's date of birth; iv) indicators for the household head's race, immigrant status, sex, and college degree attainment;[2] and v) the length of the household head's prior residential tenure in neighborhood $n(i)$. We define all variables in our base year, 2010.

Neighborhood-Level Controls Our vector of neighborhood-level controls, $X_{n(i)}$, is designed to capture characteristics of the household's origin neighborhood that are potentially correlated with our measure of neighborhood change and our set of outcome variables, Δy_i and $h(t|i)$. We first include neighborhood-level controls defined in 2010, which include i) the share of adults in the neighborhood with a college degree; ii) the share of adults in the neighborhood that identify as non-Hispanic and white; iii) the median income among working-age adults residing in the neighborhood; iv) the median property value and rent payment in the neighborhood; v) the degree of neighborhood churn which we define as the share of adults residing in the neighborhood during 2008 who also remain in the neighborhood during 2009; vi) second-order polynomials in the distance to the metropolitan division's CBD, measured in both the physical distance as well as in the cumulative share of the metropolitan division's residents residing closer to the CBD than those in tract $n(i)$; vii) fixed effects capturing five bands equidistant from the metropolitan division's CBD; and

[1] We topcode this value to 10, as a few addresses record an infeasibly high number of adult residents.

[2] While we run our analysis separately for Black and non-Black headed households, we include finer racial distinctions as part of our controls.

viii) the 5-year lag in our measure of neighborhood change.

In addition to our neighborhood-level controls defined in 2010, we also include a few contemporaneous neighborhood-level controls to capture *changing* characteristics of the household's origin neighborhood that are potentially correlated with our measure of neighborhood change and our set of outcome variables, Δy_i and $h(t|i)$. In particular, we control for changes in job market access to *tradable* industries among low-skilled workers. In terms of equation 1.15, these changes in job market access among tradable industries for low-skilled workers are defined as $\Delta \sum_{d \in \mathcal{T}} JMA_{ndt} \tilde{\theta}_d^c$, where Δ corresponds to the difference in our JMA measure across 2010 and 2019, and $\tilde{\theta}_d^c$ is the share of non-college-educated workers employed in industry d in CBSA c's state.[3] We finally include in $X_{n(i)}$ a third-order polynomial in the change in neighborhood $n(i)$'s total population.

Recall that all specifications include CBSA-level fixed effects α_{CBSA}, ensuring that variation in gentrification comes from across neighborhoods in the same CBSA.

A.3.2 Identification

A causal interpretation of our coefficients of interest, β_{NC}^{Cox} and β_{NC}^{LP}, is based on the conditionally random assignment of neighborhood change across neighborhoods between 2010 and 2019. That is, conditional on our control variables, we assume there are no unobserved neighborhood- or household-level characteristics that are correlated with our measure of neighborhood change. Given our setup, to interpret β_{NC}^{Cox} and β_{NC}^{LP} causally, we much ensure two conditions are satisfied:

1. First, residents in observably similar neighborhoods in 2010 must not differ in unobservable ways that correlate with our outcome variables, Δy_i and $h(t|i)$. If residents moving into gentrifying neighborhoods prior to 2010 are observably similar to their incumbent residents but different in unobservable ways that affect Δy_i and $h(t|i)$ (access to familial wealth, for example), β_{NC}^{Cox} and β_{NC}^{LP} will partly reflect sample selection.

[3] Note that including non-tradable industries into this measure of job market access would induce a classic "bad controls" problem, as we can easily conceptualize changes in non-tradable employment as an outcome of neighborhood change (Angrist and Pischke, 2009).

2. Second, neighborhood-level gentrification throughout our analysis period must be uncorrelated with unobserved neighborhood characteristics that independently predict incumbents' outcomes. While changes to unobserved public and private amenities caused by increased neighborhood demand among college-educated residents constitute part of our treatment, we must purge our identifying variation of shocks to neighborhood characteristics that affect incumbent residents' outcomes independently of gentrification. These unobserved shocks may include changes in tradable job market access (Kain, 1962; Miller, 2021), changes in transportation infrastructure that precedes gentrification (LeRoy and Sonstelie, 1983; Glaeser *et al.*, 2008; Curci and Yousef, 2022), or changes in neighborhood valuations resulting from secular trends in preferences or from shifts to within-city income distributions (Brueckner, 1987; Couture *et al.*, 2023). That suburbanization has continued unabated throughout our analysis periods—particularly for Black residents—makes these concerns especially acute (Bartik and Mast, 2023; Couture and Handbury, 2023).

We take steps to help ensure conditions 1 and 2 are met. First, in addition to our rich household-level controls, we often subset our sample to those who have lived in their origin neighborhood for at least five years. While this decision focuses our analysis on longtime renters, it helps disambiguate the outcomes of gentrifiers and the outcomes of our target population of incumbent low-income renters, reducing the potential for sample bias. This is especially true when considering that we control for the five-year lag in our measure of gentrification. Indeed, for sample selection to influence our results, it must be that low-income renter households predict with some accuracy how trends in neighborhood change will vary five years in the future and base their current residential choices on these predictions in a way that is uncorrelated with both their own observable characteristics (measured in 2010) and their chosen neighborhoods' observable characteristics (also measured in 2010).

Second, we carefully control for changes in neighborhoods' characteristics that may independently predict incumbents' outcomes. By controlling for changes in access to low-skill tradable employment opportunities, we help ensure aggregate economic conditions that may independently affect neighborhood composition are not driving incumbents' outcomes. This concern is relevant given the stark evidence of differential mobility responses to aggregate labor demand shocks across low- and high-skill workers (Notowidigdo, 2020). By controlling for a cubic in neighborhood-level population change over our analysis period, we further control for unobserved neighborhood-level shocks that similarly affect incumbent residents and potential gentrifiers. To take an extreme example, consider a natural disaster during our analysis period that leads to a large depopulation of affected census tracts. In this scenario, we will observe low rates of residential tenure among incumbent residents, which we—without our population controls—would falsely attribute to a decline in our measure of neighborhood change. Note that the inclusion of our population controls implies that one should interpret the coefficients on our measure of neighborhood change as the effect of changing neighborhood *composition* on incumbents' outcomes. Finally, our rich set of geographical controls ensures that we compare outcomes for incumbent residents across origin tracts equidistant from the Metropolitan Division's CBD, mitigating bias from secular trends toward suburbanization among our target population.

While we carefully choose our controls to mitigate the impact of sample selection and changes in unobserved neighborhood-level characteristics, our estimates are robust to excluding any small subset of control variables. We finally note that we explored using our IVs detailed in Section 1.6 to estimate our reduced-form equations. Our reduced-form estimates are moderately sensitive to the instrument choice and composition (i.e., years and industries selected), indicating meaningful heterogeneity in the complier characteristics of incumbent renters' origin neighborhoods across our instruments. It is, therefore, difficult to interpret the reduced-form IV estimates without placing more structure on our reduced-form equations to understand how each relevant equilibrium object (e.g., rents, non-tradable job market access, and neighborhood socioeconomic composition) mediates the impact of gentrification on incumbent renters. For now, we report our OLS estimates, which we believe offer a more transparent depiction of the impact of neighborhood change on incumbent renters' observable outcomes.

Cox Proportional Hazards Assumption Identification in our Cox Proportional Hazards models requires that we additionally meet the proportional hazards assumption. Namely, that the impact of neighborhood change on incumbents' hazard rates is constant across each year between 2010 and 2019. We test this assumption by plotting $-\log(-\log(\text{survival probability}))$ against $\log(\text{time})$ separately for incumbent residents originally residing in neighborhoods within each decile of our measure of neighborhood change. We observe parallel lines across all deciles of neighborhood change, consistent with the proportional hazards assumption. We also test the null hypothesis that the corresponding Schoenfeld residuals for our measure of neighborhood change are not serially correlated.[4] We report the results of Stata's `phtest` command for our measure of neighborhood change in Table A.1, which indicates that we cannot reject the null at the 95 percent confidence level for our full samples of Black and non-Black incumbent renters (though we are close to doing so for non-Black renters).

Table A.1. *Schoenfeld Residuals*

Shoenfeld Residuals	Black	Non-Black
$\mathbb{P} \geq \chi^2$	0.317	0.056

Notes: Table A.1 reports the results of Stata's `phtest`, testing the null hypothesis that the Schoenfeld residuals are not serially correlated. The test statistic is distributed as χ^2 under the null hypothesis of no serial correlation. *Sources*: Estimates were disclosed by the US Census Bureau's Disclosure Review Board. Project Number 2358. Disclosure Clearance Number CBDRB-FY24-P2358-R10936.

[4]The Schoenfeld residuals correspond to the difference between the observed covariate values and the expected covariate values under the Cox proportional hazards model for incumbent renters in each year between 2010 and 2019. If the proportional hazards assumption holds, these residuals should not be serially correlated (Kleinbaum and Klein, 1996).

Table A.2. *Cox Model: Effect of Gentrification on Incumbent Renters' Tenure (2010-2019)*

Hazard Rate	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Neighborhood Change	-0.0688	-0.107*	-0.0793	0.187*	0.195	0.204*	0.369	0.323*	-0.777**	-0.264	0.165	-0.901*
	(0.0542)	(0.0439)	(0.132)	(0.0781)	(0.147)	(0.871)	(0.204)	(0.108)	(0.286)	(0.325)	(0.555)	(0.411)
Select Household-level Controls												
Household Income ($1,000s)	-0.00381***	-0.00341***	-0.00320***	-0.00277***	-0.00273***	-0.00275***	-0.00326***	-0.00257***	-0.00365***	-0.00300***	-0.00301***	-0.00476***
	(0.000155)	(0.000110)	(0.000413)	(0.000271)	(0.000666)	(0.000343)	(0.000858)	(0.000453)	(0.000526)	(0.000438)	(0.000887)	(0.000733)
Household Head DOB	0.0157***	0.0157***	0.0167***	0.143***	0.0157***	0.0145***	0.0151***	0.0152***	0.0175***	0.0144***	0.0189***	0.0130***
	(0.000216)	(0.000154)	(0.000581)	(0.000381)	(0.000948)	(0.000503)	(0.00124)	(0.000675)	(0.000734)	(0.000593)	(0.00120)	(0.000973)
Residential Tenure	-0.102***	-0.105***	-0.116***	-0.116***	-0.112***	-0.118***	-0.120***	-0.116***	-0.116***	-0.114***	-0.105***	-0.106***
	(0.000833)	(0.000555)	(0.00375)	(0.00230)	(0.00601)	(0.00311)	(0.00782)	(0.00411)	(0.00478)	(0.00343)	(0.00804)	(0.00596)
Select Neighborhood-level Controls												
Neighborhood Churn	-0.328***	-0.358***	-0.901***	-0.698***	-0.696***	-0.693***	-0.740***	-0.662***	-0.928***	-0.612***	-1.145***	-0.712***
	(0.0171)	(0.0117)	(0.102)	(0.0639)	(0.144)	(0.0825)	(0.202)	(0.112)	(0.146)	(0.101)	(0.232)	(0.179)
Rent ($1,000s)	0.272***	0.197***	0.0661	0.122***	0.135*	0.100***	0.146	0.100*	0.0749	0.198**	-0.0528	0.134
	(0.0209)	(0.0143)	(0.0452)	(0.0313)	(0.0617)	(0.0303)	(0.0844)	(0.0399)	(0.0664)	(0.0654)	(0.121)	0.0784
College Share	.141***	0.212***	0.235***	0.285***	-0.0866	0.143***	0.116	0.138*	0.923***	0.818***	0.704*	1.039***
	(0.0275)	(0.0178)	(0.0662)	(0.0350)	(0.108)	(0.0494)	(0.156)	(0.0676)	(0.207)	(0.163)	(0.352)	(0.252)
White Share	0.0975***	-0.113***	0.221***	-0.155***	0.215***	-0.0654	0.0348	-0.0666	0.245***	-0.195***	0.140*	-0.123
	(0.0150)	(0.0133)	(0.0321)	(0.0246)	(0.0578)	(0.0367)	(0.0807)	(0.0461)	(0.0383)	(0.0347)	(0.0586)	(0.0629)
Sample Restrictions												
Race	Black	Non-Black	Black	Non-Black	Black	Non-Black	Black	Non-Black	Black	Non-Black	Black	Non-Black
Longtime Renters			✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Initial College Share					High	High	High	High	Low	Low	Low	Low
Fraction Developed							High	High			Low	Low
N (1,000s)	314	688	56	156	21	89	13.5	54.5	35	67	10.5	20.5

Notes: Coefficients correspond to the percent change in the hazard rate from a one-unit increase in the corresponding independent variable. A one-unit increase in our measure of neighborhood change corresponds to a one-hundred percentage-point increase. Every specification includes the full set of controls listed and detailed in Appendix A.3. We report standard errors clustered at the origin census tract level in parentheses. Longtime renters have resided in their origin Census tract since at least 2005. Tracts with a "High" ("Low") initial college share are tracts whose share of college-educated adults in 2010 is above (below) the population-weighted median in our full sample. Tracts with a "High" ("Low") fraction developed are tracts whose developed land area in 2011 is above (below) the population-weighted median in our full sample. *Sources*: ACS (2005-2021), LEHD (2010), CoreLogic (2006-2017), MAF-X (2019), and MAF-ARF (2010). Estimates were disclosed by the US Census Bureau's Disclosure Review Board. Project Number 2358. Disclosure Clearance Number CBDRB-FY24-P2358-R10936.

Table A.3. *Linear Regression Model: Effect of Gentrification on Incumbent Renters (2010-2019)*

Outcome Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: Household Outcomes												
Leave Tract	-0.0176	-0.0263	-0.016	0.07*	0.104	0.0668*	0.129	0.111***	-0.305***	-0.0888	-0.0458	-0.303
	(0.0223)	(0.0174)	(0.0509)	(0.0278)	(0.0571)	(0.0308)	(0.0738)	(0.0368)	(0.106)	(0.0894)	(0.254)	(0.158)
Moved > 5 Miles	0.038	0.0167	0.0066	0.0803**	0.0438	0.0709*	0.023	0.0926**	-0.0324	-0.0049	0.386	-0.205
	(0.0217)	(0.0171)	(0.0406)	(0.0261)	(0.0505)	(0.0294)	(0.0649)	(0.0339)	(0.0783)	(0.0802)	(0.223)	(0.154)
Leave CBSA	-0.0519***	-0.0375**	-0.0158	0.0087	-0.0194	0.0191	-0.0200	0.0179	0.0012	-0.0194	-0.0278	-0.1500
	(0.013)	(0.0122)	(0.0238)	(0.0160)	(0.0300)	(0.0178)	(0.0387)	(0.0205)	(0.0451)	(0.0432)	(0.1130)	(0.089)
Income	2,872**	2,268**	-932	-86.05	-1,458	-926.6	-1,875	311.5	-9,138*	-551.8	-23,950**	-3,874
	(948.8)	(736.4)	(2,170)	(1,375)	(2,748)	(1,547)	(3,498)	(1,816)	(3,854)	(3,427)	(9,225)	(6,444)
Commute Distance	-0.2280	-1.905***	0.637	0.5900	3.106	-0.234	3.651	1.0800	-1.854	3.146	-9.659	-3.379
	(0.776)	(0.531)	(1.609)	(1.002)	(2.059)	(1.146)	(2.529)	(1.352)	(3.132)	(2.775)	(6.284)	(5.352)
Panel B: Experienced Tract Characteristics												
Rent	0.271***	0.0960***	0.389***	0.222***	0.159**	0.132***	0.171*	0.153***	0.863***	0.607***	0.759**	0.485***
	(0.0472)	(0.0288)	(0.0620)	(0.0329)	(0.0546)	(0.0335)	(0.0701)	(0.0434)	(0.145)	(0.096)	(0.241)	(0.131)
College Share	0.726***	0.299***	1.079***	0.567***	0.539***	0.392***	0.570***	0.379***	3.726***	4.108***	3.559***	3.897***
	(0.104)	(0.0598)	(0.131)	(0.0788)	(0.0714)	(0.0525)	(0.0845)	(0.0689)	(0.250)	(0.236)	(0.565)	(0.357)
White College Share	-0.298	0.635***	0.277	1.269***	0.822*	0.449***	0.806	0.358***	7.758*	7.567***	2.832	5.206***
	(0.765)	(0.156)	(1.393)	(0.226)	(0.384)	(0.0863)	(0.419)	(0.098)	(3.951)	(1.167)	(7.965)	(1.169)
Sample Restrictions												
Race	Black	Non-Black	Black	Non-Black	Black	Non-Black	Black	Non-Black	Black	Non-Black	Black	Non-Black
Longtime Renters			✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Initial College Share					High	High	High	High	Low	Low	Low	Low
Fraction Developed							High	High			Low	Low
N (1,000s)	277	599	50	138	19	78.5	12	48.5	31	59.5	9.3	18

Notes: Coefficients correspond to the impact of a 100 percentage point increase in our measure of gentrification on a change in the associated outcome variable from 2010 to 2019. The dependent variables "Leave Tract", "Moved > 5 Miles", and "Leave CBSA" are all indicator variables equal to one if the corresponding condition is satisfied. We measure incomes and rents in 2010 dollars and commute distance in miles. We measure experienced neighborhood characteristics in percent changes. Every specification includes the full set of controls listed and detailed in Appendix A.3. We report standard errors clustered at the origin census tract level in parentheses. Longtime renters have resided in their origin Census tract since at least 2005. Tracts with a "High" ("Low") initial college share are tracts whose share of college-educated adults in 2010 is above (below) the population-weighted median in our full sample. Tracts with a "High" ("Low") fraction developed are tracts whose developed land area in 2011 is above (below) the population-weighted median in our full sample. *Sources*: ACS (2005-2021), LEHD (2010), CoreLogic (2006-2017), MAF-X (2019), and MAF-ARF (2010). Estimates were disclosed by the US Census Bureau's Disclosure Review Board. Project Number 2358. Disclosure Clearance Number CBDRB-FY24-P2358-R10936.

A.4 Microfoundations and Structural Estimation Details

A.4.1 Deriving our Estimating Equation

Consider the set of residential choices detailed in Section 1.5.2. Given these residential choices, we start the derivation of our moment restrictions with an application of the Hotz-Miller inversion, which amounts to differencing equation 1.7 across the two neighborhood choices in period t, n and n':

$$\begin{aligned}\ln\left(\frac{p^k_{xnt}}{p^k_{xn't}}\right) &= v^k_{xnt} - v^k_{xn't} \\ &= \bar{u}^k_{xnt} - \bar{u}^k_{xn't} + \delta\left(\mathbb{E}\left[\bar{V}^k_{xnt}\right] - \mathbb{E}\left[\bar{V}^k_{xn't}\right]\right),\end{aligned} \tag{A.1}$$

where the expectation operator is with respect to both the observable household-level and all the city-specific state variables. By assumption 1, we can write these expectations as

$$\begin{aligned}\mathbb{E}\left[\bar{V}^k_{xnt}\right] &= \sum_{x'}\int_{\bar{\omega}'}\bar{V}^k_{xnt}dF^{\bar{\omega}}\left(\bar{\omega}'|\bar{\omega}\right)f^x\left(x'|n, x_{nt}, \bar{\omega}^{kc}_t\right) \\ &= \sum_{x'}\mathbb{E}_{\bar{\omega}'|\bar{\omega}^{kc}_t}\left[\bar{V}(x', \bar{\omega}')|\bar{\omega}^{kc}_t\right]f^x\left(x'|n, x_{nt}, \bar{\omega}^{kc}_t\right)\end{aligned}$$

where x' and $\bar{\omega}'$ denote the next period values for x and $\bar{\omega}$. We can also replace the expectation of the ex-ante continuation values with respect to the city-specific state variables with their realized counterparts and an expectational error defined in equation 1.9:

$$\mathbb{E}\left[\bar{V}^k_{xnt}\right] = \sum_{x'}\bar{V}(x', \bar{\omega}^{kc}_{t+1})f^x\left(x'|n, x_{nt}, \bar{\omega}^{kc}_t\right) + e^{\bar{V}}(n, x_{nt}, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t+1}), \tag{A.2}$$

with

$$e^{\bar{V}}\left(n, x_t, \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t+1}\right) \equiv \sum_{x'}e^{\bar{V}}\left(x', \bar{\omega}^{kc}_t, \bar{\omega}^{kc}_{t+1}\right)f^x\left(x'|n, x_{nt}, \bar{\omega}^{kc}_t\right).$$

These realized continuation values permit minimal assumptions about households' beliefs over the evolution of the city-specific state variables. Imputing our expression for households' expected continuation values conditional on their household-level state variables in A.2 to our expression for the difference in conditional choice probabilities in A.1 yields

$$\ln\left(\frac{p^k_{xnt}}{p^k_{xn't}}\right) = \bar{u}^k_{xnt} - \bar{u}^k_{xn't}$$

$$+\,\delta\left(\sum_{x'}\bar{V}(x',\bar{\omega}^{kc}_{t+1})\,f^x\left(x'|n,x_t,\bar{\omega}^{kc}_t\right)-\sum_{x'}\bar{V}(x',\bar{\omega}^{kc}_{t+1})\,f^x\left(x'|n',x_t,\bar{\omega}^{kc}_t\right)\right)$$
$$+\,\delta\cdot\tilde{e}(x_t,\bar{\omega}^{kc}_t,\bar{\omega}^{kc}_{t+1})$$

where $\tilde{e}(x_t,\bar{\omega}^{kc}_t,\bar{\omega}^{kc}_{t+1})$ is the difference between the expectational errors when residing in neighborhood n relative to neighborhood n' in period $t+1$,

$$\tilde{e}(x_t,\bar{\omega}^{kc}_t,\bar{\omega}^{kc}_{t+1})\equiv e^{\bar{V}}(n,x_t,\bar{\omega}^{kc}_t,\bar{\omega}^{kc}_{t+1})-e^{\bar{V}}(n',x_t,\bar{\omega}^{kc}_t,\bar{\omega}^{kc}_{t+1}).$$

Next, using equation 1.8 to substitute in for $\bar{V}(x',\bar{\omega}^{kc}_{t+1})$, we obtain,

$$\begin{aligned}
&\ln\left(\frac{p^k_{xnt}}{p^k_{xn't}}\right)-\delta\left(\sum_{x'}\ln\left(p^k_{x\tilde{n}t}\right)f^x\left(x'|n,x_t,\bar{\omega}^{kc}_t\right)-\sum_{x'}\ln\left(p^k_{x\tilde{n}t}\right)f^x\left(x'|n',x_t,\bar{\omega}^{kc}_t\right)\right)\\
&=\ \bar{u}^k_{xnt}-\bar{u}^k_{xn't}+\delta\left(\sum_{x'}v^k_{\tilde{n}}\left(x',\bar{\omega}^{kc}_t\right)\,f^x\left(x'|n,x_t,\bar{\omega}^{kc}_t\right)-\sum_{x'}v^k_{\tilde{n}}\left(x',\bar{\omega}^{kc}_t\right)\,f^x\left(x'|n',x_t,\bar{\omega}^{kc}_t\right)\right)\\
&+\,\delta\cdot\tilde{e}(x_t,\bar{\omega}^{kc}_t,\bar{\omega}^{kc}_{t+1})
\end{aligned}$$

Recall that $\tilde{n}$ is a renewal action for both households. Therefore, the household-level state variables are set to the same values for both households regardless of their values in period t. This yields identical continuation values in period $t+1$ for both households. The above expression therefore simplifies to,

$$\begin{aligned}
&\ln\left(\frac{p^k_{xnt}}{p^k_{xn't}}\right)-\delta\left(\sum_{x'}\ln\left(p^k_{x\tilde{n}t}\right)f^x\left(x'|n,x_t,\bar{\omega}^{kc}_t\right)-\sum_{x'}\ln\left(p^k_{x\tilde{n}t}\right)f^x\left(x'|n',x_t,\bar{\omega}^{kc}_t\right)\right)\\
&=\ \bar{u}^k_{xnt}-\bar{u}^k_{xn't}+\delta\left(MC^k_t(\tilde{n},n)-MC^k_t(\tilde{n},n')+\tilde{e}(x_t,\bar{\omega}^{kc}_t,\bar{\omega}^{kc}_{t+1})\right)
\end{aligned}$$

Choosing n' as the city's outside option, applying assumption 2, and substituting in for the neighborhoods' flow utilities provides an equation linear in our model parameters,

$$\begin{aligned}
&\ln\left(\frac{p^k_{xnt}}{p^k_{xn't}}\right)-\delta\left(\sum_{x'}\ln\left(p^k_{x\tilde{n}t}\right)f^x\left(x'|n,x_t,\bar{\omega}^{kc}_t\right)-\sum_{x'}\ln\left(p^k_{x\tilde{n}t}\right)f^x\left(x'|n',x_t,\bar{\omega}^{kc}_t\right)\right)\\
&\qquad=\ \tilde{\alpha}^k_n+\alpha^k_t+\beta^k_w\ln(\bar{I}_{n,t})-\beta^k_r\ln(r_{n,t})+\beta^k_A\ln\left(\frac{Coll_{nt}}{Pop_{nt}}\right)\\
&\qquad+\,\beta^k_\tau\left(\sum_{x'}\ln(\tau_{xt}(x'))f^x\left(x'|n,x_{xt},\bar{\omega}^{kc}_t\right)-\sum_{x'}\ln(\tau_{xt}(x'))f^x\left(x'|n',x_t,\bar{\omega}^{kc}_t\right)\right)\\
&\qquad-\,MC^k_t(n,n_{t-1})+MC^k_t(n',n_{t-1})+\delta\left(MC^k_t(\tilde{n},n)-MC^k_t(\tilde{n},n')\right)
\end{aligned}$$

$$+ \tilde{e}(x_t, \bar{\omega}_t^{kc}, \bar{\omega}_{t+1}^{kc}) + \tilde{\xi}_{nt}^k,$$

where $\tilde{\alpha}_n^k = \alpha_n^k - \alpha^{ck}$. To condense notation, we write this equation as

$$Y_{xnn'\tilde{n}t}^k = \tilde{\alpha}_n^k + \alpha_t^k + \beta_w^k \ln(\bar{I}_{n,t}) - \beta_r^k \ln(r_{n,t}) + \beta_A^k \ln\left(\frac{Coll_{nt}}{Pop_{nt}}\right) + \beta_\tau^k \tilde{\tau}_x - \widetilde{MC}_t^k + \nu_{xnn'\tilde{n}t}^k,$$

where

$$\begin{aligned}
Y_{xnn'\tilde{n}t}^k &\equiv \ln\left(\frac{p_{xnt}^k}{p_{xn't}^k}\right) - \delta\left(\sum_{x'} \ln\left(p_{x\tilde{n}t}^k\right) f^x\left(x'|n, x_t, \bar{\omega}_t^{kc}\right) - \sum_{x'} \ln\left(p_{x\tilde{n}t}^k\right) f^x\left(x'|n', x_t, \bar{\omega}_t^{kc}\right)\right) \\
\tilde{\tau}_x &\equiv \sum_{x'} \ln(\tau_{xt}(x')) f^x\left(x'|n, x_t, \bar{\omega}_t^{kc}\right) - \sum_{x'} \ln(\tau_{xt}(x')) f^x\left(x'|n', x_t, \bar{\omega}_t^{kc}\right) \\
\widetilde{MC}_t^k &\equiv MC_t^k(n, n_{xt-1}) - MC_t^k(n', n_{xt-1}) - \delta\Big(MC_t^k(\tilde{n}, n) - MC_t^k(\tilde{n}, n')\Big) \\
\nu_{xnn'\tilde{n}t}^k &\equiv \tilde{e}(x_t, \bar{\omega}_t^{kc}, \bar{\omega}_{t+1}^{kc}) + \tilde{\xi}_{nt}^k.
\end{aligned}$$

This is the same equation reported in 1.10.

A.4.2 First-Step Conditional Choice Probabilities

To see how equation 1.11 approximates the neighborhood choice problem of households in our dynamic model, start by considering our specification for a household's conditional value function defined in 1.6,

$$\begin{aligned}
v_n^k\left(x_{it}, \bar{\omega}_t^{kc}\right) &\equiv \bar{u}_n^k\left(x_{it}, \bar{\omega}_t^{kc}\right) + \delta \mathbb{E}_t\left[\bar{V}^k\left(x_{it+1}, \bar{\omega}_{t+1}^{kc}\right)|n, x_{it}, \bar{\omega}_t^{kc}\right] \\
&= \alpha_n^k + \alpha_t^{kc} + \beta_w^k \ln(\bar{I}_{n,t}) - \beta_r^k \ln(r_{n,t}) + \beta_A^k \ln\left(\frac{Coll_{nt}}{Pop_{nt}}\right) - MC_t^k(n_t, n_{it-1}) + \tilde{\xi}_{nt}^k \\
&\quad + \sum_{\bar{\tau}=1,2}\Bigg(\delta \mathbb{E}_t\left[\bar{V}^k\left(x_{it+1}(n_t, \bar{\tau}), \bar{\omega}_{t+1}^{kc}\right)|n, x(n_{t-1}, \bar{\tau}_{t-1}), \bar{\omega}_t^{kc}\right] \\
&\quad\quad + \beta_\tau^k \ln(\bar{\tau})\Bigg) \cdot f^{\bar{\tau}}(\bar{\tau}|n_t, x(n_{t-1}, \bar{\tau}_{t-1}), \bar{\omega}^{kc}).
\end{aligned} \tag{A.3}$$

Assume for now that $\bar{\tau}_{t-1} = 1$. Then, we can re-write A.3 for some neighborhood choice n as,

$$v_n^k\left(x_{it}, \bar{\omega}_t^{kc}\right) = \gamma_{nt}^k + \mu_{\bar{\tau}} \cdot \mathbb{1}\{n = n_{t-1})\} - MC_t^k(n, n_{t-1}) \tag{A.4}$$

where,

$$
\begin{aligned}
\gamma_{nt}^{k} =& \alpha_n^k + \alpha_t^{kc} + \beta_w^k \ln\left(\bar{I}_{n,t}\right) - \beta_r^k \ln\left(r_{n,t}\right) + \beta_A^k \ln\left(\frac{Coll_{nt}}{Pop_{nt}}\right) + \tilde{\xi}_{nt}^k \\
&+ \delta\, \mathbb{E}_t\left[\bar{V}^k\left(x_{it+1}(n_t, \bar{\tau}=1), \bar{\omega}_{t+1}^{kc}\right) | n, x(n_{t-1},1), \bar{\omega}_t^{kc}\right] \\
\mu_{\bar{\tau}} =& f^{\bar{\tau}}(2|n_t, x(n_{t-1},1), \bar{\omega}^{kc})\Bigg[\beta_\tau^k \ln(2) \\
&- \delta \cdot \Bigg[\mathbb{E}_t\left[\bar{V}^k\left(x_{it+1}(n_t,\bar{\tau}=1),\bar{\omega}_{t+1}^{kc}\right)|n_t = n_{t-1}, x(n_{t-1},1), \bar{\omega}_t^{kc}\right] \\
&\underbrace{- \mathbb{E}_t\left[\bar{V}^k\left(x_{it+1}(n_t,\bar{\tau}=2),\bar{\omega}_{t+1}^{kc}\right)|n_t = n_{t-1}, x(n_{t-1},1), \bar{\omega}_t^{kc}\right]\Bigg]}_{\text{Difference in continuation values w.r.t. } \bar{\tau}_{t-1}}
\end{aligned}
$$

and where $MC_t^k(n, n_{t-1})$ is defined as usual. Equation A.4 shows that if the difference in conditional values across tenure status (τ_{t-1}) from staying in one's origin neighborhood is independent of one's origin neighborhood, n_{t-1}, then we could estimate our first-step multinomial choice model using A.4.[5] It is, however, unlikely that the difference in continuation values with respect to neighborhood tenure is independent of one's neighborhood origin. To understand why, consider two neighborhoods, n and n', in the same CBSA. If neighborhood n provides lower utility than neighborhood n', then *all* residents of neighborhood n in period t are less likely to remain there in period $t+1$ relative to residents in neighborhood n'. When all residents are unlikely to stay in a given neighborhood, the difference in continuation values for incumbent residents with different residential tenures will be small. The converse is true for neighborhoods offering higher utilities to residents.

Equation 1.11 must, therefore, capture how the difference in continuation values across neighborhood tenures varies with neighborhoods' mean utilities. We, therefore, augment equation A.4 by incorporating an interaction term between the value of residential tenure and time-varying neighborhood utilities:

$$
v_n^k\left(x_{it}, \bar{\omega}_t^{kc}\right) = \gamma_{nt}^k + \mu_{\bar{\tau}} \cdot \mathbb{1}\{n = n_{t-1})\} + \gamma_{nt}^k + \lambda_{\bar{\tau}} \cdot \mathbb{1}\{n = n_{t-1})\} - MC_t^k(n, n_{t-1}).
$$

[5]When $\bar{\tau}_{t-1} = 2$, the form of γ_{nt}^k remains the same, but $\mu_{\bar{\tau}}$ is no longer scaled by $f^{\bar{\tau}}(2|n_t, x(n_{t-1},1), \bar{\omega}^{kc})$. The current argument, therefore, remains unchanged when considering $\bar{\tau}_{t-1} = 2$.

This is the expression that appears in equation 1.11 and we believe approximates the data generating process implied by our dynamic model well.

A.4.3 Microfounding Job Market Access and Expected Income

Setup The workplace choice problem for college graduates is identical to the problem for each type-k household, as reported in Section 1.4.2. In this section, the index k also includes college graduates. For now, we assume that the labor market for each k-type is segmented.[6] Recall that in Section 1.4.2, each household's workplace choice problem is to choose which neighborhood to work in to maximize their income net of commute costs:

$$\begin{aligned}\bar{I}^k_{n,t} &\equiv b^k_t \cdot \max_m \frac{z^i_{m,t}}{d_{n,m}} w_{mt},\\ &= \max_m \frac{z^i_{m,t}}{d_{n,m}} w^k_{m,t},\end{aligned}$$

To construct our JMA instrument, we simply amend this workplace choice problem by differentiating wages across industries, I, so that the household's conditional workplace choice problem becomes,

$$\begin{aligned}\bar{I}^k_{n,t} &\equiv \max_{m,I} \frac{\zeta^k_{t,I} z^i_{m,t}}{d_{h,m}} w^k_{m,t},\\ &= \max_{m,I} \frac{z^i_{m,t}}{d_{h,m}} w^k_{m,t,I},\end{aligned}$$

where $z^i_{m,t,I}$ is distributed iid Frechet, $F(z^i_{m,t,I}) = \exp\left(-\left(z^i_{m,t,I}\right)^{-\epsilon^k_c}\right)$ with $\epsilon^k_c > 1$, for each workplace-tract-industry in the city and for each worker i of education group k. $\zeta^k_{t,I}$ is an industry- and type-specific productivity shock that captures the comparative advantage of workers of each type across industries. The probability of worker i of education group k

[6] Baum-Snow *et al.* (2019) show how to extend the model to an integrated labor market with multiple types. The resulting expressions are identical to those we derive when assuming markets are segmented. The derivations in this section follow an established literature microfounding measures of "market access" in economic fundamentals (Donaldson and Hornbeck, 2016; Tsivanidis, 2022; Baum-Snow and Han, 2023).

living in tract n taking a job in tract m is then given by,

$$\pi^k_{m|n,t} = \frac{\sum_I (w^k_{m,t,I}/d^k_{n,m})^{\epsilon^k_c}}{\sum_I \sum_{m'} (w^k_{m',t,I}/d^k_{n,m'})^{\epsilon^k_c}},$$
$$= \frac{\sum_I (w^k_{m,t,I}/d^k_{n,m})^{\epsilon^k_c}}{RMA^k_{n,t}},$$

where $RMA^k_{n,t} \equiv \sum_I RMA^k_{n,t,I} \equiv \sum_I \sum_{m'} (w^k_{m',t,I}/d^k_{h,m'})^{\epsilon^k_c}$.

Define labor supply to tract m in time t by $\ell^k_{m,t} = \sum\limits_I \left[\left(w^k_{m,I} \right)^{\epsilon^k_c} \right] FMA^k_{m,t}$, where $FMA^k_{m,t}$, represents the access firms in tract m have to k-type workers. Equating labor supply of k-type workers to tract m in period t to workers' choice probabilities yields an expression for FMA in terms of RMA:

$$\begin{aligned}
\ell^k_{m,t} &= \sum_n \pi^k_{m|n,t} \cdot \pi^k_{n,t} \cdot N^k \\
&= \sum_n \frac{\sum_I (w^k_{m,t,I}/d^k_{n,m})^{\epsilon^k_c}}{RMA^k_{n,t}} \cdot \pi^k_{n,t} \cdot N^k \\
&= \sum_I \left[\left(w^k_{m,t,I} \right)^{\epsilon^k_c} \right] N^k \sum_n \frac{\left(\left(\pi^k_{n,t} \right)^{1/\epsilon^k_c} / d^k_{h,m} \right)^{\epsilon^k_c}}{RMA^k_{n,t}} \\
&\equiv \sum_I \left[\left(w^k_{m,t,I} \right)^{\epsilon^k_c} \right] FMA^k_{m,t} \qquad \text{(A.5)}
\end{aligned}$$

The penultimate equality obtains,

$$FMA^k_{m,t} = N^k \sum_n \frac{\left(\left(\pi^k_{n,t} \right)^{1/\epsilon^k_c} / d^k_{h,m} \right)^{\epsilon^k_c}}{RMA^k_{n,t}}.$$

Furthermore, dividing both sides of A.5 by $\left(d^k_{h,m} \right)^{\epsilon^k_c}$ and summing over m yields an expression for $RMA^k_{n,t}$ in terms of $FMA^k_{m,t}$. We subsequently obtain the following system of equations for RMA and FMA:

$$FMA^k_{m,t} = N^k \sum_n \frac{e^{-\kappa \epsilon^k_c \tau^k_{n,m}} \pi^k_{n,t}}{RMA^k_{n,t}}$$

$$RMA^{k}_{n,t} = \sum_{m} \frac{e^{-\kappa\epsilon^{k}_{c}\tau^{k}_{n,m}}\ell^{k}_{m,t}}{FMA^{k}_{m,t}}, \tag{A.6}$$

where we have defined $d^{k}_{n,m} \equiv e^{\kappa\tau^{k}_{n,m}}$.[7]

Job Market Access Instrument To transparently relate our job market access instruments to the recent advances in the quasi-experimental shift-share literature, we take linear approximations of equation A.6 to obtain our measure of job market access that appears in expression 1.15:

$$JMA_{ndt} = \sum_{m \in \mathcal{N}^{c} \setminus n} e^{-\eta^{c}\tau_{nm}} l_{mdt},$$

This linear approximation of equation A.6 relates to earlier notions of market potential in international and regional trade theory, which conceptualize the demand for goods in a given region as the sum of demands in surrounding regions, weighted by bilateral transportation costs (Harris, 1954; Hanson, 2005). It is common in the regional economics literature to take linear approximations of structural measures of market access (e.g., Donaldson and Hornbeck (2016) and Herzog (2021)).[8]

Expected Income Our workplace choice model also implies that expected income discounted by commuting costs for type-k households prior to drawing the vector of neighborhood- and period-specific productivity shocks is given by,

$$\bar{I}^{k}_{nt} = \Gamma\left(1 - \frac{1}{\epsilon^{k}_{c}}\right)\left(RMA^{k}_{h}\right)^{1/\epsilon^{k}_{c}}, \quad \forall i \in k,$$

where we directly use the identify $RMA^{k}_{n,t} \equiv \sum_{I}\sum_{m'}(w^{k}_{m',t,I}/d^{k}_{h,m'})^{\epsilon^{k}_{c}}$ in this measure's construction. This is the expression that enters into the households' flow utilities.

We use the observed distribution of workplace wages among our sample of low-income renter households to construct $RMA^{k}_{n,t}$. We then regress observed workplace wages dis-

[7] $1 - e^{-\kappa\tau^{k}_{n,m}}$ represents the portion of time that type-k workers in tract n spend commuting to tract m.

[8] An alternative approach to taking a linear approximation of A.6 would be to solve for the fixed point of $RMA^{k}_{n,t}$ and $FMA^{k}_{m,t}$. However, it is unclear exactly how to relate the findings of Goldsmith-Pinkham *et al.* (2020) and Borusyak *et al.* (2022) on linear shift-share instrumentation to such a setting.

counted by commute costs on our measure of $RMA^k_{n,t}$ and predict $\bar{I}^k_{nt}$ solely using variation in $RMA_{n,t}$. This is to account for the fact that larger labor markets have mechanically higher levels of $RMA^k_{n,t}$.

Gravity and Forecasting Equations Our workplace choice model yields gravity equations we can use to estimate $\kappa\epsilon^k_c$. We follow Baum-Snow *et al.* (2019) and estimate $\kappa\epsilon^k_c$ separately for each type-k household in each city, c. Estimating $\kappa\epsilon^k_c$ separately for each type-k household in each CBSA increases the accuracy of our $RMA^k_{n,t}$ measures and thus also the power of our job market access instruments. CBSA-specific estimates allow labor demand shocks to impact job market access in neighborhoods accessible by longer commutes more in CBSAs where ϵ^k_c is lower. To obtain our estimates of $\eta^{kc} \equiv \kappa\epsilon^k_c$ we take the log of $\pi^k_{m|n,t}$ to obtain the following gravity equation,

$$\begin{aligned} \ln(\pi^k_{m|n,t}) &= \ln(RMA^k_{n,t}) + \ln\left(\sum_I \left(w^k_{m,t,I}\right)^{\epsilon^k_c}\right) - \kappa\epsilon^k_c\tau^k_{n,m} \\ &= \alpha^k_{n,t} + \rho^k_{m,t} - \underbrace{(\kappa\epsilon^k_c)}_{\eta^{kc}}\tau^k_{n,m}, \end{aligned} \tag{A.7}$$

which we estimate separately for each city and type-k households (including separately for college graduates) using 2010 commute flows constructed using the LEHD.

We obtain estimates of $\tau^k_{n,m}$ using the median tract commute time for each type-k worker in the 2005-2015 ACS surveys between tracts n and m for all tract pairs reported to have positive commute flows for any k-type worker. Since there are far more potential tract-to-tract commuting routes than ACS survey respondents, many commute routes are missing observed commute times. To estimate the remaining commute times, we follow Baum-Snow *et al.* (2019) and construct an empirical forecasting model to predict commute times between all neighborhood pairs using the distance between neighborhood centroids and the corresponding city's CBD. We estimate the following forecasting model separately for each type-k worker,

$$\ln \tau^k_{n,m} = \alpha^k_d \ln \text{Distance}_{n,m} + \alpha^k_r \ln(\text{Home CBD Dis})_n + \alpha^k_w \ln(\text{Work CBD Dis})_m + v_c + u^k_{h,m,c}.$$

With these predicted commute times in hand, we use the observed tract-to-tract commute flows for each type-k household observed in the LEHD to estimate equation A.7.[9] As per Census Bureau Disclosure Guidance, we cannot release our CBSA-specific estimates but report summary statistics of our CBSA-level estimates in Table A.4. We also report our estimates using distance—as opposed to commute time—between commuting tract pairs as a comparison. Our estimates suggest that a one-minute increase in commute time leads to an 8.8 or 8.9 percent reduction in the flow of Black and Non-Black commuters, respectively. While on the higher side, these semi-elasticities are consistent with the magnitudes found in the existing literature (e.g. Ahlfeldt *et al.* (2015)).

Table A.4. *Commute Elasticities*

	Time		Distance	
	Black	Non-Black	Black	Non-Black
Median CBSA Estimate	0.0876***	0.0885***	0.0807***	0.0782***
	(0.0000)	(0.0000)	(0.0001)	(0.0001)
Standard Deviation of CBSA Estimates	0.01353	0.01528	0.01667	0.07824
N (1,000s)	966	3,752	966	3,752

Notes: Table A.4 reports results from the gravity equations of equation A.7. We use Poisson Pseudo Maximum Likelihood to estimate these models given the prevalence of tract-to-tract potential commuting routes with zero observed commute flows (Silva and Tenreyro, 2006). We weight observations by the number of commuters in origin census tracts. To increase the number of non-zero commute flows we observe in the data, we estimate these gravity equations for all low-income households regardless of their rental status. We measure time in minutes and distance in miles. The first row reports the (pseudo) median CBSA-level estimate and its accompanying standard error in parentheses. The second row reports the standard deviation of our 50 CBSA-level estimates. *N* reports the underlying number of commuters used to construct the commute flows in the LEHD. *Sources*: LEHD (2010). Estimates were disclosed by the US Census Bureau's Disclosure Review Board. Project Number 2358. Disclosure Clearance Number CBDRB-FY24-P2358-R10957.

[9] To maximize the size of the sample used to construct our commute flows, we do not restrict our sample to renter households.

A.4.4 Welfare Analysis

Recall that our baseline measure of change in expected welfare for incumbent renters is given by,

$$\Delta W^k(n,\bar{\tau}) = \frac{\bar{V}^k(x=(n,\bar{\tau}),\bar{\omega}^{kc}_{2000}) - V^k_{ss}(n,\bar{\tau})}{\beta^k_r} \quad \forall\, k,n.$$

We now unpack $\Delta W^k(n,\bar{\tau})$, starting with $V^k_{ss}(n,\bar{\tau})$. $V^k_{ss}(n,\bar{\tau})$ is the expected welfare for an incumbent renter with residential tenure $\bar{\tau}$ residing in neighborhood n during 2000 under the assumption that the economy is in steady state. We calculate $V^k_{ss}(n,\bar{\tau})$ by finding time-invariant continuation values and levels of exogenous amenities, $\{\xi^k_n\}_{n,k}$, that induce a *stationary distribution* among our sample of low-income renter households when holding all the observed state variables fixed at their 2000 levels, i.e. $\omega^{kc}_t = \omega^{kc}_{2000}\,\forall t$.

Stationary Distribution Denote $q^k_n(\bar{\tau})$ as the model-implied share of type-k households with tenure $\bar{\tau}$ living in neighborhood n. Stacking $q^k_n(\bar{\tau})$ yields the $2(N^c+1)-$dimensional vector $\boldsymbol{q}^k$. We define a stationary equilibrium among our sample of low-income renter households as $\boldsymbol{q}^k = \boldsymbol{\Lambda}^k\boldsymbol{q}^k$, where $\boldsymbol{\Lambda}^k$ is a transition probability matrix constructed so the distribution of type-k households' locations evolve as,

$$q^k_n(\bar{\tau}') = \begin{cases} \sum\limits_{\bar{\tau}'}\sum\limits_{\bar{n}\neq n}\left[q^k_{\bar{n}}(\bar{\tau}')p^k_{n,t}\left(x(\bar{n},\bar{\tau}')\right)\right] + q^k_{\bar{n}}(\bar{\tau}=1)p^k_{n,t}\left(x(n,\bar{\tau}=1)\right)(1-g^k_n) & \text{if} \quad \bar{\tau}'=1 \\ q^k_{\bar{n}}(\bar{\tau}=1)p^k_{n,t}\left(x(n,\bar{\tau}=1)\right)g^k_n + q^k_{\bar{n}}(\bar{\tau}=2)p^k_{n,t}\left(x(n,\bar{\tau}=2)\right) & \text{if} \quad \bar{\tau}'=2. \end{cases}$$

Steady-State Equilibrium We define a steady-state equilibrium in the year 2000 as a stationary distribution among our sample of low-income renter households, where continuation values are time-invariant and state variables are fixed at their 2000 level: $\bar{\omega}^{kc}_t = \bar{\omega}^{kc}_{2000}\,\forall t$. The time-invariant continuation values in our steady-state equilibrium are precisely our steady-state measures of expected welfare, $\{V^k_{ss}(n,\bar{\tau})\}_{n,\bar{\tau}}$. We outline the procedure to compute these measures as well as the vector of exogenous amenities in Algorithm 1.[10]

[10] In step 6 of Algorithm 1, we exploit the fact that the values of exogenous amenities are the same for households regardless of their length of tenure. While we cannot guarantee Algorithm 1 induces a stationary equilibrium among type-k households with a high level of neighborhood tenure, $\bar{\tau}=2$, we find that in practice

Algorithm 1 Compute Steady State Expected Welfare

1: Guess value of unobserved neighborhood amenities: $\{\xi_n^k\}_n$.

Compute continuation values:

2: Guess $\{V_{ss}^k(n,\bar{\tau})\}_{n,\bar{\tau}}$.

3: Compute $V_{ss}'^k(n,\bar{\tau}) = \ln\left(\sum_{n'\in\mathcal{N}^c} \exp\left(\bar{u}_{n'}^k\left(x(n,\bar{\tau}),\bar{\omega}_{2000}^{kc}\right) + \delta V_{ss}^k(n',\bar{\tau})\right)\right) + \gamma, \quad \forall\, n,\bar{\tau}$

4: Set $V_{ss}^k(n,\bar{\tau}) = V_{ss}'^k(n,\bar{\tau}), \quad \forall\, n,\bar{\tau}$

Repeat steps 2 - 4 until $\max_{n,\bar{\tau}}\{|V_{ss}^k(n,\bar{\tau}) - V_{ss}'^k(n,\bar{\tau})|\} < \epsilon_V$ for some $\epsilon_V > 0$.

Compute exogenous amenities:

5: Compute the probability a type-k household chooses n given $\{V_{ss}^k(n,\bar{\tau})\}_{n,\bar{\tau}}$ and $\{\xi_n^k\}_n$

$$p_n^k\left(x(\bar{n},\bar{\tau}),\bar{\omega}_{2000}^{kc}\right) = \frac{\exp\left(v_n^k\left(x(\bar{n},\bar{\tau}),\bar{\omega}_{2000}^{kc}\right)\right)}{\sum_{n'\in\mathcal{N}^c}\exp\left(v_{n'}^k\left(x(\bar{n},\bar{\tau}),\bar{\omega}_{2000}^{kc}\right)\right)}, \quad \forall\,\bar{\tau},\bar{n}$$

6: Update exogenous amenities using observed neighborhood shares of households with low tenure, $q_n^k(\bar{\tau}=1)$:

$$\begin{aligned}\xi_n'^k = \xi_n^k + \ln\left(\hat{q}_n^k(\bar{\tau}=1)\right) - \ln\Bigg(&\sum_{\bar{\tau}'}\sum_{\bar{n}\neq n}\left[q_{\bar{n}}^k(\bar{\tau}')p_{n,t}^k\left(x(\bar{n},\bar{\tau}')\right)\right] \\ &+ q_{\bar{n}}^k(\bar{\tau}=1)p_{n,t}^k\left(x(n,\bar{\tau}=1)\right)(1-g_n^k)\Bigg)\end{aligned} \tag{A.8}$$

Set $\xi_n^k = \xi_n'^k$.

Repeat steps 2 - 6 until $\max_{n,\bar{\tau}}\{|\xi_n^k - \xi_n'^k|\} < \epsilon_\xi$ for some $\epsilon_\xi > 0$.

Take $\{V_{ss}^k(n,\bar{\tau})\}_{n,\bar{\tau}}$ as our measures of steady-state expected welfare.

We now turn to $\bar{V}^k(x=(n,\bar{\tau}),\bar{\omega}_{2000}^{kc})$, which is the expected welfare of a type-k household with tenure $\bar{\tau}$ residing in neighborhood n in the year 2000. In contrast to $V_{ss}^k(n,\bar{\tau})$, we do not assume the economy is in a steady-state equilibrium when constructing this measure. We instead compute $\bar{V}^k(x=(n,\bar{\tau}),\bar{\omega}_{2000}^{kc})$ via backward induction starting in the year 2019. We do, though, hold exogenous neighborhood amenities fixed at the values found in Algorithm 1. By holding the values of exogenous amenities fixed at their 2000 levels, we can attribute variation in $\Delta W^k(n,\bar{\tau})$ to neighborhood-level changes in rents and college shares.

it closely approximates a stationary distribution for these households. We define $\hat{q}_n^k(\bar{\tau}=1)$ as the empirical counterpart to $q_n^k(\bar{\tau}=1)$.

To initiate our backward induction solution method, we assume that the economy is again in a steady state in 2019 and calculate steady-state continuation values for 2019 as in steps 2-4 in Algorithm 1. Since we are solely concerned with expected welfare changes from the perspective of the year 2000, this assumption is innocuous given a discount rate of $\delta = 0.85$.

To compute $\bar{V}^k(x = (n, \bar{\tau}), \bar{\omega}^{kc}_{2000})$, we must also take a stance on how households form their expectations over the market-level state variables, $\bar{\omega}^{kc}_t$. One option is to assume that households have perfect foresight. We experimented with this assumption but found that households quickly reoptimize their location choices given perfect knowledge of future states. Instead, we assume that households' expectations are a weighted average of i) the future true states and ii) neighborhoods' current states multiplied by the CBSA-wide average growth rate of each state variable:

$$\mathbb{E}[\bar{\omega}^n_t | \mathcal{I}_{i,2000}] = \bar{\omega}^n_{2000} \cdot \left(1 + \frac{\sum_{n' \in \mathcal{N}^c} \left(\bar{\omega}^{n'}_t - \bar{\omega}^{n'}_{2000} \right)}{\sum_{n' \in \mathcal{N}^c} \bar{\omega}^{n'}_{2000}} \right) \cdot \mu + \bar{\omega}^n_t \cdot (1 - \mu). \tag{A.9}$$

$\mu \in [0, 1]$ determines how accurate households' beliefs are. With $\mu = 0$, households have perfect foresight, and with $\mu = 1$, households believe all neighborhoods' state variables change at an identical rate. When $\mu = 1$, households are less mobile as they expect the relative levels of neighborhoods' flow utilities to remain constant over time. We calibrate μ to match the observed average neighborhood out-migration rates in our sample of low-income renter households (we assume μ is constant across household types).[11]

Throughout our welfare analyses, we use publicly available 2010-delineated tract-level Census Survey data from IPUMS National Historical GIS (NHGIS) to compute annual rents and the distribution of households across Census tracts (Manson *et al.*, 2022). We linearly interpolate these data between survey years (2000, 2007-2019). From 2007 onward, we used the 5-year ACS aggregates to compute the tract-level data. As the publicly available tract-level data are not disaggregated enough to compute the exact shares each type-k low-income renter household across census tracts and aggregated tenure states, $\bar{\tau}$, we approximate the

[11]Note that equation A.9 is consistent with our rational expectations assumption (Assumption 3) in so far as the true data generating process renders equation A.9 unbiased.

empirical neighborhood shares used in Algorithm 1, $q_n^k(\bar{\tau} = 1)$, with the share of Black (non-Black) non-college-educated renter households living in each tract who have a tenure of less than three years.

Appendix B

Appendix to Chapter 3

B.0.1 Policy Equilibrium Uniqueness

Unique fixed point with observed heterogeneity. (Lemma 1) Assume the the model-predicted share of workers of type (ω, a) choosing to reside in district n can be written as follows,

$$\pi_n^{\omega,a} \equiv \frac{1}{I^{\omega,a}} \sum_{i \in I^{\omega,a}} \frac{\exp\left(\beta^i\, \delta_n^{\omega,a} + x_n^i\right)}{1 + \sum\limits_{m \in N^{\omega,a}} \exp\left(\beta^i\, \delta_m^{\omega,a} + x_m^i\right)},$$

where $\beta^i \in (0, 1]\ \forall i$.

Consider the function $f(\delta^{\omega,a}) = \delta^{\omega,a} + \ln\left(\boldsymbol{\Omega}/I^{\omega}\right) - \ln(\boldsymbol{\pi}^{\omega,a}(\delta^{\omega,a}))$, where $\boldsymbol{\Omega}/I^{\omega}$ is the vector of shares of ω-type workers in allocated to each district in the proposed policy equilibrium. It follows that there is a unique fixed point, $\delta^{*\omega,a}$, to f in $\mathbb{R}^{N^{\omega,a}}$.

Proof: This setup satisfies the three conditions in the contraction mapping proof in Berry, Levinsohn, and Pakes (1995):

1. *Monotonicity*: Making use of the fact that $\beta^i \in (0, 1]\ \forall i$, some simple calculus and algebra yields,

$$\frac{\partial f_m}{\partial \delta_m^{\omega,a}} \geq 0$$

and

$$\sum_m \frac{\partial f_m}{\partial \delta_m^{\omega,a}} < 1$$

2. *Lower Bound*: Note that when $\delta^{\omega,a} << 0, \forall n$

$$\frac{1}{I^{\omega,a}} \int_{i \in I^{\omega,a}} \frac{\exp\left(\delta_n^{\omega,a} + x_n^i\right)}{1 + \sum\limits_{m \in N^{\omega,a}} \exp\left(\beta^i \, \delta_m^{\omega,a} + x_m^i\right)} \leq \frac{1}{I^{\omega,a}} \sum_{i \in I^{\omega,a}} \frac{\exp\left(\beta^i \, \delta_n^{\omega,a} + x_n^i\right)}{1 + \sum\limits_{m \in N^{\omega,a}} \exp\left(\beta^i \, \delta_m^{\omega,a} + x_m^i\right)},$$

since $\beta^i \leq 1 \quad \forall i$. One can rewrite the left-hand-side of this expression as

$$\pi_n^{\omega,a}(\delta_n^{\omega,a}) = \exp(\delta_n^{\omega,a})\tilde{D}_n(\delta_n^{\omega,a})$$

where,

$$\tilde{D}_n(\delta_n^{\omega,a}) \equiv \frac{1}{I^{\omega,a}} \sum_{i \in I^{\omega,a}} \frac{\exp\left(x_n^i\right)}{1 + \sum\limits_{m \in N^{\omega,a}} \exp\left(\beta^i \, \delta_m^{\omega,a} + x_m^i\right)}.$$

Plugging this into the definition of f gives

$$f(\delta_n^{\omega,a}) = \ln\left(\Omega_n / I^{\omega_L}\right) - \ln(\tilde{D}_n(\delta_n^{\omega,a})).$$

Note that $\tilde{D}_n(\delta_n^{\omega,a})$ is declining in all the $\delta_n^{\omega,a}$'s. As all of the $\delta_n^{\omega,a}$'s approach $-\infty$, the above inequality is satisfied and $\tilde{D}_n(\delta_n^{\omega,a})$ goes to $\frac{1}{I^{\omega,a}} \sum\limits_{i \in I^{\omega,a}} \exp\left(x_n^i\right)$. Thus a lower bound for f_n is $\underline{\delta}_n^{\omega,a} \equiv \ln\left(\Omega_n / I^{\omega}\right) - \ln\left(\frac{1}{I^{\omega,a}} \sum\limits_{i \in I^{\omega,a}} \exp\left(x_n^i\right)\right) > -\infty$.

3. *Upper Bound*: Berry (1994) shows that an appropriate upper bound, $\bar{\delta}^{\omega,a}$ is found as follows. For district n, define $\bar{\delta}_n^{\omega,a}$ as the value of $\delta_n^{\omega,a}$ that would explain the share of workers of type (ω, a) choosing the outside option, $\pi_{oo}^{\omega,a}$, when $\delta_{oo}^{\omega,a} = 0$ and all the other $\delta_n^{\omega,a} = -\infty$. Then set $\bar{\delta} > \max\limits_m \bar{\delta}_m^{\omega,a}$

Thus, by the contraction mapping proof in Berry, Levinshon, and Pakes (1995), there is a unique fixed point, $\delta^{*\omega,a}$, to f in $\mathbb{R}^{N^{\omega,a}}$ ■

Proof of Conditional Uniqueness I structure the proof in five main steps. Each step yields unique equilibrium objects and policy instruments, thus demonstrating the uniqueness of any policy equilibrium. I first assume that the proposed policy equilibrium is interior ($\{H_n, L_n\}_n \in [(0, I^{\omega_H}) \times (0, I^{\omega_L})]^N$), and discuss the case where $\{H_n, L_n\}_n \notin [(0, I^{\omega_H}) \times (0, I^{\omega_L})]^N$ separately afterwards.

1. **Equilibrium Prices:** Conditional on $\boldsymbol{\theta}$, $\{H_n, L_n\}_n$, and $T(w)$, firm optimization and perfect competition in the housing market yield closed-form solutions for both $\boldsymbol{r}$ and $\boldsymbol{w}$.

2. **Non-parents' Choice Probabilities:** Conditional on $\boldsymbol{\theta}$, $\{H_n, L_n\}_n$, $T(w)$, $\boldsymbol{r}$ and $\boldsymbol{w}$, one has closed-form expressions for non-parents' district choice probabilities,

$$\pi_n^{\omega} = \frac{1}{I^{\omega,np}} \int_{i \in I^{\omega,np}} \frac{\exp\left(v_n^i\right)}{1 + \sum_{m \in N^{\omega,np}} \exp\left(v_m^i\right)},$$

where $N^{\omega,np}$ is the set of districts that are affordable for (ω, np)-type workers. There may indeed by districts that are unaffordable for low-income workers; i.e., $C_n^{\omega^L,np} < 0$ for some n.

3. **Parents' Choice Probabilities:** The total mass of each college- and non-college-educated worker in each school district can be expressed as a sum of the corresponding choice probabilities of parents and non-parents,

$$\Omega_n = \underbrace{\int_{i \in I^{np,\omega}} \pi_n^i}_{\text{Mass of Non-Parents in } n} + \underbrace{\int_{i \in I^{p,\omega}} \pi_n^i}_{\text{Mass of Parents in } n}, \text{ for } \Omega_n \in \{H_n, L_n\}$$

where $\pi_n^i = 0$ if $C_n^{\omega(i),a(i)} \leq 0$ and $\frac{exp(v_n^i)}{1+\sum_{m \in N^{\omega,a}} v_m^i}$ otherwise. Thus, one can write parents' choice probabilities that are consistent with the proposed policy equilibrium as a function of Ω_n and non-parents' school district shares, $\pi_n^{\omega,np}$. That is, $\tilde{\pi}_n^{\omega,p} = \frac{1}{I^{\omega,p}} \int_{i \in I^{\omega,p}} \pi_n^i$, where $\tilde{\pi}_n^{\omega,p} \equiv \frac{\Omega_n}{I^{\omega,p}} - \frac{1}{I^{\omega,p}} \int_{i \in I^{\omega,np}} \pi^i$.

4. **School Spending:** College-educated parents' choice probabilities can be written as,[1,2]

$$\tilde{\pi}_n^{\omega_H,p} = \frac{1}{I^{\omega_H,p}} \int_{i \in I^{\omega_H,p}} \frac{\exp\left(\tilde{\bar{\beta}}^i_{\mathcal{N}}\, \delta_n^{\omega_H,p} + y^i_n\right)}{1 + \sum\limits_{m \in N_{\omega_H,p}} \exp\left(\tilde{\bar{\beta}}^i_H\, \delta_m^{\omega_L,p} + y^i_m\right)},$$

where $\tilde{\bar{\beta}}^i_{\mathcal{N}} \equiv \frac{\bar{\beta}^i_{\mathcal{N}}}{\max\limits_{i \in I^{\omega_H,p}}\{\bar{\beta}^i_{\mathcal{N}}\}}$, $\delta_n^{\omega_H,p} \equiv \max\limits_{i \in I^{\omega_H,p}} \{\beta^i_c\} \cdot \log(\mathcal{N}_n^{\omega_H})$, and, $y^i_n \equiv \bar{\beta}^i_c \log(C_n^{\omega_H}) + \beta^i_A \log(A^i_n)$.[3] Then, since $\tilde{\bar{\beta}}^i_{\mathcal{N}} \in (0,1)$, by the Lemma there exists a unique vector $\boldsymbol{\delta}^{*,\omega_H,p}$ that matches the model-predicted shares of college-educated parents to $\tilde{\pi}_n^{\omega_H,p}$. Upon obtaining the unique vector $\boldsymbol{\delta}^{*,\omega_H,p}$ recursively via a contraction mapping, one can equate the vector to its model equivalent to obtain the following set of equations:

$$\delta_n^{*\omega_H,p} = \max_{i \in I^{\omega_H,p}} \{\bar{\beta}^i_{\mathcal{N}}\} \cdot \log\left((\gamma_s^{\omega_H} S_n^{\rho^{\omega_H}} + (1-\gamma_s^{\omega_H})\bar{q}_n^{\rho^{\omega_H}})^{1/\rho^{\omega_H}} \cdot \exp\left(\Xi_n^{\omega_H}\right)\right) \; \forall n.$$

Since each $\delta_n^{*,\omega_H,p}$ is monotonic in the level of district-specific school spending, s_n, inverting $\boldsymbol{\delta}^{*,\omega_H,p}$ yields the unique levels of school-spending that implement the proposed allocation of college-educated workers across school districts:[4]

$$s^*_n = \left[\left[\left[\frac{\exp\left(\delta_n^{*\omega_H,p} / \max\limits_{i \in I^{\omega_H,p}}\{\bar{\beta}^i_{\mathcal{N}}\}\right)}{\exp\left(\Xi_n^{\omega_H}\right)}\right]^{\rho^{\omega_H}} - (1-\gamma_s^{\omega_H})\bar{q}_n^{\rho^{\omega_H}}\right] \frac{1}{\gamma_s^{\omega_H}}\right]^{1/\rho^{\omega_H}} - \frac{H_n + L_n}{I_{\omega_H,p} \cdot \tilde{\pi}_n^{\omega_H,p} + I_{\omega_L,p} \cdot \tilde{\pi}_n^{\omega_L,p}} \tau^s_n r_n \; \forall n.$$

5. **Housing Assistance:** Observe that one can write non-college-educated parents' choice

---

[1] If $\tilde{\pi}_n^{\omega,p} \notin (0,1)$, then no such policy equilibrium exists given $\boldsymbol{\theta}$, $\{H_n, L_n\}_n$, and $T(w)$.

[2] $\tilde{\pi}_n^{\omega,p}$ can never equal to (or be greater than) 1 given that workers' idiosyncratic preferences have an unbounded support and the outside option is always available to all workers. Moreover, if $\tilde{\pi}_n^{\omega,p} \leq 0$, this would imply that non-parents choose to live in school district n, but parents cannot afford to do so. Since parents' incomes are assumed to be at least as high as non-parents (given the non-negativity constraint on housing housing), such an outcome is impossible with an unbounded support for idiosyncratic district preferences.

[3] Note that all the components of y^i_n are known conditional on $\boldsymbol{\theta}$, $\{H_n, L_n\}_n$, $T(w)$, and $\{\boldsymbol{r}, \boldsymbol{w}\}$.

[4] Note also that $\tilde{\pi}_n^{\omega,p}$ is obtained in the previous step.

probabilities as

$$\tilde{\pi}_n^{\omega_L,p} = \frac{1}{I^{\omega_L,p}} \int_{i \in I^{\omega_L,p}} \frac{\exp\left(\tilde{\bar{\beta}}_c^i\, \delta_n^{\omega_L,p} + x_n^i\right)}{1 + \sum_{m \in N^{\omega_L,p}} \exp\left(\tilde{\bar{\beta}}_c^i\, \delta_m^{\omega_L,p} + x_m^i\right)},$$

where $\tilde{\bar{\beta}}_c^i \equiv \frac{\bar{\beta}_c^i}{\max_{i \in I^{\omega_L,p}}\{\bar{\beta}_c^i\}}$, $\delta_n^{\omega_L,p} \equiv \max_{i \in I^{\omega_L,p}} \{\beta_c^i\} \cdot \log(c_{\omega_L,p,n})$, and, $x_n^i \equiv \bar{\beta}_{\mathcal{N}}^i \log(\mathcal{N}_n^{\omega}) + \beta_A^i \log(A_n^i)$.[5] Then, since $\tilde{\bar{\beta}}_c^i \in (0,1)$, by the Lemma there exists a unique vector $\delta^{*\omega_L,p}$ that matches the model-predicted shares of non-college-educated parents to $\tilde{\pi}_n^{\omega_L,p}$. Upon obtaining the unique vector $\delta^{*\omega_L,p}$ recursively via a contraction mapping, one can equate the vector to its model equivalent to obtain the following set of equations:

$$\delta_n^{*\omega_L,p} = \max_{i \in I^{\omega_L,p}} \{\bar{\beta}_c^i\} \cdot \log\left(w_{\omega_L} - T(w_{\omega_L}) - r_n(1 + \tau_n^s - \tau_n) + \Pi^{\omega_L} R\right)\ \forall n.$$

Since each $\delta_n^{*\omega_L,p}$ is monotonic in the level of district-specific housing assistance, τ_n, inverting $\delta^{*\omega_L,p}$ yields the unique levels of housing assistance that implements the proposed allocation of non-college-educated workers across school districts,

$$\tau_n^* = \left[\exp\left(\frac{\delta_n^{*\omega_L,p}}{\max_{i \in I^{\omega_L,p}}\{\bar{\beta}_c^i\}}\right) - w_{\omega_L} + T(w_{\omega_L}) + r_n(1 + \tau_n^s) - \Pi^{\omega_L} R\right] \frac{1}{r_n}\ \forall n.$$

The proof so far has assumed that the proposed policy equilibrium is interior: $\{H_n, L_n\}_n \in [(0, I^{\omega_H}) \times (0, I^{\omega_L})]^N$. I now discuss how the proof changes when $\{H_n, L_n\}_n \notin [(0, I^{\omega_H}) \times (0, I^{\omega_L})]^N$. First note that $H_n, L_n \geq I^{\omega}$ for some n is never feasible since all workers' isiosyncratic preferences for school districts (including their idiosyncratic preference for the outside option) have unbounded support, and so there will always be some positive mass of each type of worker choosing the outside option. It is also infeasible for $H_n = 0$ for any n, since, again, workers' idiosyncratic school district preferences have unbounded support. Moreover, the policy equilibrium requires that $C_n^{\omega_H,a} > C_n^{\omega_L,a}\ \forall\, n, a$. Therefore, if it is unaffordable for college-educated workers to live in district n, then it is also unaffordable for

[5]Note that the value of x_n^i is known conditional on $\boldsymbol{\theta}$, $\{H_n, L_n\}_n$, $T(w)$, $\{\boldsymbol{r}, \boldsymbol{w}\}$, and $\{S_n\}_n$

non-college-educated workers to live in district n, but then rents would equal 0, contradicting the fact that it is unaffordable for any type to live in the school district.

It may, however, be the case that $L_n = 0$ for some (and potentially many) n. This is because of the positive consumption constraint; non-college-educated workers (both parents and non-parents) may simply be unable to afford to live in a given school district. When this is the case, the proof is identical to the interior case, aside from setting $L_n = 0$ for the unaffordable districts in the proposed policy equilibrium, and taking these school districts out of the non-college-educated workers' choice sets. When evaluating a proposed equilibrium with $L_n = 0$ one must check ex-post that the districts deemed unaffordable to the non-college-educated workers are indeed so. This completes the proof ■