Essays on Places and Economic Inequality

by

Mike Zabek

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2018

Doctoral Committee:

Professor John Bound, Co-Chair
Professor Matthew Shapiro, Co-Chair
Assistant Professor Dominick Bartelme
Assistant Professor Joshua Hausman

Mike Zabek

zabek@umich.edu

ORCID iD: 0000-0002-0900-971X

DEDICATION

To Jessica Ott and to my parents. Thank you for all of your love, support, and patience.

ACKNOWLEDGEMENTS

I owe particular thanks to my advisors, Dominick Bartelme, John Bound, Joshua Hausman, and Matthew Shapiro, for patience, perspective, knowledge, and encouragement.

Several of these chapters are co-authored, and I am extremely grateful to these people for all of their work and insight. These include Pawel Krolikowski, Patrick Coate, Aditya Aladangady, and David Albouy.

Helpful academic comments from a number of people were also invaluable. These people include, but are not limited to: Martha Bailey, C. Hoyt Bleakley, Lawrence Blume, Charles Brown, Ben Farber, Pamela Giustinelli, Chris House, Chang Tai Hsieh, John Laitner, David Lam, Andrei Levchenko, Corinne Low, Lawrence Katz, Matthew Notowidigdo, Jessica Ott, Linda Tesar, Michel Serafinelli, Robert Schoeni, Christopher Smith, Melvin Stephens, Bryan Stuart, Robert Willis, Abigail Wozniak, Owen Zidar, and seminar participants at the University of Michigan, the University of Pennsylvania SSPF, H2D2 research day at the University of Michigan, the Population Association of America, the Midwest Economics Association, the Midwest Macro Association, the Urban Economics Association, the US Census Bureau, the 2017 SOLE meetings, the Barcelona GSB Summer Forum, The Federal Reserve Bank of Cleveland, Cleveland State University, and various potential employers.

Timothy Tollope and Adam Karabatakis provided excellent research assistance on various projects.

I owe very much to people who have sustatined the university in other ways. These include, but are not limited to: Laura Flack, Julie Heintz, Mary Braun, Olga Mustada, Heather MacFarland, Miriam Rahl, Lisa Neidert, Mark Champe, Mark Sandstrom, Ricardo Rodriguez, Carol Bowen, Cassandra Kelly, MaryMangum, Marc Sorace, Dan Green, Vinnie Veeraraghavan, and Nan Flood. These people have made my research possible and my time here much more enjoyable. I am grateful that they put effort into improving my time at the university.

I also am grateful to my former academic mentors. These include Nola Forbes, Sarah Calanan, Judy Kelly, Brad Hartlaub, Carol Schumacher, David Harington, Cathy Krynski, Jaret Trieber, Will Melick, Galina An, Lisa Sanbonmatsu, and Scott Schuh. Without their help, I would have never had the opportunity to write a dissertation.

The research was supported in part by NICHD center grant (R24 HD041028) to the Population Studies Center at the University of Michigan, through a Rackham Pre-Doctoral Fellowship, through a Weinberg Fellowship, through a SYLFF Fellowship, and through computational resources and services provided by Advanced Research Computing at the University of Michigan.

I also am grateful to David Dorn for sharing data and code both via his website and through personal correspondence.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

This dissertation uses changes across space to understand economic inequality. Chapter one examines people's ties to places. If people are tied to places by family and experiences, then welfare will be less equal across space. It presents evidence that this is the case, with formulae that show how this makes local policies less wasteful in declining areas. Chapter two shows that young adults, ages 25 to 35, who live in the same neighborhoods as their parents experience stronger earnings recoveries after a job displacement than those who live farther away. It presents some evidence that these differences are driven by parental help with grandchildren and perhaps help from parents in identifying new jobs. Chapter three uses changes in housing prices and rents to study inequality in standards of living. It suggests that changes in income inequality since 1930 have caused similar changes in housing expenditures, mostly through changes in the value of particular neighborhoods.

CHAPTER I

Local Ties in Spatial Equilibrium

In spatial equilibrium, the marginal person will be equally well off, no matter where she lives (Rosen (1979) and Roback (1982)). Most people are not marginal, however. Kennan and Walker (2011), for example, find that people are willing to forego around $20 thousand per year to live in the place where they were born.[1] Here, I use reduced form regressions, a parametric model, and a sufficient statistics approach to examine the consequences of people's local ties.

The main consequence is that many people are willing to stay in their birth places, even if wages are low and rents are high. So locals, or people who are locally born, make up most of the population of areas where population is declining.

People's willingness to stay in their birth places implies that equilibrium migration elasticities are lower in declining areas, where locals make up a larger share of the population. This happens for two related reasons. The first is because many locals find it worthwhile to stay in a place where they have local ties, despite declines in real wages. The second is that flows into declining areas tend to be smaller, since real wages are depressed in declining areas. Since real wages are unusually low, a smaller group of workers will find a declining place to be desirable, relative to where they are living.

The distribution of welfare depends on migration elasticities in spatial equilibrium, and migration elasticities are a sufficient statistic for evaluating a policy's effect on aggregate welfare. If migration elasticities are high, workers will move away from areas with lower real wages, and movement away from declining areas tends to increase real wages for people

[1] Large migration frictions are a common feature in the literature. Bound and Holzer (2000), Notowidigdo (2011), Molloy et al. (2011), Saks and Wozniak (2011), Ganong and Shoag (2012), Kaplan and Schulhofer-Wohl (2013), Yagan (2013), Chetty et al. (2014), and Huttunen and Salvanes (2015) (among others) have emphasized how migration appears to be surprisingly limited, particularly among people without college degrees. Several other structural models of individual migration decisions, including Coate (2017), Kleemans (2015), Oswald (2015), and Diamond (2016) have also found that people tend to have strong preferences for living close to where they were born.

who stay.[2] High migration elasticities, however, also imply that a place-based policy induce more people to live in a declining area that is receiving a place-based policy, potentially undoing any positive impacts. Understanding the magnitude of migration elasticities, then, is important for policy analysis.

To test if migration elasticities are lower in areas where a larger share of the population was born locally, I examine the impact of two labor demand shifters. The first, developed by Bartik (1991), projects national changes in employment onto local areas using employment shares in various industries based on the assumption that the initial industry shares are unrelated to changes in local labor supply (Goldsmith-Pinkham et al. (2017)). The second, developed by Autor et al. (2013), projects the impact of competition with Chinese firms onto local areas, also based on initial industry shares. After these labor demand shocks, I find that wages, labor force participation, and unemployment change by more in areas with stronger local ties. In areas with weaker ties, however, population adjusts by more, presumably through migration.

To specify how local ties affect spatial equilibrium, I present a parametric model of local ties in spatial equilibrium. In the model, local ties are a utility benefit that a worker receives if he lives where he was born. This utility benefit could represent many things, but the most obvious is a return from living close to family and friends. People may find it desirable to live near family and friends because of the benefits of close personal relationships and because economists have found that social networks can be useful for finding a job (Topa, 2011) and smoothing consumption (Kaplan, 2012), among other things.[3]

To illustrate the dynamics surrounding people's local ties, I show how two successive, identical negative shocks can have different impacts on the same area. The first shock leads to a large decline in population and a small decrease in real wages as outsiders leave the area. The second shock, however, leads to a much smaller population exodus and a bigger decrease in real wages, since the area now has a larger share of locals and a lower equilibrium migration elasticity.

In the parametric model, all areas converge to a steady state equilibrium where areas contain residents with the same levels of local ties, but only after several generations. The transitions to the steady state proceeds as people die and as children are born in areas where

[2]This is a common feature in neo-classical models of spatial equilibrium. Howard (2017), however, shows that immigration an stimulate residential construction enough to increase wages, at least in the short run.

[3]There are several relevant literatures in economics. Topa (2011) provides an overview of the literature on local networks and job referrals, which emphasizes the importance of both large social networks and geographic proximity in finding a job. Several papers examine how proximity can help to facilitate intergenerational transfers, either from parents to children, or children to parents. Examples include Konrad and Kunemund (2002), Hank (2007), Rainer and Siedler (2009), Kaplan (2012), Huttunen and Salvanes (2015), and Coate (2017).

population has grown. The dynamic of children being born disproportionately in growing areas means that people's local ties will eventually shift toward places where population has increased. Throughout the transition, however, real wages will be below their steady state values in areas where population and local ties are declining.[4]

I use a sufficient statistics approach to estimate the impact of a place-based policy on aggregate welfare. In a general model of spatial equilibrium, a cash subsidy to an individual area will decrease welfare in proportion to its impact on where people live, which is measured by the migration elasticity. This generalizes several previous discussions of place-based policies, since it provides general conditions that imply that the distortion will be a function of the size of the subsidy times the migration elasticity.[5] In addition to place-based policies, the dynamic applies to other policies that are more generous to particular areas. The most obvious is the bias towards low productivity, high amenity, often declining areas that Albouy (2009) argues is inherent in the US income tax code.

The connection between migration elasticities, decline, and the welfare implications of local policies is particularly relevant for policymakers interested in reducing geographic inequality by implementing place-based policies. A major concern about place-based policies, highlighted by Glaeser and Gottlieb (2008), Kline and Moretti (2014b), and Neumark and Simpson (2015), among others, is that they offer an incentive for people to stay in declining places. This incentive can lead to inefficiency, particularly given findings that growing up in a declining area can lead to worse outcomes in adulthood.[6] If people who live in declining areas are unlikely to move, however, then there is much less concern about the incentive effects of subsidizing declining areas. Investing in disadvantaged communities would be a more viable strategy, if the investments are well designed and well managed.

This paper also furthers the literature on internal migration by integrating microeconomic findings with aggregate spatial equilibrium.[7] Several models, following from Kennan and Walker (2011), have begun to give a much more intricate picture of various factors affecting migration, including the importance of gross flows of migrants. One fact that has emerged from these papers has been that people appear to value living in particular places for seemingly idiosyncratic reasons. For example, Gregory (2013) finds that home owners

[4]Rappaport (2004) finds a similar dynamic, where even very small migration frictions can imply that adjustments take a long time in models of migration. Empirical studies by Bartik (1993) and Beaudry et al. (2014) also find that migration adjustments appear to be similarly slow.

[5]The formula are similar to those in Kline and Moretti (2014b).

[6]Several studies find relationships in terms of a number of different outcomes. Studies of the Moving to Opportunity experiment, for example, find that young children earn more in adulthood if they live in lower poverty areas as children (Chetty et al. (2015)), and that adults are in better health after they move (Ludwig et al. (2013)). There also appears to be very different levels of intergenerational income mobility across areas (Chetty et al. (2014)), and there is some evidence that living in particular areas leads to worse outcomes in adulthood (Chetty and Hendren (2017)).

were willing to pay significant amounts to continue living in New Orleans after Hurricane Katrina. Cadena and Kovak (2016) find that immigrants to the United States are much more likely to migrate for economic reasons than people who were born in the US. This highlights the importance of a relatively small group of migrants, as opposed to a larger group of people living closer to their homes, in establishing spatial equilibrium.

Several papers have also incorporated a more micro founded view of migration into models of spatial equilibrium within countries, though none have focused on local ties. Many of these have focused on modeling advances. For example, Coen-Pirani (2010), Davis et al. (2013), and Monras (2015) model gross flows of migrants, a distinction that is important for thinking about the importance of local ties. Another advance was the estimation of a structural demand system, including endogenous amenities, by Diamond (2016). Like Diamond (2016), I use a mixed logit model with random coefficients, but unlike that paper, I do not attempt to fully estimate it. Instead, I rely on indirect inference to identify important parameters and a sufficient statistics approach.

A complementary literature has focused on the measurement and implications of frictions in the supply of housing in particular areas. It can be difficult to supply new housing because of geography (Saiz (2010)) and zoning (Gyourko et al. (2008a)). Ganong and Shoag (2012) and Hsieh and Moretti (2015), for example, argue that this limits migration into more productive areas. Similarly, the durability of the existing housing stock can make the supply of housing quite inelastic in the short term (Glaeser and Gyourko (2005) and Notowidigdo (2011)).[8]

Local ties influence migration in a way that is distinct from the durability of housing, despite the fact that both can lead to small migration elasticities in places where population is declining.[9] The main piece of evidence for this is that labor demand shocks do not appear to have different effects on the price of housing in areas with different levels of local ties; this suggests that areas with stronger local ties do not necessarily have lower housing supply elasticities. Similarly, when I include several controls for the level of house prices, I still find

[7]Seminal papers in the literature on internal migration and its effects on labor markets include Sjaastad (1962), Blanchard and Katz (1992), Bound and Holzer (2000), Kennan and Walker (2011), and Moretti (2013). Another literature examines the determination of wages and rents in response to differences in productivity and amenities in spatial equilibrium, including Rosen (1979), Roback (1982), Topel (1986), and Albouy (2016). Yet another literature focuses on economic convergence within countries, including Barro et al. (1991), Alesina and Barro (2002), and Ganong and Shoag (2012).

[7]I do not model endogenous amenities.

[8]Yet another related literature examines if homeowners holding negative equity on mortgages, after episodes of falling housing prices, discourage people from moving after negative shocks in particular areas. This literature includes Henley (1998), Ferreira et al. (2010), Modestino and Dennett (2012), and Valletta (2013), among others. Several papers find some effects, but these are generally fairly modest, especially for labor market outcomes (where the actions of renters can undo the actions of home owners).

that areas with higher levels of local ties have lower levels of migration after a shock.

The remainder of the chapter proceeds as follows: Section one provides some background on the data and examines the relationship between population growth and residents' local ties, which motivate the following sections. Section two presents reduced form regressions that measure how areas with different levels of local ties respond to labor demand shocks. Section three presents a model that incorporates local ties into spatial equilibrium and that allows them to endogenously change over time. Section four presents a sufficient statistics analysis that links migration and welfare in a wide class of models of spatial equilibrium.

1.1 How Local Ties Vary Across Areas

This section establishes two stylized facts. The first is that areas have residents with very different levels of local ties. The second is that these local ties are primarily due to different levels of population growth, since areas grow by attracting people with weaker local ties. In later sections, I show how these differences are empirically important, and I introduce a model to trace out their implications over the long term.

To show how these differences can play out in specific cases, Figure 1.1 shows population changes in two fairly typical commuting zones – Minneapolis and Atlanta.[10] In 1970, each had a similar population, but since 1980 Atlanta has grown much faster, and it is now almost 50 percent larger. Initially, the share of people who were born locally (in the same state) was about the same across areas, roughly two thirds. Since 1980, however, people born outside of Georgia have increasingly moved to Atlanta. In 2008, less than half of Atlanta's population was born in Georgia, while Minneapolis still contains mostly people who were born in Minnesota. Presumably, differences in people's origins have had important cultural and economic effects on the two cities.

The rest of this section shows how these patterns apply to places besides just Atlanta and Minneapolis. First, I briefly describe the data that I use in this and in following sections. Then I show how changes in the number of people born in an area are much slower than changes in the number of outsiders. Since growth happens by outsiders moving in, growing places have many more people who were born outside the local area. This is intuitive and unsurprising, except for the magnitude of the effect. There are large differences across

[9]Ramey and Shapiro (2001) also show how business capital can be quite durable, which would lead to similar dynamics on some, but not all, dimensions as a friction due to people's reluctance to move from their homes. For example, it is not clear why the durability of local business capital would lead to people leaving the labor force in larger numbers. Rappaport (2004), however, shows how this friction can lead to very slow population adjustments.

[10]Throughout this paper, I use commuting zones, as defined by Tolbert and Sizer (1996), as my unit of observation. These are described in the following subsection and Appendix A.1.

Figure 1.1: Population Changes in two Commuting Zones

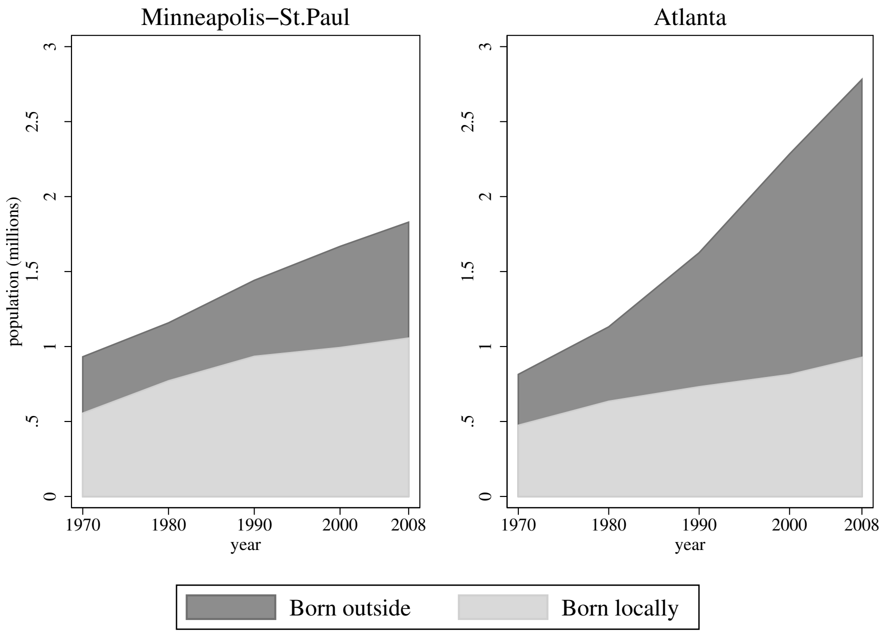

Notes: Data are from the long form decennial census and the ACS 3 year estimates (2006-2008) and are weighted to be nationally representative. Minneapolis-St.Paul and Atlanta are 1990 commuting zones 21501 and 9100. Locals are people who are born in the state they are living in (Minnesota or Georgia), while outsiders are born in other states or countries.

areas in the proportion of people who were born in the same state and in the average time a householder has lived in his house. In later sections, I show how these fairly striking differences translate into differences in how local economies respond to both policies and local shocks.

Data

Data come, predominantly, from the decennial census and ACS (via Ruggles et al. (2010)). In addition to the ACS, I use a measure of the impacts of international trade on local labor markets, coming from Autor et al. (2013) and the NBER vital statistics database. In the remainder of this section, I briefly lay out my justification for using census data, I describe the unit of analysis, and give a brief description of how I processed the data. The data and processing are more fully described in Appendix A.1.

I use data from the US Census because they provide information about migration, labor supply, wages, and housing rents for a large sample in each year. From 1980 through 2000, I use five percent samples of the population, and for the American Community Survey I use 2006-2008 (three year) estimates that include roughly one percent of the population in each year. Large sample sizes are important for measuring outcomes accurately in small areas.

Nonetheless, my preferred specifications are weighted by initial population to place more emphasis on larger commuting zones where things are more precisely measured (as in Bound and Holzer (2000)). The use of the census means that adjustments occur over a minimum of approximately 10 years (8 years for 2000-2008), matching my focus on longer term processes.

The geographic entity that I use as a unit of observation is a 1990 Commuting Zone, as described in Tolbert and Sizer (1996). Using a procedure developed by David Dorn (Autor and Dorn (2013)), I map publicly available geographic identifiers for each year in the IPUMS to the geographic boundaries of Commuting Zones (CZs). CZs are desirable here because they encompass places where people both live and work, according to 1990 commuting data, and because they cover the entire United States. I restrict to CZs in the continental United States for comparability with prior studies and because migration processes are more comparable within the continental US.

I use a person's place of birth as my primary measure of their local ties. In the data, respondents are asked to report their state of birth (or country, if they were born outside the US), which is the measure that I use. The coarseness of this measure, relative to my unit of observation, does not appear to be a big concern. Other studies have found that people value living close to their birth places according to different geographies. For example, Diamond (2016) finds that people have attachments to census divisions as well as states, and Bartik (2009) reports similar results for MSAs using the PSID. No matter the geographic detail, a person's place of birth is still only a proxy for their local ties. Some people quickly moved away from their places of birth, and some did not develop strong connections. Robustness checks, using an alternative measure of local ties, reassuringly give similar results.

I compute statistics at the level of CZs using a sample of adults. My sample includes 22-64 year old adults not living in group quarters (barracks and dorms). In computing wages, I exclude unpaid family workers and workers who did not work for pay in the past year. I report prices in 2007 dollars using the PCE deflator, and I weight all wages using labor supply weights. Appendix A.1 provides more details.

Breaking changes into locals and outsiders

To show the effects of people's local ties, I break changes in population into changes in the population of people who were born nearby and the population of people who were born elsewhere, including people born in other states and countries. The population of outsiders increases by more than the population of people born locally in places that are growing. Very few areas are actually losing people; some areas appear to be unappealing to outsiders, however, which leads their populations to stagnate.

To compare the importance of outsiders moving in against locals staying, or additional

Figure 1.2: Changes in Population Due to Outsiders Moving in

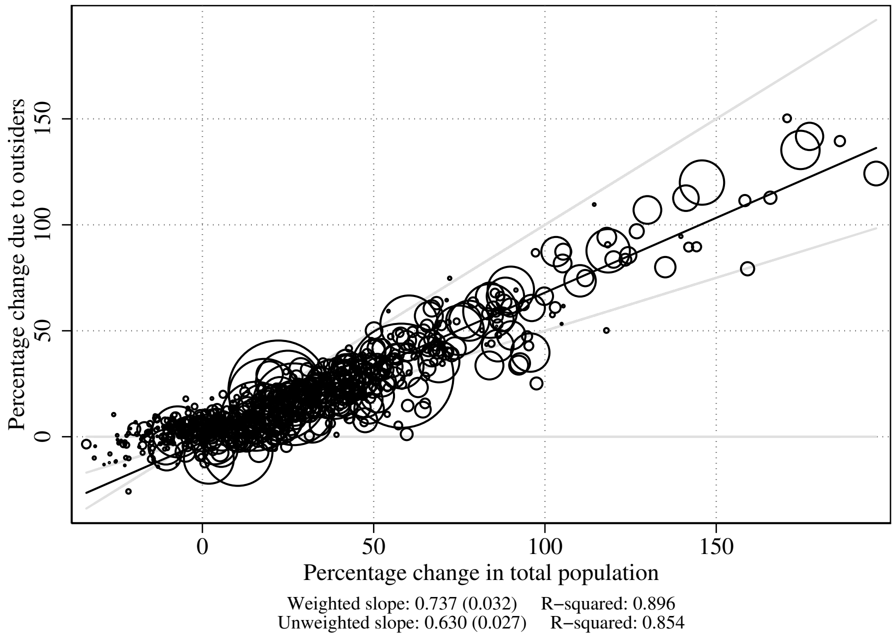

Notes: Plotted are the percentage change in population from 1980 to 2008 (x-axis), and the change in the number of people in the commuting zone who live outside their state of birth as a percentage of the total population (y-axis). Data are from the 1980 decennial census and the ACS 3 year estimates (2006-2008) and are weighted to be nationally representative. The unit of observation is a commuting zone within the continental United States. To make the figure easier to read, it does not include a small number of commuting zones where the total change in population was over 200 percent. The reported coefficients include them, however, with robust standard errors in parenthesis.

children being born, Figure 1.2 plots the changes in total population and changes in the population born somewhere else, each expressed as a percentage of the initial population. Each variable covers the period from 1980 to 2008, and the graph includes commuting zones in the continental US. The graph shows how much of the increase in population (on the x axis) is due to increases in outsiders moving in (on the y axis). Mechanically, if the only reason population changes was because more people were born, then each dot would be on the light grey line on the x axis, and all of the population change would be due to changes in the excluded population group of locals. Conversely, if there was no variation in locals staying, then all changes in population would be due to outsiders, and each dot would be on the light grey 45 degree line. If the two contributed equally, then points would be centered on the middle (22.5 degree) light grey line.

Figure 1.2 shows that outsiders drive population changes in many different areas. Dots on the graph are much closer to the 45 degree line; most are above the middle line. According to a regression with population weights, the slope is 0.74, which implies that any increase in population will be accompanied by an increase in the number of migrants equal to about three quarters of that amount. The unweighted number is lower, but still well above one half. The dots, additionally, are within a relatively narrow cone, suggesting that there is a stable relationship between the changes in the number of outsiders and the number of locals. Outsiders, taken as a whole, appear to be more responsive to changes in local economies.

The importance of outsiders in population changes is important in two ways. First, it suggests that areas have much different levels of outsiders, a fact that I establish in the remainder of this section. Second and more importantly, it suggests that the preferences of outsiders, or people choosing to live in locations that are unfamiliar to them, drive spatial equilibrium. This distinction will be an important element of my modeling strategy, which I lay out in Section 1.3.

Connecting population growth and local ties

There are substantial differences both in the percentage of residents who were born near where they currently reside and the amount of time that people have spent in their houses, which suggests that local ties vary quite substantially across the United States. These differences in residents' experience in an area are the result of the large and persistent differences in population growth rates across the United States that have been documented by Blanchard and Katz (1992), among others. Areas grow by attracting outsiders, so growing areas have many outsiders. Declining areas retain a similar percentage of local children, regardless of local conditions. In some areas, fewer than 20 percent of residents were born in the same state, while more than 80 percent of residents were in others.

Figure 1.3 shows the empirical relationship between of net changes in population from 1980 to 2008 and the amount of experience residents have in the areas where they live, as of 2008. Panel A shows population growth on the horizontal axis and the percentage of residents who were born in the state where they currently live on the vertical axis. There is a robust negative relationship between the two. On average, in a commuting zone whose population increased 100 percent between 1980 to 2008, about 30 percent less of the population will have been born in the same state. The scale of the differences are quite large. Several commuting zones have doubled in size, or more. Fast growing areas have less than a quarter of their populations born in the same state. Most commuting zones, though, have similar populations to 1980 and have more than half of their populations born in the same state. Panel B shows a similar trend for the amount of time people have lived in their houses.[11] Since the majority of moves are local, this statistic shows some supporting evidence that people in growing areas have lived in the same neighborhoods, in addition to the same areas, for less time. In a commuting zone that has grown by 100 percent more, people have lived

Figure 1.3: Population Changes and Measures of Local Ties

Panel A: Percent Born Locally

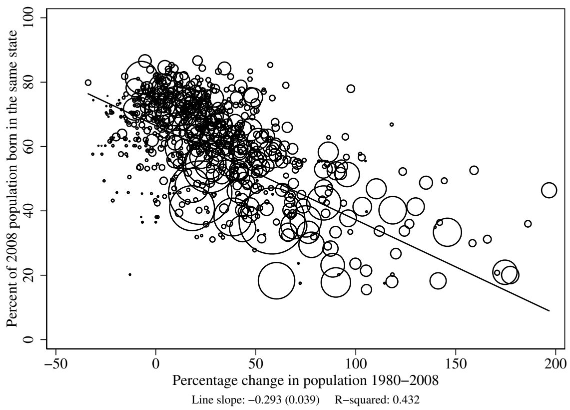

Panel B: Average Time Living in House

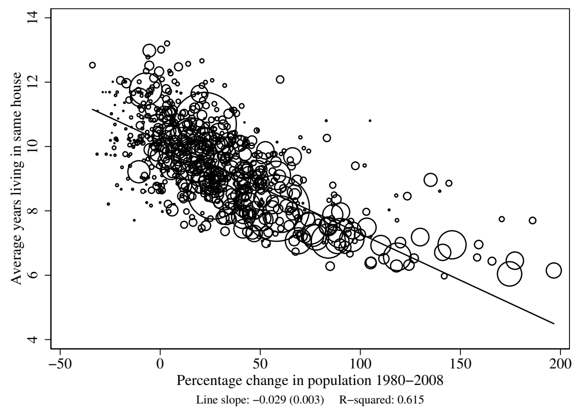

Notes: Data are from the 1980 census and 2006-2008 ACS. Each circle is a commuting zone and its radius is proportional to its population in 1980. The line is a weighted least squares regression, using the population weight. The standard error is clustered by state (a CZ is in a state if the plurality of its population resides there). Share variables are multiplied by 100. The figures omit the few commuting zones that grew more than 200 percent over the period, for visual clarity. These commuting zones are included in the regressions in table 1.1.

in their houses for about 3 fewer years.

Table 1.1 shows that these relationships are robust to omitting weights and including controls. According to each specification, areas that have grown more contain residents with less experience in the area. The magnitude of the main effects are somewhat smaller than in the figure, but they are still quite meaningful. Areas whose population have doubled have 20 percent fewer locals, as a proportion of their population, and have people who have lived in their houses for approximately 2 years less (the last columns are scaled by 100 for readability). Since population growth is quite persistent, it is difficult to disentangle if the effects are due to more or less recent population growth, but separating out growth in different time periods suggests that there is a relationship even with growth over longer time frames. For example, growth from 1980 to 1990 appears to have an effect that is stronger than growth from 2000 to 2008, at least in terms of point estimates. The (adjusted) coefficients are not different in terms of statistical significance, however.

A map of the percentage of residents born in the same state, shown in Figure 1.4, shows interesting geographic patterns. Broadly, the share of people born in the same state is much smaller in the West, particularly the Southwest. This is despite western states having higher populations and larger geographic areas. Areas with the highest percentage of residents

[11] This statistic comes from the census question asking how long the "householder," in whose name the residence is owned/rented, has been living at the residence. The statistic reports this number for all people 16-65, using person weights, so it does not necessarily reflect how long the specific individual has lived at that address.

Table 1.1: Population Changes and Measures of Local Ties

	Percent of CZ born in same state					Average years living in same house				
1980-2008 pct chg in population	-0.21	-0.18	-0.28	-0.23		-2.31	-1.94	-2.71	-2.12	
	(0.03)	(0.03)	(0.03)	(0.03)		(0.19)	(0.21)	(0.25)	(0.22)	
2000-2008 pct chg in population					-0.16					-3.18
					(0.18)					(1.71)
1990-2000 pct chg in population					-0.32					-2.98
					(0.15)					(1.09)
1980-1990 pct chg in population					-0.44					-3.10
					(0.11)					(0.94)
Weighted	No	No	Yes	Yes	Yes	No	No	Yes	Yes	Yes
Controls	No	Yes	No	Yes	Yes	No	Yes	No	Yes	Yes
Observations			722					722		
R^2	0.319	0.517	0.438	0.731	0.746	0.425	0.526	0.609	0.732	0.752

Notes: Data are from the decennial census and ACS and cover the continental United States. Regressions are weighted by initial population and standard errors in parenthesis are clustered by state (a CZ is in a state if the plurality of its population resides there). All share variables are multiplied by 100 to make them into percentage points. Controls, measured in 1980, are share college educated, share employed, share foreign born, share born in Mexico, and log population.

Figure 1.4: Percentage of Residents Born in the Same State

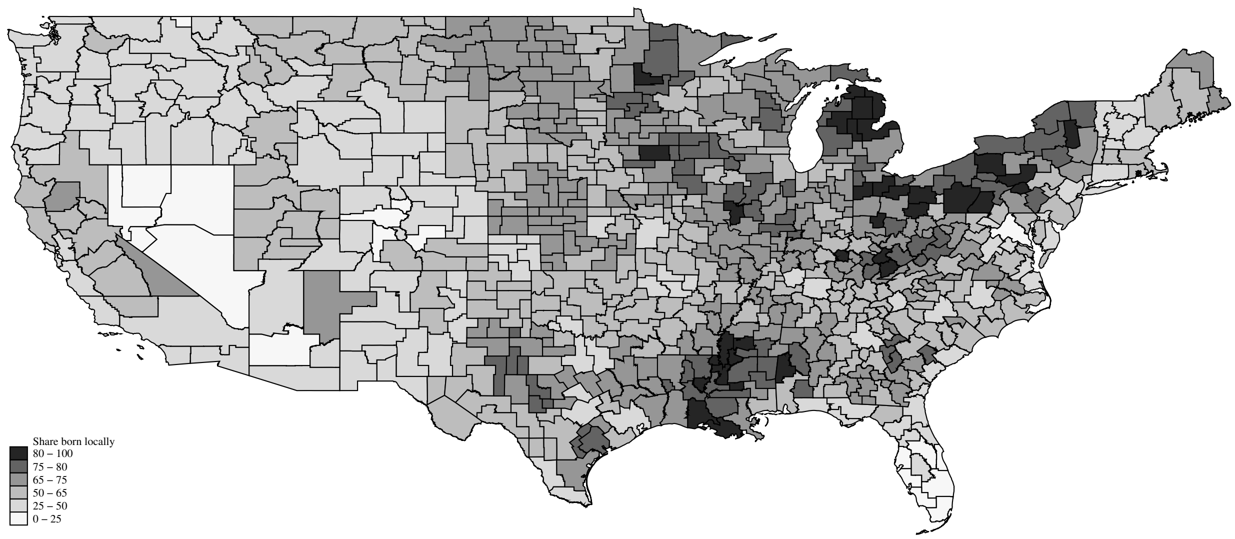

Notes: The 722 commuting zones in the continental US are shaded according to the percentage of residents who were born in their current state. Darker shades mean more were born locally. Data are from responses to the 2006-2008 ACS, via IPUMS. The statistics include people aged 22 to 64 and exclude people living in group quarters.

born in the same state tend to be rural, and they are concentrated in the Deep South, Appalachia, the Upper Midwest, and the Rust Belt. The map's scale shows, once again, that the differences are quite large. For example, parts of Michigan, Louisiana, and other states have more than 80 percent of their populations born in the same state.[12] In other areas, including commuting zones surrounding Denver, Colorado and Phoenix, Arizona, fewer than a quarter of residents were born in the state where they live. Appendix Figure A5 shows a similar pattern for the amount of time people have spent in their residences.

1.2 Reduced Form Results

To test if local ties influence outcomes in spatial equilibrium, this section examines the impact of changes in labor demand in areas with different levels of local ties using a series of reduced form regressions. I quantify impacts by decomposing the demand shift into impacts on population, residents' labor supply, wages, and rents. I use two plausibly exogenous shift-share instruments to isolate impacts on labor demand. The first, developed by Bartik (1991), uses changes in total industry employment at the national level during the 1980s.

[12] This includes the Isle de Jean Charles in Louisiana, which is notable for a $48 million (Jackson (2016), 2016 dollars) grant to resettle approximately 65 residents. A 2012 movie, Beasts of the Southern Wild, was filmed nearby (Arons (2012)) and set in a similar community. It depicts a forced migration of residents from a fictional island, presumably in southern Louisiana. The movie, and presumably the experiences of the residents involved, present an argument about the importance of residents' local ties for their migration decisions.

The second, developed by Autor et al. (2013), uses changes in industry level demand for final goods due to increased trade with China in the 1990s and early 2000s.

The results suggest that areas with different levels of local ties adjusted to labor demand shocks over different margins. Areas with lower levels of local ties adjusted their populations, in keeping with standard models of spatial equilibrium. Areas with higher levels of local ties, however, adjusted the size of their labor force, their wages, and their unemployment rates, as would be expected if people had limited geographic mobility. Rents changed by similar amounts in each area, so the differences are not driven by much larger changes in rents in declining areas.

Outcomes

To understand how areas adjust to labor demand shocks, I decompose the impacts of labor demand shocks between prices and labor supply.[13] This allows me to distinguish between population re-allocation, the standard mechanism in Rosen (1979) and Roback (1982) style models, and other possible adjustments, including people moving out of the labor force and wage changes. To compare the importance of each, I scale changes in different labor supply margins so that each represents percent changes in the number of employed workers. To illustrate, the number of people outside the labor force is about six times as large as the number of people who are unemployed. This means that if one percent of people outside the labor force started to work, and the number of jobs was constant, then the number of unemployed people would have to increase by about six percent.

Figure 1.5 illustrates the effects of a labor demand shock on employment and wages.[14] Initially, the local labor market is at equilibrium at point A, the intersection of the initial labor demand curve, L_{D1}, and the labor supply curve, incorporating all margins of adjustment, L_S. A labor demand shock of size A-B affects the local labor market, however, and shifts labor demand to L_{D2}. This leads to an increase in wages and employment at the new equilibrium, point F at the intersection of L_S and L_{D2}. The size of the wage increase, from A to D, will depend on both the elasticity of labor demand and labor supply, coming from all margins. So, for a constant labor demand elasticity, an area with a more elastic supply

[13]For these regressions, and throughout the paper, I do not distinguish between increases and decreases in any of my outcomes. This is because the reallocations that I study involve large gross flows of people into and out of employment, unemployment, the labor force, and specific local areas. So, even if moving into and out of areas involves separate concerns, and most residents are reluctant to move, there will still be no discontinuity as the net flow of population becomes negative. This is a well known fact about employment dynamics, and Monras (2015) documents this fact for migration, noting that most population changes are driven by the behavior of people moving in, while a roughly constant proportion of people move out. Responses in the model, which are continuous but vary depending on the net migration of people from their homes, illustrate this dynamic.

Figure 1.5: Effects of a Labor Demand Shock Along Multiple Margins

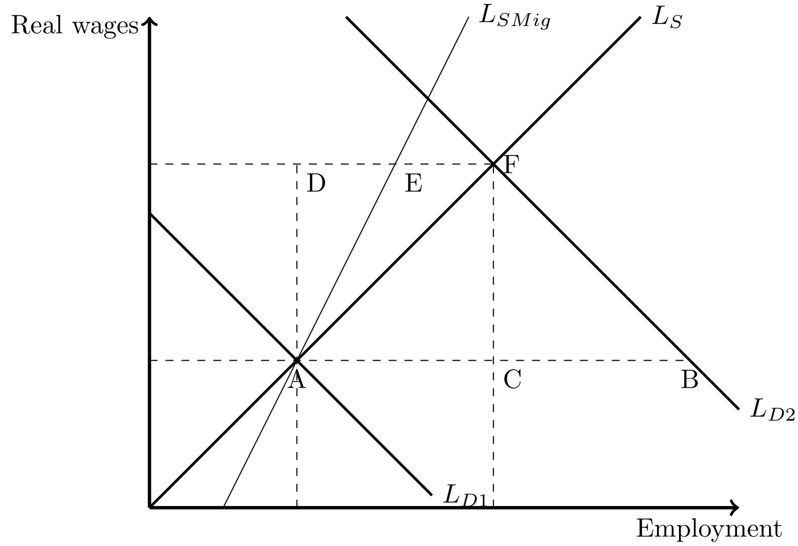

of labor will have a smaller increase in wages. This smaller increase in wages is because employment can change by more in places where labor supply is more elastic.

In addition to effects solely on employment and wages, Figure 1.5 also separates out the equilibrium change due to migration. The curve $L_{S\text{Mig}}$ illustrates changes in employment, around the initial equilibrium at point A, due to people moving in and out, while the horizontal distance from between $L_{S\text{Mig}}$ and L_S shows adjustments in all other margins. The equilibrium size of the migration adjustment, in terms of employment, is the horizontal distance from D to E in the diagram, since this is the change in employment due to the migration response after the equilibrium change in wages. So, by seeing how big the distance from D to E is relative to the total change in employment, D to F, it is possible to decompose the importance of migration, labor force participation, and possibly other distinct processes that will adjust the labor supply.

Formally, if we assume a constant elasticity of labor demand (η_D), labor supply due to migration ($\eta_{S\text{Mig}}$), and labor supply due to other adjustments ($\eta_{S\text{Other}}$), then the size of the equilibrium changes will be simple functions of the three elasticities and the size of the labor demand shock, $B - A$. The change in wages $(F - C)$, will be $\frac{B-A}{\eta_D+\eta_{S\text{Mig}}+\eta_{S\text{Other}}}$, while the change due to migration $(E - F)$ will be $\eta_{S\text{Mig}}\frac{B-A}{\eta_D+\eta_{S\text{Mig}}+\eta_{S\text{Other}}}$ and the total change in employment $(F - D)$ will be $(\eta_{S\text{Mig}} + \eta_{S\text{Other}})\frac{B-A}{\eta_D+\eta_{S\text{Mig}}+\eta_{S\text{Other}}}$. If changes in employment due to population are large relative to total changes in employment, then migration is more elastic in particular areas. This result applies regardless of the size of the labor demand

[14] For conceptual clarity, I am omitting effects on rents. Including rents would complicate the analysis, but should deliver a similar intuition.

shock, or the labor demand elasticity.

To measure the size of changes in employment due to migration, labor force participation and unemployment, I log linearize the accounting identity that the number of employees in a place is equal to the population, minus the number of people not in the labor force and the number of people who are unemployed. This allows me to compare changes in population (mainly due to migration), changes in labor force participation, and changes in unemployment rates in terms of their effect on total employment.

$$\begin{aligned} \mathrm{E} = &\quad \mathrm{P} - \mathrm{NILF} - \mathrm{U} \\ \Delta E = &\quad s_p \Delta p - (s_n \Delta n + s_u \Delta u) \end{aligned}$$

I include estimates of the effect of labor demand shocks on employment due to changes in population (Pop), the people not in the labor force (NILF), unemployment, and three other outcomes. The first is the effect on wages, which decreases with more elastic labor supply and labor demand.[15] The second is the effect on local rents. Rents play an important role in the Rosen (1979) and Roback (1982) analysis of spatial equilibrium and in the model that I present later. One concern is that some areas will be unable to build housing to accommodate additional population, so rents will rise and lower changes in population. I show rents as a rough gauge of how much changes in housing prices might affect the equilibrium. Third, as an additional measure of the size of the labor force participation response, I include the labor force participation ratio, entered as a percentage. Unlike the other measure of labor force participation (NILF, which is the scaled log change in people outside the labor force) the labor force participation rate controls for total population in the denominator, so it will not mechanically decrease if people move to the area.

Changes in Labor Demand

I use two separate shift share approaches to isolate plausibly exogenous labor demand shocks. Each works on the assumption that changes at the national level will affect a CZ proportionate to its pre-existing industrial structure, measured by its employment shares in particular industries at the beginning of the period. The idea is that whatever drives the national changes is presumably not due to supply factors within the CZs that are affected. The first, Bartik (1991), shifter is perhaps the most straightforward in that it simply takes

[15] In my specifications I do not attempt to compute real wages. Instead I present separate results for nominal wages and rents (for housing). These could be combined to compute a proxy; Albouy (2016) suggests that local rents can proxy for 1/2 of local consumption, while national accounts suggest that about 1/3 of consumption is spent on housing and utilities.

changes in industrial employment at the national level (excluding the CZ in my case) and projects them onto CZs. The second, Autor et al. (2013), instrument isolates the effects of Chinese manufacturing competition.

Bartik Shifters

The commonly used Bartik (1991) shifter projects industry level employment changes outside of a CZ onto it using the CZ's share of employment in each industry at the beginning of the period. For area j from period $t-1$ to t it can be written as follows:

$$\hat{\Delta L}_{j,t} = \sum_{i \in \text{ind}} \left(\frac{L_{i,-j,t} - L_{i,-j,t-1}}{L_{i,-j,t-1}} \right) \frac{L_{i,j,t-1}}{L_{j,t-1}}$$

The Bartik shifter is a weighted average of changes in industry level employment outside the CZ (the term in brackets). The weights are the area's share of employment (the second term) in period $t-1$.[16]

The Bartik shifter is a good choice in that it has enough variation to ensure some power, and in that it also can be thought of as plausibly exogenous. To make the case for exogeneity stronger, I only use the instrument in the 1980s because the instrument's logic of projecting industry level trends onto local areas is particularly compelling in the 1980s. Many of the changes in the 1980s are due to national changes that led to a decline in manufacturing employment. The instrument has a good amount of power since manufacturing tends to be spatially concentrated. At the same time, few industries are so concentrated that a single CZ makes up an excessively large proportion of total employment. I rely on the more specific trade shifter that I describe below to provide some evidence from later periods.

Trade instrument

The trade instrument uses a similar shift share strategy, but focuses on a very specific process – increased competition with Chinese manufacturers. Autor et al. (2013) document that imports from China to the United States increased significantly over the 1990s and early 2000s as China entered the World Trade Organization and emphasized an export-led development strategy.

[16] Bartik (1991), Blanchard and Katz (1992), and Bound and Holzer (2000) include changes in employment within the region in question in their calculation of industry wide changes in national employment, which simplifies the construction of the variables. I follow more recent practice, however, and calculate "leave one out" Bartik instruments by excluding each local labor market in question from the nationwide changes used to project employment changes in each industry.

$$\hat{\Delta L}_{j,t} = \sum_{i \in \text{ind}} \frac{-\Delta M_{i,t}}{L_{i,t-1}} \frac{L_{i,j,t-1}}{L_{j,t-1}}$$

Again, the equation computes a weighted average using an area's share of employment in a particular industry $\left(\frac{L_{i,j,t-1}}{L_{j,t-1}}\right)$ as weights. In this case, however, the quantity in the parentheses is different. Instead I measure the size of Chinese import competition in a particular industry, modified by a negative sign to make it have the same sign as above. Specifically, $\Delta M_{i,t}$ measures the dollar value increase in imports coming from China in industry i in thousands of dollars. The results of this instrument, then, can be interpreted as the effect of an increase in imports from China equal to one thousand dollars per worker.[17]

The exclusion restriction for this instrument is more credible than a Bartik specification since it isolates changes due to a single change that was driven by factors outside the United States. Businesses in China increased exports to the US for reasons that are likely to be unrelated to supply shifts in parts of the US. To bolster the case even further, I follow Autor et al. (2013) in instrumenting for Chinese import penetration in the US using Chinese import penetration in other countries.

Specification

I examine differences in responses to labor demand shocks using two different specifications. I estimate each at the CZ level, removing time invariant characteristics of CZs by first differencing all variables. The bins specification separates CZs into two bins, those with low and high levels of local ties, and estimates effects for each bin. This is my preferred specification, since it allows an easy interpretation of the magnitudes involved. To allow for more straightforward hypothesis testing and to show that the effect is not dependent on the cutoff between two bins, however, I also present a triple differences specification. The triple difference specification allows the impact on individual CZs to vary linearly, but continuously, with differences in the CZ's local ties.

Bins

My preferred specification estimates the reduced form effect of labor demand shocks separately for areas with high and low levels of local ties by separating them into two bins. The first bin contains labor markets where less than 60 percent of workers were born in the

[17] Autor and Dorn (2013) present their regressors using different notation and with a different ordering of terms. For this exercise I use the variables from their published dataset, so I am mechanically using the same variation. I differ from their notation and ordering for presentational reasons, to maintain continuity with the Bartik formula.

same state ($\mathbf{1}_L = 1$), and the second contains areas where more than 60 percent of workers were born locally ($\mathbf{1}_H = 1$).[18] Roughly 10 year changes in the outcomes are linear functions of these shifters and an extensive series of controls:

$$\Delta y_{j,t} \quad = \quad \alpha_t + (\beta_L \mathbf{1}_L + \beta_H \mathbf{1}_H)\Delta \hat{L}_{j,t} + \gamma_L \mathbf{1}_L + \gamma_H \mathbf{1}_H + \gamma_X X_{j,t-1} + \epsilon_{j,t} \tag{1.1}$$

Here $\Delta \hat{L}_{j,t}$ is the labor demand shifter[19] and the β coefficients show the effect of these shocks for the specified subset of local labor markets. In addition, α_t is a dummy for the time period where the regressions encompass multiple time periods, and X are the controls.[20] I follow much of the literature by estimating this equation in first differences, which controls for time invariant effects. In the cases with only two periods, this is exactly equivalent to using fixed effects, but in cases with more than two periods, it relies on slightly different assumptions.[21] In this and the triple difference specification, I report standard errors that are clustered by the state the CZ had the plurality of its population within.

Triple difference

To allow for more straightforward hypothesis testing and to show that the results are robust to different cutoffs, I also use a triple difference specification. The triple difference specification implies that that the effect of the labor demand shock varies linearly with the shore born locally:

$$\Delta y_{j,t} \quad = \quad \alpha_t + \beta_{\text{Main}} \Delta \hat{L}_{j,t} + \beta_{\text{Inter}} \Delta \hat{L}_{j,t} \text{ShLocal} + \gamma_{ShLocal} \text{ShLocal} + \gamma_X X_{j,t-1} + \epsilon_{j,t} \tag{1.2}$$

I regress an outcome ($\Delta y_{j,t}$) on a labor demand shifter ($\Delta \hat{L}_{j,t}$) multiplied by the demeaned share of local workers (ShLocal, which measures average levels of local ties), the direct effect

[18]I chose 60 percent because it creates roughly two equal sized groups in most years. Earlier versions used 50 percent and the second specification does not rely on a specific cutoff. The cutoff is mainly designed to produce precise coefficient estimates.

[19]The labor demand shifters are either the Bartik projection or the projected dollar value of imports competition from China, instrumented for by lagged Chinese exports to other countries. I do not use these shifters as instruments for employment. Instead I estimate the reduced form effect of these changes in labor demand.

[18]In addition to dummy variables for each bin of local labor markets, I control for the share of working age adults outside the labor force, unemployed, foreign born, having entered the state in the past five years, and the share of adults who are under 35 and 50 to 64. Generally, specifications are not sensitive to the choice of controls.

[19]Wooldridge (2002) notes that first differencing is preferred when the outcome is a random walk, while fixed effects is preferred when the outcome has serially uncorrelated errors.

of both, and controls for the time period α_t, if there are multiple time periods. The coefficient of interest is β_{Inter}, which represents how the effect varies with changes in residents' average level of local ties. In this framework, tests that effects vary across areas are tests that β_{Inter} is different from zero. Since the share local term is demeaned, the coefficient on the labor demand shifter (β_{Main}) represents the effect for an area with an average share of workers born locally.

Results

Summary statistics

To show some basic characteristics of the sample, Table 1.2 reports summary statistics for the major outcomes, some covariates that I use as controls, and the plausibly exogenous labor demand shocks. Panel A reports statistics about the levels of variables among all 722 continental CZs in 1980, while Panel B shows outcomes and labor demand shocks in the form they enter the regression equations. First I show the scaling factor (if applicable) and the mean and standard deviation among all CZs (unweighted). Next, I show the mean broken out by areas with high and low levels of local ties. Unless otherwise noted, the statistics are computed either in 1980 or from 1980 to 1990 for the scaled log changes.

The first two columns of Panel A show that the average CZ has a modest population, most people 22-65 are employed, and that most people were born in the same place. The average unweighted population of a continental CZ was 162 thousand people in the sample, but the standard deviation is quite large, in accordance with Zipf's law. Most people are employed, but about 30 percent of adults 22-64 were outside of the labor force. The average CZ had about 66 percent of its residents living in the state of their birth as of 1980, but the standard deviation of this number was relatively large in keeping with the previous discussion.

Areas with different levels of local ties also differ in terms of other covariates. Panel A shows that CZs high ties tend to be smaller, in keeping with the relationship between population stagnation and higher levels of local ties. CZs with high ties tend to have slightly older, less educated populations who earn lower wages and pay less money in rent. Somewhat surprisingly, they also have higher labor force participation. Differences in these covariates suggest that it is important to control for level differences across areas, which the first differences specification does.

Panel B shows that population changes have large impacts on employment, much more so than other categories. The standard deviation of scaled population changes are roughly three times larger than changes in people in the labor force and nine times larger than

Table 1.2: Summary Statistics

Panel A: Levels of Covariates

	Mean	StD	Low Ties	High Ties
Population (thous)	162.2	453.9	282.6	117.4
NILF (thous)	43.1	116.3	73.7	31.7
Unemployed (thous)	6.5	18.9	10.6	5.0
Real wages (hourly)	15.1	1.8	16.1	14.8
Real rents (monthly)	475.8	77.5	539.5	452.1
Percent in labor force	71.5	4.4	71.2	71.6
Percent unemployed	3.9	1.6	4.1	3.9
Percent locals	66.4	16.2	44.0	74.8
Average time in house	8.5	1.4	7.0	9.1
Percent college edu	33.5	8.3	40.5	30.9
Percent foreign born	2.7	3.4	5.1	1.9
Percent under 35	42.7	3.5	44.8	41.9
Percent over 50	27.1	3.0	25.4	27.7

Panel B: Outcomes and Regressors

	Scaling	Mean	StD
Population	145	8.11	18.38
NILF	39	-5.86	5.92
Unemployed	6	0.97	2.04
Real wages	100	-2.84	6.19
Real rents	100	7.36	11.21
Pct in labor force		5.13	1.91
Bartik shock (80-90)		12.94	4.24
Trade shock (90-00)		-1.18	1.78
Trade shock (00-08)		-2.64	3.02

Notes: The tables show unweighted summary statistics for the sample of 722 continental CZs. Unless otherwise specified, the statistics refer to values in 1980. The first columns show the mean and standard deviation among all CZs, the next two show means for areas with low and high ties (above or below 60 percent locals), and the final three show the scaling parameter and the scaled log change in the variable, except for the percent in the labor force where the value is simply the percent change. Note that the shock variables are themselves scaled log changes, but these statistics appear in the first columns instead. The variables are grouped into outcomes, controls, and regressors that are used in the reduced form regressions.

changes in unemployment. The table also shows the impact of women entering the labor force in greater numbers from 1980 to 1990, since the labor force participation rate grew by five percent on average.

Bartik shifter

Table 1.3: Bartik Shocks by Share Born Locally

Panel A: Bins Specification

	Pop	NILF	Unemp	Emp	Wages	Rents	LFP
Bartik: Low ties	2.11	0.48	-0.05	1.62	0.26	0.25	0.05
	(0.56)	(0.17)	(0.05)	(0.39)	(0.24)	(0.33)	(0.03)
Bartik: High ties	0.53	0.00	0.05	0.46	0.29	0.29	0.08
	(0.32)	(0.08)	(0.03)	(0.25)	(0.21)	(0.25)	(0.03)
P-val: No diff	0.01	0.02	0.08	0.01	0.92	0.93	0.28
R^2	0.58	0.52	0.67	0.52	0.35	0.54	0.36
Observations				722			

Panel B: Triple Difference Specification

	Pop	NILF	Unemp	Emp	Wages	Rents	LFP
Interaction	-4.24	-1.10	0.04	-3.08	0.95	0.48	0.02
	(1.34)	(0.44)	(0.12)	(0.91)	(0.67)	(0.98)	(0.07)
Main effect	1.24	0.20	0.02	0.99	0.26	0.20	0.07
	(0.28)	(0.08)	(0.03)	(0.21)	(0.16)	(0.22)	(0.02)
Percent locals	0.32	0.05	0.00	0.25	-0.02	-0.23	0.01
	(0.25)	(0.08)	(0.02)	(0.18)	(0.15)	(0.17)	(0.02)
R^2	0.60	0.55	0.66	0.54	0.37	0.55	0.36
Observations				722			

Notes: OLS regression coefficients, weighted by initial population, are plotted for either the main effect plus a linear interaction term with the demeaned share locally born, or the coefficient separately estimated for less than or greater than 50 percent locally born CZs. Controls, measured in 1980, are: the birth share variable used in the interaction term, the share of working age adults outside the labor force, unemployed, with a college education, foreign born, having entered the state in the past five years, and the share of adults who are under 35 and 50 to 64. Wald tests are presented for the hypothesis that the effect is constant across states with high and low in state birth shares. Data are from the decennial census from 1980 to 1990 including all CZs in the continental US. Variables are in percentage changes, except for the linear interaction terms, which are proportions (divided by 100). Log numbers of people (unemployment, labor force exits, and log population) are scaled by their ratio to the number of employed workers and wages are residualized according to the text. Standard errors in parentheses are clustered by state (determined by the state where the plurality of residents live).

Specifications using Bartik shifters from 1980 to 1990, shown in Table 1.3, show strong migration responses in areas with low levels of local ties and smaller responses in areas with higher levels of local ties. Differences in the coefficients on population are statistically and economically significant in both specifications, as are differences in the number of people outside the labor force. There is little evidence of differences in other coefficients. Put together, the changes suggest that migration is less responsive to labor demand shocks in places with high local ties, in keeping with the intuition that local ties are a barrier to migration.

The most striking difference in Table 1.3 is the response of total population in each area. The bins specification for areas with low levels of local ties, shown in Panel A, show that population changes add two percentage points to the stock of potential workers after a one percent increase in predicted local labor demand. Alongside this, however, the number of people outside of the labor force increases by about 0.5 percent of the initial workforce, leaving a 1.5 percentage point increase in employment. For areas with higher levels of local ties, the population response is muted, equal to only about 0.5 percent of the workforce, and there is no discernible change in the number of people outside the labor force. The triple difference specification confirms these results. The interaction term is scaled by 100 for readability, so the roughly 30 percent difference between the average high and low ties area implies a 1.3 percentage point smaller change due to population changes, and a 0.3 percentage point smaller change due to people entering the labor force, in an area with higher ties. The two specifications appear to match quite well, as expected.

Differences between high and low ties areas in Table 1.3 suggest that adjustments are different in areas where people have strong local ties. The population changes are one quarter as large in areas with stronger ties, and the difference is statistically significant at the one percent level. This suggests that areas with lower ties can adjust after changes in labor demand by absorbing additional population, as in Blanchard and Katz (1992). Areas with higher ties, however, adjust along other margins. In the bins specification, areas with high levels of local ties have a statistically significant increase in labor force participation rates after a positive shock, mirroring Bartik (1993), who finds that residents benefit from local labor demand shocks.[20]

Differences in age structure, differences in educational attainment, or developments in the housing market do not appear to drive these differences between areas with high and low levels of local ties. Appendix Table A9 includes separate interactions with the local age structure, local educational attainment, the initial percentage of residents employed, and

[20] One somewhat puzzling result is that an increase in labor demand increases the number of people outside the labor force in areas with low ties. Appendix Tables A6 and A7 separate the effect out for men and women to show that the effect is driven by women. According to Table A7, areas with few locally born workers experience increases in the number of adult women outside of the labor force that are roughly 1/3 the size of the increase in population. This relative size suggests that women migrated in and remained outside of the labor force, likely because of their partner. This supports a literature on "tied migration," including Sandell (1977) and McKinnish (2008), that finds that women often drop out of the labor force after moves. The 1/3 figure is also consistent with average labor force participation rates among women in the 1980s, which are between 50 and 60 percent. Impacts on the labor force participation rates, which control for changes in population due to migration, are never significantly negative in Tables 1.3, A6, and A7.

several measures of initial rents. It shows that the main results I outlined above are robust to including these, and appear to actually grow if other interactions are included. Another piece of evidence that suggests that these effects are driven by local ties themselves is the near equal sized impacts on wages and rents. If, for example, differences between growing and declining areas were due to housing being inelastically supplied in declining areas because housing is durable, then rents should increase by much more in areas with higher levels of ties.[21] While the estimates for wages and rents in Table 1.3 are imprecise, there is little evidence that these differences are very large in this context, at least.

Local ties can have meaningfully different implications than durable housing stocks. Even if housing is durable, it still depreciates at least 3 percent per year, and it is built in very specific lots. For example, it seems apparent that local ties should be more important than durable housing in the city of Detroit, whose population has declined by more than half since 1950. In the process of that decline, much of the housing stock has been destroyed either by nature, by vandalism, or through demolitions by the city.[22] The city's decline in population has not been accompanied by a decline in the population of the broader metropolitan statistical area, which has actually gained population since 1950. One explanation for the relative stability of the population of the MSA, and of many other MSAs, is that people's local ties kept them close to family and friends in the area, but not necessarily the city itself.

Trade shifter

Regressions using the trade instruments, shown in Table 1.4, also show that areas with low ties adjust in terms of population, while places with higher ties adjust along other margins. The bins specification shows that a $1,000 per worker decrease in import competition from Chinese firms leads to an increase in population equal to 1 percent of the initial workforce.[23] The number of workers outside the labor force increases by about one quarter as much, but this difference is statistically indistinguishable from zero. The effect on labor force

[21] Glaeser and Gyourko (2005) and Notowidigdo (2011) emphasize the importance of durable housing in declining areas. Durable housing can lead to an inelastic housing supply in areas where population is declining. Since houses depreciate slowly, decreases in labor demand (and amenities) can lead to very low rents, which keep people in the area. This is particularly true for poorer residents, since poor people spend a higher share of their incomes on housing. This mechanism is most important over short time frames and fine geographies, however. Over longer time frames, like the time frames associated with changes in residents' local ties, even small depreciation rates can lead to substantial decreases in the housing stock. Another factor to consider is that the size of households has been declining (Albouy and Zabek (2016)) so housing stocks would have to expand to house a population that remained constant (few areas have consistently declined in population). Previous examinations by Rappaport (2004) and Davis et al. (2013) have also found that durable housing plays a modest role.

[22] The city has actually paid to destroy vacant housing, since many argue that it increases crime and poses a danger to public safety (e.g. Kurth and MacDonald (2015)). This destruction has left much of the city composed of green spaces, with occasional single family homes placed in oddly compact lots.

Table 1.4: Import Shocks by Share Born Locally

Panel A: Bins Specification

	Pop	NILF	Unemp	Emp	Wages	Rents	LFP
Trade: Low ties	1.06	0.30	0.03	0.74	0.09	1.37	-0.01
	(0.52)	(0.21)	(0.06)	(0.37)	(0.25)	(0.28)	(0.10)
Trade: High ties	-0.13	-1.10	-0.20	1.20	0.64	1.19	0.78
	(0.41)	(0.22)	(0.06)	(0.45)	(0.18)	(0.57)	(0.17)
P-val: No diff	0.02	0.00	0.00	0.28	0.08	0.78	0.00
R^2	0.47	0.62	0.63	0.34	0.12	0.18	0.54
Observations				1444			

Panel B: Triple Difference Specification

	Pop	NILF	Unemp	Emp	Wages	Rents	LFP
Interaction	-2.01	-3.84	-0.70	2.64	2.11	-0.79	2.45
	(2.50)	(0.97)	(0.25)	(1.87)	(1.07)	(2.20)	(0.54)
Main effect	0.43	-0.34	-0.07	0.86	0.30	1.29	0.33
	(0.43)	(0.16)	(0.04)	(0.40)	(0.15)	(0.28)	(0.11)
Percent locals	-0.39	-0.21	-0.03	-0.16	-0.04	-0.10	0.08
	(0.10)	(0.04)	(0.01)	(0.07)	(0.04)	(0.09)	(0.02)
R^2	0.49	0.63	0.63	0.35	0.14	0.18	0.55
Observations				1444			

Notes: Two stage least squares estimates using Chinese trade with other countries in each industry to instrument for trade with the US only. Coefficients are plotted for either the main effect plus a linear interaction term with the demeaned share locally born, or the coefficient separately estimated for less than or greater than 50 percent locally born CZs. Controls, measured in the beginning of each period, are: the birth share variable used in the interaction term, the share of working age adults outside the labor force, unemployed, foreign born, having entered the state in the past five years, the share of adults who are under 35 and 50 to 64, and a year fixed effect. Wald tests are presented for the hypothesis that the effect is constant across states with high and low in state birth shares. Data are from the decennial census (1990 and 2000) and ACS (2008) including all CZs in the continental US. Variables are in percentage changes, except for the linear interaction terms, which are proportions (divided by 100). Log numbers of people (unemployment, labor force exits, and log population) are scaled by their ratio to the number of employed workers, and wages are residualized according to the text. Standard errors are in parentheses and clustered by state (determined by the plurality of the population is in that state) and all results are weighted by population.

participation, which controls for population, is a fairly precise zero. Interestingly, wages appear to be barely affected, but rents jump substantially. In places with high levels of local ties, however, people enter the labor force. Population changes are negative, though small and statistically insignificant, and the stock of workers outside the labor force decreases by about 1 percent of the initial workforce. Changes in the number of unemployed workers are also meaningful, at 0.2 percent of the workforce, particularly if one considers that the value is scaled. Putting these together, the effect on the labor force participation ratio are substantial. The $1,000 decrease in competition leads to a roughly 0.75 increase in the percent of workers in the labor force. In addition, wages increase substantially in places with high local ties – by 0.8 percent in response to the $1,000 per worker decrease in competition.

The results in Table 1.4 are robust to using the triple difference specification and the magnitudes of each are also in line with the bins specification. The estimated interaction terms are negative for population, the number of people outside the labor force, and unemployment, though the interaction with population is imprecisely estimated and statistically indistinguishable from zero. Effects on wages and labor force participation increase with higher ties, in keeping with the limited population response. Each suggests substantial losses for the local population as Chinese firms entered the market in the 1990s and early 2000s, keeping with Autor et al. (2013) and Feler and Senses (2015). Interestingly, these losses appear to be highly concentrated in areas where workers have higher levels of local ties and were either unwilling or unable to migrate.

The findings are also robust to including other possible differences as interactions; if anything, they appear to be strengthened by them. Appendix Table A10 shows that including the same interactions with age structure, educational attainment, the percent employed, and several measures of rents does not change the main findings. The results are often stronger with additional interactions. The estimated impacts on rents are also quite similar, which suggests that the housing markets in each area respond in similar ways.

Summary

Areas where people have higher levels of ties respond to labor demand shocks in different ways. In areas where people have higher levels of ties, population changes are smaller, people move into or out of the labor force, and wages change by more. All of these suggest that there are forces that make migration slower in places where people have higher ties, and

[23] Since the median exposure to Chinese import competition is roughly $1,000 per worker, the coefficients show the effect in terms of the exposure of a fairly typical area. The mean, shown in Table 1.2, is higher because the measure is particularly large in some rural areas. It is more than $15,000 per worker in some areas and maxes out at $43,000 per worker in Murray, Kentucky. The results are weighted by population and are robust to dropping several of these areas.

that these have implications for welfare.[24] To see how these differences can emerge in spatial equilibrium, the next section presents a model that incorporates a flexible specification of people's ties to their homes. It shows how declining areas are unappealing to most people born outside, but still are appealing to a considerable fraction of people who have strong local ties. Thus, high average levels of local ties emerge endogenously. I use the model, and generalize it using a sufficient statistics approach, to examine the welfare implications of different migration elasticities. I also use the model's structure, combined with these estimates, to calculate migration elasticities after Bartik shocks in the 1980s and increases in Chinese trade in the 1990s and early 2000s.

1.3 Model

Here I present a parametric model to formalize and extend the intuitions that I laid out above. The model specifies how areas develop different levels of local ties, how different levels of local ties lead to different migration elasticities, and how local ties may be reallocated over time. The model also allows me to distinguish between equilibrium in the steady state and in the short and medium run. Previous shocks have no relevance for the model's steady state, since ties are able to be completely reallocated, as with models that do not include local ties. It takes a long time to reach the steady state, however, and ties are extremely relevant in the interim.

In the model, different levels of local ties are a consequence of population growth, and local ties lead to different migration elasticities across areas. In the short run, negatively shocked areas can become unattractive to outsiders, while still retaining workers with local ties. This makes their migration elasticities low; outsiders are reluctant to move in and locals

[24] This appears to be true in terms of population increases and decreases, since the specifications do not distinguish between these two margins. It would be difficult to credibly identify asymmetric effects in this context, since my instruments are not the only factors leading to population growth and decline. For example, a positive shock to labor demand might be in a place that is shrinking. Also, in the data, gross flows of people into and out of CZs appear to be much larger than net flows, and people are still moving to places that are declining. This implies that there is not a discontinuous change in the profile of the marginal migrant as population changes turn negative. In my model, presented in Section 1.3, effects of labor demand shocks are actually symmetric locally. This is because people born in an area might reside elsewhere, but wish to return. In declining areas these "exiles" are a larger proportion of people who might consider moving in, and the preferences of this group of marginal people changes continuously. Kennan and Walker (2011) also emphasize that return migration makes up a large percentage of gross migration flows, so this mechanism is supported by the data.

are reluctant to move out.[25]

Local ties can be reallocated, but the reallocation takes several generations.[26] People move to areas after favorable shocks, and they have children who become tied to these areas. In this way, children's local ties reflect these changes in local productivity and amenities. Permanent subsidies are undesirable, since they keep population from being fully reallocated. Subsidies that last for a generation, however, can subsidize areas without affecting migration by very much. Similarly, if changes in amenities or productivity are only temporary, then it may be socially desirable to retain the previous distribution of local ties.

The remainder of this section presents the model, its calibration, and its implications in the short and long run. First, I present the model's setting, the various agents who interact in the model, the problems agents solve, and the resulting equations that describe the model's equilibrium. Next, I show how migration elasticities incorporate both the traditional mechanism of rents rising to discourage additional migration, and also local ties affecting migration, through a term connected to how selected the average resident is. This illustrates how local ties form and how they impact migration elasticities. The following section presents the model's calibration, using indirect inference and off the shelf parameters. The final two subsections present the short and long run implications of including local ties in the model. The short and long run implications link the model to my reduced form results and provide longer term implications of including local ties in a model of spatial equilibrium.

Setting, agents, and equilibrium

The model is in spatial equilibrium, with a large number of areas, indexed by j, that individual workers, indexed by i, are free to move between in each period. Workers have ties to their home areas, indexed by k. I model these ties using a parametric distribution of amenity values that workers enjoy if they live in their homes (where $j = k$). Each area also has an amenity value, A_j, that applies to both locals and outsiders. Local firms offer wages, w_j, based on levels of local productivity, and the prices a national firm pays to combine local goods into a tradeable consumption good. Landlords charge rents, r_j, based on the scarcity of land in an area. The interesting counterfactual relates to the impact of government subsidies, g_j.

[25]Since there are always gross flows of locals moving out and outsiders moving in, migration elasticities are symmetric across small positive and negative shocks. So, a small increase and a small decrease in wages and/or amenities will have effects that are approximately equal in magnitude. Population changes have a non-linear relationship with the size of the underlying changes, however. In general, the first half of a decrease in productivity will change population by more than the second half, since the migration elasticity will be smaller for the second half.

[26]Rappaport (2004) finds that reallocation can occur very slowly in models with relatively small frictions to capital, or labor, reallocation in the form of convex adjustment costs.

I distinguish between equilibrium in the short and long run. In the short run, people's local ties are fixed, so population adjustments are limited by people's local ties. In keeping with my empirical approach, I assume that these are proportional to where people were born, and that the initial distribution of population determines these birth places. In the long run, however, I allow local ties to adjust with population changes, and I show that this leads to a steady state.

Worker's problem

For a worker of type i, living in area j, and with home area k, utility is Cobb-Douglas in a final consumption good c_j (with a price normalized to one) and housing h_j.[27]

$$u_{ijk} = (1 - \alpha^H) \ln(c_j) + \alpha^H \ln(h_j) + A_j + \mathbb{1}(k = j)\mu_i + \xi_{ij}$$

With a corresponding budget constraint (taking into account government subsidies, (g_j):

$$c_j + r_j h_j = w_j + g_j$$

The μ_i term is a random coefficient describing a worker's preference for their home. μ_i depends on the worker's unobserved type i; it allows me to specify the distribution of preferences for residing in one's home, or the distribution of people's local ties. I assume that all areas have the same initial distribution of local ties. A sorting process, similar to the one I document in Section 1.1, will make it so that some areas will have residents who have stronger local ties, however.[28]

The government provides a net subsidy, g_j, to workers. People also gain utility from general local amenities A_j and an area specific error term ξ_{ij}. Following much of the literature estimating Rosen (1979) and Roback (1982) style models, I assume workers inelastically supply one unit of labor once they choose their location.

The above implies a log linear indirect utility function.

$$u_{ijk} = \ln(w_j + g_j) - \alpha^H \ln(r_j) + A_j + \mu_i \mathbb{1}(k = j) + \xi_{ij}$$

If ξ_{ij} is distributed type one extreme value, then the likelihood that a person of type i, with home k, locates in area j, which I denote with ψ_{ijk}, takes on a very convenient form.

[27] I have omitted time subscripts for parsimony, because the problem is static once one considers workers' local ties.

[28] Details of my empirical implementation of the μ_i term are provided in the calibration section. I use a normal distribution, with a mean and variance that I calibrate using indirect inference. Train (2009) provides an introduction to logit and mixed logit (random coefficient) models, including details of their development in describing substitution patterns in consumer demand for products.

$$\psi_{ijk} = \frac{\exp(\omega_{ijk}/\sigma_\xi)}{\sum_{j' \in J} \exp(\omega_{ij'k}/\sigma_\xi)} \tag{1.3}$$

where ω_{ijk} is the worker's utility in area j, excluding ξ_{ij}, and the σ_ξ term is a measure of the variance of ξ_{ij}.

Local goods firms

Local good varieties for each area are produced by (a representative) perfectly competitive firm in each area, called the local goods producer. The local goods producer combines capital, K_j, which is supplied in a national financial market at interest rate ρ, with local labor, N_j, to produce Y_j of the local good. The production function is parameterized as:

$$\begin{aligned} Y_j &= f(\theta_j, K_j, N_j) \\ &= \theta_j K_j^{\alpha^Y} N_j^{1-\alpha^Y} \end{aligned}$$

θ_j is a area specific productivity term.

National firm

A perfectly competitive national firm produces the tradeable consumption good out of each local good. It buys each local good at a price of p_j, and sells the tradeabale good at a price of one.

$$Y = \left(\sum_{j' \in J} \phi_{j'}^{\frac{1}{\eta^Y}} (Y_{j'})^{\frac{\eta^Y - 1}{\eta^Y}} \right)^{\frac{\eta^Y}{\eta^Y - 1}}$$

Where j' indexes the goods produced in each area, η^Y is the Armington elasticity (of substitution) between the local goods, and ϕ_j is a demand shifter for each local good.

Government

The government can provide subsidies (net of taxes) to workers living in each area, g_j. The subsidy represents programs that are either directly or indirectly targeted to certain areas, as well as the taxes that pay for them. The largest subsidies are likely to be due to a lack of adjustments for local prices in both the tax code (Albouy (2016)) and in means based programs (Notowidigdo (2011)).[29] Spending on transportation infrastructure, grants to provide utilities, state and federal grants for education, and explicitly place-based policies, like federal neighborhood improvement grants, could also be included in g_j. For simplicity, I assume that these programs are valued by residents at the cost it takes to provide them.

Since the government has to balance its budget, the following equation must hold.

$$\sum_j g_j N_j = 0$$

Housing market

The housing good h_j in a local area represents both housing and non-tradeable local goods and services. The price of the housing good, r_j is determined by supply and demand in the local area.

Housing is supplied by landlords, assumed to be absentee, who develop plots of land that can be turned into housing at monotonically increasing marginal costs.[30] I parameterize these using the following supply function, which gives the cumulative sum of all housing units that would be rented out for rent r_j:

$$H_j^S = F(r_j) = r_j^{\eta^H}$$

Where η^H is the local housing supply elasticity. The consumer problem implies (suppressing tax terms) that workers will demand housing as a fixed proportion of their income:

$$H_j^D = \frac{\alpha^H w_j N_j}{r_j}$$

Law of motion for local ties

The law of motion for the number of people born in area k, specified below, implies that ties are formed proportionate to the population of each area. The first term specifies that a fixed proportion, s_D, of the population will die and have their ties reallocated in each period. The additional s_D ties are allocated according to the current population. For area k, current

[29]The subsidies are likely the most policy relevant in declining areas. Albouy (2016) notes that progressive tax rates subsidize people living in less productive places, since they do not adjust for local prices. A similar argument applies to other means tested programs. This bias is often by design. For example, the governing documents of the European Union explicitly allow national governments to pursue place-based policies to "promote the economic development of areas where the standard of living is abnormally low" (Article 107 of EU (2012)).

[30]Another interpretation of the absentee landlord assumption is that households own their houses but effectively rent their houses to themselves each period. Most notably, this abstracts from the investment motive for owning a home. Zhang (2016) gives some analysis of the investment role of housing, separating it out from consumption choices, to some extent. There is a more applied literature on a specific implication of the financial role of housing – housing lock in. In the appendix, I give a brief overview of the literature on housing lock in.

population is $\sum_{k'} N_{kk'}$. The $\frac{s_D}{1+s_F}$ term specifies the s_D ties are reallocated so that there is a constant population where a constant proportion is foreign born.

$$N'_k = N_k(1 - s_D) + \sum_{k'} N_{kk'} \frac{s_D}{1 + s_F}$$

To keep the model simple, I assume that the destruction of local ties is random, and that workers do not have preferences about the distribution of other people's local ties. $s_D N$ workers die in each period, to be replaced with children who were born according to the distribution of deaths, and of total population. This is very helpful in terms of making the model tractable. Since the deaths are random, there is no impact on the shape of the distribution of μ_i. Since workers do not attempt to influence the distribution of local ties, I can continue to solve the model according to the earlier equations, which are specified conditional on the current distribution of local ties.

While it is a strong assumption, there are reasons to suspect that workers lack strong preferences about the distribution of local ties after s_D has arrived. One way of rationalizing this is to assert that parents act as if their children will have the same preferences that they do, or that their choices about where to live have a very small impact on their children's local ties. Another is to assert that parents think about forming local ties in a way that is myopic, behaving as if either they cannot influence their family's local ties, or they cannot forecast future conditions well enough to make such a long term decision.

The model will be in a steady state if the following equation is satisfied. Essentially, the number of people who were born in the area is equal to the population over one plus the share of foreign born people in the country.[31]

$$N_k = \frac{\sum_{k'} N_{kk'}}{1 + s_F}$$

Equilibrium

Equilibrium is a set of prices and quantities (p_j, w_j, r_j, N_j) where markets clear, and agents have solved their individual problems. It obeys the following equations:

Labor market supply and demand locally:

$$N_j = \sum_{j' \in J} \sum_{k' \in K} \psi_{i'jk'} N_{i'k'} \tag{1.4}$$

$$w_j = (1 - \alpha^Y)(p_j \theta_j)^{1/(1-\alpha^Y)} \left(\frac{\alpha^Y}{\rho} \right)^{\alpha^Y/(1-\alpha^Y)} \tag{1.5}$$

[31] It is straightforward to show that this steady state is stable and unique.

Equilibrium in the housing market and equilibrium in the market for the local good:

$$r_j = \left[\alpha^H w_j N_j\right]^{\frac{1}{1+\eta^H}} \tag{1.6}$$

$$\theta_j N_j \left(\frac{p_j \theta_j \alpha^Y}{\rho}\right)^{1/(1-\alpha^Y)} = Y \frac{\phi_j}{p_j^{\eta^Y}} \tag{1.7}$$

Migration responses

The migration elasticity, analytically derived from the model and shown below, reveals the influence of local ties as well as the more typical influences of housing supply in a Rosen-Roback framework.

$$\frac{d\ln(N_j)}{d\ln(w_j)} = \frac{[1+\eta^H - \alpha^H]}{1+\eta^H + \alpha^H(1-\bar{\psi}_j)/\sigma_\xi} \frac{(1-\bar{\psi}_j)}{\sigma_\xi} \tag{1.8}$$

There are two distinct factors that determine the elasticity: The first term is the traditional Rosen-Roback mechanism of increases in wages causing increases in rents. Rents increase because higher wages lead to higher demand for housing, both from new residents and from current residents who spend some of their increased wages on housing. The second term measures the influence of worker preferences, and it is inversely proportional to residents' local ties. To see this, strip away the effects of prices by considering the impact of an increase in log wages, with rents unchanged.

$$\frac{\partial\ln(N_j)}{\partial\ln(w_j)} = \frac{(1-\bar{\psi}_j)}{\sigma_\xi} \tag{1.9}$$

The denominator shows a dynamic that is typical in models of spatial equilibrium; migration elasticities decrease when people have stronger idiosyncratic preferences. The spread of the person by area is given by σ_ξ, so as the spread widens, the migration elasticity decreases.

The numerator shows how local ties influence migration elasticities, since it is a decreasing function of $\bar{\psi}_j$. Equation 1.3 defines $\bar{\psi}_j$: it is the ratio of ω_{ijk} terms in area j over the sum of $\omega_{ij'k}$ terms in every area, averaged over all residents of area j. Intuitively $\bar{\psi}_j$ will be higher if more residents are inframarginal, or if they greatly prefer area j over all other areas. This implies that areas with many locally born residents will have lower migration elasticities, since these residents tend to have higher values of μ_i, so they have idiosyncratically high preferences to reside in area j. I show how this is the case both for the calibrated model and in empirical regressions later in the paper.[32]

Calibration

To calibrate the model, I use a mixture of off the shelf parameters as well as an indirect inference procedure. The most important parameters are the two variance terms for the idiosyncratic error distributions (μ_i and ξ). These determine how important idiosyncratic factors are in workers' preferences for their birth areas and for other areas. I set the variance of non-local term, ξ, based on estimates in the literature provided by Suárez Serrato and Zidar (2014). For μ, I use an indirect inference procedure to match a regression coefficient relating workers' decisions to stay in the areas of their birth with the change in log population in that area. This regression is shown in Appendix Table A4. It is very much related to the estimates in Figure 1.2, in that it is driven by the relationship between people staying and outsiders moving in. I set the mean of μ so that roughly 60 percent of workers stay in the area where they are born. This roughly matches the national share of workers who stay in their state of birth.[33]

I implement the indirect inference procedure in two steps. First, I set the share of people who are foreign born, as well as the spread of ξ. These two parameters, along with the two parameters defining the distribution of μ_i, are the only parameters in the model that will affect the moments that I target. I set σ_ξ to 0.4, which produces average elasticities close to those found by Suárez Serrato and Zidar (2014), and I assume that 13 percent of workers are foreign born, based on current population statistics.[34] The second step is to simulate a series of productivity draws and match the estimated relationship of locals staying and outsiders moving in, shown in Appendix Table A4, as well as approximately 60 percent of workers staying in their homes. Table 1.5 shows the implied parameters.

I set the other parameters, reported in Table 1.5, according to national averages and to

[32]The functional form of ξ_{ijk} tends to work against this result. For individuals, the logit error structure implies that workers with a 50 percent probability of residing in an area (based on observables) will be the most responsive to changes in real wages. This is at odds with the findings in the first section that workers are unresponsive to changes in general amenities of their home areas, and that workers have a roughly 50 percent chance of living in their home state. Since the structure of the logit model seems to make an unrealistic functional form assumption about individual migration elasticities, it is particularly important that the μ_i terms have a distribution of values. If the μ_i terms all had the same value, then this would actually make the migration elasticities of locals higher than those of outsiders.

[33]Empirically I calculate the normal distribution of μ_i using Gaussian quadrature with 100 nodes, which provides a good approximation with moderate computational effort.

[34]Within the model, foreign born workers have $\mu_i = 0$ for all areas, since these workers were not born in any area of the US. The model does not feature migration across international borders, however, so they only choose across local labor markets in the United States. The assumption underlying this choice is that changes in any one individual area will affect the attractiveness of the country as a whole by a very small amount.

Table 1.5: Parameter Values

	Description	Value	Reason
σ_ξ	Preference spread	0.4	Suárez Serrato and Zidar (2014)
μ_{μ_i}	Preference for home	3.23	Indirect inference
σ_{μ_i}	Preference for home spread	3.7	Indirect inference
η^Y	Armington elast	4	Feenstra et al. (2014)
α^Y	Capital share	0.33	Standard
ρ	Real interest rate	0.05	Standard
η^H	Housing supply elasticity	15	Green et al. (2005a)
α^H	Non-tradeable share of cons	0.33	Albouy (2009)
J	Number of areas	722	Number of CZ's
s_F	Share foreign	0.13	US population
s_D	Share dying	0.02	60 year avg lifespan

Note: These are the parameter values used for the model's calibration, including a text description and a note describing the reasoning behind each value. Variables noted with "indirect inference" were computed using an indirect inference procedure that seeks to match the share of people who live in their home areas (states in the data) as well as a coefficient of a regression of the proportion of people who stay in their home state on log changes in the population of that state.

the literature. Three important parameters are the Armington elasticity (η^Y), the share of housing/local goods in overall consumption (α^H), and the elasticity of supply for housing (η^H). For η^Y I follow Feenstra et al. (2014) and choose an Armington elasticity of 4. I set α^H (the share of locally produced / housing goods in consumption) based on Albouy (2016). I set η^H to be equal to roughly the middle of the estimates in Green et al. (2005a).[35]

For the exercises below, I focus on a situation where all areas are initially the same. I set the location specific terms (A_j, θ_j, and ϕ_j) to be identically equal to one. These represent demand and supply shifters that have impacts on the levels of wages, rents, and amenities but not on responses to them, which are the focus of the exercises below. I set J so that there are 722 areas, which matches with the number of local labor markets in my data.[36]

[35] There is a substantial literature showing that housing supply elasticities vary significantly based on local geography and zoning (Most prominently Saiz (2010) and Glaeser and Gyourko (2002)). A concern with η^H is the functional form of the supply curve. Glaeser and Gyourko (2005) and Notowidigdo (2011) argue that this is an important concern, particularly for rapid declines in productivity. Other studies, including Davis et al. (2013), find that it plays a limited role. Empirically, I find little evidence that housing supply varies greatly between areas with high and low ties, which supports the idea that local ties are built up over long periods where housing can depreciate more easily.

[36] Albouy (2016) and Diamond (2016) do attempt to allow areas to vary in terms of local productivity and local quality of life, and they present estimates for various areas. I do not attempt this exercise here, however, because it is extremely data intensive and because the focus of this paper is not on estimating the relative merits of various areas, in terms of productivity and amenities. Instead, I use the model as a way of exploring how integrating people's local ties impacts the dynamics of reaching spatial equilibrium after a series of shocks, in a simple and tractable way.

Mechanisms and dynamics

This section shows how local ties form, how they impact migration and local wages, and how they can be reallocated over time. Local ties are stronger in areas where productivity has declined, since people with strong local ties are much more willing to endure the low real wages that come with a decline in productivity.[37] These higher levels of local ties imply that new shocks will have smaller impacts on population, and larger impacts on real wages, than in areas that are growing. People's local ties can be reallocated, but the process can be slow – particularly if it involves new people being born and growing up in different areas.

Figure 1.6: Effects of two Consecutive Decreases in Productivity

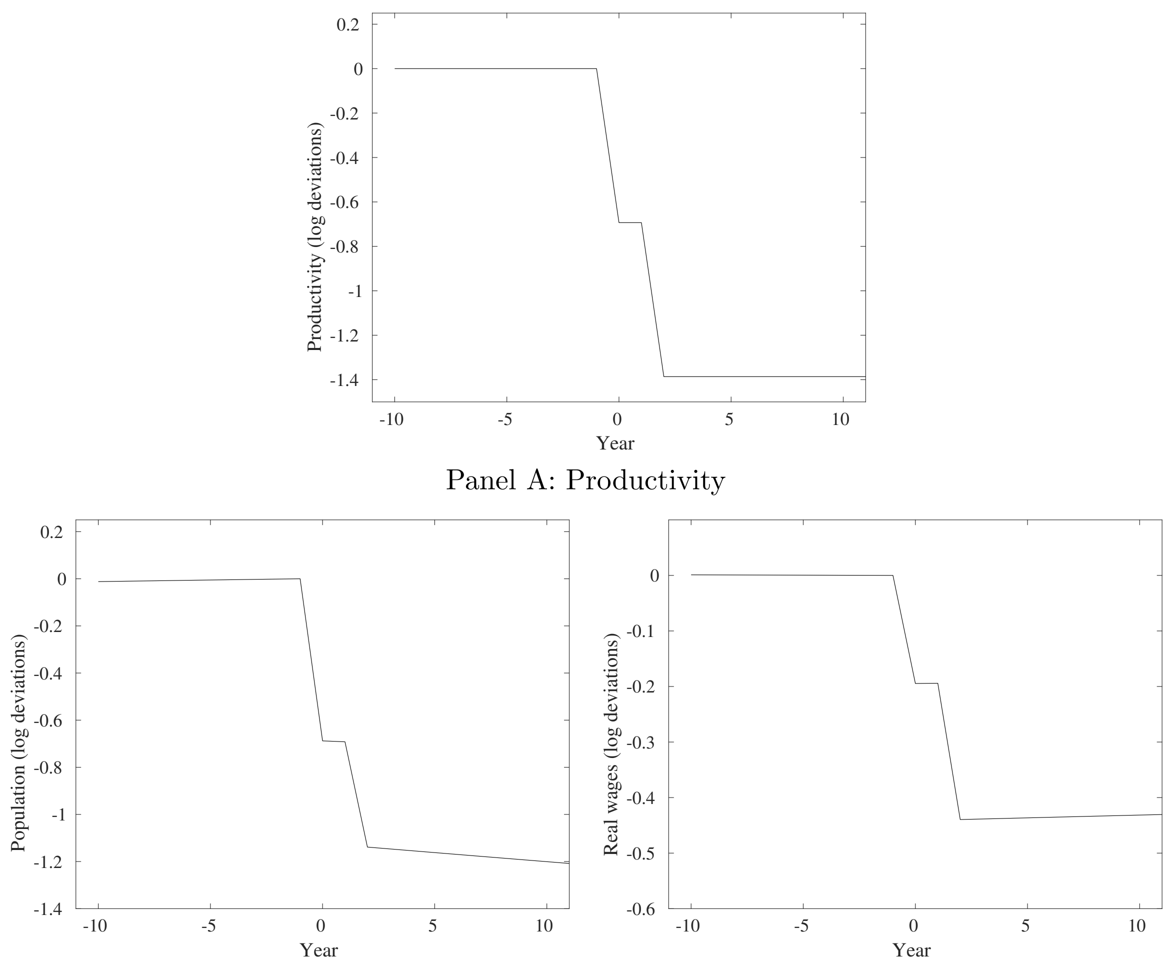

Panel A: Productivity

Panel B: Population

Panel C: Real wages

Note: This plots the population and wages change after two exogenous shocks to productivity, with each variable measured in log deviations. Each shock decreases productivity by 50 percent, about 0.69 log points.

Local ties change after population changes, so two identical changes in productivity can

[37] A similar dynamic applies to changes in amenities.

have different effects if they happen in succession. Figure 1.6 shows this dynamic for two permanent 50 percent decreases in local productivity (θ_j), which is plotted in Panel A.[38] The change in productivity decreases both population, plotted in Panel B, and real wages, plotted in Panel C. The size of these drops varies, however; population drops by more initially, and real wages drop by more in the second period. Population drops by .58 log points initially, and .43 log points in the second period (44 and 19 percent). Real wages, however, drop by .19 and .26 log points, or 18 and then 23 percent. The second decrease in productivity affect local real wages by more, because there are fewer people who are willing to adjust to it by moving to other places.

Areas have different responses after successive changes in productivity because each change leads to different types of residents. Growing places attract outsiders because they offer higher real wages. The lower real wages available in declining places mean that their populations have stronger local ties, however, because people will only be willing to locate there if they have local ties that make up for the low real wages.

Figure 1.7: Share Locals and Population Changes

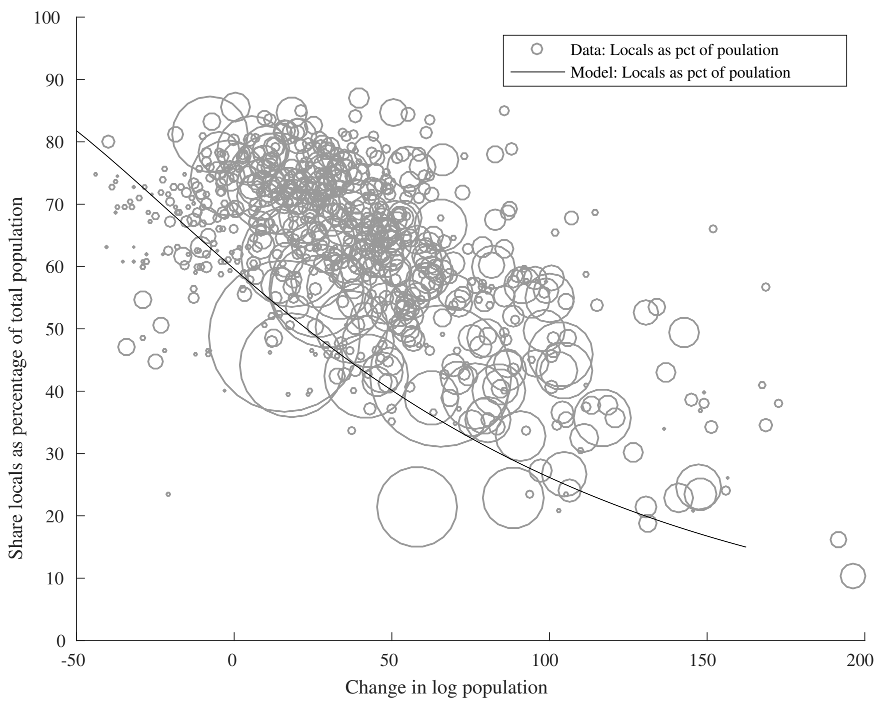

Note: The black line is model predicted share of local workers after an increase in productivity that generates the population change charted. The circles are individual CZs, based on population changes from 1970 to 2008 and the share of locals in 2008.

[38]In addition to productivity, θ_j, I also decrease demand for the local good from the national goods firm, given by θ_j. This separates the impact of local ties from the imperfect substitutability of the various local goods.

The inverse relationship between population growth and people's level of local ties is very apparent in the data, and the model is able to match it quite well. Figure 1.7 plots how the share of people who were born nearby, my empirical measure of local ties, changes as places have bigger increases in population. It plots both values from the data, where each circle represents a commuting zone, and the implied relationship in the model as the solid line, based on changing levels of productivity in the model. Each relationship has a pronounced negative slope.

Figure 1.8: Migration Elasticities After a Shock to Local Productivity

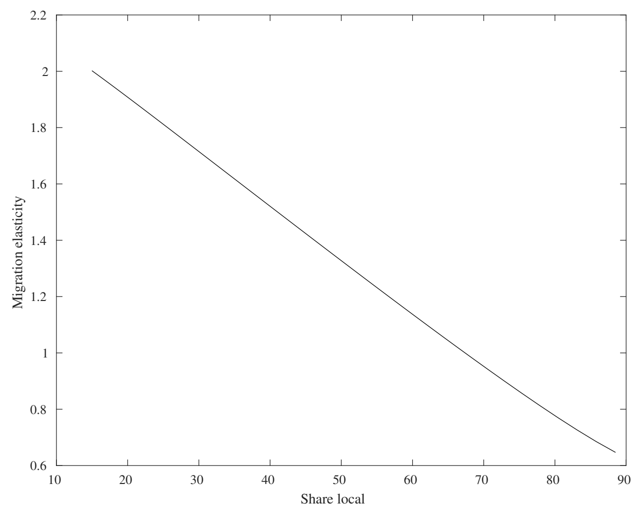

Note: The y axis represents migration elasticities as derived in equation 1.8 while the x axis shows the share of locals in total population. Differences in share of locals are the result of exogenous changes in local productivity.

These different levels of local ties lead to differences in migration elasticities, so the amount of migration that happens in response to a change in productivity will vary depending on an area's level of local ties. Figure 1.8 quantifies this by comparing places with different changes in productivity, and hence different levels of local ties. Areas that have had increases in productivity, and hence have residents with fewer local ties, have smaller migration elasticities. A place where a bit less than 30 percent of the population was born locally, like Denver, will have a migration elasticity of around 1.8. Conversely, a place where 70 percent of the population was born locally, like Detroit, will have a migration elasticity that is one half as large, at 0.9. Compared to the estimates coming from responses to labor demand shocks, these differences are fairly modest.

Another departure from more traditional models of spatial equilibrium is that local ties lead to very slow adjustment processes after local shocks, since it can take a long time for local ties to be reallocated. Figure 1.9 shows an example of a slow adjustment process after a

Figure 1.9: Effects of a Decrease in Productivity Over Time

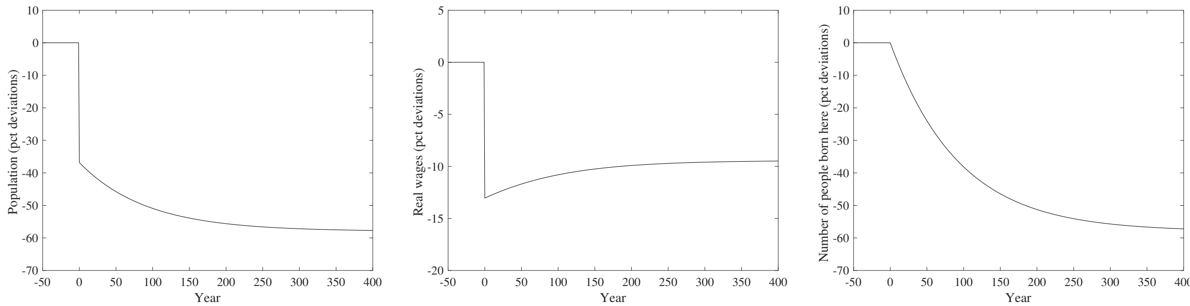

Panel A: Population　　Panel B: Real wages　　Panel C: People born here

Note: This plots the population, real wages, and the number of people born in an area after a permanent 50 percent decline in productivity. Real wages are the total of the wage, rent, and amenity terms in the utility function and the number of people born in the area is the total number of people whose birth place is the area, N_k

50 percent permanent decline in productivity. On impact, the drop in productivity leads to a 37 percent decrease in population and a 13 percent decrease in real wages. The impact is different in the long term, however, once people's ties are reallocated to other places. Panel C shows how the drop in population leads fewer people to be born in the area over time, as people are born with ties to different areas. This process occurs only very gradually, however. Local ties are still being reallocated 300 years after the change in productivity.

The gradual decline in people's local ties leads to a gradual decline in population and a gradual increase in real wages. Intuitively, since people are less enthusiastic about living in the area, employers have to offer higher wages, and landlords have to charge lower rents, to attract workers. The eventual population is about one third lower, and wages are about one third higher, than the initial drop.[39] The larger initial drop in real wages illustrates a form of hysteresis in the model. Previous declines in productivity can lead to lower levels of real wages that can persist for a long time.

The same general patterns apply to different sized declines in both local productivity and local amenities. Table 1.6 summarizes these results by showing the size of the initial and the eventual drop in both population and real wages, as well as the time that it takes for population to be within 1 percent of its steady state value. Even small changes in productivity and amenities can take generations to be fully reflected in local populations. For example, population is more than one percent away from its steady state value 90 years

[39]Intuitively, real wages can decline in this context because of the other, idiosyncratic term in people's preferences, ξ_{ijk}. The structure of this term means that some people will have a preference for living in this area despite the low real wages, even though they have no local ties in the area. More canonical Rosen-Roback models lack this term. In one of these models, real wages would be equalized across areas.

Table 1.6: Time to Convergence After Various Shocks

	Population decrease		Years to
Productivity decrease	Initial	Total	steady state
0.75	-55	-82	417
0.50	-37	-58	351
0.25	-19	-30	264
0.10	-8	-12	163
0.05	-4	-6	91
Amenities decrease			
0.75	-45	-69	381
0.50	-34	-54	341
0.25	-20	-32	272
0.10	-9	-14	180
0.05	-5	-8	110

Note: The table shows the initial and total drop in population (measured in percent), as well as the time it takes to be within one percent of the steady state value after permanent decreases in productivity and amenities of the percent indicated. Initial drop means the drop in population immediately after the shock, while total drop gives the drop from the initial to the steady state level of population. Both are measured as a percentage of the initial population. The next three columns give the time it takes to be within 1 percent of the steady state values. These are separated by the time without a subsidy, the time when there is a subsidy to the area and the difference between the two. The subsidy is assumed to be a 10 percent subsidy that declines by 4 percent each year. For the smallest shocks, the subsidy can be larger than the effect of the shock, since the subsidy does not vary with the size of the initial shock.

after a five percent drop in productivity or amenities. Adjustments are even slower after larger changes in productivity, with a 50 percent drop taking 400 years to be fully reflected in local population counts. After a change in productivity or amenities, population falls by about 1/3 less than its total decline.

1.4 Migration and Welfare After Local Subsidies

Migration elasticities determine local levels of welfare in spatial equilibrium. Higher migration elasticities imply that equal sized changes in labor demand will have smaller effects on wages and other margins of adjustment, like people dropping out of the labor force. Migration also impacts welfare after a subsidy to a particular area. If migration is large enough, then it will undo the effects, in terms of welfare, of a local subsidy. The literature on place-based policies, surveyed by Glaeser and Gottlieb (2008) and Neumark and Simpson (2015), very much emphasizes this dynamic.

Here I use a sufficient statistics approach, applicable to most models of spatial equilibrium, to show how migration determines the welfare implications of spending in local areas. The logic is straightforward, and it relates to arguments originally made by Harberger (1964); a subsidy to a local area induces misallocations in the form of people moving into that area, when they would prefer to live other places. This misallocation, as it turns out, is actually the only effect on total welfare, since the envelope theorem will apply to any model that begins in undistorted equilibrium.[40] The application of the envelope theorem, and the summing of total welfare, mean that this is actually a fairly straightforward application of the sufficient statistics techniques (Chetty (2009)).

It is possible to measure migration elasticities using a two stage least squares, since they are the responses of log population to an increase in log incomes, allowing local prices to change. To instrument for local incomes, I use the two plausibly exogenous changes in labor demand that I identified before, a Bartik shifter in the 1980s and the impact of Chinese imports in the 1990s and early 2000s. I find that migration elasticities are quite low in areas where people have high levels of local ties. Even very large subsidies will cause only small distortions in the local area. In areas with lower levels of local ties, migration elasticities are bigger, but they still are fairly small. Local subsidies in these areas will be more distortionary.

Welfare impacts of local subsidies

To show that migration elasticities are sufficient statistics for welfare, consider the effects of a local stimulus in a general model of spatial equilibrium. For simplicity, I focus on a cash transfer to an area.[41] In the model, I have five actors: workers, landlords, local firms, a national firm, and a national government. The model is in undistorted spatial equilibrium, before a subsidy is enacted for a particular area, paid for with lump sum taxes levied on other areas.

I add up the welfare of each agent, each converted to a money metric, to provide a measure of total welfare in the economy. For many of the agents, utility can already be thought of in terms of a money metric – profits. For workers, however, I need to convert utility into a money metric. To do this, I use an indirect compensation function, as defined in Varian (1984). Using this construct gives me a measure of equivalent variation due to a change in prices.

[40] To my knowledge, the first paper to note that the envelope theorem applied in spatial equilibrium was Busso et al. (2013). Kline and Moretti (2014b) present the idea in the context of a particular parametric model, that provides a good amount of intuition about the effects. Some models, such as the one contained in Suárez Serrato and Zidar (2014), do contain distortions that make the envelope theorem inapplicable. Presumably, however, the literature on analyzing dead weight losses in the presence of other distortions, enjoyably reviewed in Hines (1999), could also give some guidance.

[41] Other programs should have similar dynamics in terms of migration, so I have not focused on them here. The most obvious distinction is that other consumption focused programs would have costs that might not exactly equal their benefit to residents. E.g. the public goods that they provide may be worth more, or less, than the cost of providing them. The welfare implications of polcies that focus on increasing local firm productivity, or drawing local firms, will similarly depend on migration elasticities, as I show in the context of my parametric model below. I do not attempt to expand my argument to them here, but it should be a straightforward exercise.

For worker i, the indirect compensation function, $m_i(w_1, r_1, g_1; w_0, r_0, g_0)$, is a function of an initial menu of wages (w_0), rents (r_0), and governmental subsidies (g_0) across areas, as well as a counterfactual menu of new wages (w_1), rents (r_1), and subsidies (g_1). $m_i(w_1, r_1, g_1; w_0, r_0, g_0)$ measures the additional income needed, at the initial wages, rents, and subsidies, for worker to have the same utility she has at a new set of wages, rents, and subsidies. In other words, $m_i(w_1, r_1, g_1; w_0, r_0, g_0)$ measures equivalent variation.[42]

For local landlords, local firms, and national firms, profits will vary depending on the price of inputs and outputs. Landlords' profits are an increasing function of the level of rents, $\pi_j^H(r_{1,j})$, where j denotes a local area. The assumption is that landlords will develop all parcels of land where the cost is lower than $r_{1,j}$ and gain positive profits off the cheapest parcels of land to develop. In keeping with much of the literature, they do not use labor in the production process. Local firm profits depend on the cost of labor locally, as well as the price of the local good that they sell, $\pi_j^Y((w_{1,j}, p_{1,j})$. Similarly, the national firm's profits depend on the price of the tradeable consumption good (normalized to one), relative to the local goods used to make it (p_1).

Taking all of these components, and adding in the cost of providing subsidies ($\sum_j g_{1,j} N_{1,j}$, which ignores distortions in raising the tax revenues for a subsidy), gives the total amount of welfare in the model. Note that I have allowed all of these sub-components to have very general functional forms:

$$W(w_1, r_1, p_1, g_1) = \sum_i m_i(w_1, r_1, g_1; w_0, r_0, g_0) + \sum_j \left[\pi_j^H(r_{1,j}) + \pi_j^Y(w_{1,j}, p_{1,j}) - g_{1,j} N_{1,j}\right] + \pi^Y(p_1) \tag{1.10}$$

In this setup, the relevant comparative static is the effect of an increase in $g_{1,j}$, relative to $g_{0,j} \equiv 0$, on total welfare. An increase in $g_{1,j}$ is literally a cash subsidy to a local area, like the implied subsidy inherent in the income tax system (Albouy (2016)), but it

[42]Compensation functions, also called money-metric utility functions, are described in a more general context in Varian (1984) and more formally in Chipman and Moore (1980) . In this case, I use what Chipman and Moore (1980) calls "generalized equivalent variation" to measure welfare, since I am interested in changes from a fixed equilibrium. I follow the literature on spatial equilibrium in performing this money metric aggregation while describing distributional effects by showing impacts on workers, landlords, and firms separately. Note that converting to a money metric and aggregating can be more problematic that it may at first appear. For example, converting to a money metric treats money going to all people equally, ignoring differences in marginal utilities of income and preferences about equality. These different marginal utilities are highlighted in Glaeser et al. (2008) and distinguishes their setup from later work. Blackorby and Donaldson (1990) formalize this point and summarizes a large body of work in welfare economics that identifies other problems with the approach, even where equalizing transfers are feasible. Allowing different weights, based on either equality concerns or some idea of marginal utility, might allow local subsidies to increase social welfare, which they never do in this framework. Determining appropriate social welfare weights, of course, is problematic (e.g. Arrow (1950)).

could also be thought of as a program that benefits residents exactly as much as it costs to implement. In this way, they could be programs to improve public spaces, investments in public schools, investments in residential streets, or even programs to pay the college tuitions of the area's children (e.g. the Kalamzoo Promise program in Kalamazoo, Michigan). These programs may have higher or lower values to residents than the cost of implementing them. Nonetheless, I assume that these costs and benefits balance out, since it is a reasonable starting point for studying a general program, without basing too many conclusions on the specific details of its implementation.

The effect of an increase in $g_{1,j}$ on total welfare is the total derivative of equation 1.10 with respect to this increase in $g_{1,j}$, so it takes into account migration, as well as readjustments due to changes in equilibrium prices. As originally noted by Busso et al. (2013), however, the envelope theorem will apply in this case. Each agent will only want to make very small adjustments in their behavior, since the small change in prices brought on by the program will still leave them quite close to an optimum. This means that I can ignore the effects of re-optimization, since these will have minimal impacts on welfare.[43]

The first equation below shows the other effects, which are due to two different types of changes. The first change is in the number of people in the area, which will increase the subsidy payment, since now there are more people to compensate. The envelope theorem, however, means that we do not need to consider differences in people's welfare as they move, since people who chose to migrate feel that the new area is only marginally better. The second change is in various prices. The changes in prices will mean that some actors will do better than others. For example, if wages fall in response to increased population, then local firms will have lower costs that are exactly equal to the lower level of compensation for local workers. I break out formulas for each actor in the appendix.

$$\begin{aligned}\frac{dW}{dg_{1,j}} &= \sum_{j'} \left[(H_{j'} - N_{j'}h_{j'}) \frac{dr_{1,j'}}{dg_{1,j'}} + (N_{j'} - N_{j'}) \frac{dw_{1,j'}}{dg_{1,j'}} + (Y_{j'} - Y_{j'}) \frac{dp_{1,j'}}{dg_{1,j'}} \right] + (N_{j'} - N_{j'}) - \frac{dN_j}{dg_{1,j}} g_{1,j} \\ &= -\frac{dN_j}{dg_{1,j}} g_{1,j}\end{aligned}$$

where $H_{j'}$ is the total amount of housing, $N_{j'}$ is the total population, $Y_{j'}$ is the total production of the local good, and $h_{j'}$ is the total amount of housing purchased per worker in area j'. Each of these terms measures the total amount of a good purchased (housing, labor, and the final good) in area j' and multiplies it by the change in its price. Note that I have

[43] While this model is quite general, it does not include some other dynamics, like frictional unemployment. Schmieder and von Wachter (2016) show how this logic can be extended to frictional unemployment and how adjustments in terms of search effort (and job finding rates) would similarly drop out of the calculation.

omitted the before and after subscripts for simplicity.

Since both the buyer and the seller are on equal terms here, however, the transfers all cancel out and the welfare impact only depends on a single term. This term, $(N_{j'} - N_{j'})$, gives the subsidy payment to workers in the area, minus the cost of that payment for people already residing in the area.

The fact that changes in prices only entail transfers from one party to another means that all of these effects will cancel, leaving only the effect of population changes on the tax necessary to finance the subsidy. This happens in the second line, and it makes the connection to Harberger (1964). The welfare impact of the subsidy is negative and equal to the size of the distortion that it induces in people's choices about where they would like to live, since some of the subsidy goes to paying to relocate people to places they would not otherwise live.

Figure 1.10: Dead Weight Loss Due to a Location Specific Subsidy

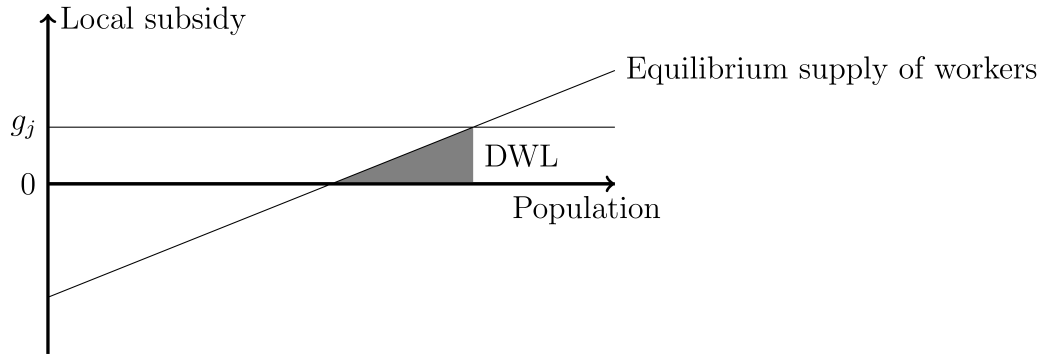

Figure 1.10 gives a graphical interpretation of this distortion. It plots the total population of an area, on the x axis, against the subsidy paid to residents, on the y axis. The shaded triangle, labeled DWL, gives the size of the welfare loss, or the dead weight loss. This triangle is the integral of the additional population attracted by the subsidy times the subsidy payment that goes to making them indifferent towards living there. Since the migrants do not gain utility from this portion of this subsidy, this portion is effectively wasted.

In a wide class of models of spatial equilibrium, the welfare implications of a local subsidy are proportional to the subsidy's effect on migration. Even in models where there are other distortions, or where the programs are more complicated, this will be an important effect. Noting this regularity gives a simple way of thinking about how various public policies impact the distribution of welfare across areas. It formalizes the idea that we should be concerned about population reallocation in response to local subsidies. In a world where areas have different migration elasticities, the formula shows whether it is feasible to change welfare through governmental action in a particular area.

Measuring migration elasticities

It is possible to measure migration elasticities with minimal modeling assumptions by measuring changes in population after exogenous changes in local incomes. So long as one uses an exogenous change in local incomes, and one does not control for changes in local prices, the estimate will embed the effects of people's preferences about the local area as well as equilibrium changes in prices.

The equation below shows the basic empirical specification that I use to recover the migration elasticity, $\eta_{\text{Mig},j}$. The migration elasticity, $\eta_{\text{Mig},j}$, measures the effect of an increase in log incomes on log population, including the endogenous responses of other local prices. Since I intend to include the effects of these other local prices, like housing prices, I do not attempt to control for them. Following the reduced form results, I do control for decade fixed effects, γ_t, and the standard set of controls from the reduced form regressions, βX_{jt}. These ensure that the regressions are not being driven by different trends for areas where people are of different ages, different education levels, or places where more people are foreign born, for example. Following my earlier regressions, I allow heterogeneity across areas, j, by splitting areas into bins based on their levels of local ties, and also by including a continuous interaction with the level of local ties in each area. I present these variations in Section 1.2.

$$\Delta \log \text{pop}_{jt} = \eta_{\text{Mig},j} \Delta \text{income}_{jt} + \gamma_t + \beta X_{jt} + \epsilon_{jt} \tag{1.11}$$

To isolate plausibly exogenous changes in local incomes, I use both the Bartik shifters in the 1980s and the Chinese import measures in the 1990s and early 2000s. As I discussed in Section 1.2, each has a good claim at exogeneity in this context.[44] To maximize power, I stack the data for each of the three decades and estimate one set of parameters in the second stage. I allow the Bartik instruments to have different first stage effects from the trade instruments, but I assume the impact of the trade instruments is the same in each decade. [45]

I measure changes in incomes by combining information about changes in wages with information about the availability of jobs, as measured by the employment to population ratio. Wages are an imperfect measure of labor incomes because there appear to be significant

[45]Another point about the instruments is that the use of labor incomes abstracts from people's labor leisure choices. In my model, and in much of the literature on spatial equilibrium, an increase in labor incomes has an identical effect as an equivalent increase in local subsidies, because people work for a fixed number of hours in the place where they live. By the logic of the sufficient statistics derivation, however, the impacts on people's labor leisure choice should fall out from the first order welfare impacts of a local subsidy. Intuitively, people are roughly indifferent about working more or searching harder for a job. The most serious limitation for my empirical work appears to in terms of attracting population; local subsidies may be more or less appealing to migrants than increases in wages.

frictions to their adjustments, particularly in periods when labor demand is falling. Workers and employees may be reluctant to accept declines in nominal wages, for example, and search frictions could also play a role.[46]

In my empirical setup, labor incomes are the product of wages once one is employed times one's probability of being employed, as in Harris and Todaro (1970). Potential migrants consider not only wages, but also the difficulty of finding and keeping a job. I use the employment to population ratio as a measure of this probability.

Changes in log labor income, then, are changes in log wages, Δwage_{jt}, plus changes in the local employment to population ratio, $\Delta \text{emp ratio}_{jt}$.

$$\Delta \text{income}_{jt} = \Delta \text{wage}_{jt} + \Delta \text{emp ratio}_{jt}$$

Estimated migration elasticities

The estimated migration elasticities are significantly lower in areas where people have higher levels of local ties, but they are still fairly modest in areas where people have relatively low levels of local ties. The implication of these relatively low migration elasticities is that local subsidies are the least distortionary in areas where people have many local ties. Distortions are also relatively modest, however, in areas where people have few local ties.

Migration elasticities, reported in Table 1.7 are lower in areas with higher levels of local ties. According to the first column, the migration elasticity is around one in an area with a lower level of local ties, and around one tenth in an area with a higher level of local ties. The linear interaction in the second column implies that a 10 percent increase in the share of people who were born locally leads to a decrease in the migration elasticity of around 0.35. Figure 1.11 plots the estimated migration elasticities from this linear interaction and shows that the implication is that areas with very low levels of local ties, e.g. in Miami, where 15 percent of the population was born locally, the migration elasticity is around two. In areas with very high levels of local ties, the estimated migration elasticity reaches zero.[47]

[46]It also is possible to only use wages. When I only use wages as a measure of labor incomes my results are similar, but much less precise. These results are available upon request.

[47]Appendix Table A11 separates out estimates from the 1980s using the Bartik instruments from later estimates using the trade instruments. It produces qualitatively similar results, but there is some suggestion that migration elasticities were uniformly lower in the later periods. This seems to support several literatures that find that Americans are becoming less and less mobile over time. Note that the estimates are quite imprecise and that there may be some concerns about weak instruments, however. Another robustness exercise, available upon request, is using only wages to measure labor incomes. Using only wages results in migration elasticities that are slightly higher, but much more imprecisely measured, and results in some concerns about the strength of the first stage. None of the specifications give evidence to overturn the finding that migration elasticities are lower in places where people have stronger local ties.

Table 1.7: Instrumental Variables Estimates of Migration Elasticities

	(1)	(2)	(3)
Low ties: Wages	0.99		
	(0.39)		
High ties: Wages	0.08		
	(0.24)		
High ties indicator	-0.95		
	(1.14)		
Main effect of wages		2.40	1.35
		(0.34)	(0.58)
Interaction (x100)		-3.44	
		(0.64)	
Main effect of ties		-0.14	
		(0.07)	
Year fixed effects	Y	Y	Y
Controls	Y	Y	Y
P-val: No diff	0.03	0.00	
First stage F: Wald	35	37	46
First stage F: K-P	12	11	14
Observations		2166	

Note: This table shows results from IV regression results, where I use Bartik (for the 1980s) and trade (for the 1990s and early 2000s) shocks as instruments for hourly (residualized) wages. The coefficients on wages are the relevant migration elasticities in each type of area, and the interaction term is multiplied by 100 so it displays more significant digits. Standard errors, in parentheses, are clustered at the state level. P-val: No diff reports the p value against the null hypothesis that there are no differences between areas with different levels of local ties, and the two first state F statistics are the standard Cragg-Donald Wald (partial) F statistic as well as the robust Kleibergen-Paap Wald rank F statistic, in that order. The regression contains 722 commuting zones that are first differenced across four years (so there are three observations per commuting zone). The controls and weights are the same as in the earlier reduced form regressions. Areas with low levels of local ties are areas where less than 60 percent of the population was born in the same state they are currently living in.

Table 1.8: Impacts of Local Subsidies

	Migration elasticity					
Subsidy	5	2.5	1	0.5	0.1	0
20 %	50.0	25.0	10.0	5.0	1.0	0.0
10 %	25.0	10.0	**5.0**	2.5	**0.5**	0.0
5 %	10.0	5.0	2.5	1.3	0.3	0.0
1 %	2.5	1.3	0.5	0.3	0.1	0.0

Note: The table shows the size of deadweight losses brought on by a subsidy for differently sized migration elasticities (columns) and per person subsidies. The deadweight losses are measured as a percentage of the cost of the subsidy, if it were to apply to only the initial population. The size of the subsidy is measured as a percentage of the initial level of income. The highlighted values are the dead weight losses that would apply for a 10 percent subsidy to an area with low ties, and the same subsidy to an area with high ties, according to the first column of Table 1.7.

Figure 1.11: Estimated Migration Elasticities

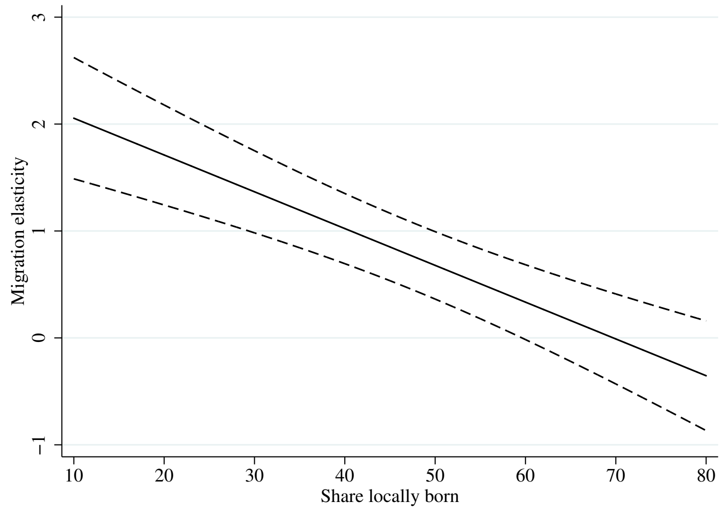

Note: This plots the estimated elasticities from the regression reported in the second column of Table 1.7, for areas with different levels of local ties. The dotted line is a 95 percent confidence interval.

Table 1.8 quantifies the welfare implications of local subsidies in areas with different migration elasticities. It measures the welfare cost of the distortion as a percentage of the benefit of the subsidy to the initial population. According to Table 1.8, a subsidy to an area with a low level of local ties results in a small but meaningful distortion. The same subsidy in an area with a high level of local ties, however, results in a distortion that is economically insignificant. A subsidy equal to 10 percent of an area's initial wages, for example, leads to a dead weight loss equal to five percent of the subsidy's payments to the initial population in an area with a migration elasticity of one. In an area with an elasticity of 0.1, however, the distortion is one tenth that size.

In general, the parameter estimates imply that local subsidies will create distortions that are fairly small, particularly in declining areas with high levels of local ties. In fast growing areas, with very low levels of local ties, the distortions could be meaningful after large subsidies, with estimates implying that these distortions could equal about 20 percent of the size of a decently sized local subsidy. For declining areas, however, there will be very small distortions, even after very large subsidies.

1.5 Conclusion

In this paper, I examined how people's ties to declining areas impact migration elasticities and the processes of growth and decline more broadly. I confirm previous findings that people feel strongly about their birth places, which I use as a proxy for most workers' homes. Using both reduced form regressions and a model of spatial equilibrium, I show how these connections can lead to lower migration elasticities in declining areas. The model shows how heterogeneous migration elasticities have important implications for welfare and for place-based policies. The results suggest that declining areas have lower levels of welfare than otherwise might be expected in spatial equilibrium and that targeted place-based policies may be particularly effective in declining areas. These results differ substantially from the story of utility equalization in the benchmark model of Rosen (1979) and Roback (1982).

Low migration elasticities in declining areas suggest that migration should not equalize utility across areas within countries, leaving room for governmental interventions. In this paper, I have focused on place-based policies, and I have modeled them in a somewhat reduced form way, abstracting from details specific to their implementation. This is not to say that these details are unimportant. Several excellent papers have rigorously evaluated individual place-based policies that show promise for equalizing welfare in the absence of migration. Indirect place-based policies, which Neumark and Simpson (2015) define as policies that seek to move people out of declining areas, may also play a role. My findings suggest that these policies face a difficult battle in convincing people to move from declining areas, but they are also consistent with some barrier to certain people's migration that may be possible to eliminate. One example might be the lack of social networks in places with more opportunities, which Yannay Spitzer (2015) emphasizes in the context of international migration.[48]

A better understanding of why people have local ties would also be of much use to policymakers and to researchers. If, for example, people have strong local ties because they need to care for elderly relatives, then it might make sense for the government to subsidize care in old age. If people have emotional connections to places, on the other hand, then it may not be as reasonable to expect people to move.

[48]Examples of studies of direct policies include Busso et al. (2013) (Enterprise Zones), Kline and Moretti (2014a) (The Tennessee Valley Authority), and Bartik and Sotherland (2015) (The Kalamazoo Promise Scholarship program). The Gautreaux (Rubinowitz and Rosenbaum (2000)) and Moving to Opportunity (Ludwig et al. (2013), Chetty et al. (2015)) programs are prominent examples of indirect place-based policies, though they move people at a level of geography below the level that I focus on here.

CHAPTER II

Parental Proximity and Earnings After Job Displacements

with Patrick Coate and Pawel Krolikowsi

2.1 Introduction

Americans live very close to their parents. In the Panel Study of Income Dynamics (2017), the median household head, ages 25 to 35, lives just over five miles from their parents and about one-fourth of these people live in the same neighborhoods as their parents.[1,2] In modeling location choice, economists often explain this proximity to parents with a substantial amenity value of being close to home.[3] We find that young adult children also experience labor market benefits when they live close to their parents.

Our focus on the labor market outcomes of young adults is motivated by two literatures. First, Wilson (1987) and Munshi and Rosenzweig (2016) show how parents can provide informal social insurance; Kaplan (2012) and Dalton (2013) link parental insurance to the labor market outcomes of young adults. Second, Granovetter (1995) and Ioannides and Loury (2004) present evidence that the use of friends and relatives is prevalent and productive during the job search process, potentially improving the match quality between workers and firms. More specifically, Corak and Piraino (2011) and Kramarz and Skans (2014) show how parents can use their professional networks to help ease children's entry into the workforce.

[1] As in Hellerstein et al. (2015), for example, we use census tracts as a measure of neighborhoods.

[2] Researchers find similar patterns using other data sources. For example, Compton and Pollak (2015) use the National Survey of Families and Households to find that most Americans live within 25 miles of their mothers, and Bui and Miller (2015) use the Health and Retirement Survey to show that median older Americans live 18 miles from their mothers.

[3] Kennan and Walker (2011), for example, describe young men's utility gain from living in the state where they grew up as equivalent to a $20,000 wage increase per year. Other migration models with many locations (Bishop, 2008; Gemici, 2011; Diamond, 2016; Coate, 2017) find similar magnitudes in the US using slightly different combinations of datasets and definitions of home location.

To quantify these parental benefits, we focus on displaced workers, or workers who involuntarily lose their stable jobs through no fault of their own (e.g. being laid off).[4] We focus on job displacements because of their plausible exogeneity (von Wachter et al., 2009) and because these events are followed by large and persistent earnings losses, on average. We follow the standard difference-in-difference methodology of Jacobson et al. (1993) to document the earnings losses of displaced workers in the PSID, but we analyze separately these losses by workers who lived in and out of their parents' neighborhoods at the time of displacement.

We show that young adults, ages 25 to 35, who lived in their parents' neighborhoods prior to a displacement experience a remarkably strong earnings recovery, catching up to a control group of non-displaced workers five years after the displacement. Those living farther from their parents experience a large, permanent decline in earnings, similar in magnitude to the losses found by previous researchers. Stronger earnings recoveries for young adults living closer to their parents are driven by smaller on-impact wage reductions and stronger post-displacement recoveries in both hours and wages. We find no benefit from parental proximity after job loss for older workers.

The earnings results for young workers do not seem to be driven by either selection or differential migration between groups. People who live in the same neighborhoods as their parents tend to earn less, be less educated, and to live in areas with lower employment-to-population ratios. We show that our results are robust to using a propensity score reweighting technique that constructs a sample of people who live farther from their parents, but who are similar in terms of these observable characteristics. In particular, the reweighting procedure controls for the quality of the jobs that both groups lose, since it uses several detailed measures of job quality. We suspect that any differences that remain after the reweighting will lead us to understate the true effect of parental proximity, since workers who have difficulty coping with changes should tend to remain closer to their parents. We also find similar earnings differentials for young workers when we restrict to a subsample of workers who do not migrate after the displacement.

We present three facts that highlight mechanisms besides parental housing transfers acting as insurance after job losses, a channel that Kaplan (2012) emphasizes. First, we find that the benefit of parental proximity extends to people who live in the same neighborhood as their parents, but who do not coreside with them.[5] Second, children who live closer to

[4]See Jacobson et al. (1993) and Stevens (1997). Early literature reviews include Hamermesh (1989), Fallick (1996), and Kletzer (1998). More recent work includes von Wachter et al. (2009) and Davis and von Wachter (2011).

[5]We cannot reject the hypothesis that the earnings recovery of workers coresiding with their parents prior to displacement is the same as those living in their parents' neighborhoods but not coresiding. We show, however, that for workers living outside of the neighborhood – in the same commuting zone and beyond – the benefit dissipates.

their parents spend similar amounts of time unemployed and work a similar number of hours when they are employed after a displacement. Third, although we do find some support for increases in housing transfers around displacement, these increases are small and they do not appear to move differentially for those living close to and farther away from their parents.

A simple theory shows that our results are consistent with parents facilitating higher-wage job offers for their children and inconsistent with selection based on children's desire to live at home. Parents could facilitate higher-wage jobs for their children by either finding jobs for them, or supporting them when they accept new jobs. Regarding the former, we find weak evidence that young adults who live closer to their parents are more likely to be employed in their parents' industries after a job displacement. Regarding the latter, young adults who have children of their own drive the stronger earnings recoveries among workers closer to their parents; this suggests that grandparents' help in caring for their grandchildren is important. The theory also shows that a preference of some children to live closer to their parents would lead these children to have weaker wage recoveries after a job displacement.

The rest of the paper proceeds as follows. Section 2.2 describes our data, describes our sample, and presents our main earnings results using simple averages. Section 2.3 presents these results with a more sophisticated econometric setup, decomposes them into hours and wages, and presents additional results by age, proximity to parents, and geographic mobility. Section 2.4 shows that the earnings results are robust to reweighting and including interactions with additional characteristics. Section 2.5 investigates possible mechanisms related to parental employment networks and the role of housing transfers. Section 2.6 discusses selection and unobserved heterogeneity in the context of our results and presents some back-of-the-envelope calculations that use our empirical results to calculate the value of living close to one's parents. Section 2.7 discusses broader implications of our work and avenues for future research.

2.2 Analysis Data and Sample Averages

2.2.1 Dataset and Sample Construction

In choosing the appropriate dataset for our analysis, we are confined by three major restrictions: 1) The data need to include inter-generational linkages; 2) the data need to include job-history information to identify worker displacements and to measure labor earnings; and 3) the data need to provide (preferably many) repeated observations on workers to implement a difference-in-difference approach. To our knowledge, the PSID is one of few datasets that meets all three requirements.[6] The PSID began in 1968 with an interview of approximately 5,000 families, and follows any new families formed from the original group. We use the 1968 to 2013 waves of the PSID and we employ both the Survey Research Center (SRC) and Survey of Economic Opportunity (SEO) samples with longitudinal weights. Our conclusions, however, are unchanged when we drop the SEO sample and do not use weights. In our analysis, we include observations on people between the ages of 18 and 62 and we restrict our attention to what the PSID refers to as "household heads," who are at least 16 years old with the most financial responsibility for the family. The PSID generally defines this as the male in a husband-wife pair or an unmarried couple who has been co-residing for at least one year. Our results are similar when we include both heads and wives (Appendix B.1).

Due to the genealogical nature of the PSID we have the location of adult children and their parents in each wave if they choose to respond. At the time of the survey, the PSID also collects information about people's labor market experience, including their earnings during the previous calendar year.

Our main results are for job displacements that occur when workers are between 25 and 35 years old, though we also present results for job displacements that occur when workers are between 36 and 55. Job displacements are determined from questions that are asked to employed and nonemployed individuals. Employed individuals who have less than a year of tenure with their present employer (and in some survey years, individuals who started their current job after January 1 of the previous calendar year) are asked: "What happened to the job you had before?" Nonemployed individuals are asked: "What happened to that employer (job)?" (the individual's previous job). The two categories of responses used to identify displacements are "plant closed/employer moved" and "laid off/fired." As is standard in the displaced worker literature, we also impose that workers had at least two years with their employer and were working full-time before the displacement so that our workers have a strong connection to the labor market. Our results are qualitatively similar with different definitions of attachment.[7]

We construct the analysis dataset in the following way. For a given age (the "base age") we include heads that were displaced between the date of their last survey and their current survey and heads that were not displaced. This is the "treatment" and "control" group for

[6]Other possible sources of publicly available data are the National Longitudinal Survey of Youth (NLSY), or the Survey of Income and Program Participation (SIPP). The advantage of the PSID relative to each of them in this context is the wealth of data it collects about parents and children over many years. This is particularly true relative to the SIPP, which only contains information about respondents' birth states.

[7]In the baseline approach we follow the job displacement literature in imposing a positive tenure cutoff, but setting this too high (like six years in Jacobson et al., 1993, for example) causes small sample sizes in our context.

this base age. We include heads who were and were not living in the same neighborhoods as their parents at the time of the previous interview.[8] We repeat this procedure for every base age between 25 and 55 and stack all the samples to create the final dataset.[9,10] To track when workers are displaced or not, let the *relative year* be zero in the base age, one in the year after, etc. For example, for the base age 40, the relative year is -8 when workers are 32, zero when workers are 40, and 6 when workers are 46.[11]

Table 2.1 shows the summary statistics for the final sample, which we restrict to observations with non-missing parents' location information.[12] Since we are using the PSID's poverty (SEO) oversample, and hence sampling is likely endogenous to our outcome (earnings), we take the suggestion of Solon et al. (2015) and use the longitudinal weights provided by the PSID throughout the analysis. The dataset consists of around 50,000 records, with an average of 20 years of observations for each, yielding roughly 1,000,000 person-year observations. The final dataset contains about 1,460 displacements, of which 320 took place while a worker resided in their parents' neighborhoods and approximately 1,140 occurred while a worker was not in their parents' neighborhoods. The average annual displacement probability is around three percent in our sample.[13] Before displacement, the displaced workers

[8]Because of the genealogical nature of the PSID data, we typically observe the parents of single adults or of one set of parents of a married couple. We treat cases in which we have the location of the husband or wife's parents symmetrically, although sometimes this means we are using the household head's parents and sometimes parents-in-law. In some cases, we will observe multiple parents' locations (typically due to divorce of an original PSID household head); in these cases an adult child is coded as same neighborhood if they live in the same census tract as any parent or in-law.

[9]Note that workers may appear more than once in the final dataset because they may be in the control group several times, or in the treatment group at one base age, but in the control group at another base age, etc. In our results, we will cluster standard errors at the worker level to account for these multiple observations.

[10]We use an unbalanced panel of data that was collected about workers when they were 18 to 62. Displacements are only included if they are reported in surveys from 1969 to 1997, while respondents are from ages 25 to 55. Our restriction to displacements from 1969 to 1997 is to preserve the interpretation that someone was displaced in the previous year, since the 1968 survey asks about displacements in the last ten years, and surveys after 1997 ask about the previous two years (the PSID is a biannual survey after 1997). We restrict to base ages between 25 to 55 in order to insure that workers have an attachment to the labor force.

[11]Due to the survey design of the PSID, the location of household heads is only observed if they have previously moved out of their parents' house. Therefore, adult children who have never moved out of their parents' home are outside the scope of our analysis. The United States Census Bureau (2016, Table AD-1) reports that 50 to 60 percent of 18 to 24 year olds live with their parents (including college students living in dorms during the academic year), but only 10 to 20 percent of 25 to 34 year olds do. Thus, beginning our analysis at age 25 mitigates this sample selection issue.

[12]Common reasons we do not know parents' locations are because the parents are deceased and because they were never interviewed by the PSID. The structure of the PSID means that we are much more likely to see parents who were interviewed by the PSID in later years, since parents of original respondents are not interviewed.

[13]This is consistent with displacements rates in previous work. See, for example, Davis and von Wachter (2011) where it is between three and four percent and Kuhn (2002), where it is between four and five percent.

are slightly younger, less educated, earn less, and have been with their employer for a less time than their non-displaced counterparts. Around 16 percent of adults live in the same neighborhoods as their parents. We analyze the data separately for younger workers (ages 25 to 35) and older workers (ages 36 to 55); Table 2.1 presents summary statistics separately for this younger group of workers as well.

2.2.2 Some Preliminary Evidence

Figure 2.1 plots the average earnings before and after displacements. The top panel in Figure 2.1 presents the average earnings of workers who were displaced (dashed) and not displaced (solid) averaged over the base ages 25 to 35. All earnings are measured in 2007 dollars (CPI-U-X1). These lines highlight the dramatic earnings consequences of worker displacement. The figure delivers three messages, which have been documented in many prior studies. First, displacement leads to a large initial drop in earnings. Annual earnings drop by around \$10,000, or around 20 percent of pre-displacement earnings.[14] Second, while earnings for these displaced workers recover, this recovery does not exceed the earnings gains experienced by the control group of non-displaced workers. Even 10 years after the displacement, the earnings of displaced workers have not caught up with the earnings of non-displaced workers, despite earnings recovering to their pre-displacement levels after around six years. Finally, there do not appear to be differences in the trends of earnings prior to the displacement. There is a difference in the level of earnings before displacement between the control and treatment groups. We address these level differences in the empirical exercises that follow.

The bottom panel of Figure 2.1 decomposes the average earnings into those that were in their parents' neighborhoods in relative year -1 (light gray), and those that were not in their parents' neighborhoods (dark gray). Many people in our sample aged 25 to 35 are not in their parents' neighborhoods, so the average earnings of those workers (displaced or not) are close to the average earnings presented in the top panel of the figure. Before describing the effects of displacement for these two groups, it is worth pointing out that workers who live in their parents' neighborhoods have significantly lower earnings, even before the displacement.

The bottom panel of Figure 2.1 shows that displaced workers who were not in the same neighborhoods as their parents see large earnings losses relative to a group of workers who

[14] The earnings question in the PSID refers to the earnings during the last calendar year. The displacements have been coded to have happened between the previous survey date and the current survey date. Since most PSID interviews happen in April and May, most of our displacements are referring to displacements that happen at the end of the previous calendar year. As such, the earnings on-impact, although they fall, may not reflect the entirety of the displacement event as the earnings from the last calendar year were largely unaffected by the displacement. Rather, in the year following the displacement the largest reductions may be documented. As such, in the top panel of Figure 2.1 the declines at year '1' are larger than at year '0'.

were not displaced and not in the same neighborhoods. This gap persists over the next 10 years. In stark contrast, those workers who were in the same neighborhoods as their parents in the year prior to the displacement see a much healthier recovery in earnings. Prior to the displacement the difference in the earnings of the displaced and non-displaced who live in their parents' neighborhoods is around \$3,000 and the earnings of the displaced workers recover to this pre-displacement difference around six years after the displacement. The gap in earnings between these displaced workers and the non-displaced group closes entirely within nine years of the displacement. Appendix Figure B1 presents a similar figure for the natural logarithm of earnings with the same conclusions.

In the next two sections we verify these results with a standard displaced worker specification (Section 3), which controls for, among other things, worker fixed effects, and a propensity score reweighting exercise (Section 4), which also controls for observable differences between those in their parents' neighborhood and those farther away. The preliminary results presented in this section are robust to these more sophisticated methods.

2.3 Regression Results

2.3.1 Earnings Losses by Geographic Proximity to Parents

To control for differences between workers who are displaced and not displaced we follow a standard difference-in-difference methodology and estimate the following equation:

$$e_{iat} = \alpha_{ia} + \gamma_t + X_{iat}(\beta^A + \beta^H H_{ia}) + \sum_{k=-4}^{10^+} \left(D^k_{iat}\delta^k + D^k_{iat} H_{ia}\zeta^k \right) + \epsilon_{iat} \tag{2.1}$$

Here e_{iat} is the annual earnings of worker i in calendar year t when the base age is a, α_{ia} represents a worker-base-age dummy, γ_t controls for calender year fixed effects, the X_{iat} terms control for an age quartic, and H_{ia} is a dummy variable indicating whether worker i was neighbors with their parents in the year prior to age a. This dummy is interacted with the age quartic in X_{iat} to allow for different age-earnings profiles for those living near their parents and farther away, which captures the different counterfactual (non-displaced) earnings trajectories for these two groups that we observe in Figure 2.1. The variable D^k_{iat} captures whether worker i at time period t and base age a was displaced k periods ago. We pool the -4 dummy and the $+10$ dummy and omit the -2 dummy so all results are relative to two years before the displacement.[15] As a result, the coefficient δ^k captures the change in earnings for a worker who was displaced k periods ago and was not living in their parents' neighborhoods prior to the displacement relative to other workers who were not neighbors with their parents and were not displaced.[16] The coefficient ζ^k picks up the additional effect of being neighbors with your parents on the earnings outcomes of displaced workers. This approach does not consider how other factors, correlated with living in the same neighborhood as one's parents, might explain the differential impact of displacement on earnings.[17] The propensity score reweighting in Section 2.4 addresses these concerns.

Figure 2.2 presents the effect of displacement on earnings for workers farther away from their family, $\hat{\delta}_k$, and the effect of displacement for workers living in the same neighborhoods as their parents, $\hat{\delta}_k + \hat{\zeta}_k$, for workers experiencing a displacement between ages 25 to 35. These results tell the same story as the simple averages presented in Figure 2.1. At the time of displacement, workers experience large declines in earnings; around \$10,000 for those living in their parents' neighborhoods and around \$15,000 for those living farther away. With the average pre-displacement earnings of these groups being around \$35,000 and \$45,000, respectively, this represents a 30 percent decline in earnings at the time of displacement. The post-displacement recovery, however, is quite different for the two groups. The group living farther away from their parents experiences a small recovery in the short- to medium-run but still has earnings losses of around 25 percent even 10 years after the displacement. In contrast, the group that was living in the same neighborhoods as their parents prior to the displacement experiences a steady recovery in the years following the displacement, with earnings losses indistinguishable from a full recovery after four years. The difference between the earnings of the two groups is statistically significant at longer horizons. The results are

[15]We omit the dummy for two years before the displacement because it is the most recent period that is still before both the displacement and the period that we use to define whether someone is living close to their parents, which is one year before the displacement. Choosing a period before when we define workers' proximity to their parents eliminates any mechanical difference in the baseline level of earnings due to either the displacement, or conditioning on someone living in the same neighborhoods as their parents in that period. This choice also allows us to recover a more precise estimate of our base period, since relatively young workers do not have earnings information going back very far. Kletzer and Fairlie (2003) make a similar choice.

[16]This approach is most closely related to the approaches taken by Davis and von Wachter (2011) and Huttunen and Salvanes (2015). See Krolikowski (2017a) for a more thorough discussion of choosing a control group for displaced workers.

[17]We follow Kletzer and Fairlie (2003), who also focus on young workers, and estimate the earnings model in equation (2.1) without worker-specific time trends. For young workers there exist relatively few pre-displacement earnings observations so we think that worker-specific trends are unlikely to be well estimated. Also, as we shall see, the post-displacement effect of being in the same neighborhood as one's parents on earnings is gradual and builds over time so it might be incorrectly attributed to worker-specific trends if these were included in the estimating equation. Finally, our specification includes worker-base-age dummies, α_{ia}, which vary within workers by base age and therefore likely already pick up worker-specific trends in earnings.

similar if one drops observations that have zero annual earnings or if one uses the log of annual earnings on the left hand side in equation (2.1) as opposed to the level of earnings (Appendix Figures B2 and B3, respectively).

2.3.2 Employment, Hours, Wages, and Unemployment Duration

Figure 2.3 presents the results from estimating equation (2.1) with three different outcomes: an indicator for whether the person worked positive hours in the previous calendar year, the number of hours worked during the previous calendar year (conditional on positive hours), and earnings per hour. The top panel shows the probability of positive hours last year. This falls during the survey after the displacement, as some workers experience an entire year out of work. The graph suggests that displaced workers, regardless of location, are around 4 percentage points (pp) less likely to have employment in the year after the displacement than non-displaced workers. Although the recovery appears slightly stronger a few years after the displacement for those living near their parents, it is difficult to tell the two groups apart with the large standard errors. As such, the two groups seem to have similar post-displacement employment outcomes.

The middle panel of Figure 2.3 shows the results from estimating equation (2.1) with the hours worked last calendar year as the outcome, where we condition on positive hours. On-impact the reduction in hours for the two groups is similar, around 350 hours (approximately 18 percent of the 2,000 hours prior to displacement). Although those near their parents see a larger fall in hours upon displacement, this difference is not statistically significant. The recovery in hours, however, appears stronger for those living in their parents' neighborhoods. In particular, from two to ten years after the displacement, there is a statistically positive increase in the hours of those living in their parents' neighborhoods, whereas those living farther away see their hours plateau. We cannot reject, however, the null hypothesis that the hours recoveries are the same for the two groups.

The bottom panel of Figure 2.3 shows how hourly earnings, conditional on positive hours, move around displacement. At the time of displacement, those living in their parents' neighborhoods experience a significantly smaller wage reduction (around \$1.50/hr) than those living farther away (around \$4/hr). Moreover, workers who lived in the same neighborhoods as their parents at the time of displacement see their wages recover fully, and workers who lived farther away see no recovery. In Section 2.4 we use propensity score reweighting to account for observable differences between the two groups, including pre-displacement wages. The results presented in Figure 2.3 are qualitatively similar.

Although the intensive and extensive margins plotted in Figure 2.3 suggest that the two groups of workers were unemployed for similar amounts of time, the PSID allows us to look directly at the number of weeks a worker spent unemployed in the previous calendar year. Figure 2.4 presents these results separately for those living close to their parents and those living farther away. Not surprisingly, in the year of displacement, the time spent unemployed rises sharply by around seven weeks, but the increase is remarkably similar for the two groups. Over the next few years, the decline in weeks spent unemployed is also very similar. We see these duration results as evidence that longer job search is unlikely to be an important explanation for the differing post-displacement earnings outcomes of the two groups.

2.3.3 Heterogeneous Effects of Parental Proximity

We present a series of stylized facts in this section that give some indirect evidence about the mechanisms that lead to stronger earnings recoveries after job displacements among young workers who live closer to their parents. While none of these facts can conclusively identify a particular mechanism, they still are useful in identifying specific properties of the mechanisms involved. In particular, it appears that the mechanism is strongest when children live quite close to their parents, that the mechanism operates for people 25 to 35, but not for people 36 to 55, and that the mechanism operates both for workers who move and for workers who do not move after displacement. We also find evidence that, among those workers who live close to their parents, the mechanism is only operative for those who have children in their household at the time of displacement. In Section 2.5 we investigate possible mechanisms for our earnings results more directly, and in Section 2.6 we present a simple theoretical framework to interpret our findings.

Figure 2.5 shows the results of estimating equation (2.1) for older workers, ages 36 to 55. As with the earnings outcomes for the young, the earnings losses associated with displacement are large and persistent. This figure, however, suggests that for older workers, living in their parents' neighborhoods does not help post-displacement labor outcomes in the same way that it assists younger workers. If anything, it might be detrimental, but we cannot reject the null hypothesis that the post-displacement earnings effects are the same for the two groups. The discrepancy between the results for the young and the old likely reflects a change in the direction of resource flows, since older adults often live close to their parents in order to care for them in their old age. For example, Chari et al. (2015) estimate the opportunity cost of informal elder care in the US at $522 billion annually and Lin and Wu (2010) find that among people 65 and older who had difficulties with instrumental activities of daily living, about 35% report that a child is a source of informal support.

Figure 2.6 shows the results of estimating equation (2.1), where we look at young workers who are living very close to their parents (same neighborhood), close to their parents (same commuting zone, but not same neighborhood), and farther away from their parents (outside of the commuting zone) at the time of displacement.[18] This figure suggests that those living close to their parents, but not in the same neighborhoods, experience similar post-displacement earnings outcomes than those who live farther away. This evidence leans against the idea that parental employment networks, which should operate within the commuting zone, are an important part of the explanation. We present some evidence on parental networks in Section 2.5.2.

As in Huttunen and Salvanes (2015) and Cao and Stafford (2017), we document a large impact of job displacement on regional mobility (Appendix Figure B4 shows that switching commuting zones increases sharply at the time of displacement).[19] To abstract from mobility, we check that the post-displacement earnings trajectories are not driven by "movers" by restricting the sample to workers who do not change geographic location after the displacement. For this analysis we take the county as the relevant measure of geography because restricting to young workers who never change tracts after a displacement reduces the sample sizes dramatically.

Figure 2.7 presents the results with this restricted sample together with the original results from Figure 2.2. Perhaps not surprisingly, the earnings outcomes of the sample that are restricted to no mobility after a displacement are almost always worse than for the unrestricted sample. However, the differences between those who, prior to the displacement, resided close to their parents and farther away are equally pronounced for this restricted sample. Therefore, post-displacement mobility patterns are unlikely to account for our baseline earnings results.

To estimate the differential impact of parental proximity by whether children are present in the household prior to the displacement, we estimate the following equation:

$$e_{iat} = \alpha_{ia} + \gamma_t + X_{iat}(\beta^A + \beta^H H_{ia}) + \sum_{k=-4}^{10^+} \left(D^k_{iat}\delta^k + D^k_{iat}H_{ia}\zeta^k + D^k_{iat}C_{ia}\nu_k + D^k_{iat}H_{ia}C_{ia}\eta_k\right) + \epsilon_{iat}, \quad (2.2)$$

which is similar to our baseline equation (2.1) but with two additions. First, we add terms, $D^k_{iat}C_{ia}$, and the coefficients on these terms, ν_k, capture the additional effect of children's presence on the post-displacement earnings trajectories of workers. Second, we include the

[18] We augment equation (2.1) by including an interaction of the age quartic with an indicator for whether a worker is in the same commuting zone, but not the same tract, and for whether a worker is in the same tract as their parents. Additionally, we interact the displacement dummies with whether a worker lives in the same commuting zone as their parents and the same tract as their parents. This approach allows for mutually exclusive age quartics for the three different groups (same tract, same commuting zone but not same tract, and different commuting zone) while testing for a distinct effect of displacement for those in the same tract as opposed to just the same commuting zone as their parents.

[19] See Mincer (1978) for early work on family ties and migration decisions. Molloy et al. (2011) and Kaplan and Schulhofer-Wohl (2017) provide a more extensive analysis of recent trends in migration, including declining inter-state migration at different geographies.

three-way interaction between the displacement dummies, D^k_{iat}, the dummy variable indicating whether worker i was neighbors with their parents in the year prior to age a, H_{ia}, and a dummy variable indicating whether worker i had children present in their household in the year prior to age a, C_{ia}, with associated coefficient η_k. We are most interested in the differential effect of parental proximity on the earnings outcomes of displaced workers by whether they have children in the household or not, which is captured by the difference between $\hat{\delta}_k + \hat{\zeta}_k + \hat{\nu}_k + \hat{\eta}_k$ and $\hat{\delta}_k + \hat{\nu}_k$.

Figure 2.8 presents evidence that it is those workers in their parents' neighborhoods that have children in their household that drive our overall earnings results for young workers. In particular, Figure 2.8 Panel A shows that for those with children, the earnings trajectories around displacement look similar to the overall results. Panel B shows that for workers who live in their parents' neighborhoods and who do not have children in their household, the earnings recovery looks remarkably similar to those without children who live farther away from their parents. We interpret these findings as evidence that parents can help their adult children by stepping in to provide childcare after a job displacement.

Several additional findings, which we present in Appendix B.1, suggest that the earnings differences that we observe are driven by having parents nearby. The baseline earnings results are robust to not using PSID sample weights (Appendix Figure B5), to including controls for local labor market conditions (Appendix Figure B6), to a distance-based measure of proximity (based on latitudes and longitudes of block groups; Appendix Figure B7), and are similar for men and women (Appendix Figure B8).[20] The results are also similar when we look at workers who are actually co-residing with their parents as opposed to living in the same neighborhoods as their parents (Appendix Figure B9), in the spirit of Kaplan (2012). This finding suggests that the effect of parental proximity is not limited to transfers while coresiding, a topic we discuss further in Section 2.5. We have also found that the earnings differences for the two groups persist even if one includes additional interactions of the displacement dummies in equation (2.1) with whether the worker was displaced while living in the county that they grew up in (Appendix Figure B10). These results suggest that parental proximity has an independent effect on post-displacement earnings from other factors in a worker's home county.

[20] Local labor market conditions include county-level employment-to-population ratios and unemployment rates. Employment-to-population ratios were obtained by merging in information from County Business Patterns (CBP) and population information from the National Historical Geographic Information System, where we linearly interpolate between census years. The former data are available from 1969 onwards. Using county-level unemployment rates from the Local Area Unemployment Statistics (LAUS) program delivers a similar conclusion but those data are only available after 1980, substantially reducing our sample.

2.4 Propensity Score Reweighting

In this section, we show that our results are robust to using a propensity score reweighting technique that chooses a comparison group of young adults, aged 25 to 35, who live farther from their parents, but who lose similar jobs and who have similar characteristics as those who live closer to their parents. Workers who live closer to their parents are less educated, earn less, and tend to live in places with lower employment-to-population ratios, so it is possible that these differences lead to differences in post-displacement earnings.[21] Through our reweighting, we emphasize workers who are similar in these dimensions, but who live farther away. This allows for heterogeneous effects of a job displacement on earnings owing to these observable characteristics.

The reweighting does not control for unobservable difference between the two groups, but we suspect that these are unlikely to change our results for two reasons. The first is that our results do not appear to be sensitive to either including additional characteristics, or taking characteristics away from our reweighting procedure. Appendix B.2 (Appendix Figure B11) presents the main results from the reweighting with different sets of characteristics, and the results tend to be quite stable once one controls for basic educational and demographic differences.[22] The second is because we suspect that the bias would go in the other direction. A worker who knows that they are less able to adapt to new circumstances should prefer to live closer to their parents, since moving involves adapting to new circumstances and since parents can act as supports. If this is the case, then selection on unobservables should lead to more severe effects of job displacements among workers who live in their parents' neighborhoods.

2.4.1 Methodology

Following the literature on propensity score reweighting (Rosenbaum and Rubin, 1983; Hirano et al., 2003), we reweight observations from different groups of young adults so that they have the same characteristics as young adults who are displaced while living closer to their parents. This involves reweighting three separate groups of workers so that each group

[21] Since we use worker fixed effects in our regressions, for these differences to affect our earnings regression results, they would have to have a larger effect on earnings after displacement. For example, this would be the case if workers who lost jobs with lower wages tended to have smaller earnings losses after displacements.

[22] An additional concern is that by controlling for initial wages, we are choosing a comparison group of workers who are less skilled than the group living in their parents' neighborhoods. This would occur, for example, if there was a wage penalty for living close to one's parents, and this wage penalty meant that a worker who earned a given amount in another area would actually earn slightly less if they stayed at home. To address this concern, we present results in Appendix B.2 where we exclude wages from the calculation of our reweights. Instead we include only education, demographic characteristics, and local employment-to-population ratios. Our findings are similar the baseline results presented here.

has similar observable characteristics as the group of displaced workers who lived closer to their parents. To accommodate multiple groups, we follow Imbens (2000).

We compute the weights, W_{ia}, for person i at base age a, who is in a group defined by whether they lived close to their parents (H or A) and whether they were displaced (D or N), $j_{ia} \in \{HD, AD, HN, AN\}$, using the following formula:

$$W_{ia} = \frac{\mathrm{P}(j_{ia} = HD|X_{ia})}{\mathrm{P}(j_{ia} = HD)} \frac{\mathrm{P}(j_{ia})}{\mathrm{P}(j_{ia}|X_{ia})} \tag{2.3}$$

The formula is an application of a typical reweighting scheme (DiNardo et al., 1996; Fortin et al., 2011) to multiple groups. The weight is one for the treatment group ($j_{ia} = HD$) since we are reweighting all other observations to have the same characteristics as this group.

We can recover the conditional and unconditional probabilities in a semiparametric way using sample averages and logit regressions with flexible functional forms. We estimate the probabilities conditional on X_{ia} using a multinomial logit regression, as suggested by Imbens (2000). The predictors are a linear and a square term in pre-displacement earnings, average year to year changes in pre-displacement earnings, wages, a dummy for one-digit PSID occupations, a dummy for being college educated, the number of years of schooling that the worker has completed, the worker's age, the worker's tenure in the job, the employment-to-population ratio in the worker's county, a dummy for whether the worker is male, and a dummy for whether the worker is African American.[23] The unconditional probabilities in equation (2.3) are the proportion of the sample made up by the group.

Table 2.2 shows a validation of the weights for young adults using several of the covariates in X as well as some other variables that were not included in the reweighting. It reports the mean of the variable among different groups of workers and p-values of a Wald test of equality with the group of people who were young and lost their jobs while living in the same neighborhood as their parents. In keeping with our regression analysis it includes each person separately for each year they were in the sample of people at risk for a displacement. Panel A shows these statistics using the initial PSID person weights and Panel B uses the propensity score reweights. As intended, the differences across samples disappears in Panel B where each group has similar initial earnings, ages, years of education, and a similar likelihood of having children. We also verify in Appendix B.3 that average earnings are similar between

[23]The regression is unweighted and all of the controls are the average values of the variables in the three years leading up to the event (ignoring years where they are not observed).

the reweighted groups of workers who were not displaced (Appendix Figure B12).[24]

2.4.2 Earnings Results Using Propensity Score Reweighting

We begin by showing the effects of reweighting in terms of the simple means that we began with in Section 2.2.2. Figure 2.9 shows reweighted means of earnings around potential displacements for each group of young adults. Each group has similar earnings before the potential displacement, and the groups of workers who do not suffer a displacement have very similar trajectories after, which suggests that the weights emphasize workers with similar counterfactual earnings trajectories in each group. At the time of displacement, workers have similar losses in earnings regardless of whether they live in their parents' neighborhoods, though workers who live closer do earn slightly more on average. The earnings of those who were closer to their parents begin to out-pace the earnings of workers who were farther away after the displacement, however. In the final years this difference is quite large, at around $10,000.

Figure 2.10 shows the baseline regression specification (equation 2.1) with the new weights.[25] It confirms that the inverse probability reweighting procedure produces similar qualitative results as the main specification. The differences between groups are smaller, but they remain both economically and statistically significant. As before, there are substantial drops in earnings following a displacement, though the initial losses are roughly equal and smaller at around $10,000. A steadily increasing difference in earnings emerges in the years after displacement, however, as the group who lives closer to their parents makes up much of the earnings penalty from the displacement. People living closer to their parents see no detectable earnings losses in years four through ten, and they have essentially zero estimated earnings losses from year six onwards, with some positive point estimates. Statistically, the group of workers who were living closer to their parents earns significantly more than the other group in years six and ten, at the five percent level. These differences do appear to be economically significant as well, with a difference of around $8,000, or about 10 percent of initial earnings. People who live farther away from their parents have permanent earnings losses, and their earnings are always statistically significantly different from zero after the

[24] The reweighting exercise is similar to allowing several interaction terms in the regression specification. For comparison we present results from such an interacted model separately in Appendix B.4 (Appendix Figure B13). In Appendix Figure B14, we implement the same propensity score approach, but use only a subset of observations where there is strong common support according to the selection method proposed by Crump et al. (2009) with similar results.

displacement.

2.5 Investigating Mechanisms

We investigate housing transfers and parental networks to see if they might be responsible for differences in post-displacement earnings between people who live different distances from their parents. A relatively small amount of money is transfered through gifts of in-kind housing after displacements, and we find no evidence that these transfers are larger for people who lived closer to their parents before displacement. We find that children who live close to their parents are more likely to be employed in the same industry as their parents, controlling for local industrial composition, although the estimates are noisy. Additionally, in Appendix B.5 we investigate whether search intensity (Appendix Table B3) and industry switching (Appendix Figure B15) can account for our baseline earnings results, and we find that these mechanisms are unlikely to explain our baseline findings.

2.5.1 Housing Transfers

To see if our results are due to children being able to move in with parents who live nearby, we estimate the cash value of parents allowing children to live with them, and we investigate how these housing transfers change around job displacements. Kaplan (2012), in particular, emphasizes that housing transfers can help children to earn more after job displacements by allowing them to be more selective about job offers, and by facilitating investments in their careers. If housing transfers drive our results, we would expect to see increases in the amount of in-kind housing transfers among workers who lived closer to their parents before they lost their jobs.

We use two complementary approaches to measure housing transfers, which may be under-reported in surveys like the PSID. The first records people who report that they receive all of their rent as a gift. The PSID only asks about this if people report that they pay no rent, however, so this approach misses people who pay below-market rents.[26] To address the possibility that some households pay below-market rents to live with their parents, we also construct an estimate of how much the child saves by living in their parents' household.

[25]We continue to use clustered standard errors, which address heteroskedasticity induced by the weighting, along with mechanical correlations across observations. Goldschmidt and Schmieder (2015) also use a similar approach. The main difference with their approach is we use reweighting, while they rely on nearest neighbor matching.

[26]More precisely, this includes households who report having neither owner nor rented and who then go on to say that they live rent free because of a gift, inheritance, or some other non-work related reason.

The second approach backs out how much a child saves by living with their parents in situations where a child moves in with a parent who is also a respondent in the PSID. When a respondent child moves in with a parent, the PSID classifies the household as two different families living within a single housing unit. In these situations, the interviewer will assign everyone living in the housing unit to a family unit and then conduct separate interviews, including questions about housing, with each family unit. The information about total housing costs, combined with the composition of the household allows us to construct a level of housing consumption, using an OECD equivalence scale. We then ask how much this level of consumption would cost if the family lived separately. This value gives us an amount of rent that the child would have to pay, were they to live alone and have the same level of consumption, and the difference between this hypothetical rent payment and the child's actual rent payment is the rent transfer from their parents.[27]

Table 2.3 reports the proportion of households in our baseline sample of adults ages 25 to 35 who receive housing transfers and the average value of these transfers among households who receive them. We report measures from the survey question about receiving rent as a gift as well as measures of how often children live with their parent, and the implied rent savings, according to our procedure.

According to Table 2.3, a relatively small proportion of 25 to 35 year olds receive transfers of rent, and these transfers are modest relative to both average rents and the earnings losses after a displacement. According to both measures, less than ten percent of the sample receives a transfer of housing at the date of the survey. Eight percent of the sample live with a parent and around two percent receive all of their rent as a gift. Among households who receive a transfer, the average transfer was around \$4,300 according to the implied savings, and around \$2,500 according to the survey question. Each is much smaller than the average rent of around \$6,800 and the estimated earnings losses of around \$15,000 in the year after a displacement.[28]

The regression coefficients plotted in Figure 2.11 suggests that housing transfers may spike around displacements, but that the spikes around displacement are economically small and statistically insignificant. There is no evidence that there are larger increases in housing

[27] In situations where the home is owned, we convert this housing value into a rental value, using the rough conversion factor of 0.0785 (also used in Albouy and Zabek, 2016). We provide a more detailed description of the procedure in Appendix B.6.

[28] This difference between the value according to the two estimates could be for several reasons. The most obvious is because the counterfactual is different between the question and our exercise. The counterfactual in the question is what would be the rent if the respondent's current dwelling were rented, while our question is how much it would cost for the respondent and their family to find similar accommodation. To the extent that dwellings are shared, and the market does not value living with one's parents as much as an OECD scale would suggest, these two estimates should diverge in the direction that we find. It also is possible that children are prone to under-estimating the amount of free rent that they receive.

transfers among people who lived closer to their parents before they lost their jobs. In fact, there is some suggestion that people who live closer see decreases in housing transfers after a displacement.[29] The point estimates in Panel A show that households are around four percent more likely to receive all of their rent as a gift in periods around displacements. A four percent increase is quite large relative to the two percent likelihood in the baseline sample, but it is a small slice of the overall population. Panel B shows the dollar values of the transfers involved; it also suggests that the transfers are modest at best. The implied rent savings estimates are noisy, but we can reject that there is an increase of $500 or more per year coming from a PSID parent. This is at least an order of magnitude smaller than the earnings losses after a displacement.[30]

2.5.2 Employment in Parents' Industry

Young adults living close to their parents may have more productive job search experiences and healthier earnings post displacement as a result of family networks, as documented in Kramarz and Skans (2014). The basic idea is that, after job loss, parents may be able to assist their adult children by tapping into their own employment networks to help their adult children to gain employment in their own industry at jobs with favorable wages and wage growth.[31] We can look for direct evidence for this mechanism using PSID data because we have industry codes for all workers, including parents and their adult children.

Table 2.4 presents some summary statistics on the probability of workers' working in their parents' industries. The table reports the fraction of workers employed in the same one-digit industry as their parents. We focus on employed individuals because when unemployed, the industry of an individual is the industry they were last employed in.

The table suggests that, on average, young workers living in the same neighborhood as their parents are slightly less likely to be working in their parents' industry than those living farther away. Workers in both groups have around a 25 percent probability of working in their parents' industry.

[29] Results for reported transfers of money, presented in Appendix B.6 (Appendix Figure B16), suggest that children who lived farther from their parents received larger cash transfers after a displacement, and that children who lived closer received no such transfers.

[30] Kaplan (2012) argues that the option value of moving in with parents is important, which could mean that the option value to workers was more valuable than the dollar value of the realized transfers. Another way that this exercise may be understating the value of in-kind transfers of housing from parents is by missing frequent movements of people in and out of their parent's homes, and the additional value coming from this flexibility. Our measures are only based on where people live at the time of the survey each year.

[31] There is some debate about how referral networks affect workers' wages. For example, Dustmann et al. (2017) find that hires from employment networks raise wages, while Bentolila et al. (2010) suggest that networks may reduce wages because they might assign workers to jobs in which they do not have comparative advantage. Alesina et al. (2015) (p.599) find that "individuals who inherit stronger family ties are less mobile, have lower wages and higher unemployment..."

We estimate a specification that is similar to our baseline regression (equation 2.1), where the outcome variable is an indicator of whether the young adult child is employed in the same one-digit industry as their parent. To control for local industry composition, we include employment industry shares at the county level from County Business Patterns (CBP) data.

Figure 2.12 shows the probability of working in a parent's industry rises in the years after displacement and this increase occurs only for workers living in the same neighborhood as their parents, although the results are noisy.[32] The effect is relatively large with a 10pp increase in working in a parent's industry on a base of around 25 percent (Table 2.4), although the estimates are noisy and the differences are not statistically significant from each other. Our results imply that for those who are displaced and living in the same neighborhood as their parents the probability of working in the same industry as their parents is elevated several years before the displacement. Appendix Figure B17 shows that older workers living close to their parents do not tend to move into their parents' industries after a displacement. Appendix Figure B18 shows that workers living in the same commuting zone as their parents do not experience a significant increase in their probability of working in their parent's industry after a displacement. These results are consistent with our findings that older workers and workers living farther away from their parents do not experience a post-displacement earnings benefit from parental proximity. We think our results suggest that parental employment networks may be operative for young displaced workers when they live in the same neighborhood as their parents.

2.6 Discussion

In this section, we use a simple model to show that our results are consistent with some parents improving their children's wage-offer distributions and that selection based on an unobserved preference for home would actually lead to the opposite results. We also investigate selection on ability and conclude that, on average, workers who live closer to their parents tend to be less skilled, and so this is unlikely to drive our results. Back-of-the-envelope calculations suggest that this improved wage-offer distribution is worth around $1,000 per year for a risk-neutral agent at an average risk of displacement, and this labor market advantage is one reason why children would prefer to live close to their parents.

To see the implications of heterogeneity in wage-offer distributions and preferences for home, first consider a simple economy where all workers are ex-ante homogeneous. Workers draw wages from two locations, home and away and these wage distributions are identical.

[32]We find similar increases when we use three-digit industries, but at longer horizons. We see similar results with two-digit industries as the ones presented here. We do not find any meaningful movement in the probability of working in the same occupation as parents around a displacement.

Suppose further that there are no moving costs, but that living at home is associated with positive utility payoff, b.

With homogeneous workers, people who move away are paid more because they need to be willing to forgo the utility payoff in their homes. This is one of the reasons why we pursue the propensity score reweighting exercise: even in an environment with ex-ante identical agents, selection ("luck") means that the earnings losses of those living farther away from their parents may be larger because they had higher pre-displacement earnings. Our reweighting exercise removes this selection effect because it only uses workers living farther away from their parents who have similar pre-displacement wages to those living close to their parents.

Once one controls for differences in workers' initial jobs, however, the post-displacement earnings of homogeneous workers will be identical. In order to match our finding of different post-displacement earnings outcomes, we consider three types of worker heterogeneity.

First, suppose that workers differ in their preference for living at home. In particular, suppose that some workers ("homebodies") prefer to live close to their parents and receive payoff b, while others ("explorers") have no preference for living close to parents. Notice that, on average, explorers will have higher wages because they receive no utility from being close to their parents and therefore simply seek the highest wage. Moreover, in equilibrium, those observed away from home are more likely to be explorers than homebodies. As before, the reweighting scheme will address pre-displacement selection on wages, but since workers away are more likely to be explorers they will, on average, have better wage outcomes after the displacement. As such, this sort of heterogeneity works against our empirical findings where, after a displacement, those close to their parents tend to have better earnings outcomes than those farther away.

Second, suppose that all workers receive utility b when living near parents, but workers differ in the wage-offer distribution they face. In particular, away from home, homebodies face a wage-offer distribution with mean μ, but at home the mean is $\mu + w_0$, where $w_0 > 0$. Explorers do not have this advantage and face the same distribution at home and away, with mean μ. Notice that, in equilibrium, a worker who is away is more likely to be an explorer because homebodies have a stronger preference for home as a result of the better wage-offer distribution. Also notice that homebodies will, on average, have higher wages due to the wage shifter, w_0. However, note that the expected wage of those at home could be below the expected wage of those away due to the selection on b. As before, the reweighting scheme will address pre-displacement selection on wages, but since workers away are more likely to be explorers they will, on average, have worse wage outcomes after the displacement. Therefore, our main empirical finding can be explained by differences in the wage-offer distribution.[33]

Third, suppose that all workers receive utility b when living near parents, but workers differ in their unobserved ability and that the return to this ability can be earned in both locations, home and away. In this simple framework, when we compare workers in different locations who have the same wage, as we do with our reweighting approach, the workers living at home will, on average, have higher ability. In principle, this selection on worker ability could be driving our empirical results; however, previous literature finds that this simple intuition is not supported by the evidence. In fact, Topel (1986), Bound and Holzer (2000), and Notowidigdo (2013) all find that low-skilled workers are less mobile in response to adverse labor demand conditions. As such, we think that geographic selection on unobserved ability is unlikely to explain our main findings.

To calculate the value of an improved post-displacement wage-offer distribution in our second though experiment, we take an estimate of the earnings differences after displacement, and we modify it to represent an expected value for a worker that has an average lifetime risk of being displaced. We begin with the differences in post-displacement earnings from Section 2.4, and we discount them by an annual interest rate of four percent. This simple calculation implies that the lifetime benefit (over a career lasting 35 more years) of living close to parents, conditional on a displacement, is around \$100,000. In our sample, the probability that a young worker experiences displacement is around 20 percent, so the expected total benefit of living close to home is around \$20,000 for someone at average risk of displacement. This suggests that the benefit of parental proximity after job displacement is associated with an annual value of around \$1,000. Similarly, if we perform the same calculations with our baseline estimates from Section 2.3, we obtain a value of around \$2,500. Note that this calculation only assumes the benefits of being close to home will apply after a job displacement; if workers received similar benefits after less severe labor market disruptions, our estimates would be a lower bound on the wage benefits of being close to home.

2.7 Conclusion

Young adults who live in the same neighborhoods as their parents experience stronger earnings recoveries after job displacements than those who live farther away. This result persists after we apply a reweighting scheme that controls for observable differences between the two groups. Our results are consistent with a simple theory where some parents can facilitate higher wage job offers for children who live nearby. We find some weak evidence

[33] Notice that with a positive moving cost, $c > 0$, even if *all* workers faced a better wage-offer distribution at home, we would get the desired result. This is because people who moved away have to pay the cost, c, to move back home and, as a result, they will on average have worse post-displacement wage outcomes.

that parents find jobs in their industries for their children. Young adults with children of their own drive the stronger earnings recoveries among workers closer to their parents, so grandparents' help with childcare might also be important.

Longer job searches and children's ability to move in with parents who live nearby do not appear to explain the result. Unemployment durations and changes in hours worked after displacement are similar regardless of how far children live from their parents. Children do tend to move in with their parents after displacements, but the transfers involved are relatively small and involve relatively few workers. Children who lived closer to their parents before they were displaced are not any more likely to move in with their parents.

Our results suggest that one reason workers live close to their parents might be because this helps them after a negative labor market shock. This phenomenon can explain why people appear so reluctant to move from declining areas (Ganong and Shoag, 2012; Zabek, 2017) and why migration responses have been smaller than expected after several local shocks (Bound and Holzer, 2000; Yagan, 2017). This reluctance to relocate could explain why workers appear to be less mobile in economic downturns (Molloy et al., 2011) and why immigrants appear to be more mobile than natives (Cadena and Kovak, 2016). Our results could also inform research on the decline in inter-state migration (Molloy et al., 2011; Kaplan and Schulhofer-Wohl, 2017), the recent increase in young adults living with their parents (United States Census Bureau, 2016), and the recent rise in leisure among younger men (Aguiar et al., 2017). More directly, parental resources may be an important explanation for the finding that workers place a large premium on living close to their places of origin (Kennan and Walker, 2011; Coate, 2017). Based on our empirical findings, simple calculations suggest that parental proximity after job displacement is associated with an annual value of around $1,000.

An important implication of our finding that young adults experience post-displacement labor market benefits from living close to their parents is that any government program that substitutes for parental proximity will increase children's mobility. Even if a program perfectly crowds out parents' efforts, this program would remove the incentive for workers to locate close to their parents. Without these incentives, younger workers would move to higher-wage jobs, and these movements would increase total output in equilibrium.

Going forward, we hope researchers use other data sources to verify our baseline results on parental proximity and post-displacement earnings losses. Ideally, these new data would also facilitate additional analyses that focus on the mechanisms leading to our baseline results. We think that building and estimating a model that incorporates parental location (Kennan and Walker, 2011; Coate, 2017) and matches well the earnings losses of displaced workers (Jarosch, 2015; Krolikowski, 2017b) is a particularly fruitful way to proceed.

Figure 2.1: Average Earnings for Young Displaced Workers by Proximity to Parents

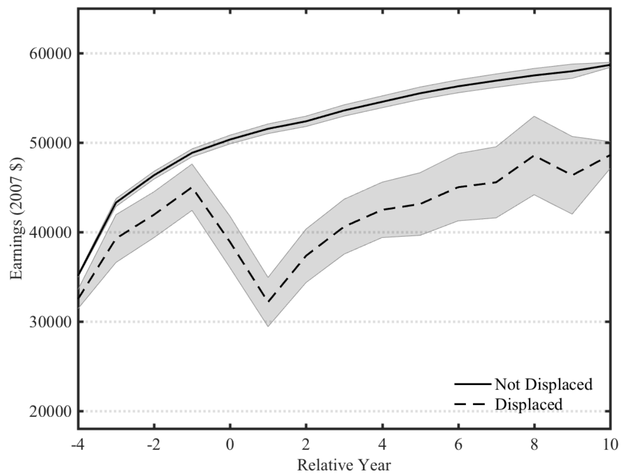

Panel A: Average Earnings for Young Displaced and Non-Displaced Workers

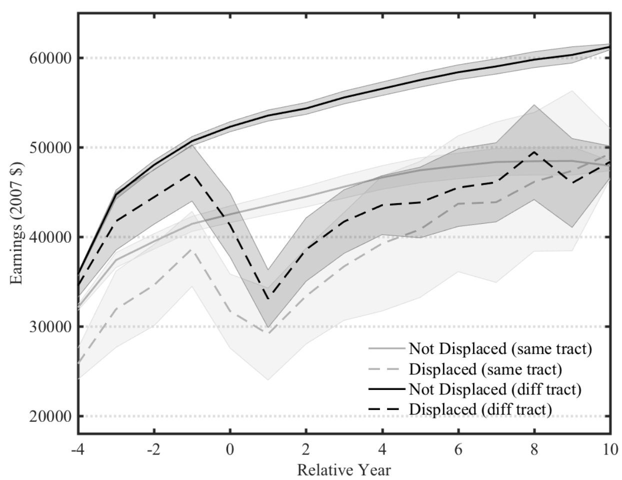

Panel B: Average Earnings for Those in Their Parents' Neighborhoods and Not

Note: Young workers who live in their parents' neighborhoods experience stronger earnings recoveries after a displacement than young workers who are not living in their parents' neighborhoods. These figures plot average earnings for displaced and not displaced young workers (aged 25 to 35 in year zero). The shading represents 95 percent confidence intervals, computed by clustering standard errors at the worker level. All of the workers were employed in a job for at least two years. Workers were displaced if they reported that they were no longer in that job because the plant closed, because they were laid off, or because they were fired. The subgroups in Panel B are defined based on how close they lived to their parents two years before they were at risk of a displacement. The same tract group lives in the same census tract as their parents, while the different tract group lives in a different census tract. See Section 2.2 for more information on the sample construction, data, and definitions.

Figure 2.2: Earnings Losses for Young Displaced Workers

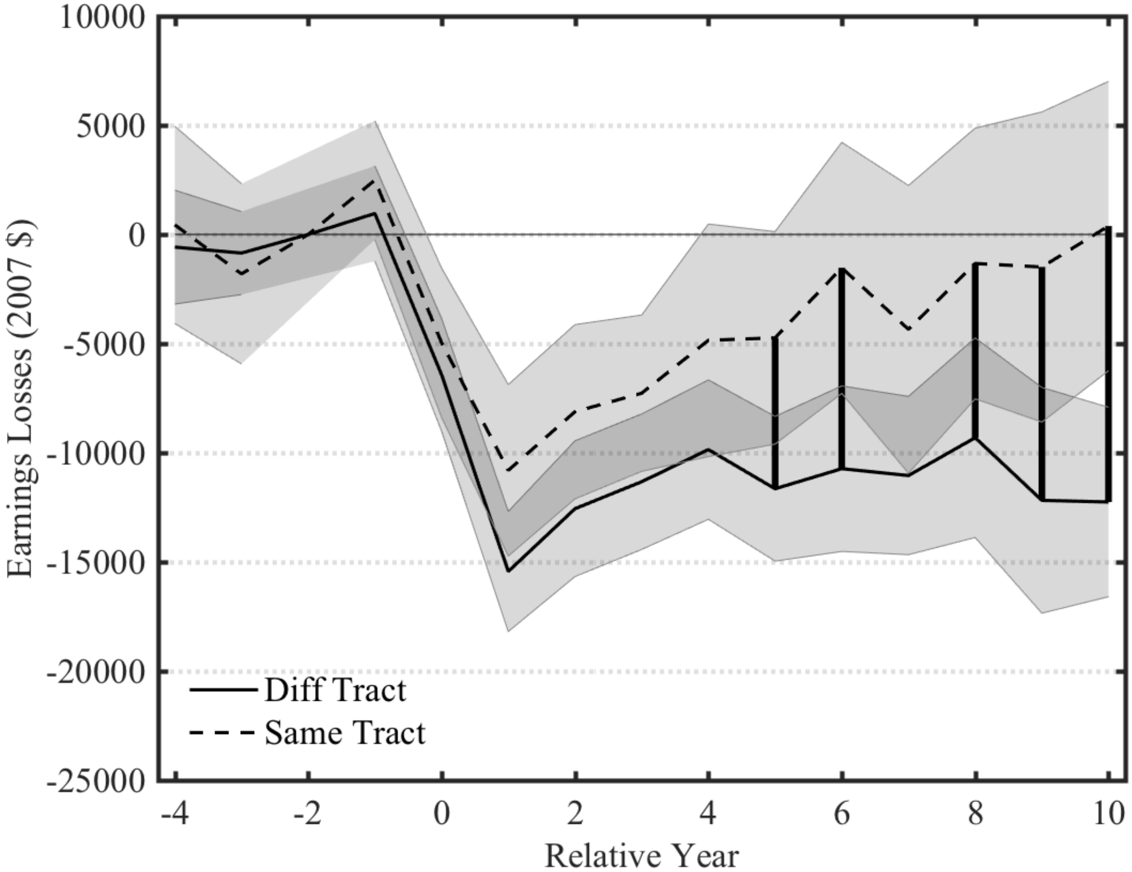

Note: In the medium- and long-run, young workers living in the same neighborhoods with their parents experience a full recovery in earnings after a displacement. Young workers not living in their parents' neighborhoods experience large and permanent earnings losses, amounting to around 30 percent of their pre-displacement earnings even 10 years after the displacement. The figure plots regression coefficients from equation (2.1) describing the impact of a job displacement on the earnings of groups of young workers, aged 25 to 35 at the time of displacement. The shading represents 95 percent confidence intervals and any vertical bars represent statistically significant differences at the five percent level. Standard errors are clustered at the worker level. The definitions of displacements and of whether workers live in the same tract as their parents follow those in Figure 2.1, and Section 2.2 contains more information on the sample construction, data, and definitions. The regression controls for worker and year fixed effects as well as a quartic term in each workers' age that we allow to differ between the different groups plotted on the figure.

Figure 2.3: Positive Hours, Hours Worked, and Wages for Young Displaced Workers

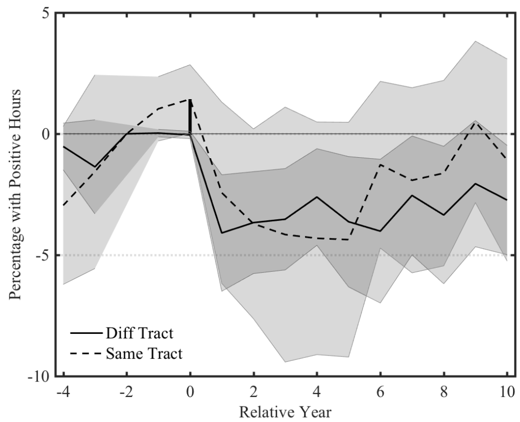

Panel A: Indicator for Positive Annual Hours

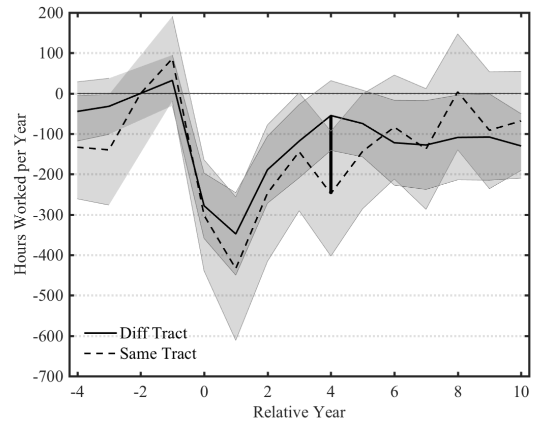

Panel B: Hours Worked

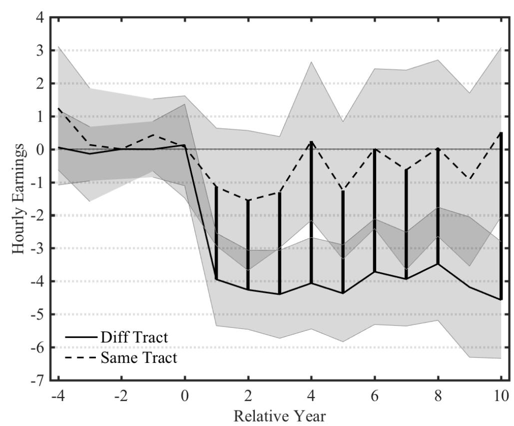

Panel C: Hourly Earnings

Note: The intensive margin and wages drive the recovery in annual earnings documented in Figure 2.2. At the time of displacement, wages fall less for workers living near their parents and hours fall slightly more, although the hours differences are not statistically significant. These figures plot regression coefficients from equation (2.1) describing the impact of a job displacement on measures of labor supply and of wages for groups of young workers, aged 25 to 35 at the time of displacement. The shading represents 95 percent confidence intervals and any vertical bars represent statistically significant differences at the five percent level. Standard errors are clustered at the worker level. The definitions of displacements and of whether workers live in the same tract as their parents are the same as in Figure 2.1 and Section 2.2 contains more information on the sample construction, data, and definitions.

Figure 2.4: Weeks Spent Unemployed

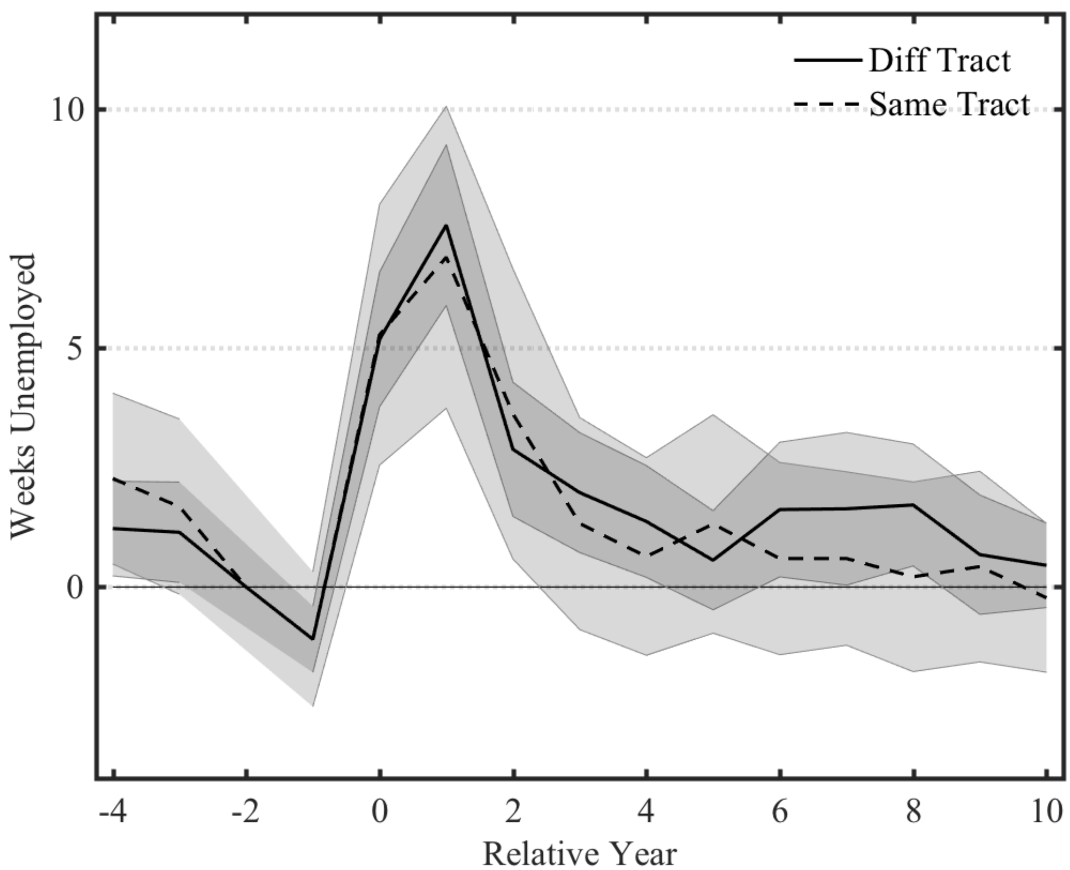

Note: Younger workers who live in their parents' neighborhoods experience very similar unemployment durations around a displacement to workers who live farther away. Both groups see an increase of around seven weeks on-impact and a steady decline over the next 10 years. The figure plots regression coefficients from equation (2.1) describing the impact of a job displacement on the number of weeks that workers were unemployed in each year. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. In this figure, the differences are not statistically significant. Shading and statistical significance are based on standard errors computed by clustering at the worker level. The definitions of displacements and of whether workers live in the same tract as their parents follow those in Figure 2.1, and Section 2.2 contains more information on the sample construction, data, and definitions.

Figure 2.5: Earnings Losses for Older Displaced Workers

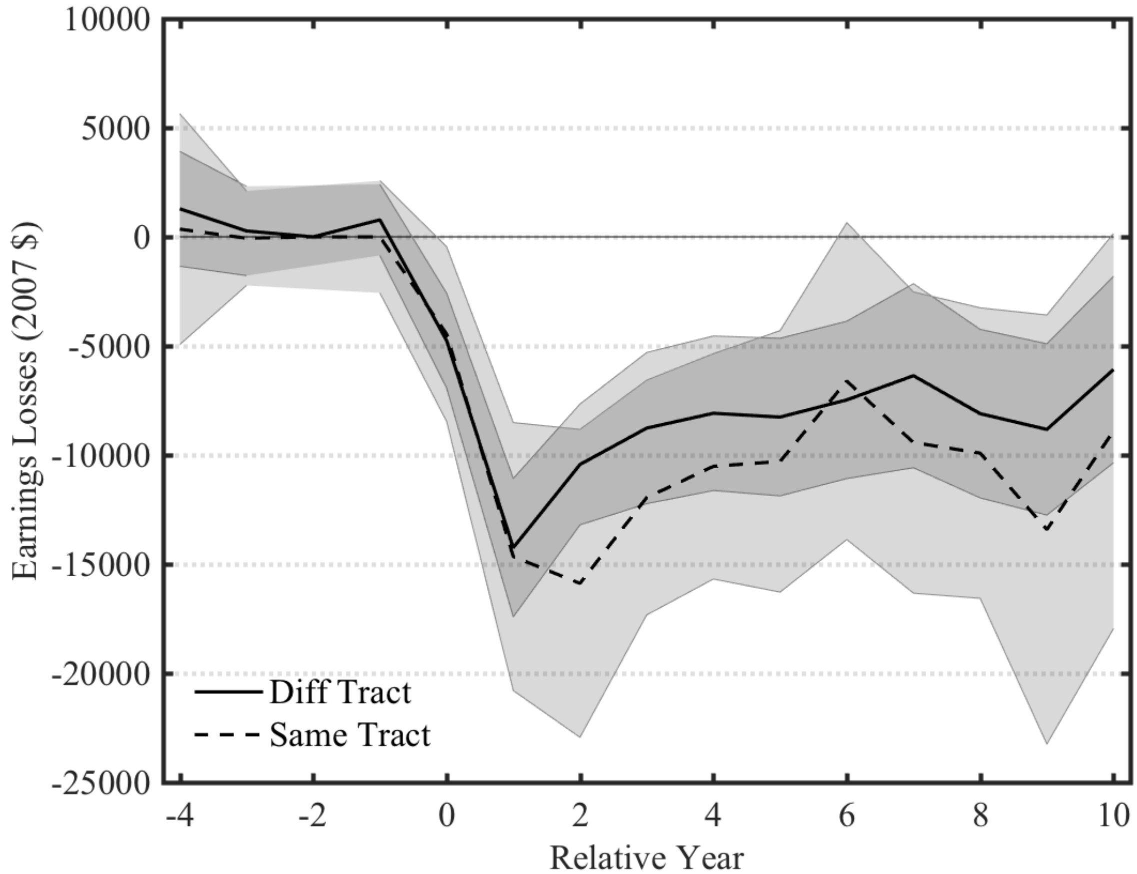

Note: Older workers (ages 36 to 55) who live in their parents' neighborhoods do not experience the same benefit of living close to their parents as young adults. If anything, living near parents prior to displacement has a detrimental effect, but these differences are not statistically significant. The figure plots regression coefficients from equation (2.1) describing the impact of a job displacement on the earnings of groups of older workers, aged 36 to 55 at the time of displacement. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. Shading and statistical significance are based on standard errors computed by clustering at the worker level. The definitions of displacements and of whether workers live in the same tract as their parents follow those in Figure 2.1, and Section 2.2 contains more information on the sample construction, data, and definitions.

Figure 2.6: Earnings Losses for Young Workers by Different Proximities to Parents

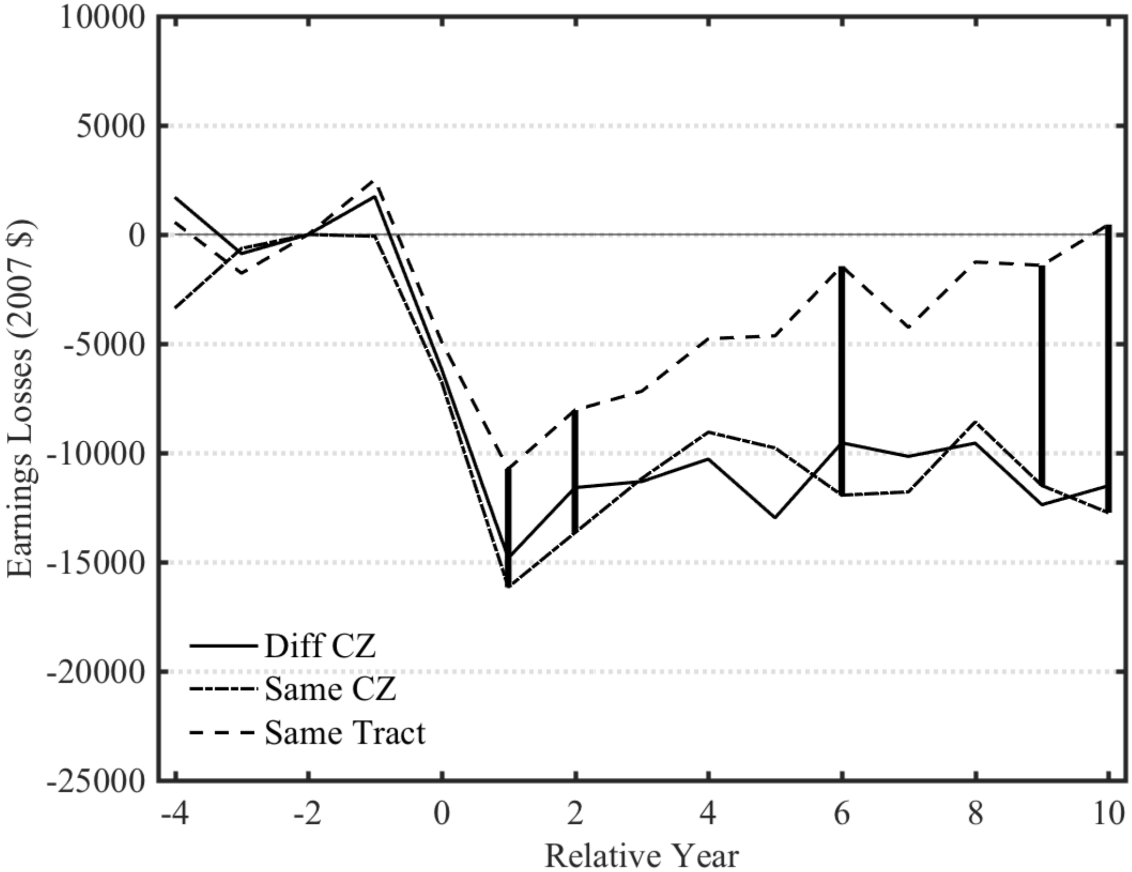

Note: Those workers living close to their parents (in the same commuting zone), but not in the same neighborhoods, do not experience significantly better post-displacement earnings outcomes than those who live farther away. The figure plots regression coefficients from a specification similar to equation (2.1) describing the impact of a job displacement on the earnings of three mutually exclusive groups of young workers. The three groups are defined as workers who lived in the same census tract as their parents, those who lived in the same commuting zone but not the same tract, and those who lived in a different commuting zone as their parents. Each group is defined based on their location two years before the displacement. The figure includes vertical bars that connect the line for workers who live in the same tract with the line for workers who live in the same commuting zone, but not the same tract. We include these when the estimates are statistically significantly different from one another at the five percent level. Statistical significance is based on standard errors computed by clustering at the worker level. The definitions of displacements follow Figure 2.1, Section 2.3.3 describes the specification in more depth, and Section 2.2 contains more information on the sample construction, data, and definitions.

Figure 2.7: Earnings Losses for Young Workers Who Do Not Move

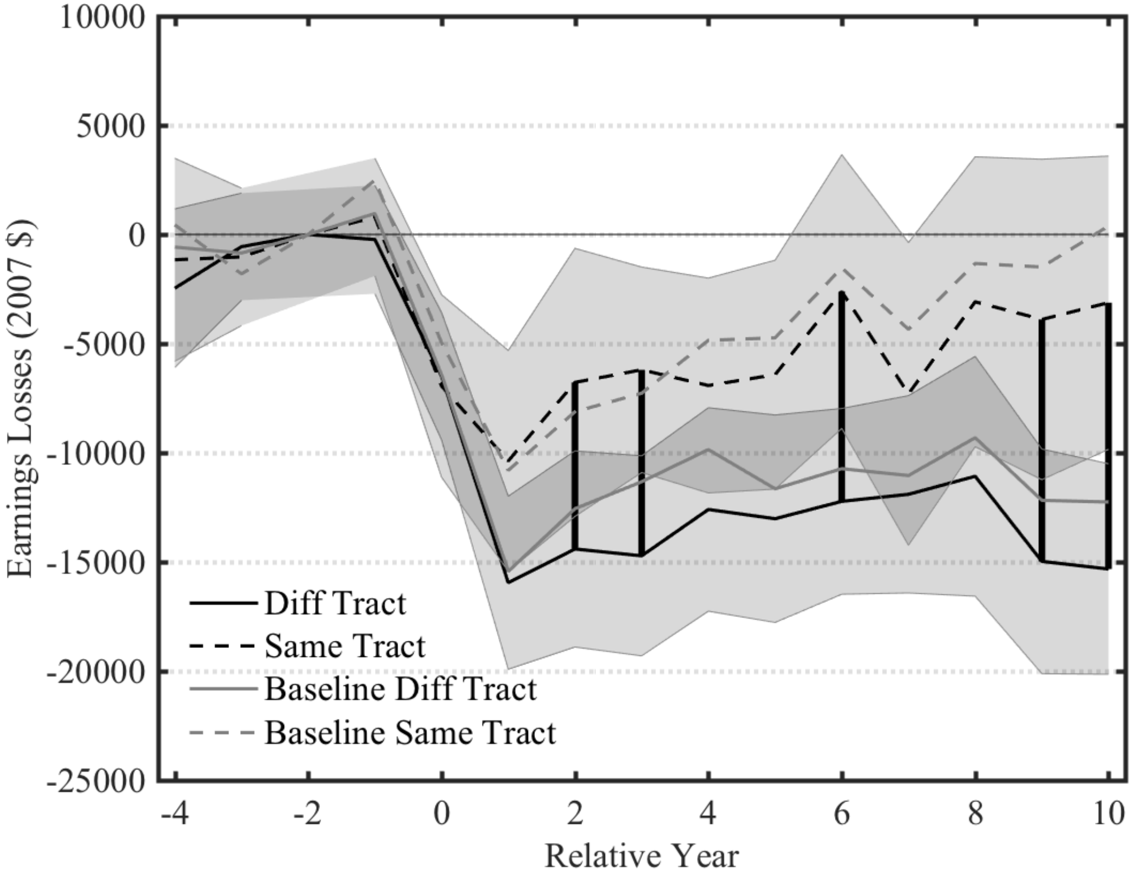

Note: Restricting the sample to workers who do not switch counties after the displacement does not affect the baseline result. Perhaps not surprisingly, the sample almost always has worse earnings outcomes (point estimates) than the unrestricted sample. The figure plots regression coefficients from equation (2.1) describing the impact of a job displacement on the earnings of young workers who do not move between counties after a job displacement. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. We only report the vertical bars for significant differences between the regression results with no mobility, since Figure 2.2 reports them for the whole sample. Shading and statistical significance are based on standard errors computed by clustering at the worker level. The definitions of displacements follow Figure 2.1, and Section 2.2 contains more information on the sample construction, data, and definitions.

Figure 2.8: Earnings Losses for Young Workers With and Without Children

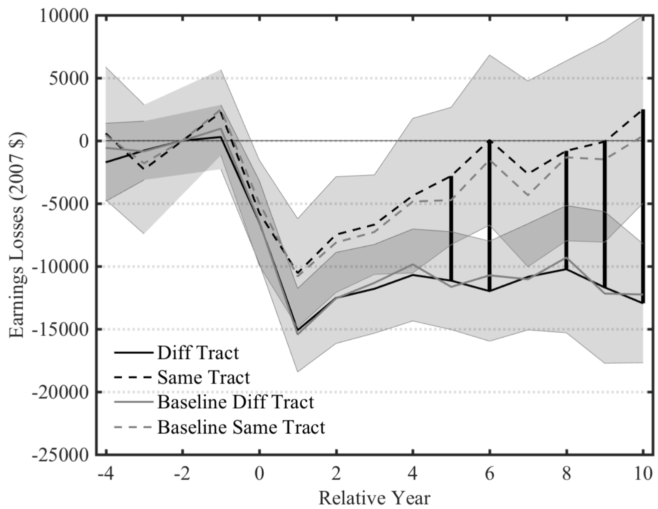

Panel A: Young Workers With Children

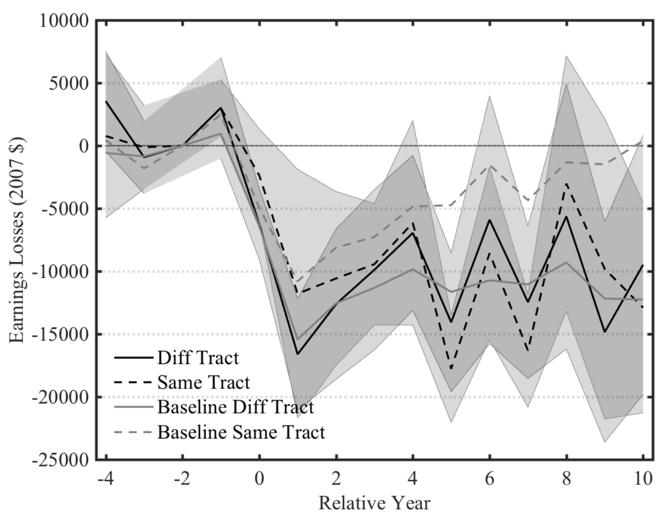

Panel B: Young Workers Without Children

Note: The earnings of young workers who have children and live close to their parents recover after a displacement, while those of workers who live farther away, or who do not have children, are permanently lower. The figure plots regression coefficients from a specification similar to equation (2.1) describing the impact of a job displacement on earnings. Panel A presents the coefficients that apply to workers with children, and Panel B to those without. All of the coefficients come from the same regression specification. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. Shading and statistical significance are based on standard errors computed by clustering at the worker level. The definitions of displacements follow Figure 2.1, and Section 2.2 contains more information on the sample construction, data, and definitions.

Figure 2.9: Means After Propensity Score Reweighting

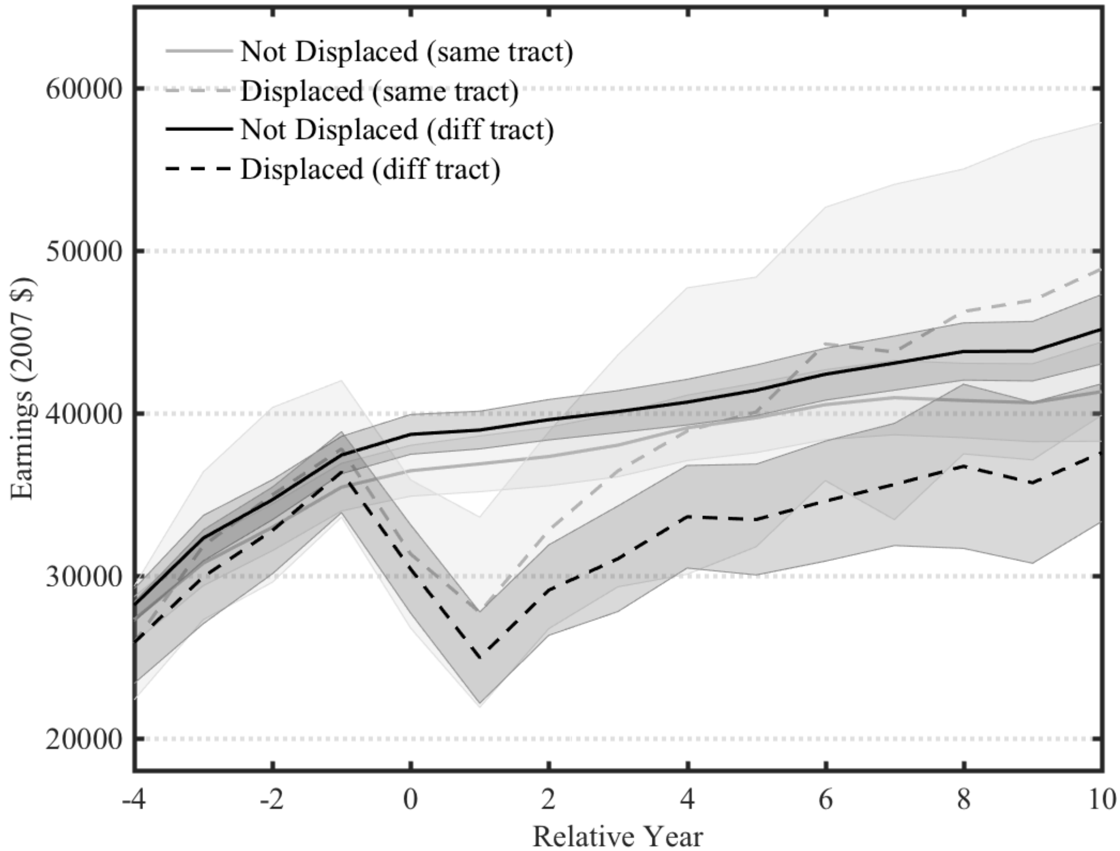

Note: Simple averages, after applying propensity score weights, still suggest that workers living close to their parents have significantly better post-displacement earnings outcomes than those who live farther away. This figures plot propensity score weighted average earnings for displaced and not displaced young workers. The shading represents 95 percent confidence intervals, computed by clustering standard errors at the worker level. We designed the weights to make each other group of workers comparable to the group of workers who live in the same census tract as their parents two years before they experience a job displacement. We include characteristics of workers' jobs, of workers' levels of education, of employment-to-population ratios where workers live, and of workers' demographics, including whether they have children or not. The definitions of displacements follow Figure 2.1, Section 2.4 describes the reweighting, and Section 2.2 contains more information on the sample construction, data, and definitions.

Figure 2.10: Earnings Losses After Propensity Score Reweighting

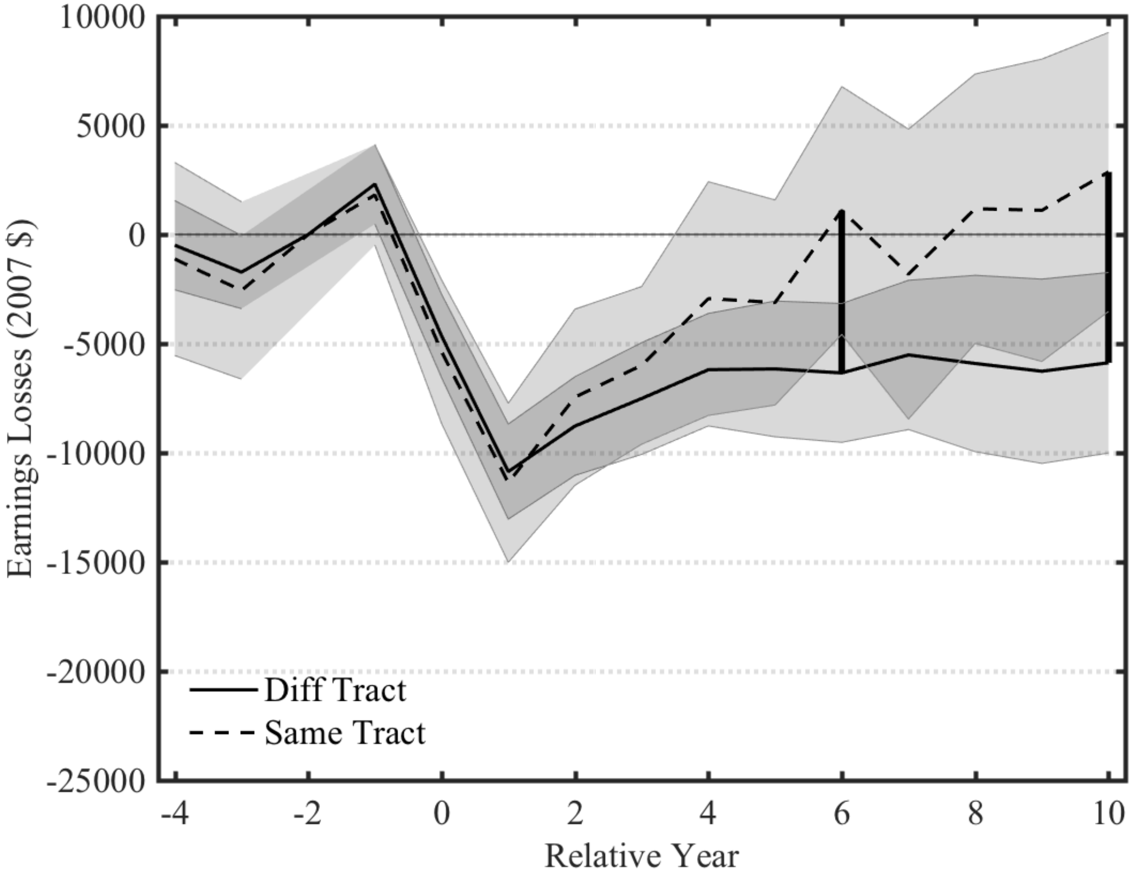

Note: Even after controlling for observable differences using propensity score reweighting, young workers living in their parents' neighborhoods at the time of displacement experience healthier earnings recoveries than those living farther away. Although this difference is quantitatively smaller than in Figure 2.2, it is still statistically significant at longer horizons. The figure plots propensity score weighted regression coefficients from equation (2.1) describing the impact of a job displacement on the earnings of young workers. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. Shading and statistical significance are based on standard errors computed by clustering at the worker level. We designed the weights to make each other group of workers comparable to the group of workers who live in the same census tract as their parents two years before they experience a job displacement. We include characteristics of workers' jobs, of workers' levels of education, of employment-to-population ratios where workers live, and of workers' demographics, including whether they have children or not. The definitions of displacements follow Figure 2.1, Section 2.4 describes the reweighting, and Section 2.2 contains more information on the sample construction, data, and definitions.

Figure 2.11: Housing Transfers Around Displacements

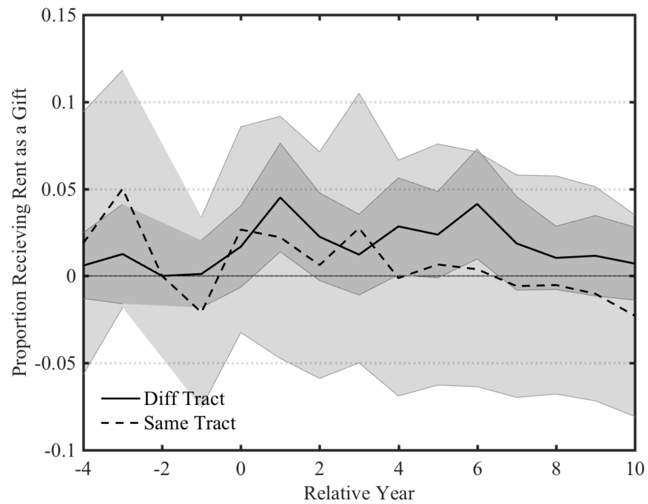

Panel A: Receiving Rent Entirely as a Gift

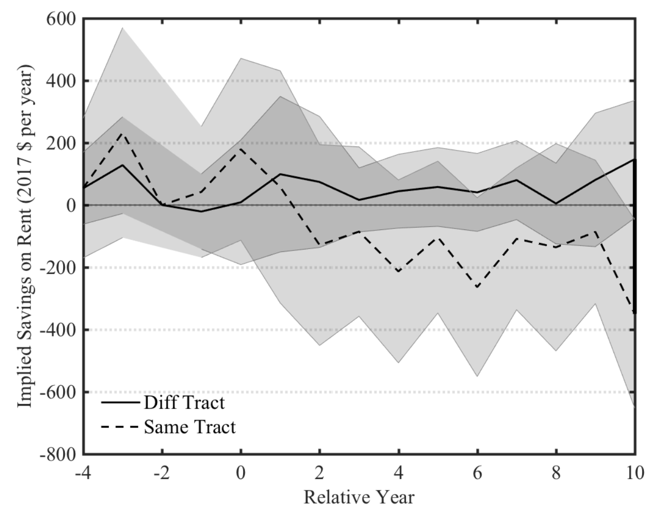

Panel B: Implied Rent Savings

Note: There is a small, but detectable, increase in housing transfers around displacement, primarily among workers lived outside their parents' neighborhoods before a displacement. These figures plot regression coefficients from equation (2.1) describing the impact of a job displacement two measures of in-kind transfers of housing to young workers. The measure in Panel A is the proportion of workers who report that they pay no rent and who then volunteer that this is because someone provided it as a gift to them. The measure in Panel B is the estimated dollar value that the worker's family unit saves in rent, based on living with another family unit. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. Standard errors are clustered at the worker level. Each measure is explained in more depth in Section 2.5.1, the definitions of displacements follow Figure 2.1, and Section 2.2 contains more information on the sample construction, data, and definitions.

Figure 2.12: Working in a Parent's Industry

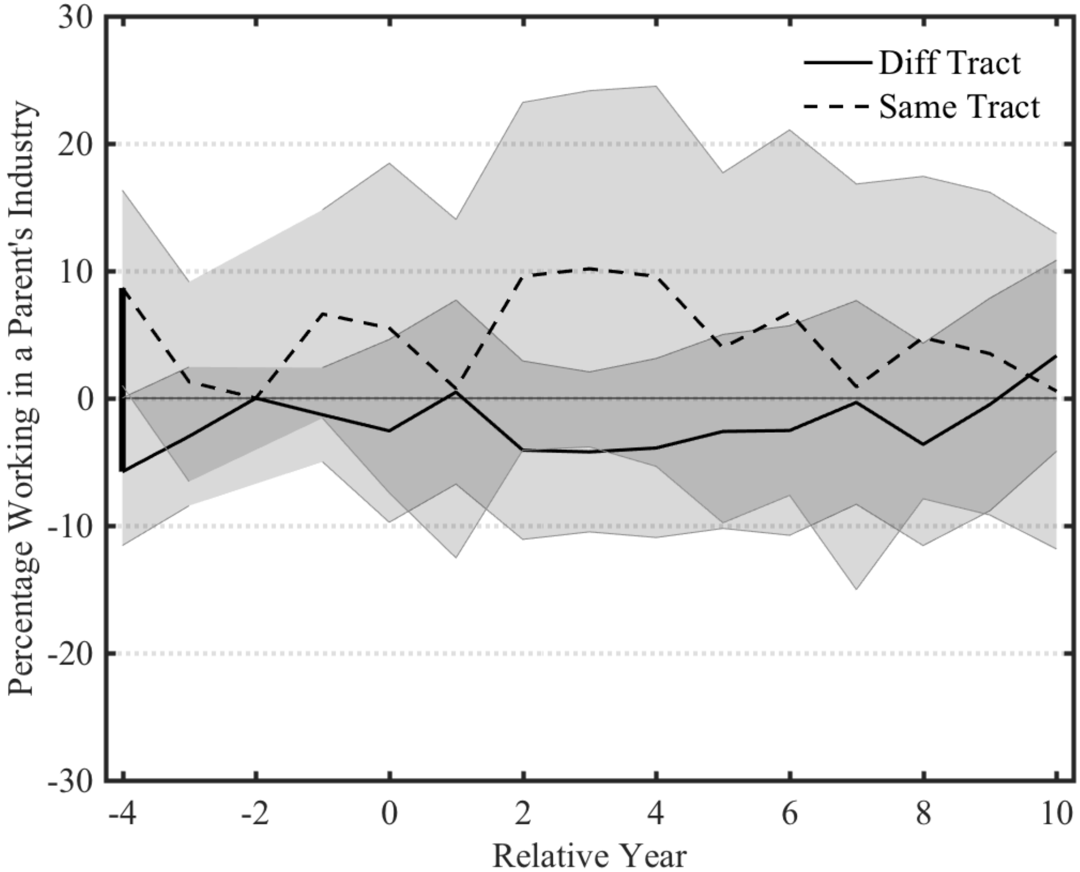

Note: Workers appear to be slightly more likely to work in their parents' industries around a displacement, though the effect can vary depending on how industries are measured. It does not appear to be due to local industrial composition, however. These figures plot regression coefficients from equation (2.1) describing the impact of a job displacement on the proportion of workers who work in the same one-digit PSID coded industry as a parent. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. Shading and statistical significance are based on standard errors computed by clustering at the worker level. The definitions of displacements follow Figure 2.1, and Section 2.2 contains more information on the sample construction, data, and definitions.

Table 2.1: Summary Statistics

Variable			Same Tract		Different Tract	
	Displaced	Not Displaced	Displaced	Not Displaced	Displaced	Not Displaced
	Panel A: All Workers Age 25 to 55					
Earnings	48,300	56,000	41,600	48,400	49,900	57,400
Age	36.4	37.9	34.0	35.5	37.0	38.3
Years of Schooling	12.9	13.6	12.5	13.1	13.0	13.6
Employer Tenure	7.5	10.6	7.1	9.8	7.5	10.7
Fraction Co-residing	0.10	0.08	0.52	0.52	0.00	0.00
Fraction in Parents' Tract	0.19	0.15	1.00	1.00	0.00	0.00
Number of Children	1.24	1.24	1.37	1.25	1.21	1.23
Fraction Male	0.84	0.83	0.84	0.78	0.84	0.84
Wages ($/hr)	21.7	24.7	18.6	21.6	22.5	25.3
Hours Worked	2,260	2,310	2,280	2,270	2,250	2,310
# of records	1,464	44,732	322	7,548	1,142	37,184
	Panel B: Young Workers Age 25 to 35					
Earnings	45,100	49,000	39,000	41,600	47,100	50,800
Age	28.9	29.4	28.4	29.2	29.1	29.4
Years of Schooling	13.1	13.8	12.5	13.1	13.3	13.9
Employer Tenure	5.3	6.5	5.5	6.7	5.2	6.4
Fraction Co-residing	0.10	0.08	0.42	0.39	0.00	0.00
Fraction in Parents' Tract	0.25	0.20	1.00	1.00	0.00	0.00
Number of Children	1.14	1.11	1.22	1.20	1.11	1.09
Fraction Male	0.84	0.82	0.83	0.79	0.85	0.83
Wages ($/hr)	20.3	21.6	17.6	18.6	21.2	22.3
Hours Worked	2,240	2,310	2,250	2,280	2,240	2,320
# of records	707	17,996	204	4,005	503	13,991

Note: Workers who were displaced and workers who live in the same tract as their parents are younger, tend to earn less, and thend to have lower soceoeconoic status than workers who were not displaced and who lived furhter from thier parents. This table presents (individual PSID) weighted averages using unbalanced data from the 1968-2013 PSID surveys. Dollar figures are in 2007 dollars using the CPI-U-X1. All variables are measured in the year before a potential dispalcement (relative year -1). The PSID sample of household heads is composed chiefly of men. We restrict to observations that appear in our baseline sample (equation 2.1). The definitions of displacements follow Figure 2.1 and Section 2.2 contains more information on the sample construction, data, and definitions.

Table 2.2: Means Before and After Reweighting

Variable	Same Tract		Different Tract		Same Tract		Different Tract	
	Displaced	Not Displaced	Displaced	Not Displaced	Displaced	Not Displaced	Displaced	Not Displaced
	Panel A: PSID Weights				Panel B: Reweighted			
Earnings	\$35,000	\$39,700	\$44,700	\$48,300	\$35,000	\$33,100	\$33,100	\$34,800
	[1.00]	[0.03]	[0.00]	[0.00]	[1.00]	[0.36]	[0.43]	[0.93]
Average Change in Earnings	\$2,900	\$2,300	\$2,900	\$3,500	\$2,900	\$2,100	\$3,300	\$2,500
	[1.00]	[0.36]	[0.97]	[0.39]	[1.00]	[0.25]	[0.64]	[0.58]
Years of Schooling	12.50	13.10	13.31	13.94	12.50	12.46	12.37	12.52
	[1.00]	[0.00]	[0.00]	[0.00]	[1.00]	[0.84]	[0.61]	[0.93]
Share in Goods Industries	0.53	0.43	0.51	0.34	0.53	0.48	0.51	0.40
	[1.00]	[0.06]	[0.81]	[0.00]	[1.00]	[0.33]	[0.86]	[0.02]
Share Manager/Professional	0.20	0.27	0.31	0.41	0.20	0.20	0.17	0.20
	[1.00]	[0.11]	[0.03]	[0.00]	[1.00]	[0.92]	[0.54]	[0.91]
Employer Tenure	5.32	6.79	5.16	6.42	5.32	5.39	5.20	5.33
	[1.00]	[0.00]	[0.65]	[0.00]	[1.00]	[0.81]	[0.76]	[0.98]
Unemp Rate in County	7.47	7.17	7.48	6.80	7.47	7.42	7.83	7.45
	[1.00]	[0.34]	[0.98]	[0.04]	[1.00]	[0.89]	[0.45]	[0.96]
Age	28.32	29.26	29.14	29.51	28.32	28.04	27.95	28.09
	[1.00]	[0.00]	[0.02]	[0.00]	[1.00]	[0.34]	[0.32]	[0.43]
Number of Children	1.26	1.21	1.11	1.11	1.26	1.16	1.23	1.18
	[1.00]	[0.72]	[0.32]	[0.26]	[1.00]	[0.42]	[0.83]	[0.56]
Fraction Male	0.82	0.81	0.86	0.85	0.82	0.81	0.77	0.83
	[1.00]	[0.82]	[0.34]	[0.39]	[1.00]	[0.89]	[0.36]	[0.73]
Number of Records	190	3,620	456	12,633	190	3,620	456	12,633

Note: After applying the propensity score weights, the sample of workers who live in the same tract as their parents and those living farther away are statistically indistinguishable in terms of many observable characteristics. This table reports means for each group using PSID weights in the first four columns and the propensity score weights in the last four columns. For each variable, we report the mean and a p-value in brackets of a Wald test that this mean is the same as the value in the first column. Standard errors and p-values adjust for clustering at the worker level. The number of records refers to the number of person by age records where we have sufficient earnings observations to include the record in the main sample, which we use for our main results. Missing data for some variables mean that some of the statistics on this table, most notably the local unemployment rates, are computed based on fewer records. The definitions of displacements follow Figure 2.1, Section 2.4 describes the reweighting, and Section 2.2 contains more information on the sample construction, data, and definitions.

Table 2.3: Measures of Housing Transfers

	Rent	Gifted Rent		Implied Savings	
	Dollar value	Proportion	Dollar value	Proportion	Dollar value
Value	\$6,800	0.02	\$2,500	0.08	\$4,300
Standard error	(90)	(0.002)	(170)	(0.005)	(430)
N	7,890	18,396	417	18,396	543

Note: Less than 10 percent of young workers recieve discounted housing, and the implied amount tends to be modest, and less than the average amount that families spend on rent when they live alone. The first column reports average annual rents for the baseline sample and the following columns report measures of housing transfers. The rows report means, standard errors of those means, and sample sizes. Gifted rent reports the proportion of households who report receiving all of their rent as a gift, and the annual value of that gift. Implied savings reports the proportion of people who live with a parent and the implied annual dollar value that they receive from that parent. Note that the dollar value is only observable when the parent is a PSID respondent themselves. We explain each measure of housing transfers in more depth in Section 2.5.1.

Table 2.4: Summary Statistics of Sharing Parent's Industry by Parental Proximity

	Same Tract	Diff Tract
$\mathbb{P}[in\ parent's\ industry]$	0.25	0.28
Standard error	(0.0018)	(0.0032)
N	89,157	32,371

Note: Simple averages suggest that employed young workers living in their parents' neighborhoods are slightly less likely to be working in their parents' industry. Results are based on large sectors but looking at finer levels of disaggregation does not alter the conclusions. This table reports the proportion of young workers who are employed in their parent's one-digit PSID industries by whether they currently live in the same census tract as their parents. The rows report means, standard errors of those means, and sample sizes.

CHAPTER III

Housing Inequality

with Aditya Aladangady and David Albouy

3.1 Introduction

In this paper, we examine a dimension of inequality that has received surprisingly scant attention – inequality in housing outcomes. We find that measures of inequality in housing prices and rents in the United States exhibit a U-shaped pattern over the last 85 years, resembling patterns of income and wealth inequality, often referred to as a "great compression" followed by a "great divergence" (Piketty and Saez (2003), Saez and Zucman (2014)). Housing-value inequality fell from 1930 to 1970 as home ownership expanded, but has subsequently risen. Rent inequality also fell, but it has risen only slightly since. Combining both measures into a rental equivalent, we again see a U-shape.

To understand the fall and subsequent rise of housing inequality, we use decomposition techniques to quantify the impacts of key variables. Changes in housing inequality have occurred primarily within cities and are not explained by *observable* changes in dwelling characteristics. Thus, changes in the desirability of particular neighborhoods, reflected in land values, appear to be the main contributor to changes in housing inequality, although we cannot rule out changes in unobserved housing quality. Using a series of simple regressions, we also find that local housing inequality is related to local income inequality at magnitudes implied by reasonable values for the income elasticity of housing consumption.

Knowledge of housing inequality sheds light on larger issues concerning consumption and wealth inequality. It informs debate over whether consumption inequality has grown as much as income inequality over the last few decades.[1] Housing accounts for a large share of consumption, and has done so more stably than other items, such as food or health care.[2] Housing may represent permanent income particularly well as it is durable (Friedman (1957), p 208). The great compression and great divergence in housing inequality are of roughly similar magnitude, commensurate with income changes. However, housing inequality has less to do with tangible dwelling characteristics — such as living space — and more to do with what people pay to live in different locations. These locations offer different "intangibles" such as access to employment and local amenities, such as schools, safety, and natural features.

There has also been recent debate over whether wealth inequality has increased as much as income inequality.[3] Housing informs this debate as it accounts for much of the overall capital stock and is the principal asset for most Americans with savings.[4] Our evidence that housing values have diverged indirectly supports the view that wealth inequality has increased, albeit to levels lower than before World War II, due to increases in home-ownership.[5]

To our knowledge, our paper is the first to document inequality in housing prices and rents over such a long period and to relate them to measures of income inequality over space and time. Our analysis of housing expenditures sheds light on consumption inequality prior to World War II – before widely available household-level consumption data – providing direct evidence of a great compression in consumption inequality contemporaneous with

[1]Krueger and Perri (2006) and Meyer and Sullivan (2010) argue that consumption inequality has increased much less than income inequality using the Consumer Expenditure Survey. Aguiar and Bils (2011) propose a correction for measurement error that results in a consumption inequality measure that mirrors income. Work using PSID consumption measures by Attanasio and Pistaferri (2014) and earlier evidence by Cutler and Katz (1992) also suggest that consumption inequality has increased in line with income inequality.

[2]As a fraction of Personal Consumption Expenditures (PCE), "Housing and utilities" rose from 16.6 to 18.1 percent between 1959 and 2014. Similar fractions for "Food and beverages purchases for off-premises consumption" are 19.4 to 7.5 percent; "Clothing and footwear," 8.0 to 3.1 percent, "Motor vehicles and parts," 5.9 to 3.7 percent; "Health care," 4.7 to 16.5 percent. The share of housing in aggregate expenditures is roughly one-sixth in PCE and one-third in the Consumer Expenditure Survey (CEX) (Albouy and Lue (2015).)

[3]Kopczuk (2004) find a great compression of wealth inequality but little divergence afterwards. Using capitalization methods, Saez and Zucman (2014) find greater divergence, although Kopczuk (2015) disputes this.

[4]Housing accounts for roughly one third of total household wealth and roughly 40 percent of the capital stock. Housing is almost two thirds of wealth for the middle three quintiles (Wolff, 2014). In 2004, 62 percent of housing wealth was held by the bottom 90 percent, and only 9.8 percent by the top 1 percent; for stocks and mutual funds, the comparable numbers are 14.6 and 44.8 percent (Wolff 2009, p. 160).

[5]Differences in wealth due to inequality in housing values have more complex implications than differences due to other types of capital. An appreciation in home prices will only increase a home-owner's permanent wealth if they have a less expensive alternative to living in their now more expensive house. Rognlie (2014) and Bonnet et al. (2014) have emphasized this point in critiquing Piketty and Zucman (2014). One interesting feature of our paper is our finding that home owners have alternatives that are observably quite similar but much less expensive. This suggests that the distinction is less important.

similar changes in income and wealth.[6] While several studies (e.g. Van Nieuwerburgh and Weill (2010), Moretti (2013), and Gyourko et al. (2013)) have examined recent changes in housing-price inequality across cities, our analysis extends for longer, covers rents, and examines dwelling characteristics and variation in prices within cities.

3.2 Data, Inequality Measures, and Empirical Techniques

3.2.1 Housing Data and Sample Selection

Our data are drawn from the full 1930 and 1940 Census, the long form of the 1960 through 2000 Censuses, and the American Community Survey (ACS) for 2009-12.[7] These surveys ask owners to report the current value of their home and renters to report their monthly rent. Renters from 1970 onward and home owners from 1980 onward also report their utility costs. Home values are recorded as intervals from 1960 through 2000 and monthly rents are reported as intervals for 1960, 1980, and 1990. In all other years the questions record a continuous measure with relatively little top coding.[8]

We restrict our sample to focus on residential homes and to maximize consistency across years. This involves three restrictions. First, we restrict the sample to the continental United States. Second, we eliminate all farms and homes used for commercial activities, such as dental and medical offices. Third, we remove owner-occupied units in multi-family structures. Appendix table C1 lists the number of houses in each category. Removing farms is our most important sample restriction. The 1930 and 1940 data do not include dwelling characteristics, nor do they indicate the presence of businesses, so we cannot remove businesses and multi-family structures in these years. These two categories account for only 6 percent of owner-occupied houses in 1960.[9]

In addition to the Census data, we use the Survey of Consumer Finances (SCF) from 1983 through 2013 to provide some insight into how movements in mortgage debt may have influenced inequality in home equity, as opposed to gross holdings of housing assets. The SCF oversamples high-income, high-welth households which tend to have lower response

[6] Goldin and Margo (1992) used methods similar to ours to highlight and examine the great compression in incomes in the 1930's and 1940's. Piketty and Saez (2003) examine the great compression of top incomes, particularly from capital.

[7] The 1950 Census does not include information about housing, and we o not have access to the 1950 Census of Housing Microdata.

[8] The home value is the owner's estimate. We discuss below how question changes impact our analysis. We take data from the ACS from 2009 (not 2008) through 2012 to reflect conditions after the latest housing boom and bust.

[9] Enumerator instructions for 1940 instruct respondents to exclude the value of units rented out and to also exclude the area of the house that is used for business purposes if this is a "considerable portion" of the house. These instructions are re-printed in the data appendix.

rates in order to provide a more representative sample throughout the wealth and income distribution, making it a useful source of data to study inequality. The survey is conducted every three years and asks respondents about their wealth and income. Specifically, homeowners report house values of both their primary residence, as well as any other real estate properties owned. In addition, households provide detailed information on financial and nonfinancial wealth and various forms of debt, which are aggregated with housing wealth data to compute net worth for each household and deflated using the CPI to reflect 2013 dollars. All charts and figures using the SCF data are computed using cross-sectional survey weights.

3.2.2 Interpolation and Extrapolation Procedures

The Census often asks questions about home values and rents in terms of intervals. Even when questions are asked in terms of continuous dollars, the data are top coded, and respondents often round. To account for these issues, we use a Pareto interpolation procedure to allocate responses within intervals. Where y are values distributed according to $F(y)$, this procedure fits $\ln[1 - F(y)] = \alpha_j[\ln(k_j) - \ln(y)]$, where $\alpha_j > 1, k_j > 0$ are parameters for an interval j. In other words, it fits a linear spline of the "tail" function $\ln[1 - F(y)]$ to $\ln(y)$. For high values of y, the α_j parameters are fairly constant, implying that the Pareto distribution provides a good fit. To improve comparability and handle rounding problems, we also interpolate in years with continuous data, using 25 artificial intervals, with 4 percent of the data in each.[10]

At the top of the distribution, we use standard Pareto extrapolation procedures (e.g. Atkinson et al. (2011)). In years that we have information about the mean of top coded values, we use the estimate of $\alpha_J = E[y|y > k]/(E[y|y > k] - k)$. In other years, we extrapolate α_J for the top coded bin from the α_j from intervals in the top 10 percent of the distribution.[11] For the bottom interval we extrapolate values using a uniform distribution to deal with bottom coding and problems associated with zero values. We set the minimum of this interval as a proportion of average household income in that year. More details are in Appendix C.2.

[10] Because of rounding, these intervals can sometimes contain only one value. In these cases, we decrease the number of groups until each interval has a top value at least 2.5 percent higher than its bottom value.

[11] The one exception to this are rent values in 1930 and 1940, where we hard code the top Pareto distribution parameter to be 2.5. In these years we have a continuous distribution of rents (and home values) with very few top codes; parameters estimated from the mean procedure lead to suspiciously high inequality statistics.

3.2.3 House Values, Gross Rents, and Consumption Equivalents

We present all of our results using the Census and ACS data separately for housing prices and rents. To integrate the two and account for changes in home ownership (which mainly changed from 1940 to 1960), we impute a rental equivalent measure for home owners, based on a constant user cost of 0.0785 times the value of the house and plus utility payments.[12] We then divide these rental measures by a household equivalence scale to provide a per-person measure of housing consumption.[13] We also provide results for home equity in both primary residences and all real estate using the SCF data.

3.2.4 Inequality Measures

To describe inequality in housing expenditures we use three scale and population invariant measures: the variance of logs, the Gini coefficient, and the Theil entropy index.[14] We decompose the variance of logs and the Theil index to show how much overall inequality is due to variation across versus within areas. Note that the variance of logs is particularly sensitive to the bottom of the distribution while the Gini coefficient and Theil mainly measure changes in the middle of the distribution (Atkinson (1970)).

In addition, we describe inequality in housing and net worth over the full distribution using Lorenz Curves. Doing this allows us to further decompose the contribution of home equity (net housing wealth) to the overall Lorenz curve for net worth.[15]

[12]This commonly-used value is based on Peiser and Smith (1985). User costs of home ownership are foregone interest, depreciation, and property taxes, minus price appreciation. User costs vary over space, time, and possibly occupant, but there are various empirical and conceptual challenges in calculating year-by-city user costs (e.g. those examined in Verbrugge (2008) and Poterba and Sinai (2008)). We choose a constant number so as to avoid the data being overly influenced by our choice of methodology, recognizing that scale invariance of the inequality measures will help to undo differences among owners.

[13]Equivalence scales are designed to take into account economies of scale in consumption, such as from the sharing of common spaces, kitchens, etc. We use the OECD equivalence scale: $1 + 0.7(A - 1) + 0.5C$ where A is the total number of adults in the household and C is the number of children (14 years old or below).

[14]The Theil index is computed by taking a weighted average of the log of the (normalized) expenditure on housing by each household. Where $\bar{y} \equiv E[y_i]$:

$$T = \frac{1}{N} \sum_{i=1}^{N} \left[\frac{y_i}{\bar{y}} \ln \left(\frac{y_i}{\bar{y}} \right) \right]$$

The Gini coefficient and Theil index satisfy the principle of transfers, so a transfer from a rich to a poor person always decreases inequality. We decompose Theil entropy indices and variances of logs by calling inequality in mean levels of housing expenditures across areas the between (areas) component, and inequality within areas the within component. As Cowell (2011) notes, the variance of logs is only decomposable with constant means of logged values, not levels.

[15]The Lorenz Curve for net worth in the p-th percentile of the population is given by:

$$L(p) = \frac{\sum_{i=0}^{p} w_i}{\sum_{i=0}^{1} w_i} = \frac{\sum_{i=0}^{p} \left[w_i^h + w_i^{nh} \right]}{\sum_{i=0}^{1} w_i}$$

3.2.5 Dwelling Characteristics and Location Measures

Dwelling characteristics are provided only from 1960 onwards. They include the age of the building, the number of rooms, the number of bedrooms, the presence of complete plumbing, and what heating system, if any, was installed (gas, oil, wood, electric, etc.).

Measures of location extend back to 1930. We consider the lower 48 states, and 722 commuting zones (CZs). CZs are defined like metropolitan areas, corresponding to local labor markets, but include rural areas, and are constant over time.[16]Some locational advantages may be permanent, such as climate and natural features. Others, such as safety or employment proximity, may have changed considerably since 1930. We remain agnostic about which amenities households value, mainly because we can measure so few of them.

3.2.6 Re-weighting Analysis

To account for how observable changes in housing characteristics have affected inequality, we use the propensity-score re-weighting approach of DiNardo et al. (1996). This approach estimates what the distribution of housing prices would be if the observed characteristics of houses had remained identical to earlier years, and those characteristics had been priced as in later years. In this way, it rules out general-equilibrium effects (such as changes in demand) affecting prices. Nonetheless, it is a useful accounting method, as changes from re-weighting tell us how well observable characteristics could affect inequality.

The re-weighting makes the observable characteristics of a sample of houses in a later "current" year, t (e.g. 2010), resemble characteristics from a previous "target" year, t' (e.g. 1970). The weights are the odds of a house being observed in the target year relative to the current year, which we determine using a logit regression on the pooled sample in years t and t'. Fortin et al. (2011) provides detailed conditions for identification and implementation details. The main requirement for identification is conditional independence – after conditioning on all characteristics, a house's price is independent of the year we observe it in.

where total net worth w_i can be decomposed into net housing wealth (home equity) w_i^h and net nonhousing wealth w_i^{nh} and the index $i \in (0, 1)$ is a percentile ordering that sorts households from lowest to highest net worth.

[16]We use 1990 Commuting zones, which are more fully explained in Tolbert and Sizer (1996). They are made up of counties and designed to be places where people both live and work. For 1930 and 1940 we are able to assign every house to one commuting zone. In 1960 they are not available in public PUMS samples. From 1970 onward some houses are identified as being in PUMA's or county groups in the IPUMS data that cut across multiple commuting zones. For these years we follow Dorn (2009), Autor et al. (2013), and others in probabilistically assigning houses to commuting zones based on the proportion of responses in their most detailed geographic category that were in one commuting zone or the other. For most of our specifications we also report statistics using states. We are able to match each house to a state for each year.

3.3 Empirical Results

3.3.1 General Trends over Time

Table 3.1 provides basic characteristics of the sample. The first row shows how home-ownership fell after the Great Depression and rose dramatically after World War II, hovering just above 60 percent for over fifty years. Household size consistently fell (from 3.8 in 1930, 3.3 in 1960, to 2.5 in 2012) and the number of rooms per house rose slowly after 1960 (from 5 to 5.75) along with the number of bedrooms. Indoor plumbing was common but not ubiquitous in 1960, but became so by 1980. Taken together, these trends imply Americans consumed much more housing per person. The fraction living in the South and West also ballooned from 35 to 60 percent, while the Midwest and Northeast lost its former predominance.

Figure 3.1 graphs Lorenz curves for various measures of housing outcomes and incomes in 1930, 1970, and 2012. We present Lorenz curves because they show how each part of the distribution has changed over time.

Panel A shows that, among home-owners, 2012 and 1930 have almost identical levels of inequality, with 50 percent of housing value accruing to the top 20 percent. Inequality is uniformly lower in 1970 with 40 percent going to the top 20 percent. To better understand the distribution of housing assets (if not equity) Panel B includes renters as zeroes. Here we see 1930 is considerably less equal, with 80 percent of owner-occupied housing wealth accruing to 20 percent of households, resembling the original "Pareto Principle," for 19th-Century landowners. In comparison, inequality in 2012 is much lower, although changes in the number of farms may influence this number.

Panel C shows that rents were quite unequal in 1930, but that they were much more equal in 1970. Since then, rent inequality has only grown slightly, primarily for costlier units. Panel D shows consumption equivalents, which summarize much of the earlier discussion: 1930 had the highest inequality, 1970 the least, and 2012 resembles 1930 at the top of the distribution but 1970 at the bottom.

Panels E and F describe household wage income (not available in 1930) and total income (not available in 1930 or 1940). Total income inequality is generally smaller, especially at the bottom (partly due to lack of zeroes), and increased less than wage inequality between 1970 and 2012.

Figure 3.2 graphs our preferred inequality statistics in all years. Generally these show U-shaped patterns from 1930 to 2012. Before 1970 there was a great compression in housing consumption, similar to trends in the inequality of household income shown in Panel D.

From 1930 to 1970 the Theil entropy index and the variance of the log of home prices roughly halved. Each statistic increases after 1970 — with a blip in 1990 that seems due to regional housing booms — and settles in 2012 at a level slightly below its value in 1930. It is especially remarkable that home-ownership waxed considerably during the great compression, while it did not wane during the great divergence.

Rents exhibit a different pattern. While inequality in rents declined dramatically between 1930 and 1960, it increased less than home values in the period starting in 1980. Since inequality changed more at the top than at the bottom, this may relate to increased levels of housing assistance for low-income households, or rent control in major cities like New York and San Francisco.

The consumption equivalent measure accounts for changes in demographics and home ownership, and exhibits a U-shaped pattern that is slightly less extreme in the latter half of the U. This measure shows the lowest levels of inequality in 1980, while the other measures bottomed-out closer to 1960. Inequality in consumption equivalents appears to be currently at its highest level since World War II and seems to be increasing.

While we only have household wage income in 1940 (and nothing in 1930) it is worth noting that inequality in our housing measures fell *more* between 1940 and 1960 than measures of income. Subsequent increases in housing inequality since 1960 have been comparatively gentler.

3.3.2 Decomposition over Space

In Figure 3.3 and Table 3.3 we use a spatial decomposition to see how much inequality is due to differences across areas as opposed to within them. Differences across cities are likely labor market driven, since CZs are designed to resemble local labor markets. Differences within cities are likely due to dwelling characteristics, local amenities, and commuting opportunities.

In all cases, within-area inequality is much larger than between-area inequality, such that the within statistics are generally between three to ten times as large. The U-shaped patterns observed at the national level for each measure are reflected both within and across states and commuting zones. However, given the relative magnitudes, changes in inequality are mainly driven by differences within metro areas. In contrast, much of the literature (examples include Van Nieuwerburgh and Weill (2010), Diamond (2016), Gyourko et al. (2013), and Moretti (2013)) focuses on growing inequality between metro areas. Between-metro changes did figure prominently in 1990 blip, which occurred in rents as well as prices, but have otherwise been dwarfed by changes within metros.

The large increases in household wage inequality within areas in the late 20th century, mirrors those in Baum-Snow and Pavan (2013), who find income inequality grew the most within the largest U.S. cities. Meanwhile, inequality between metro areas declined considerably between 1940 and 1960, and did not change considerably afterwards outside of the 1990 blip.

3.3.3 The Role of Observable Characteristics

Although the share of income devoted to housing appears to have stayed relatively constant over the 20th century, patterns of housing consumption appears to have changed considerably. As pointed out in Table 1, while households have shrunk in (human) size, their housing units have gotten larger, and they have moved south and west.

Increases in housing price inequality since 1970 could be the result of two forces. First, some households may have moved into ever larger units, while others moved into ever smaller units. Second, Americans might have moved to high and low value cities, heightening inequality.

These hypotheses are examined in Figure 3.4 and Table 3.4. They show that if households lived in housing units with dwelling characteristics, locations, and household composition observably identical to those in 1970, housing values would be only slightly more equal. This accountable change is very small in relation to the overall changes.[17]

The re-weighting results in Table 3.4 show that dwelling characteristics can explain at most 30 percent of the change in the two inequality statistics using the consumption equivalent measure. The direction varies, however, with decreases in inequality among home owners and *increases* for renters. Generally, changes in the location of houses and household composition from 1970 to 2012 have tended to make households *more equal* in their housing outcomes. It is also interesting, even surprising, that neither location nor dwelling characteristics do much to explain changes in household wage inequality.

Our results on dwelling characteristics support previous reflections on this issue. For example, Glaeser and Gyourko (2008) find that per-capita square footage consumed by rich and poor households has become more equal over time. Since construction costs vary little within cities, much of the growing inequality in housing value seems to be due to growing inequality in land values, or the right to build on such land.

[17] The re-weighting specification we use allows for depreciation by including their age as an explanatory variable. This rules out controlling for "vintages" of houses with features particular to their year of construction. The "Vintages" specification includes a house's vintage, but not its age (implicitly assuming there is no depreciation net of usual maintenance). Davis and Heathcote (2007) find that spending on structural improvements in the US (roughly 0.8 percent) is very similar to a reasonable estimate of structural depreciation (roughly 1 percent). Our conclusions using vintages reinforce the idea that changes in dwelling characteristics have not driven the changes in housing inequality that we document.

3.3.4 The Relationship between Income and Housing Inequality

The patterns documented in section 3.3.2 suggest a link between local income and housing inequality. Studies of housing demand (e.g. Polinsky and Ellwood (1979), Hanushek and Quigley (1980), and Mayo (1981)) have generally concluded that housing is a necessity, a regularity sometimes known as "Schwabe's Law." Most estimates of the income eladststicity of demand for housing are 0.3 and 1.0. So variations in income should be reflected in housing consumption. In fact, other things equal, the variance of log housing consumption should equal the square of the income elasticity of housing times the variance of log income.[18]

Patterns of income and housing inequality over time and space suggest a considerable relationship between the two. For instance, between 1970 and 2012, the variance of log income increased by 0.26, while the variance of log housing consumption increased by 0.09, consistent with an income elasticity of 0.59.

We examine this relationship spatially in Figure 3.5 and Table 3.5 by comparing how inequality within CZs for housing relate to income. Here we see that the two inequalities are strongly related, and share a relationship of roughly the same magnitude. Column three of Table 3.5 takes the square root of a regression of the variance of the logs of measure of housing consumption on the variance of logs of household income. This produces estimates of the income elasticity of housing expenditures in the range of 0.7 to 0.9, which are reasonable, if slightly higher than values from temporal variation.

3.3.5 Role of Housing in Driving Wealth Inequality

Growing inequality in house prices from 1970 onwards, documented in the previous sections, likely maps into growth in inequality in overall wealth studied in the literature. Moreover, since housing assets are used as collateral for mortgages, home equity loans, and lines of credit, households with housing assets may purchase them by increasing leverage. As such, inequality in net positions in housing may show larger inequality than those in gross positions.[19] To provide insight into the quantitative importance of housing and housing-related debt in driving wealth inequality, we use household-level data from the SCF which provides a comprehensive picture of household balance sheets throughout the wealth distribution.

[18]This exercise carries several caveats: General equilibrium effects may interact with consumer preferences to either dampen (Matlack and Vigdor (2008)) or amplify (Van Nieuwerburgh and Weill (2010)) the direct effects of changes in incomes. Additionally, since housing is so heterogeneous, it is difficult to quantify how much "housing service" a house of a given size or in a given neighborhood provide.

[19]When a household owns a home with a mortgage, the mortgage issuer has a claim on the gross housing wealth. These claims may, in turn, be sold or further collateralized and eventually held by other household, likely higher up the wealth distribution. We focus here on home equity rather than indirect claims on housing assets for two reasons. First, data on indirect claims are not readily available. Second, since houses are primarily purchased using debt, capital gains always accrue to homeowner alone.

Specifically, the data allow us to both net out housing debt held against housing assets and compute the fraction of overall net worth comprised of housing wealth at various points in the distribution.

Using data from the SCF, Figure 3.6 plots the Gini coefficient, Thiel index, and the standard deviation of logs for gross housing wealth, home equity, and net worth.[20] As with the Census data, results from the SCF suggest inequality in gross housing wealth has risen since 1983. The increase in inequality based on home equity appears to have been somewhat larger than the increase based on gross housing wealth, suggesting that homeowners with lower overall wealth also have higher leverage. While inequality in both net and gross housing wealth has risen, the change in inequality in overall net worth appears somewhat larger.

To study the role of housing in driving rising wealth inequality, Figure 3.7 plots Lorenz Curves for overall net worth along with the contributions from net equity in the household's primary home and net equity in other real estate. The overall Lorenz curves in each year (solid lines) are computed as the fraction of aggregate net worth held by the lowest *pth* percentile of the net worth distribution. The shaded regions show the fraction of this group's net worth comprised of home equity in their primary residence and other real estate respectively. At the 100th percentile, the charts show home equity makes up just over 20 percent of overall net worth in the SCF, broadly consistent with aggregate national statistics such as the Flow of Funds Accounts.

In the cross-section, home equity explains much of the inequality in housing wealth at the middle of the wealth distribution, where housing wealth comprises a larger fraction of overall net worth. Households at the top of the distribution are more likely to hold larger shares of their overall wealth in other non-real-estate assets. Looking over time, housing wealth's contribution to overall wealth inequality appears to have diminished somewhat since 1983. This is evident from the fact that the shaded region below the dashed line has fallen by more than the solid line at each quantile. Moreover, the Gini coefficient for overall net worth can be decomposed into its contribution from home equity as in Lerman and Yitzhaki (1985). Doing so suggests that home equity explained roughly 20 percent of inequality in net worth in 1983, but has fallen to explaining only 16 percent in 2013. The decline occurs because the decline in the share of net worth coming from home equity outweighs the increase in inequality within home equity over this period.

[20] All housing statistics apply to primary residence only. Gross housing wealth is the self-reported house price for owner-occupied homes. Home equity is gross housing wealth less mortgages, home equity loans, and home equity lines of credit secured with the respondent's primary residence. Net worth is home equity plus all financial and nonfinancial assets, less all household debts. Imputed values of pensions are included in all waves except 1983 and 1986.

3.4 Conclusion

Our results provide some refinements to the debate on inequality, particularly in terms of consumption and wealth. The similarity of housing and income inequality over space and time according to plausible income elasticities appears to support those arguing that consumption inequality does reflect income inequality. Additionally, changes in dwelling characteristics and differences between cities explain only a small fraction of recent increases in housing inequality. This suggests that the value of land plays an important role, even within cities.[21]

Several studies — e.g., Green et al. (2005b), Gyourko et al. (2008b), and Saiz (2010) — have emphasized that regulatory and geographic constraints on housing supply may play an important role in price differences across cities. Our findings suggest constraints may play a role within cities, as new housing in the most desirable neighborhoods may be the most constrained. These may interact with findings by Rossi-Hansberg et al. (2010), Guerrieri et al. (2013), and Autor et al. (2014) on how local externalities that can lead to substantial income sorting within cities. Generally, our findings suggest that researchers would do well to more closely examine differences in land prices across neighborhoods. An interesting research project would be to determine how much they are driven by local externalities, dwelling characteristics, and fixed neighborhood amenities.

The growing inequality in housing prices that we document also indirectly supports findings that wealth inequality has increased. High home-ownership levels do imply that inequality in housing wealth is still smaller than it was in 1930. Nevertheless, the windfall gains from unequal housing price changes, which benefited some homeowners relative to others, may help stir once-popular Georgeist concerns about unequal land-ownership from the beginning of the 20th Century. Moreover, it appears that housing inequality will continue to grow in line with any further increases in income inequality.

[21] Our result also supports findings by Watson (2009) and Reardon and Bischoff (2011) that segregation by income has increased since 1970. This contrasts with findings by Glaeser and Vigdor (2012) and others that black-white segregation peaked in 1970 and has since declined.

Figures

Figure 3.1: Lorenz Curves for Home Values, Rents, Housing Consumption, and Household Income

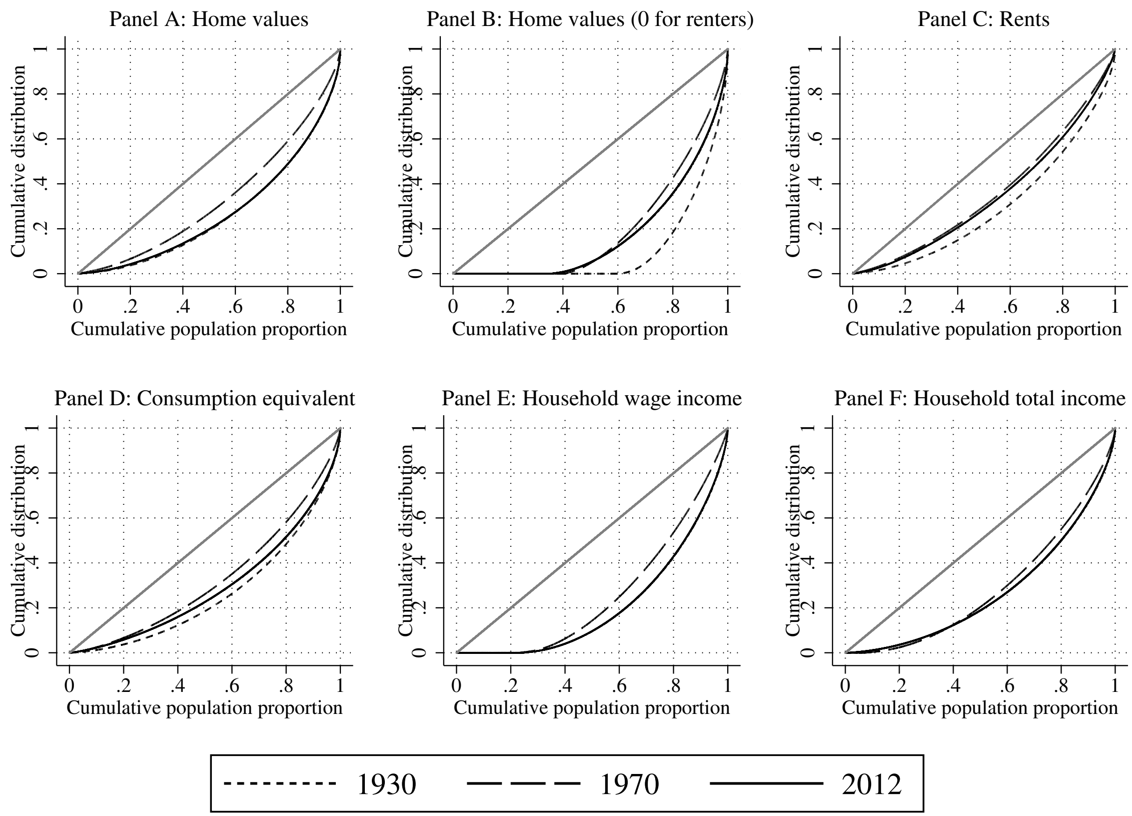

NOTE: These curves graph the cumulative percentage of housing expenditures (from lowest to highest), against the cumulative percentage of households. Panel A, is for owning households; Panel C is for renters. All others are for the sample of home owners and renters combined. Panel B includes renters with an implied home value of zero. Data are interpolated within intervals as described in the text. Data are from the Decennial Census and the 2009-2012 ACS. See the text and data appendix for more detail.

Figure 3.2: Inequality over Time in Home Values, Rents, Housing Consumption, and Household Income

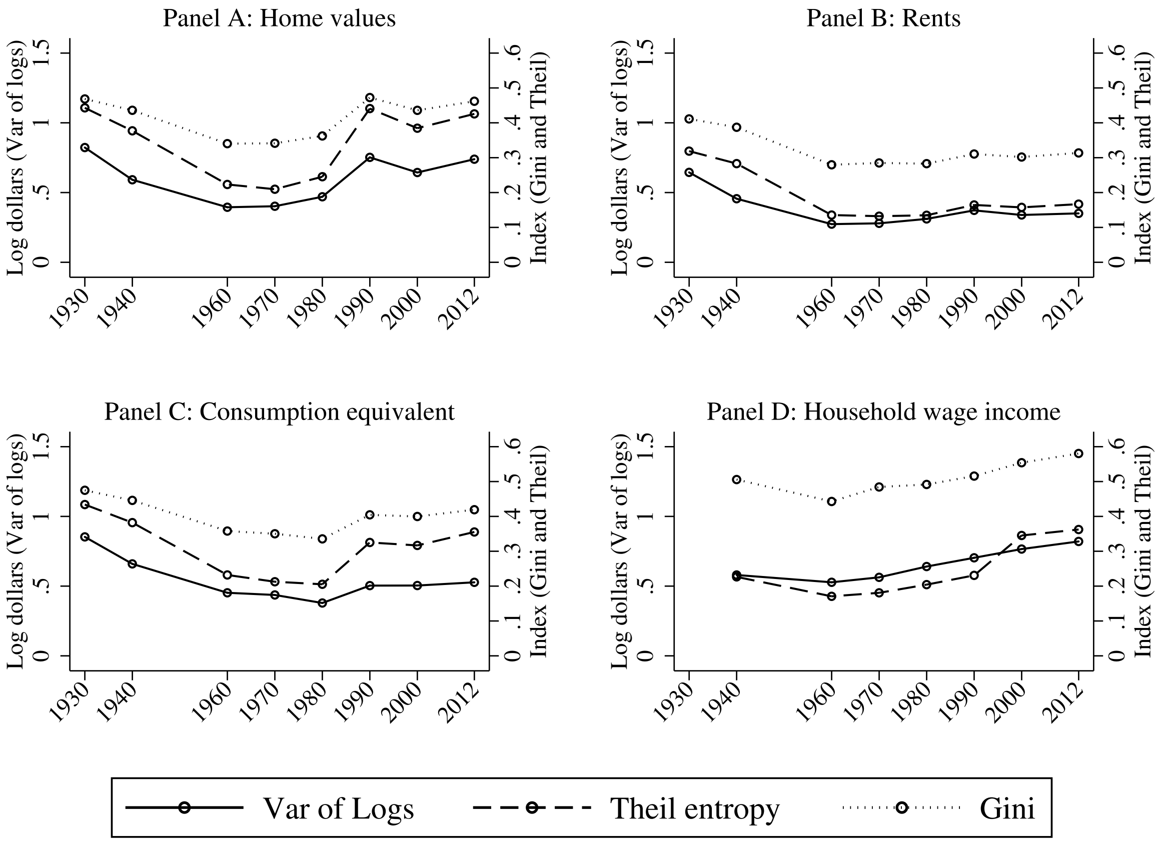

NOTE: Each presents inequality measures for separate samples. The first is for home owners and is the dollar value of their primary residence. The second is for renters and is their cash expenditure per month on rent and utilities. The final combines the two to compute an consumption measure per person, as explained in the text. Farms and houses used for business are excluded. Data are interpolated within intervals. The observation for 2012 comes from ACS for 2009-2012. See text and Data Appendix for more detail.

Figure 3.3: Decomposition of Values Between and Across Geographies.

Panel A: Theil entropy

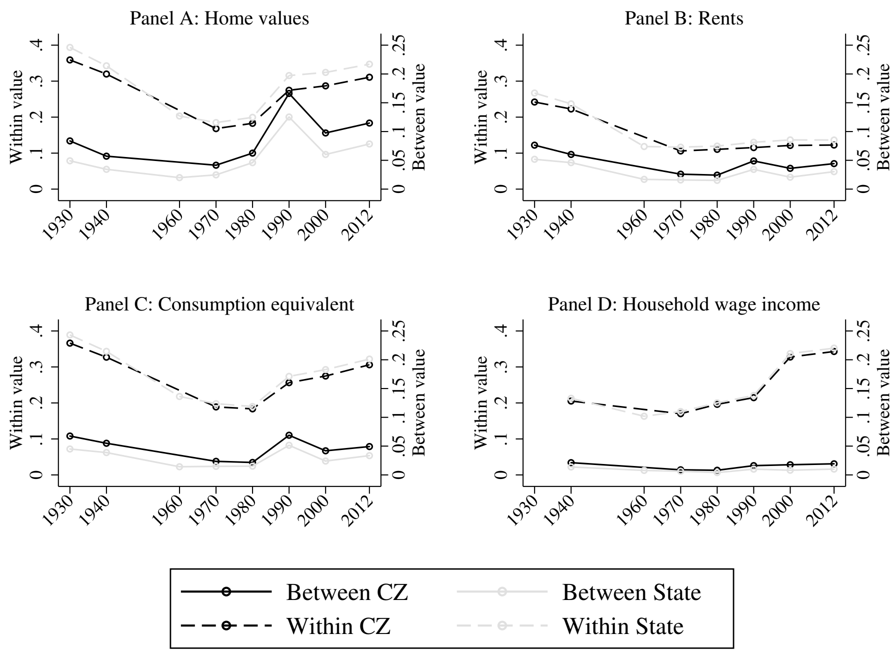

Panel B: Variance of logs

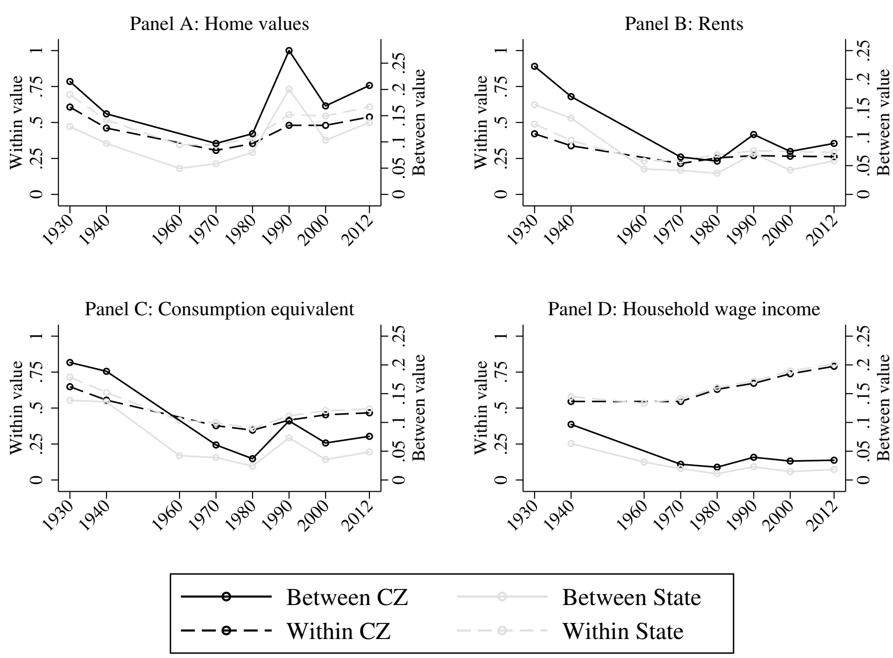

NOTE: Between and within decompositions of housing expenditures are shown for the Theil entropy index and the variance of logarithms. See earlier figure notes for details about the sample.

Figure 3.4: Explanatory Power of Observable Dwelling, Location, and Demographic Characteristics

Panel A: Theil entropy

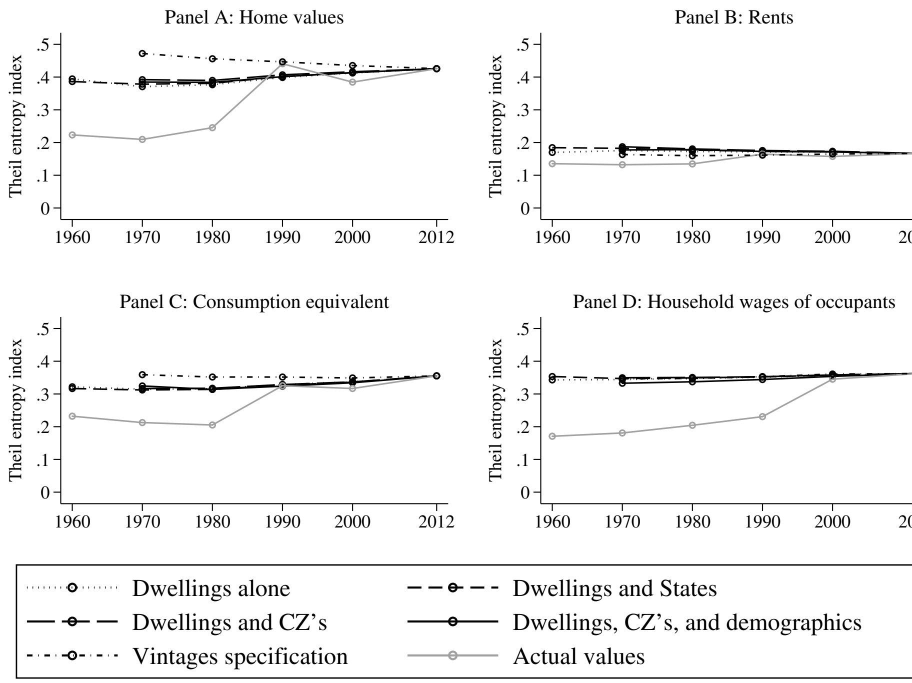

Panel B: Variance of logs

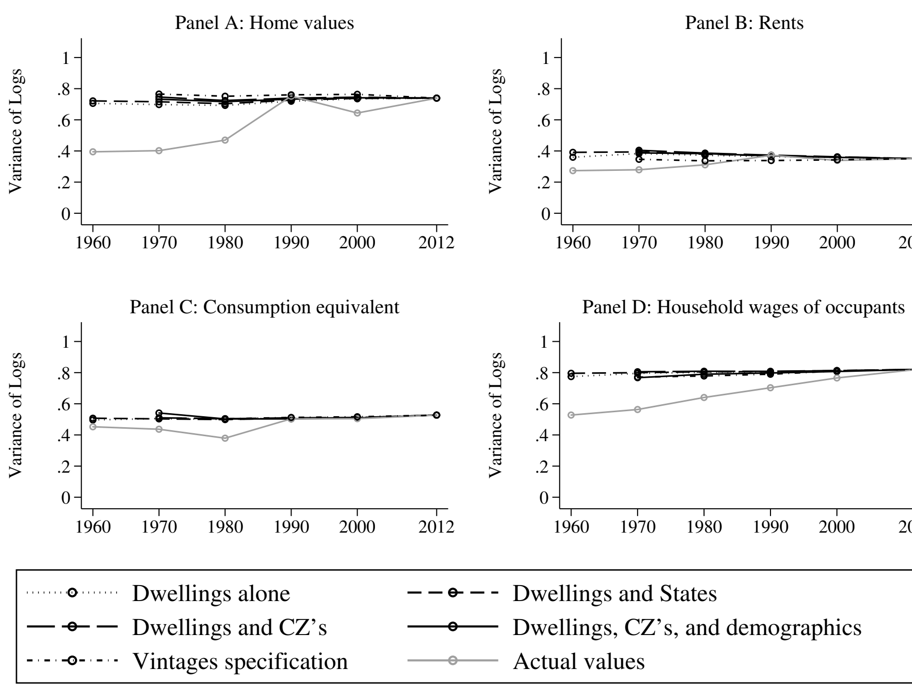

NOTE: Lines in black refer to re-weightings of the 2012 data to be comparable in composition in terms of the observed named variables. The line in grey refers to the actual values in the specified years. See Table 4 for variable descriptions. The solid line contains weights accounting for all of these which were computed as noted in the text, following DiNardo et al. (1996) and Fortin et al. (2011).

Figure 3.5: Local Inequality in Housing versus Income

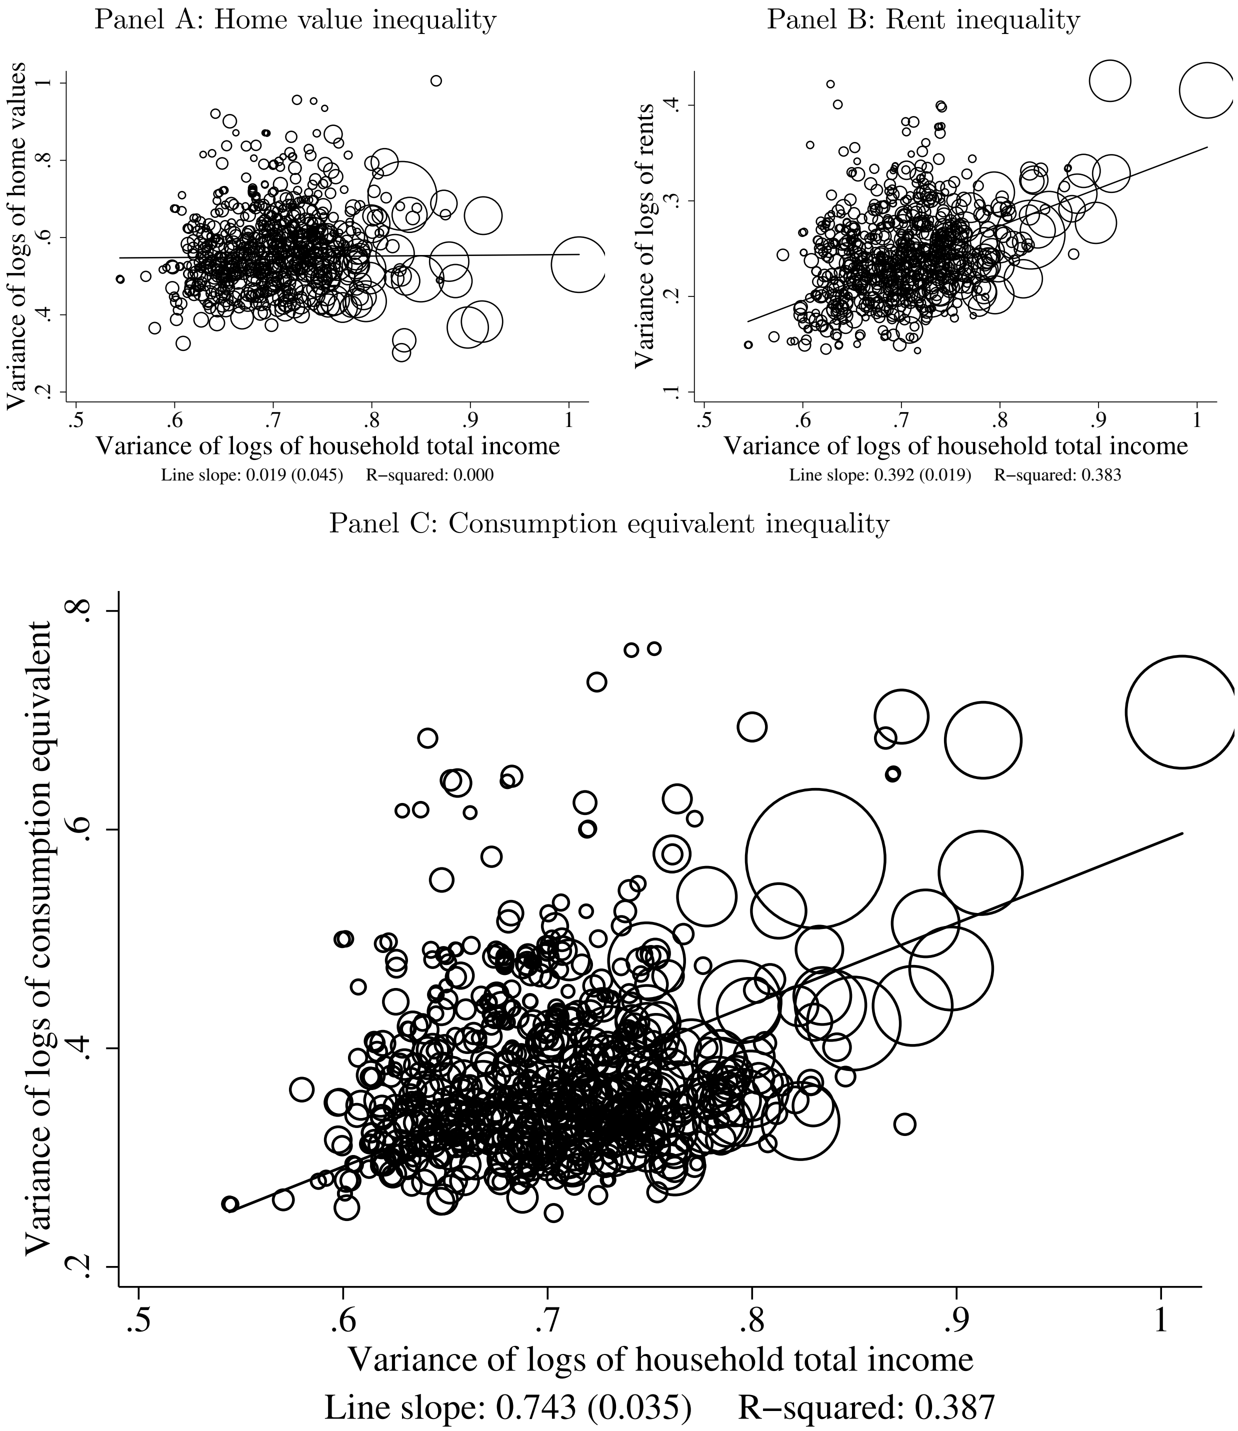

NOTE: These panels display scatter plots for commuting zones with the variance of logs of total income recorded in the 2009-2012 ACS on the x axis and each measure of housing expenditure inequality on the y axis (with zeros excluded). Each variable is computed for its relevant population with the variance of the logs of household total income including both owners and renters in each graph. The size of each circle is proportionate to the number of (weighted) households in each commuting zone. The OLS regression line plotted incorporate these weights.

Figure 3.6: Inequality in Home Equity

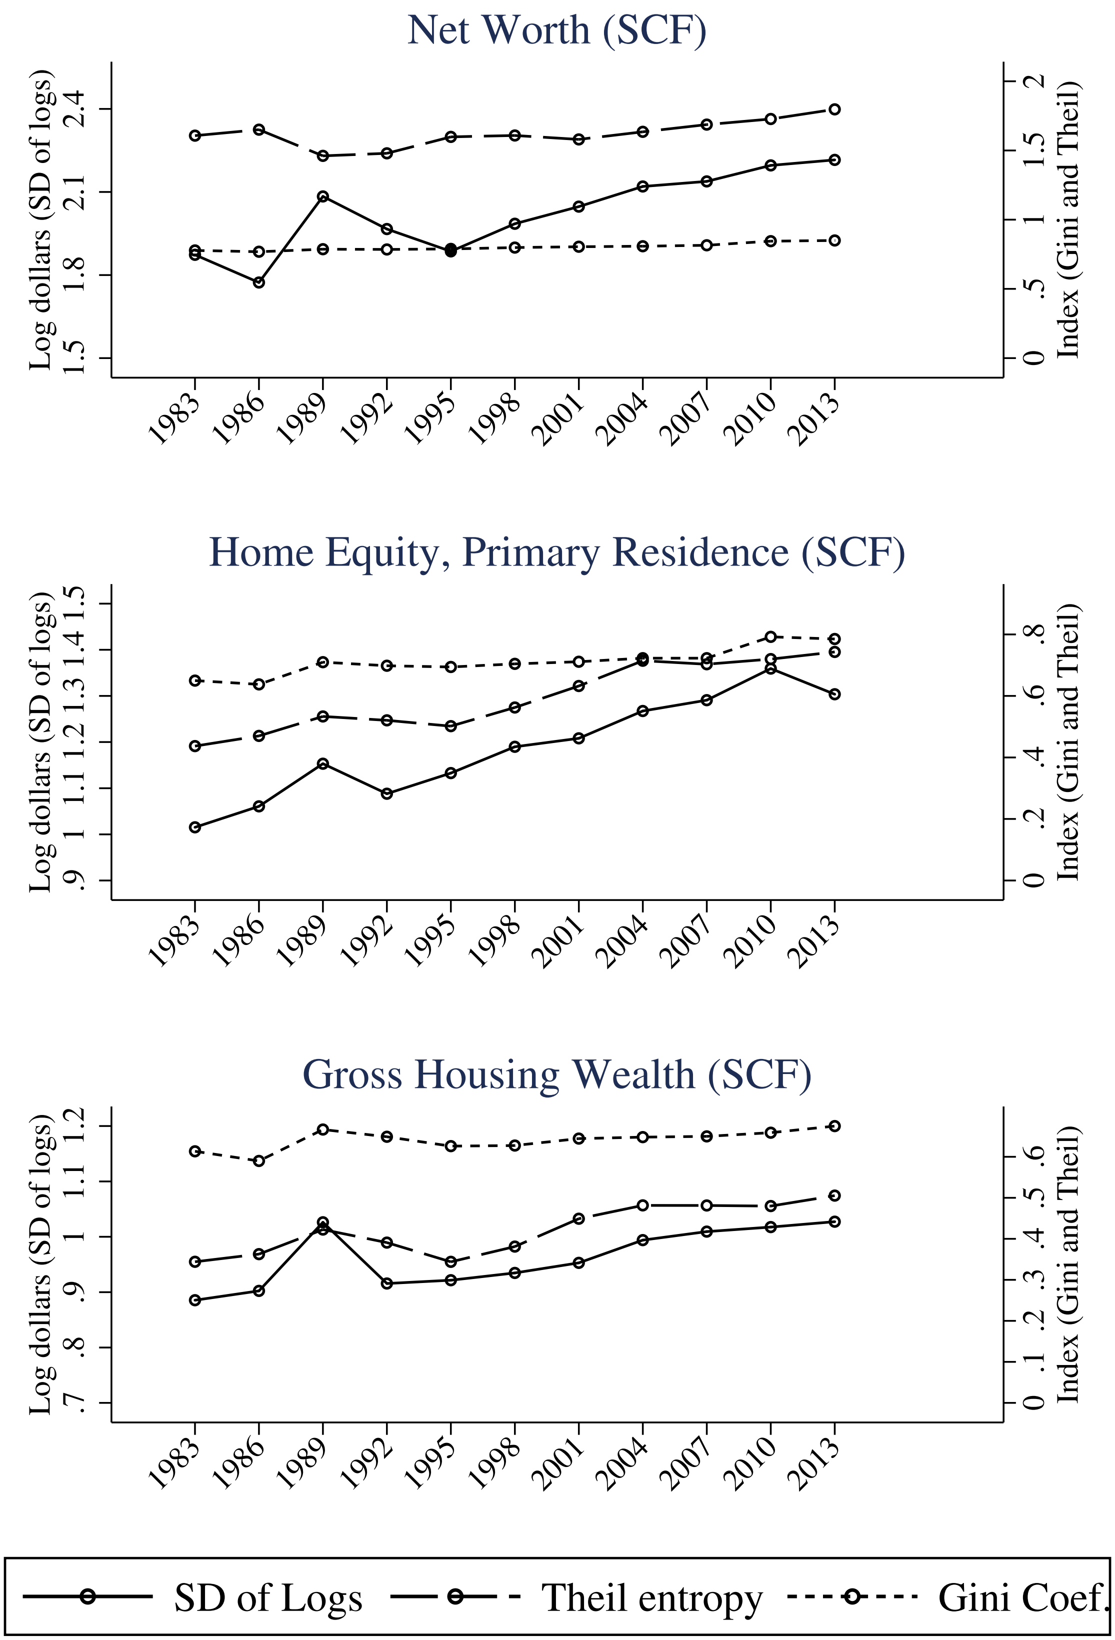

NOTE: Each panel presents inequality measures for overall net worth, home equity (net housing wealth), and gross housing wealth in the full SCF sample by wave.

Figure 3.7: Contribution of Home Equity to Lorenz Curves for Net Worth

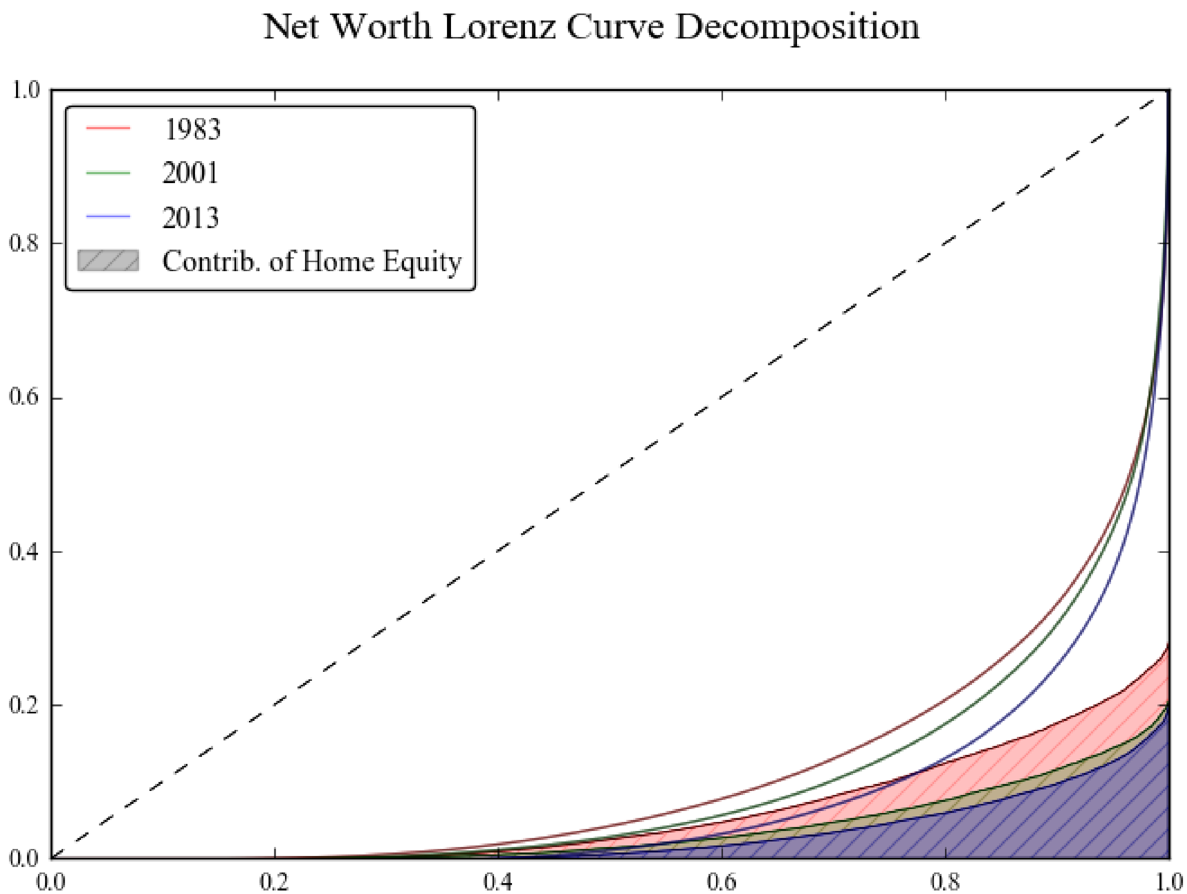

NOTE: These curves graph the cumulative percentage of total net worth held by households in various percentiles of the wealth distribution. The area under the curve is further decomposed into the contribution from home equity in primary homes (shaded hatched region) for each of the three waves drawn.

Tables

Table 3.1: Descriptive Statistics

	1930	1940	1960	1970	1980	1990	2000	2012
Sample size	1,146,534	1,436,399	451,066	2,237,327	3,589,168	4,172,556	4,749,600	4,410,4
Owner occupied	0.45	0.41	0.61	0.62	0.64	0.63	0.65	0.65
Median home value (thous)	53	40	74	83	113	111	130	154
Median monthly rent	312	292	361	433	472	572	609.18	649
Median annual consumption equivalent ($/person)	1,523	1,396	2,955	3,588	4,962	5,465	6,064	6,950
Median annual household wage income (thous)		14	29	35	32	35	37	33
Median annual household total income (thous)		14	34	41	40	46	50	47
Positive wage income		0.79	0.82	0.77	0.79	0.77	0.78	0.76
Household size	3.84	3.44	3.27	3.11	2.75	2.64	2.60	2.54
Number of rooms			4.92	5.09	5.31	5.37	5.45	5.74
Number of bedrooms				2.41	2.48	2.53	2.57	2.71
Plumbing facilities			0.88	0.95	0.99	0.99	0.99	1.00
Midwest	0.32	0.31	0.28	0.27	0.26	0.24	0.24	0.23
South	0.23	0.24	0.28	0.31	0.33	0.35	0.37	0.38
West	0.11	0.13	0.17	0.18	0.19	0.21	0.21	0.22

NOTE: Sample includes renting or home-owning households (when at least one resident is the owner-occupier). Data comes from the Decennial Census and 2009-2012 ACS. Blank cells indicate that the relevant statistic is not available. Values are means, and dollar values area deflated by the CPI excluding structures, unless otherwise specified.

Table 3.2: Inequality Statistics

Home Values	1930	1970	2012
Variance of logs	0.823	0.402	0.739
Theil entropy	0.443	0.210	0.426
Gini coefficient	0.468	0.341	0.462
Ratio of 90 to 10th percentiles	11.159	5.399	9.295
Rent			
Variance of logs	0.645	0.279	0.351
Theil entropy	0.319	0.132	0.167
Gini coefficient	0.411	0.285	0.314
Ratio of 90 to 10th percentiles	8.066	3.989	4.619
Consumption equivalent			
Variance of logs	0.853	0.437	0.527
Theil entropy	0.434	0.213	0.355
Gini coefficient	0.475	0.350	0.419
Ratio of 90 to 10th percentiles	10.790	5.283	6.146
Household wage income	1940	1970	2012
Variance of logs	0.580	0.563	0.820
Theil entropy	0.227	0.181	0.362
Gini coefficient	0.506	0.484	0.580
Ratio of 90 to 10th percentiles			

NOTE: Home values refer to owner-occupied homes. Rents are gross and include the cost of utilities. Both are self reported. Consumption equivalents combine gross rents with imputed rents based on a percentage of the home value plus utility costs, divided by the equivalence scale $1 + 0.7(A - 1) + 0.5C$ where A is the number of adults and C is the number of children under 14. Household wage income refers is the sum of wages and salaries from all household members. Values are interpolated within intervals using standard procedures described in the text.

Table 3.3: Between-Within decomposition

Panel A: Home Values

	1930	1970	2012
Variance Overall	0.823	0.403	0.746
Variance Between	0.215	0.097	0.208
Variance Within	0.607	0.306	0.538
Theil Overall	0.443	0.210	0.426
Theil Between	0.084	0.042	0.115
Theil Within	0.359	0.168	0.311

Panel B: Rents

	1930	1970	2012
Variance Overall	0.645	0.280	0.351
Variance Between	0.223	0.065	0.089
Variance Within	0.422	0.215	0.262
Theil Overall	0.318	0.132	0.167
Theil Between	0.076	0.026	0.044
Theil Within	0.242	0.106	0.122

Panel C: Consumption equivalents

	1930	1970	2012
Variance Overall	0.853	0.438	0.543
Variance Between	0.204	0.061	0.076
Variance Within	0.648	0.377	0.467
Theil Overall	0.434	0.213	0.355
Theil Between	0.068	0.024	0.049
Theil Within	0.366	0.189	0.306

Panel D: Household wage income

	1940	1970	2012
Variance Overall	0.643	0.586	0.825
Variance Between	0.097	0.027	0.034
Variance Within	0.546	0.546	0.790
Theil Overall	0.227	0.179	0.362
Theil Between	0.021	0.009	0.019
Theil Within	0.205	0.170	0.343

NOTE: Between and within decompositions refer to the decomposition of each statistics to variation within states or commuting zones (CZs) versus between them. The sum corresponds to values in table 3.2.

Table 3.4: Re-weighting

Panel A: Theil entropy

	Consumption equiv	Rents	Home values	HH wages
2012	0.355	0.167	0.426	0.362
Dwellings	0.315	0.176	0.371	0.344
Dwellings and state	0.312	0.182	0.378	0.347
Dwellings and CZ	0.317	0.187	0.392	0.350
Dwellings, CZ, and demographics	0.325	0.177	0.385	0.333
1970	0.213	0.132	0.210	0.181

Panel B: Variance of logs

	Consumption equiv	Rents	Home values	HH wages
2012	0.527	0.351	0.739	0.820
Dwellings	0.505	0.383	0.698	0.796
Dwellings and state	0.503	0.394	0.716	0.800
Dwellings and CZ	0.510	0.404	0.746	0.806
Dwellings, CZ, and demographics	0.541	0.388	0.733	0.768
1970	0.437	0.279	0.402	0.563

NOTE: Statistics are taken of 2012 data, re-weighted to emulate the distribution of observable dwelling and location characteristics of previous years. The top and bottom row refers to the actual values in the specified years. Dwellings refers to: Indicators for the number of rooms, bedrooms, the decade of construction, plumbing facilities, and the heating system type. States and CZs include indicators for the geographic entities, and demographics is the interactions of the household type (single, a couple, a single parent, or non-related individuals) with indicators for the number of children and adults (separately). The solid line contains weights accounting for all of these which were computed as noted in the text, following DiNardo et al. (1996) and Fortin et al. (2011).

Table 3.5: Implied Relationships with Income Inequality

	Home values	Rents	Consumption equivalent
1970	0.720	0.629	0.755
	(0.133)	(0.048)	(0.054)
1980	0.764	0.690	0.715
	(0.157)	(0.072)	(0.039)
1990	0.365	0.654	0.782
	(0.167)	(0.031)	(0.128)
2000	0.444	0.620	0.789
	(0.124)	(0.037)	(0.074)
2012	0.138	0.626	0.862
	(0.568)	(0.057)	(0.094)

NOTE: Each coefficient is the (re-signed) square root of the absolute value of the regression coefficient relating the variance of the log of household wage income within each CZ with the variance of the log of the variable within each CZ. Standard errors clustered by the state the plurality of the CZ's population lives in are in parenthesis. The unit of observation for each regression is a CZ year and each coefficients is from a different regression. Each is computed using the relevant sample for that variable while income inequality is computed using the universe of all home owners and renters. The third column represents an estimate of the income elasticity of demand for housing services. See tables 3.1 and 3.2 and figure 3.5 for more detail.

APPENDICES

APPENDIX A

Local Ties in Spatial Equilibrium

A.1 Data

A.1

The data comes primarily from the decennial census and ACS as collected by IPUMS at the University of Minnesota (Ruggles et al. (2010)). Data on the impact of trade on individual local labor markets comes from Autor et al. (2013), and the vital statistics data comes from the NBER. I restrict my sample to prime-aged (16-64 inclusive) people not living in group quarters (barracks and dorms). In computing wages I exclude unpaid family workers and workers who did not work for pay last year. Generally, I aggregate these data up to the Commuting Zone (Tolbert and Sizer (1996)) level and perform my analyses at this level, except in some cases where I focused on states to better match Vital Statistics and migration data.

IPUMS

The data from the US Census comes via the IPUMS sample detailed in Ruggles et al. (2010). I use several PUMS samples: For 1970 I use the form 1 one percent sample at either the state or metro level, depending on whether the analysis uses states or commuting zones. For 1980, 1990, and 2000 I use the five percent samples. For 2008 I use the ACS 3 year estimates from 2006 to 2008. For the bulk of specifications I exclude people residing in group quarters, such as military barracks or dormitories. The only exception is the growth accounting by state that I performed in section two. In that case I include people residing in group quarters because this exclusion might cause me to lose young adults born 16 to

21 years earlier. For worker wages I exclude unpaid family workers and only include people who worked last year. In regressions using commuting zone data in 1970 I exclude 1990 commuting zone number 24600 because I suspect its geographic definition was mis-coded.[1]

I also compute "labor supply weights" following Autor et al. (2013) that weight each worker by their total hours worked last year, and I exclude the top and bottom 1 percent of wages from the computation. All wages are deflated using the personal consumption expenditure chain type price index available from the Federal Reserve Bank of St. Louis via their FRED service. The reference year is 2007.

Vital Statistics

I use datasets containing the data from the US vital statistics (National Center for Health Statistics (2014)) that were created by Jean Roth at the NBER and are available publicly in the NBER website. The only cleaning that I perform on the data is to collapse it (weighting by whether it is a 1 of 2 or full sample for the state year combination I am concerned with) at the state level and convert the alphabetic numbering of states to standard fips codes. To exactly match 16 to 21 year olds as of the census data in 1990, I also exclude entries for certain years where the person born would report being 15 or 22 at the time of the survey date.

Population Changes

To investigate demographic changes (births, deaths, aging, and migration) underlying the growth of local areas, I combined data from the US vital statistics with data from the US Census IPUMS to compare the variance of birth rates, aging, and gross migration rates across US states.[2] I focus on people over 16 and under 65.

I compute four quantities that show the main drivers of population changes: Births, people aging out of the population 16 to 65, gross in-migration (immigration), and gross out-migration (emigration). Since the census long form asks about migration relative to where a worker was living five years ago, I focus on each over the past five years to make them comparable. Thus, I focus on births 16 to 21 years before the census date, the total number of people in the state aged 55 to 60 in the previous census, the number of current residents living elsewhere five years ago, and the number of people living in the state five years ago, but living elsewhere now. I focus on changes in population between 1980 and 1990,

[1] Using the commuting zone crosswalks from David Dorn suggests that its population in 1970 was ten times larger than its population in 1980.

[2] This analysis is for US states based on the ease of matching Vital Statistics geographic identifiers to states. In principle this analysis could be done with local labor markets using publicly available data.

which is the first period where I can use readily available data from the Vital Statistics.

Birth Locations

The best available data I have access to concerns workers' states of birth. Unfortunately, this is the most detailed geography that the census bureau asks for, so it is impossible to determine precisely what local labor market a respondent was born in without using an outside data source. Consequently, I tally the proportion of residents of a local labor market who are living in the state of their birth. For large states with many local labor markets (California and Texas are examples) this should lead me to overshoot the proportion of residents living in the area of their birth. For labor markets that cut across state lines (New York for example) I would be understating the proportion of residents living in the same area they were born in since a resident could be in a different state, but the same commuting zone they were born in. On the one hand, imprecision in the measure of the proportion of residents born in the same local labor market is a concern. It is important to note that in a world where areas are not unique islands, the ideal geographic construct may be different in terms of work, family, and other considerations. For example, a worker may prefer to live further away from her parents compared with her work and a worker living in the same state but a different commuting zone as they were born might be almost as constrained as a worker living in the same commuting zone.

Another other issue with the variable is that most births are in hospitals and sometimes children will be born in a hospital in a different state from where their mother lives. Bartik (2009), for example, documents this using data from the PSID. In this situation, the question appears to ask for the state of the hospital, which is a poor proxy for the concepts I am examining. While this variable is far from perfect, its concordance with other measures of a respondent's local "ties" such as their tenure in their home should suggest that it is still meaningful for this application.

For all of these reasons, I include alternative specifications that use alternative measures of local ties. Generally these results are quite similar.

Local Labor Markets

I define a local area for this project as a Commuting Zone (CZ) defined by Tolbert and Sizer (1996). Commuting Zones are designed to reflect local labor markets where workers live and work, based on commuting data collected in the 1990 Census.[3] A given CZ can contain

[3]Different Commuting Zones exist following the 2000 census, however I keep with Autor et al. (2013) and use the 1990 definitions. I do this to keep CZ definitions constant and I use 1990 because it reflects local areas at the beginning of the sample.

multiple states and states can contain multiple CZs. CZs are quite similar to Metropolitan Statistical Areas (MSAs) that are more commonly used, but CZs also include rural areas, covering the entire area of each of the 50 states. They are constructed to be an ideal analogue to the areas in traditional models of migration where workers live and work in the same area. To merge the IPUMS data I use in my specification I use the crosswalks created by David Dorn and available via his academic website. For historical charts, I exclude commuting zone 24600, which I believe may be improperly coded in 1970.

A.2 Growth due to Migration and Natural Changes

As a first step, I decompose of changes in local population into migration and natural changes (births and aging) across states. Table A2 calculates the components of working age (16-64) population changes by continental US states over the period from 1980 to 1990. It uses estimates of migration from 1985 to 1990 from the 1990 census, age structure information from 1980, and birth data from the vital statistics to compare flows due to migration and natural changes.[4]

The main implication of the decomposition in Table A2 is that migration is much more important than natural changes in terms of changes in population across the United States. The standard deviation of net migration across continental US states is nearly four times as large as that of natural changes (13.9 against 3.6 percent). Gross migration into areas also varies much more than any other component, suggesting large differences in areas' abilities to attract workers born elsewhere. This is despite fertility and mortality driving changes in aggregate population. A similar exercise by Berry and Dahmann (1977) produced similar results.

To provide more evidence about the influence of migration relative to natural changes I plot the ratio of gross changes due to migration over gross changes due to natural causes in Figure A2. The figure has two panels: the first plotting flows that increase population, and the second plotting flows that decrease population. In each case a higher value of the ratio means that migration has a larger contribution to a given state's population dynamics. On the x axis I plot the state's log change in population, to show how the importance of migration varies with a state's population growth over the period.

The plots in Figure A2, supporting the earlier evidence, show that migration is more

[4] I perform the decomposition across states for convenience and because it coincides with my current measure of a person's place of birth. I hope to perform it at the level of commuting zones in a later version based on more detailed data about people's birth places. Additionally, I focus on an age range with relatively low mortality rates in the United States, so differences in morality rates across states should not be large enough to affect the results. More information about the datasets and methodology is contained in the data appendix.

important than natural changes for states that grew. On the left panel, higher population growth is clearly associated with a higher ratio of immigration to births. The relationship is weaker for decreases in population, though it does appear that out migration is important growing states as well as declining ones. For decreases, there is almost no relationship, though a slight positive association emerges if Wyoming is removed from the analysis.

The result that migration drives local growth is consistent with economists understandings of both of migration and fertility. Standard models of migration (e.g. Rosen (1979) and Roback (1982)) predict that areas with attractive amenities will gain population. On the other hand, models of fertility have little to say about fertility in one place or the other. Perhaps the closest connection is through a possible income effect, where richer parents will choose to have more children, so long as children are a normal good. A problem with this argument is the emphasis by Becker (1960) on the quality-quantity trade of in child rearing. For example, Willis (1973) and Becker and Lewis (1973) argue that while parents in areas with higher wages will tend to spend more on children, this may lead them to invest more in the "quality" than the quantity of their children. Previous studies, such as Lindo (2010) and Black et al. (2013), have found that local shocks have relatively small (positive) impacts on fertility.

A.3 Sufficient Statistics Derivation

The following are the equilibrium conditions and the effect of a change in g_j on welfre for each actor in the model presented in the main paper. Note that I ommit the 1 and 0 subscripts, except for equations that involve the indirect compensation function. This is for simplicity, and since the equations would apply in either scenario.

Household

Household i maximizes utility subject ot a budget constraint:

$$\max_{j,c_j,h_j} u(c_j, h_j, a_j) + \xi_{ij}$$
$$\text{s.t. } g_j + w_j = c_j + r_j h_j$$

Where c_j is the level of tradable consumption in area j, h_j is its housing (non tradable) consumption in j, a_j is the local amenity level, w_j is the wage, r_j is the rent, and g_j is the net governmental transfers. ξ_{ij} is an arbitrary distribution of areas specific preferences household i for area j. Workers inelastically provide labor, though this can also be relaxed.

The first order conditions are:

$$\frac{\partial u}{\partial c_j} = \lambda_j$$
$$\frac{\partial u}{\partial h_j} = \lambda_j r_j$$

Where λ_j is the marginal utility of consumption in area j.

To have comparability between households and other ators, I measure households' welfare using an indirect compensation function (Varian (1984)), using initial prices. Chipman and Moore (1980) calls this "generalized equivalent variation." The compensation function is defined in the model as $m_i(w_1, r_1, g_1; c_{0,j}, h_{0,j}, g_{0,j}, j_0) \equiv e(w_1, r_1, g_1; u(c_{0,j}, h_{0,j}, g_{0,j}) + \xi_{ij_0})$ where $e(\cdot;\cdot)$ is the more common expenditure function, giving the expenditure necessary to equal the initial level of utility, and the subscripts 0 and 1 denote initial values (0) and values after some change (1).[5] To get to the indirect compensation function, one simply replaces the initial level of utility with the indirect utility function, at initial prices:

$$m_i(w_1, r_1, g_1; w_0, r_0, g_0) \equiv e(w_1, r_1, g_1; v(w_0, r_0, g_0))$$

After a change in g_j, each of the arguments in the function change. Luckily the envelope theorem applies (Small and Rosen (1981), Kline and Moretti (2014b), and Chetty (2009)) so the change in utility is equal to the change in expenditures at the initial levels of consumption, holding location (j_0) fixed:

$$\frac{dm_i(w_1, r_1, g_1; w_0, r_0, g_0)}{dg_{1,j}} = \mathbb{1}(j_0 = j) + \frac{dw_{1,j_0}}{dg_{1,j}} - h_{0,j_0}\frac{dr_{1,j_0}}{dg_{1,j}}$$

The increase in the subsidy is the first term, but it only appears if the subsidy would apply in the area the household was living in initially. The second term is the change in earnings, and the final is the effect of the change in rents, given the initial level of housing consumption. There is no effect on tradeable good consumption, since it is the numeraire.

Landlords

Landlords make the difference between the rent that they charge and their cost of providing housing, $c_j(H_j)$ where H_j is the total amount of housing in areas j. $c_j(H_j)$ is (quasi) monotonically increasing, since land becomes increasingly costly to develop into good housing

[5]If utility is linear in income, as in Busso et al. (2013) and Kline and Moretti (2014b), then the compensation function, at current prices, is the utility.

as there is less and less available land.

$$\pi_j^H = \max_{H_j} r_j H_j - \int_0^{H_j} c_j(x)dx$$

This gives a simple FOC:

$$r_j = c_j(H_j)$$

Landlord welfare is π_j^H. After a change in subsidies, total profits will change in the following way (based on the envelope theorem):

$$\begin{aligned}
\frac{\partial \sum_{j'} \pi_{j'}^H}{\partial g_j} &= \textstyle\sum_{j'} \frac{\partial \pi_{j'}^H}{\partial r_{j'}} \frac{\partial r_{j'}}{\partial g_j} \\
&= \textstyle\sum_{j'} H_{j'} \frac{\partial r_{j'}}{\partial g_j}
\end{aligned}$$

Local Firms

A local firm produces local output Y_j and sells it at price p_j. It employs local labor and buys capital on a national market at interest rate ρ. The firm's profit, which it maximizes, is:

$$\pi_j^Y = \max_{N_j, K_j} p_j Y_j(N_j, K_j) - w_j N_j - \rho K_j$$

So, the first order conditions are:

$$\frac{\partial Y_j}{\partial L_j} = \frac{w_j}{p_j}$$
$$\frac{\partial Y_j}{\partial K_j} = \frac{\rho}{p_j}$$

The welfare of firms are their profits, π_j^Y. Again, because of the envelope theorem, they are only affected in terms of prices after a subsidy into the local area:

$$\begin{aligned}
\frac{\partial \sum_{j'} \pi_{j'}^Y}{\partial g_j} &= \textstyle\sum_{j'} \frac{\partial \pi_{j'}^Y}{\partial w_{j'}} \frac{\partial w_{j'}}{\partial g_j} + \frac{\partial \pi_{j'}^Y}{\partial p_{j'}} \frac{\partial p_{j'}}{\partial g_j} \\
&= Y_j \frac{\partial p_{j'}}{\partial g_j} - N_{j'} \frac{\partial w_{j'}}{\partial g_j}
\end{aligned}$$

The Final Goods Firm

There is a national firm that takes local output Y_j and produces the numeraire tradeable consumption good (Y).

$$\pi^Y = \max_{\text{all } Y_j} Y(Y_1, Y_2, \ldots, Y_J) - \sum_{j'} p_{j'} Y_{j'}$$

So, the first order conditions are, for all j areas' goods:

$$\frac{\partial Y}{\partial Y_j} = p_j$$

The welfare of these firms are their profits, π^Y. π^Y may be affected by the subsidy to the areas, since these firms have to pay for their inputs:

$$\begin{aligned} \frac{\partial \pi^Y}{\partial g_j} &= \textstyle\sum_{j'} \frac{\partial \pi^Y}{\partial p_{j'}} \frac{\partial p_{j'}}{\partial g_j} \\ &= -Y_j \frac{\partial p_{j'}}{\partial g_j} \end{aligned}$$

Aggregation

I aggregate the welfare of households, landlords, and each type of firm by measuring each in monetary terms and adding each up. There is a continuum of households, and $N_{0,j}$ of them live in area j initially, I need to cumulate the effect across all of these households. There is one (representative) landlord and local firm per area, and only one national firm. Adding these up gives the formulation in the main paper. The welfare result, similarly, can be obtained by simply adding up all of the effects on the welfare of various actors.

A.4 Factors that Lead to Local Ties

Residents' preferences for their places of birth, much like people's preferences about living in any location, represent many factors. Disentangling these factors is an active and interesting area with many recent contributions. For example, Kennan and Walker (2011), Coate (2017), and Diamond (2016) use structural micro-economic models to estimate residents' preferences for different areas, including their birth places. Somewhat atypical to this literature, I have remained mostly agnostic about the specific factors that lead workers to have the preferences that I measure, since my basic results do not rely on these distinctions.

My reliance on cross sectional snapshots of the population also makes if difficult to credibly disentangle different factors.

The underlying reasons for these measured preferences have implications for possible policy responses, including alternatives to place-based policies. One possibility is that the preferences that I measure are the result of other frictions that policies may remove. In this way it may be possible to change their magnitude and move workers out of declining areas. For example, if workers reside in their birth places because of mobility costs, then it may be cost effective to pay workers to migrate. It may also be possible to move a given community, including most of its members, to a new place with higher productivity in keeping with the spirit of the literature on "dynamic mobility."

In this section I briefly discuss several underlying factors that may lead to these measured preferences. I discuss literatures on social networks, literal migration costs, frictions in the housing market, information frictions, and endogenous human capital formation. I also include a brief discussion of endogenous preference formation. My tentative conclusion is many different stories are consistent with the preference for home areas that I and other researchers have measured. It seems unlikely that inexpensive interventions will induce workers to move out of declining areas much more than they already do. A constructive path for future research would be to directly study these pathways using more detailed micro data.

Networks of People

Some of the most valuable connections people have in a place is the collection of people that they know. These ties are especially strong with parents and members of a person's nuclear family at various points in their lives. Aging parents, in particular, represent a strong tie to particular local areas, as shown by Konrad and Kunemund (2002), Hank (2007), Rainer and Siedler (2009), and others. Friends may also have a meaningful influence. For example, Topa (2011) summarizes a literature on networks and job referrals. Even if workers do not obtain jobs from their contacts, they may rely on friends and family for informal insurance in difficult times (e.g. Kaplan (2012), Huttunen and Salvanes (2015)), or for information about particular services and opportunities that are available locally. Forming new friendships involves significant effort and even if such effort is expended, the returns are uncertain and it can be difficult to form relationships that are as viable as relationships that people gave up by moving.

The most relevant economics literature on this phenomenon is the literature examining "dynamic mobility." Carrington et al. (1996) shows that large scale migration tends to follow pattern where there tend to be trailblazers who establish links between a sending and destination community. Significant migration occurs only after these links have been formed. They show that this is apparent in the great migration of African Americans from the south to the north of the United States and Yannay Spitzer (2015) shows a similar result for Eastern European Jews migrating to the United States. This pattern emphasizes that network links between sending and destination communities are important, and that people tend to migrate in ways that leverage their social networks.[6]

If networks are important, then a lack of migration might be because residents lack connections in desirable destinations. It is unclear how this limitation may change with the population dynamics I describe. People may be cutoff in areas where they primarily encounter people born in the same place, since this will tend to make their networks more locally focused. However, if links are formed by trailblazers from sending communities, then the opposite may be the case. An area may reach a tipping point where there are enough people who have migrated to establish a second community in another place.

Another explanation is that people form networks with different numbers of local and non-local links over time. Some people's networks will be local and others will naturally have many links that come from elsewhere. If the proportion of people with local and non-local links is more or less fixed in growing and declining areas, possibly because the two factors above tend to balance out, then this could generate the patterns I find. People with external links may be more likely to migrate, but people with only local links will be unlikely to. Another scenario might be that some people are highly reliant on networks while others are not, or are able to easily form new ones in destination communities. These differences may be correlated with socioeconomic status, but there may also be other important factors.

Migration Costs

One explanation is that workers face different mobility costs in different areas that I am not modelling. These different mobility costs may be due to literal differences in the cost of hiring movers or selling a house because of different wage rates, regulation, capital costs, and market thickness in different areas. Credit may also be more or less available, or needed, in

[6] As this literature shows, racial and ethnic segregation plays an important role in migration. It is difficult, however, to disentangle how segregation would impact people's local ties or my setup. Minority groups may have a smaller choice set of locations that are open to them – either because of discrimination or because of personal preferences. Bound and Holzer (2000) finds some evidence that less educated blacks are less likely to migrate after a local demand shock. They also note that gross migration rates are lower for blacks, even within education groups. Cadena and Kovak (2016) find that immigrants, who often live in segregated communities, migrate in greater numbers than natives. In my model, a lack of immigrants appears to be lessening migration rates in declining areas since immigrants have weaker local ties. Another consideration is that specific minority groups may have preferences about locations that are distinct from other groups. This will lead particular minority groups to be more likely to migrate if these places become desirable after shocks to their original locations.

areas with different shares of locally born workers.

There are a number of reason why literal mobility costs should play a limited role. The first is the length of the periods that I examine. In the empirical regressions I use a roughly 10 year time frame. This means both that workers have adequate time to save for a move, and also that the move is a small percentage of their consumption over the entire period, so the benefits should be quite large in comparison with the costs. Notably, Blanchard and Katz (1992), suggest that migration responses take roughly 10 years to play out.

A second reason why literal mobility costs may be less important is because gross migration is much larger than net migration. Since people often string together multiple moves, it implies that literal migration costs are a relatively small factor in their decision making. In particular, a person who lives in their state of birth is fairly likely to have lived outside for at least some time. For example, Kennan and Walker (2011) find that roughly 1/4 of all moves in the NLSY79 are back to a respondent's state of residence at age 14. Decennial census data tells a similar story. Among people living in their state of birth in 2000, nearly half had moved at some point over the past five years, and 3.5 percent had moved home from somewhere outside of the state. While the second number is modest, it is much larger than the effects that I observe over a period that is twice as long.

A third reason is that, because moving involves paying costs in two separate places, there is a limit on the extent that moving costs can differ by sending areas. Since much of the financing of the move should occur in the area that a worker moves to, not the place they are moving from, a worker's chosen destination should matter more in terms of credit availability than the area they are moving from. Also, even if mobility costs were only paid in the sending area, costs are likely to be lower in declining areas. Wages will tend to be lower, suggesting that movers will cost less, and the typical six percent commission on a real estate sale will also tend to be cheaper in levels.

A final piece of evidence on this channel is the study by Huttunen and Salvanes (2015) examining moves by displaced workers. Their finding that recently displaced (fired or laid off) workers migrate in greater numbers than a control of non-displaced workers suggests that other factors outweigh mobility costs when workers face large earnings losses. They also find that workers tend to move closer to their parents.

Housing Frictions

Several theories postulate that home ownership might tie people to local areas in ways besides the impact of durable housing, which I addressed earlier. Homeowners may be less likely to move after a decrease in local housing prices because loss aversion makes them less willing to suffer the capital loss or because it makes it harder for them to afford a new down payment. This effect may dominate the increased number of foreclosures in areas with declines in home prices, and negative economic shocks more broadly. Another theory, originally advanced by Oswald (1996), is that the transaction cost of selling a house are large enough depress migration, and therefore increase the unemployment rate.

One concern is that the local ties that I observe are entirely explained by these frictions in the housing market, perhaps as brought about by larger declines in housing prices in declining areas. Tables A9 and A10 address this concern by allowing the effect of a local labor demand shock to vary based on local ties and the level and lagged changes in rents. I find that my results are similar and in some cases stronger.

The finding that ties appear stronger than changes in housing prices is not surprising given previous literature on housing lock in. Evidence on the effects of house price declines on migration is mixed.[7] One fairly consistent finding, however, is that effects on the labor market are small. Modestino and Dennett (2012) squares this with their finding of meaningful effects on migration among homeowners by noting that only about 20 percent of migrants are homeowners. Even if home owners are much less likely to move, a somewhat controversial claim, then renters may undo the labor market impacts.

More broadly, there are many reasons to suspect that owning a house will tie people to individual areas. The most obvious is the high transaction costs that are associated with selling a house. This effect is complicated, however, by the fact that people choose to become homeowners knowing these costs. Since many people are inframarginal about migration, and migration is more common among younger households who are less likely to own, the effect may be small. In terms of determining a causal effect, this selection issue has complicated previous investigations, which have produced mixed results.[8] Housing frictions may amplify other ties that people have to places, but available evidence about them is mixed.

Information

Information about relevant alternatives is an important factor in the migration process. The recent literature on migration, following Kennan and Walker (2011), places a great deal of emphasis on information and its ability to explain repeated moves. Within a framework of net migration, limits on information may increase effective mobility costs of people who decide to move. These factors likely interact with other preferences about homes, since

[7]Farber (2011) and Bricker and Brian Bucks (2013) find small or no effects of house price decreases on migration, Molloy et al. (2011) and Donovan and Schnure (2011) find effects only for small distance moves, and Henley (1998), Ferreira et al. (2010), and Modestino and Dennett (2012) find some effects for longer distance moves. Farber (2011), Modestino and Dennett (2012), and Valletta (2013) find that homeowners' negative equity does not appear to affect the labor market.

[8]Coulson and Fisher (2009) is a recent investigation that contains a review of earlier studies.

Kennan and Walker (2011) and Gregory (2013) find that people's preferences about their homes are not completely explained by information frictions. Schmutz and Sidibe (2016) argue that frictions in job finding rates for movers can explain more, though they still find that mobility costs are substantial.

Levels of information may vary across places or among people in a given place giving rise to differences in people's choices. Particular places might have worse information because of physical geography, economic geography, migration patterns, vacation patterns, availability of information resources like the internet, or other factors. Declining areas could be more isolated than other areas, so this may influence the process. Another possibility is that different people within an area are differentially informed about outside opportunities and that this is relatively constant across areas. If this were true then it might behave in a way similar to a difference in costs across residents. It seems that most implications of information for net migration flows will be modeled in my framework, though with different interpretations.

Establishing how well information about far off alternatives circulates to particular areas appears to be an interesting target for future research. One factor that may limit information's implications in this context is that there are several incentives for "arbitrage" on the part of employers if workers are misinformed in a particular area, however.

Endogenous Human Capital

Another explanation is that workers have location specific human capital. This area specific human capital could be job related skills, or information about the local area. This could mean that workers prefer to leave their birth states, but would suffer large wage losses if they did.

There are several reasons to suspect that workers have substantial location specific human capital. Topa (2011) summarizes a large literature that emphasizes the importance of local social ties for job referrals. These may take many years to form. Local knowledge might also be particularly valuable if workers interact with specific laws, regulations, business structures, contacts, or natural features.

An emphasis on local human capital, however, is somewhat counter-intuitive given the fairly robust finding that college educated workers are significantly more likely to migrate than less skilled workers. A literature on job specific human capital, including Blatter et al. (2012) and Hudomiet (2014), that suggests that job specific human capital is higher for more skilled workers. In particular, agglomeration effects should lead workers to be less mobile if they have skills that are industry specific and the industry is highly concentrated. An example of this might be an investment banker in New York.[9]

Location specific human capital might be inversely correlated with overall human capital if it is an inferior good. Unions have historically been more important for low skilled workers and union jobs may be controlled on the basis of local connections. These jobs may be more valuable for less skilled individuals; for example, this might be true if unions tend to compress the wage distribution. Local licensing also may bind more in the low skilled labor market. Kleiner and Krueger (2013) suggest that licensing may be moving in such a way that the rise of licencing and fall of unions roughly offset each other in terms of the number of workers affected by either.

There seems to be very limited scope for policies to affect workers' levels of location specific human capital. Referral networks appear to serve a valuable role in conveying information about job applicants and location specific skills are by definition valuable for productivity in particular local areas. Some political institutions, like recent state licensing laws, may be possible to change, but political economy considerations may complicate these efforts. Voters may reward local politicians who enact policies that advantage people with local connections.[10]

Endogenous Preferences

Workers may develop preferences for particular places based upon spending time in a place. In particular, many activities involve fixed cost that may have to be paid again if one were to move to a new location. Children may participate in activities, play certain sports, cheer for particular sports teams, or eat certain foods that may only be popular in particular places. For children in particular, some of these affiliations may be malleable, since many take up different sports as adults, but preferences for food may be much more fixed in adulthood. Adults as well may develop particular local affiliations, such as membership in local clubs, knowledge about certain local features like hiking trails, local resources like bookstores, local community groups, or local activities. Patterns of dressing, tastes for home styles and decor, language differences, political affiliations, and many other cultural factors likely contribute to this.

Endogenous preference formation most likely interacts with other factors. So, for example, many social relationships may be based around particular local activities. Individual

[9] Agglomeration effects in training, for example in PhD Economists, may make skilled workers more mobile since they often spend extended periods in unfamiliar locations. Workers in these occupations experience many different areas so they might be expected to have fewer ties to individual areas.

[10] Anthropologist Scott (1998) argues that there is a fundamental political tension between local and central control. This can lead to policies that are intentionally designed to make local knowledge more valuable. He argues that the process of constructing a nation state like France involved replacing these policies. He argues that this process has had many unintended negative consequences.

participants may not enjoy specific activities more than alternatives in other places, but they value the social interactions that they get out of these activities, and these might be hard to form elsewhere. Married couples may be tied down by one spouse's like for particular local rituals. Children may find it optimal for parents to move, but parents may be tied to particular places by their preferences for particular local activities. Workers may also be more willing to invest in location specific human capital or make decisions that increase their mobility costs if they enjoy living in their current area.

Policy Implications

The main policy implication from this discussion is that it may be possible to remove some of the reasons for some people's unwillingness to migrate. If it were cost effective to induce people to migrate from declining areas without decreasing their welfare by very much, then this would obviously be a solution to a host of local economic problems.

The cheapest friction to eliminate would be a direct mobility cost driven by limited access to credit. Unfortunately, available evidence suggests that this mechanism is unlikely to drive most of the effects that I document. Access to credit could be beneficial in a number of other ways – such as encouraging education, small business formation, and other productive investments – but it appears unlikely to be much more valuable in this area than in others. A reasonable first step in this area would be to allow easier transfer of benefits between different state programs. While available evidence suggests this might have a small impact, it presumably can be done at relatively low cost and it is difficult to imagine how it would harm welfare.

Another possibility would be to create links from declining areas to other growing areas. This could involve encouraging the migration of influential "trailblazers" and their continued integration in sending communities, or by establishing agencies devoted to establishing workers in other places. A problem with this approach is that it would be difficult for any governmental agency to properly assess needs for migration services. It also is not clear that a market failure is at work – an employer in a growing area, for example, would have an incentive to hire workers from declining areas at cheaper wages if it were to recruit in those areas. Such programs might also face political pressure in sending communities since they would be designed to de-populate them and reduce their influence. More subtle policies, such as the integration of local employment agencies, or the standardization of state level credentials would be likely to help. An additional benefit is that these would also improve labor market "fluidity" across all areas, and not necessarily only declining ones.

Many of the other explanations for people's preferences involve factors that are either costly or impossible to adjust. If all preferences are formed in childhood, for example, then an area would only gradually lose its appeal as a smaller and smaller proportion of the population grows up there. The system would eventually return to a "steady state" equilibrium where population reflects the common valuations of productivity and consumption amenities that are common in a Rosen-Roback framework, but this evolution would be much slower than is commonly assumed. Compensatory place-based policies might slow this evolution, but since they do not change population by very much, they are likely to have small effects.

Appendix Tables and Figures

Table A1: Association Between Populations of Locals and Outsiders

	Pct chg in population born locally			
Pct chg in population born outside	0.36	0.28	0.22	0.16
	(0.05)	(0.04)	(0.04)	(0.03)
Weighted	No	No	Yes	Yes
Controls	No	Yes	No	Yes
Observations	722			
R^2	0.287	0.434	0.215	0.483

Notes: Coefficients are from a regression of changes in the population born locally on the change in population born outside over the period from 1980 to 2008 for the 722 commuting zones in the continental US. Each is measures as a percentage of the initial population (including all people). Data are from the long form 1980 decennial Census and the 2006-2008 ACS. Data are are weighted to be nationally representative. Locals are people who are born in the state they are living in, while outsiders are born in other states or countries.

Table A2: Components of Population Changes from 1980 to 1990

		StD	Mean	N
Net	Migration	13.78	4.02	48
	Natural changes	3.69	9.96	48
Gross	Immigration	14.82	29.73	48
	Births	2.81	24.55	48
	Emigration	8.97	25.71	48
	Aging	1.48	15.66	48

Notes: Standard deviations and means are expressed as a percentage of the initial population for all continental US states with equal weights. For example, a state with $100,000$ births and $1,000,000$ in initial population would have a value of 10 percent for births. Data are from the decennial census and vital statistics (National Center for Health Statistics (2014)) covering the continental United States. Births are from 1969 to 1974, aging is the population 55 to 60 in the 1980 decennial census, and migration statistics are from the 1990 census. Migration includes moves from abroad but not moves from the state to abroad, since the sample only includes people who are in the United States when the census was conducted. Each is multiplied by two to represent total population movements over 1980 to 1990. Net migration is immigration minus emigration and "natural" changes are births minus aging out of the age range. Immigration and emigration are relative to states, not countries, and population in this context is the population of people aged 16-64.

Table A3: Persistence of Population Changes

	All	Large	All	Large
Lagged pct chg in population	0.52	0.53	0.35	0.29
	(0.06)	(0.08)	(0.03)	(0.13)
Twice lagged pct chg in population			0.13	0.08
			(0.02)	(0.06)
Thrice lagged pct chg in population			0.04	0.13
			(0.01)	(0.07)
Observations	1444	48	721	16
R^2	0.569	0.678	0.702	0.831

Notes: Results are from an autoregression of changes in population on lags of itself. "All" denotes results using all commuting zones, "Large" denotes commuting zones that had populations of more than 1 million people initially. Data is from the decennial census and ACS. Regressions are weighted by initial population and standard errors in parenthesis are clustered by state (a CZ is in a state if the plurality of its population resides there). Year fixed effects are included for panel regressions.

Table A4: Locally Born Workers Staying and Population Changes

	Percent of people born in the state staying	
1970-2008 log change in working age population	0.09	0.12
	(0.04)	(0.02)
Controls	No	Yes
Observations	48	48
R^2	0.170	0.762

Notes: Data is from the decennial census and ACS and cover the continental United States. Regressions are weighted by initial population and robust standard errors are in parenthesis. All share variables are multiplied by 100 to make them into percentage points. Controls are share college educated, share employed, share foreign born, share born in Mexico, and log population – all measured in 1970. The share of workers born in the same state includes all adults 16-65 born in that state and living somewhere in the United States from 2006-2008 (the ACS 2008 3 year sample window).

Table A5: Associations Between ADH and Bartik Instruments

	Bartik	ADH trade exposure	ADH IV
Bartik	1.00		
ADH trade exposure	0.21	1.00	
ADH IV	0.26	0.73	1.00

Note: Correlation coefficients are shown between instrumental variables related to Chinese import competition and Bartik labor demand instruments. The table describes the correlation between a given CZs Bartik instrument for 1980 to 1990 against its Chinese import exposure for 1990-2000, then again for 2000-2008. I use population weights at the beginning of the period relevant for the Chinese import shock.

Table A6: Bartik Shocks by Share Born Locally: Men

Panel A: Bins Specification

	Pop	NILF	Unemp	Emp	Wages	LFP
Bartik: Low ties	1.85	0.02	-0.07	1.85	0.37	0.15
	(0.47)	(0.06)	(0.05)	(0.41)	(0.24)	(0.03)
Bartik: High ties	0.46	-0.07	0.05	0.51	0.32	0.10
	(0.28)	(0.03)	(0.03)	(0.28)	(0.23)	(0.04)
P-val: No diff	0.01	0.20	0.05	0.01	0.89	0.32
R^2	0.57	0.20	0.64	0.53	0.28	0.35
Observations			722			

Panel B: Triple Difference Specification

	Pop	NILF	Unemp	Emp	Wages	LFP
Interaction	-3.70	-0.17	0.10	-3.50	0.97	-0.16
	(1.15)	(0.17)	(0.12)	(1.00)	(0.74)	(0.11)
Main effect	1.10	-0.05	0.01	1.13	0.33	0.13
	(0.25)	(0.03)	(0.03)	(0.24)	(0.17)	(0.03)
Percent locals	0.30	-0.03	-0.01	0.31	0.02	0.03
	(0.23)	(0.03)	(0.02)	(0.21)	(0.16)	(0.03)
R^2	0.59	0.22	0.63	0.55	0.31	0.35
Observations			722			

Notes: CZ level results with statistics including only men aged 16-65. Weighted by initial population with clustered (by state) standard errors and controls as in table 1.3. See table 1.3 for full notes.

Table A7: Bartik Shocks by Share Born Locally: Women

Panel A: Bins Specification

	Pop	NILF	Unemp	Emp	Wages	LFP
Bartik: Low ties	2.48	1.01	-0.02	1.29	0.05	-0.07
	(0.70)	(0.35)	(0.06)	(0.36)	(0.23)	(0.05)
Bartik: High ties	0.63	0.07	0.05	0.34	0.15	0.05
	(0.37)	(0.18)	(0.03)	(0.20)	(0.17)	(0.04)
P-val: No diff	0.02	0.02	0.26	0.01	0.72	0.06
R^2	0.59	0.55	0.62	0.49	0.45	0.40
Observations			722			

Panel B: Triple Difference Specification

	Pop	NILF	Unemp	Emp	Wages	LFP
Interaction	-4.97	-2.31	-0.07	-2.63	0.79	0.20
	(1.60)	(0.87)	(0.12)	(0.82)	(0.58)	(0.15)
Main effect	1.46	0.47	0.03	0.76	0.07	0.00
	(0.33)	(0.16)	(0.03)	(0.18)	(0.15)	(0.03)
Percent locals	0.35	0.12	0.02	0.17	-0.08	-0.02
	(0.29)	(0.15)	(0.02)	(0.15)	(0.13)	(0.03)
R^2	0.62	0.57	0.61	0.52	0.46	0.39
Observations			722			

Notes: CZ level results with statistics including only women aged 16-65. Weighted by initial population with clustered (by state) standard errors and controls as in table 1.3. See table 1.3 for full notes.

Table A8: Bartik Shocks by Local Average Household Tenure

Panel A: Bins Specification

	Pop	NILF	Unemp	Emp	Wages	Rents	LFP
Bartik: Low ties	2.36	0.42	-0.07	1.88	0.28	0.88	0.10
	(0.42)	(0.14)	(0.06)	(0.31)	(0.23)	(0.34)	(0.04)
Bartik: High ties	0.43	-0.01	0.03	0.40	0.34	0.15	0.07
	(0.27)	(0.08)	(0.03)	(0.21)	(0.19)	(0.26)	(0.03)
P-val: No diff	0.00	0.01	0.11	0.00	0.83	0.12	0.52
R^2	0.63	0.60	0.68	0.56	0.37	0.55	0.38
Observations				722			

Panel B: Triple Difference Specification

	Pop	NILF	Unemp	Emp	Wages	Rents	LFP
Interaction	-0.55	-0.14	0.01	-0.40	0.10	-0.06	-0.00
	(0.11)	(0.03)	(0.01)	(0.08)	(0.09)	(0.14)	(0.01)
Main effect	1.16	0.11	-0.02	1.00	0.34	0.48	0.11
	(0.28)	(0.07)	(0.03)	(0.23)	(0.18)	(0.23)	(0.03)
Avg time in house	2.69	-0.32	-0.52	2.88	0.51	2.85	0.47
	(1.69)	(0.58)	(0.24)	(1.29)	(1.33)	(1.70)	(0.22)
R^2	0.65	0.66	0.69	0.57	0.40	0.56	0.41
Observations				722			

Notes: Regression coefficients are plotted for either the main effect plus a linear interaction term with the demeaned average household tenure in the CZ, or the coefficient separately estimated for CZs with fewer or more than 8 years of average household tenure. Controls, measured in 1980, are: the household tenure variable used in the interaction term, the share of working age adults outside the labor force, unemployed, foreign born, having entered the state in the past five years, and the share of adults who are under 35 and 50 to 65. Results are weighted by initial population with clustered (by state) standard errors and controls as in table 1.3. See table 1.3 for additional notes.

Table A9: Bartik Regressions Including Other Interactions

	Log population						Log NILF					
Ties interaction	-2.38	-4.24	-4.82	-6.90	-5.51	-5.22	-0.47	-1.10	-1.25	-1.93	-1.52	-1.35
	(1.21)	(1.34)	(1.10)	(1.13)	(1.10)	(1.84)	(0.44)	(0.44)	(0.35)	(0.29)	(0.33)	(0.53)
Bartik shock	0.91	1.24	1.02	1.02	0.91	1.27	0.10	0.20	0.14	0.13	0.09	0.20
	(0.32)	(0.28)	(0.29)	(0.29)	(0.28)	(0.26)	(0.08)	(0.08)	(0.08)	(0.08)	(0.07)	(0.06)
Percent locals	-0.10	0.32	0.45	0.60	0.49	0.51	-0.05	0.05	0.09	0.14	0.11	0.09
	(0.19)	(0.25)	(0.24)	(0.20)	(0.20)	(0.23)	(0.08)	(0.08)	(0.07)	(0.06)	(0.06)	(0.06)
Pct under 35 interaction			27.13						10.05			
			(15.44)						(4.46)			
Pct 50 to 64 interaction			38.09						12.60			
			(13.23)						(4.49)			
Pct college interaction				-9.01						-2.81		
				(2.56)						(0.75)		
Pct employed interaction					-15.23						-5.09	
					(3.68)						(1.13)	
Rents interaction						-1.49						-0.31
						(2.09)						(0.68)
Pos rent chgs interaction						5.09						1.75
						(1.89)						(0.56)
Neg rent chgs interaction						3.63						1.55
						(3.29)						(1.32)
Main Controls	No	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
Observations				722					722			
R^2	0.474	0.604	0.629	0.626	0.637	0.621	0.313	0.547	0.576	0.575	0.593	0.591

Notes: Regressions, as explained in the text, using shift share (Bartik) demand indexes and data from the decennial censuses in 1980 and 1990 for individual commuting zones in the US. In addition to the setup described in table 1.3, these also control for the main effect of the interaction term, if it is not already included.

Table A10: Trade Regressions Including Other Interactions

	Log population						Log NILF					
Ties interaction	-5.87	-2.01	-3.66	-4.91	-2.44	-5.58	-4.80	-3.84	-3.56	-4.78	-3.96	-4.61
	(2.69)	(2.50)	(2.74)	(2.58)	(2.33)	(3.28)	(1.01)	(0.97)	(1.06)	(1.18)	(0.92)	(1.28)
Import shock	0.41	0.43	0.42	0.37	0.45	0.42	-0.33	-0.34	-0.32	-0.36	-0.33	-0.31
	(0.61)	(0.43)	(0.51)	(0.47)	(0.43)	(0.54)	(0.15)	(0.16)	(0.16)	(0.16)	(0.15)	(0.14)
Percent locals	-0.39	-0.39	-0.43	-0.47	-0.41	-0.46	-0.17	-0.21	-0.20	-0.23	-0.22	-0.25
	(0.07)	(0.10)	(0.11)	(0.11)	(0.10)	(0.11)	(0.02)	(0.04)	(0.04)	(0.05)	(0.04)	(0.05)
Pct under 35 interaction			-3.47						-1.71			
			(27.94)						(7.85)			
Pct 50 to 64 interaction			10.99						-4.15			
			(27.52)						(7.72)			
Pct college interaction				-6.79						-2.19		
				(3.78)						(1.41)		
Pct employed interaction					-8.78						-2.38	
					(5.36)						(2.22)	
Rents interaction						-1.77						-1.15
						(1.56)						(0.43)
Pos rent chgs interaction						-3.68						1.38
						(2.96)						(1.40)
Neg rent chgs interaction						8.90						2.69
						(5.04)						(1.81)
Main Controls	No	Yes	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
Observations	1444						1444					
R^2	0.298	0.485	0.491	0.491	0.488	0.522	0.533	0.629	0.629	0.632	0.628	0.668

Notes: Regressions using a shift share index of trade in final goods' impact on local commuting zones in the US in first differences for the periods from 1990 to 2000 and 2000 to 2008. Data come from the decennial census and three year ACS. In addition to the setup described in table 1.4, these also control for the main effect of the interaction term, if it is not already included.

Table A11: Instrumental Variables Estimates of Migration Elasticities, Separately

Panel A: Bartik

	(1)	(2)	(3)
Low ties: Wages	2.51		
	(1.17)		
High ties: Wages	-0.18		
	(1.57)		
High ties indicator	21.52		
	(19.58)		
Main effect of wages		7.41	1.86
		(3.37)	(0.86)
Interaction (x100)		-9.75	
		(5.01)	
Main effect of ties		0.68	
		(0.57)	
Controls	Y	Y	Y
P-val: No diff	0.23	0.05	
First stage F (Wald)	5.8	13.8	24.9
First stage F (K-P)	0.6	3.8	4.6
Observations		722	

Panel B: Trade

	(1)	(2)	(3)
Low ties: Wages	1.58		
	(1.57)		
High ties: Wages	-1.40		
	(1.36)		
High ties indicator	-5.34		
	(2.49)		
Main effect of wages		4.60	0.32
		(9.89)	(0.41)
Interaction (x100)		-8.45	
		(18.55)	
Main effect of ties		-0.22	
		(0.10)	
Year fixed effects	Y	Y	Y
Controls	Y	Y	Y
P-val: No diff	0.30	0.65	
First stage F (Wald)	4.1	1.1	54.1
First stage F (K-P)	1.1	0.2	17.3
Observations		1444	

Note: This breaks out IV results using the Bartik and trade shocks in separate specifications. Otherwise, it follows the methodology in Table 1.7.

Figure A1: Changes in the Population of Locals and of Outsiders

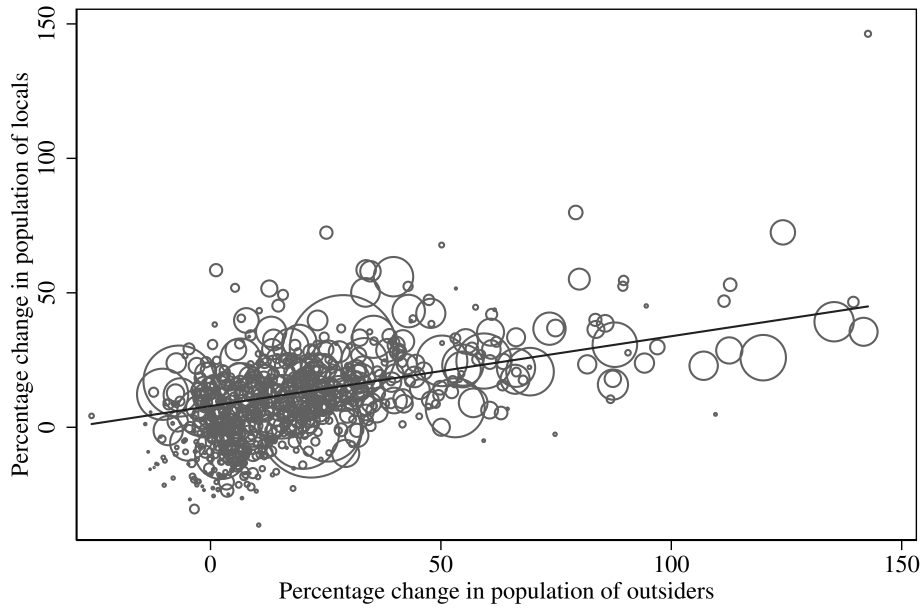

Line slope: 0.260 (0.031) R−squared: 0.247
Standard deviation of x and y variables: 28 and 14

Notes: Plotted are changes in the total population of people who were born outside of (inside) their current state, from 1980 to 2008, divided by the total population of the commuting zone in 1980 and multiplied by 100. In this way it represents the contribution of this population group to changes in the commuting zone's population. Data are from the long form decennial census and the ACS 3 year estimates (2006-2008) and are weighted to be nationally representative. The unit of observation is a commuting zone within the continental United States. The figure excludes the small number of commuting zones where the each statistic was over 150 so it is easier to read. Regressions in Table A1 include them, however.

Figure A2: Ratios of Migration and Non-migration Population Changes

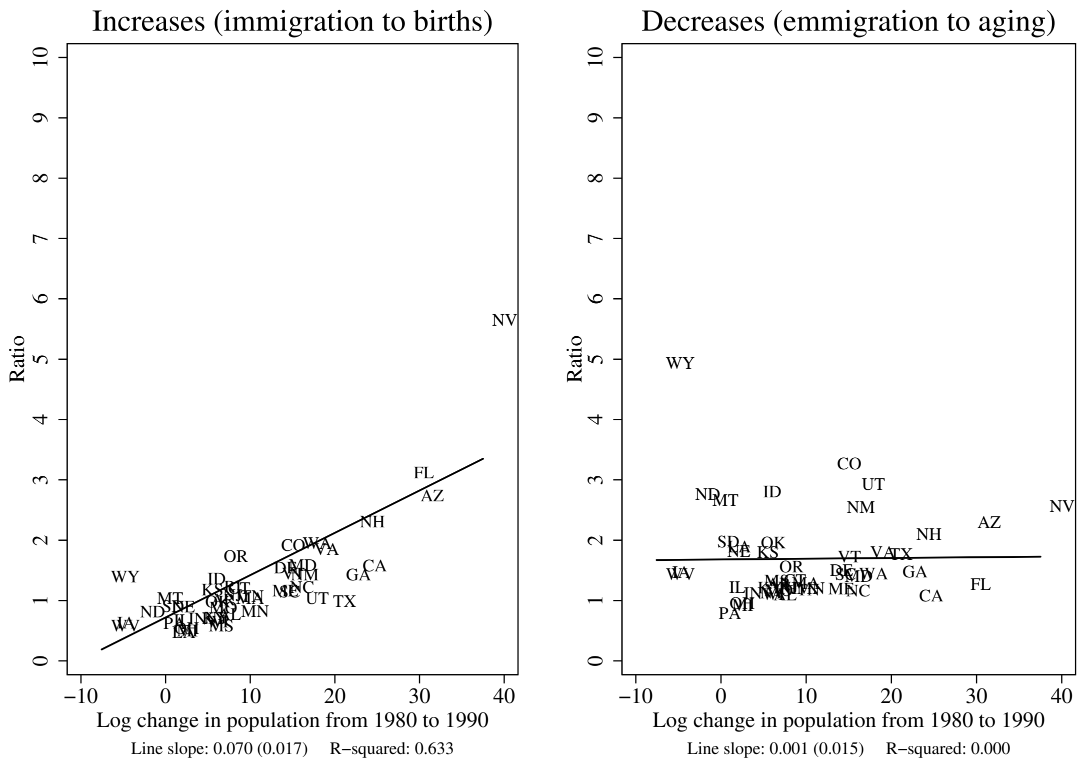

Notes: Data are from the decennial census and vital statistics (National Center for Health Statistics (2014)). Births are from 1969 to 1974, aging is the population 55 to 60 in the 1980 decennial census, and migration statistics are from the 1990 census. The regression line is an OLS regression using each state as an observation. Robust standard errors are in parenthesis.

Figure A3: Correlations Between 10 Year Changes in Working Age Population

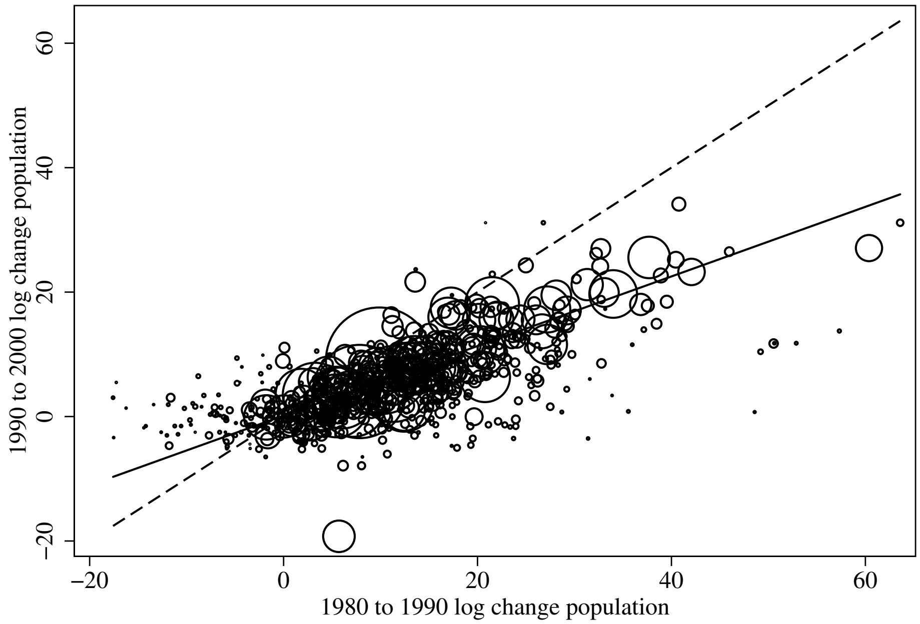

Line slope: 0.506 (0.012) R−squared: 0.443

Figure A4: Population Changes and Locally Born Workers Staying

Panel A: Unadjusted

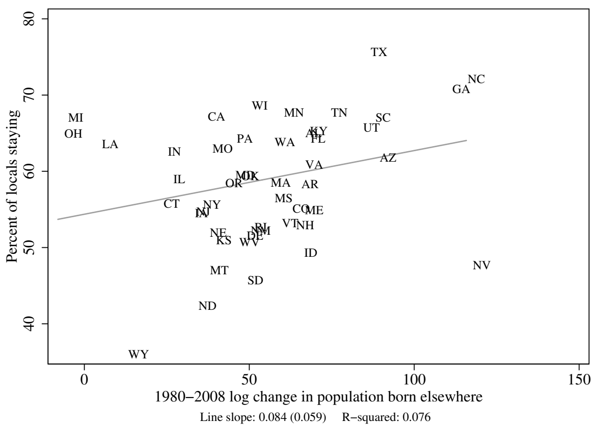

Panel B: Residualized

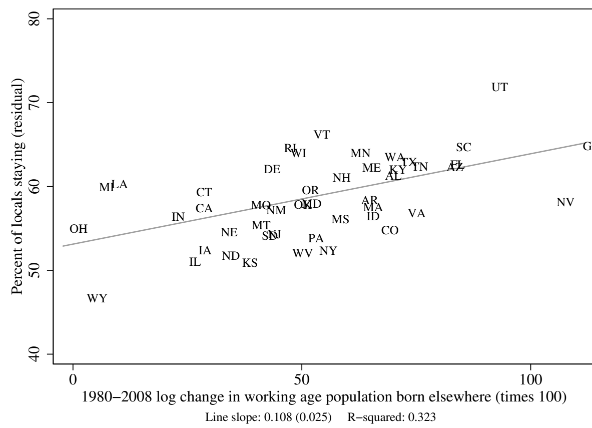

Panel C: Unadjusted and Weighted

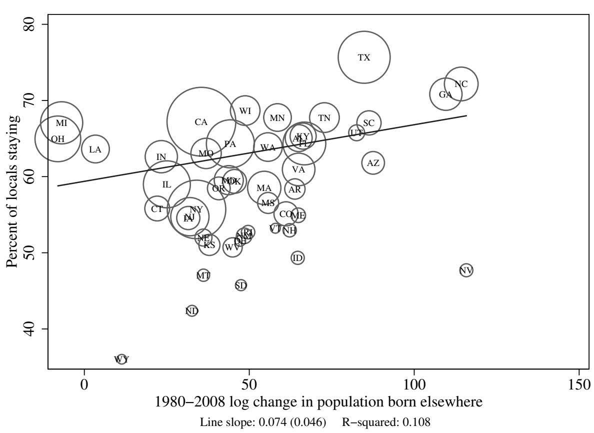

Panel D: Residualized and Weighted

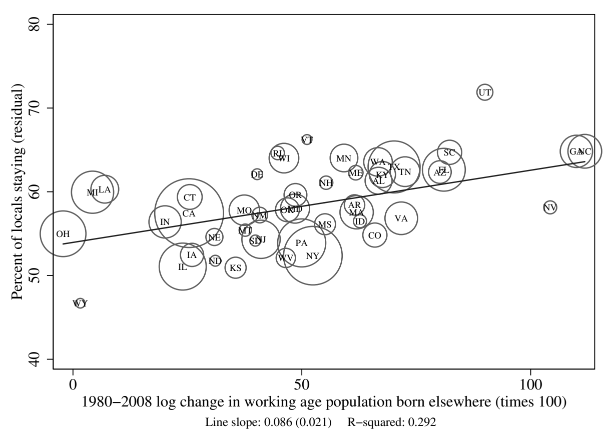

Note: Plotted are either the observed share born locally, or residuals of a regression of the share locally born residuals on a series of controls, with the constant added back in after. The controls are: share college educated, share employed, share foreign born, share born specifically in Mexico, and population 40 years previously. The line is from an OLS or WLS regression and the standard error is clustered by census division.

Figure A5: Average Time Living in the Same House as of 2000

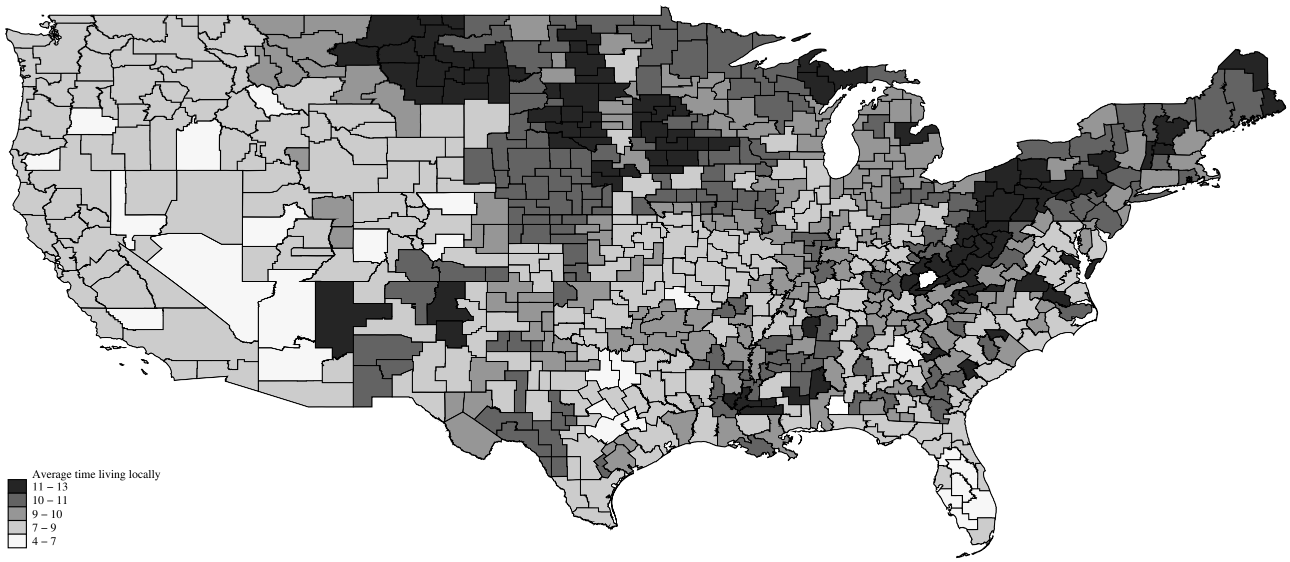

Notes: The 722 commuting zones in the continental US are shaded according to how long the average "householder," in whose name the residence is owned/rented, has been living at their current residence. The statistic is weighted according to the number of adults 16-64, fulfilling other sample restrictions, who live at that residence. Darker shades mean longer average times living in the residence. Data are from responses to the 2000 long form census via IPUMS.

Figure A6: Local Labor Demand Shocks

Panel A: Bartik Shocks: 1980 to 1990

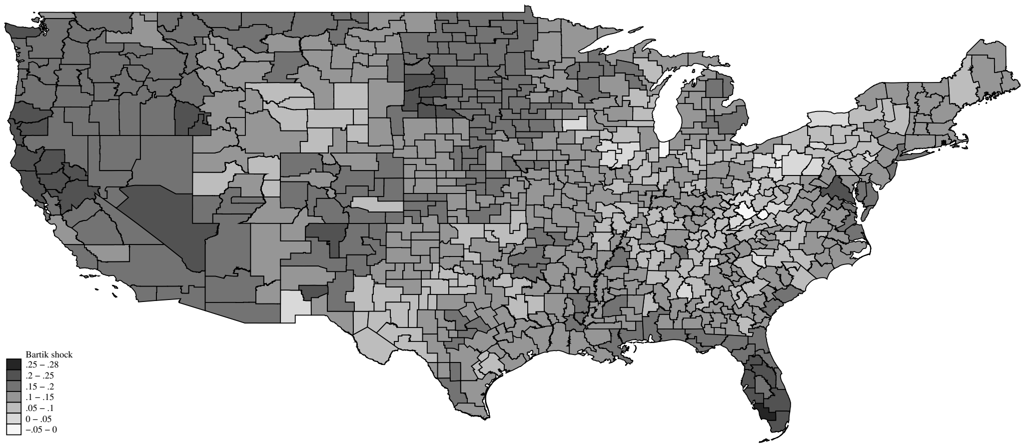

Panel B: Trade Shocks: 1990 to 2000

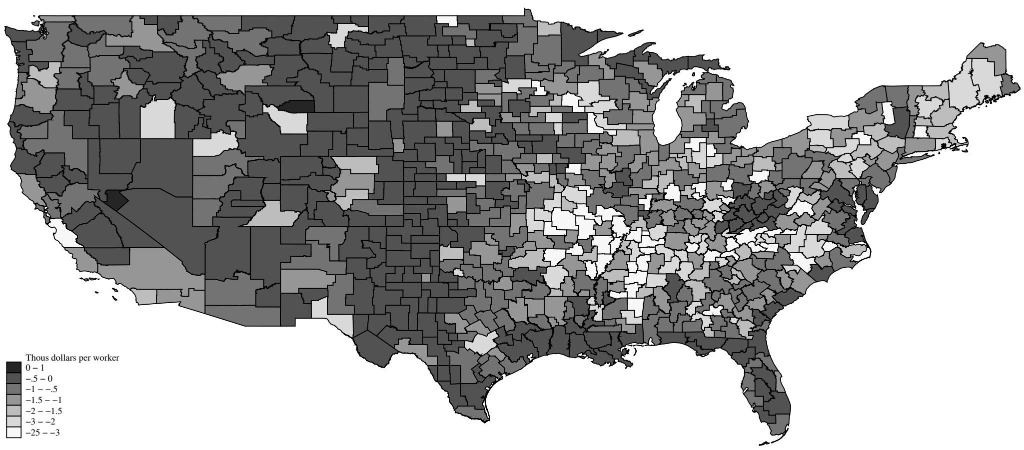

Panel C: Trade Shocks: 2000 to 2008

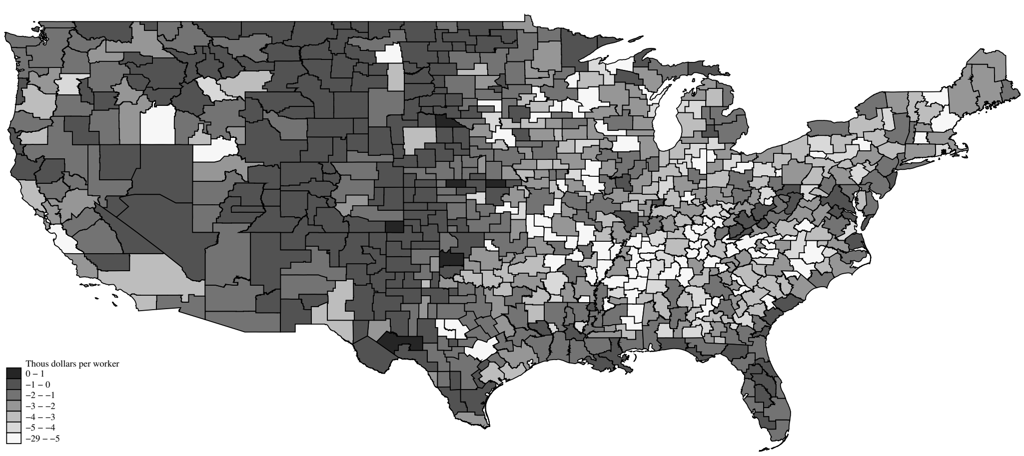

Note: The figure plots commuting zones shaded based on the severity of the local labor demand shock in the period. Data are from the Decennial Census, ACS, and Autor et al. (2013).

Figure A7: Scatterplot of ADH and Bartik Instruments

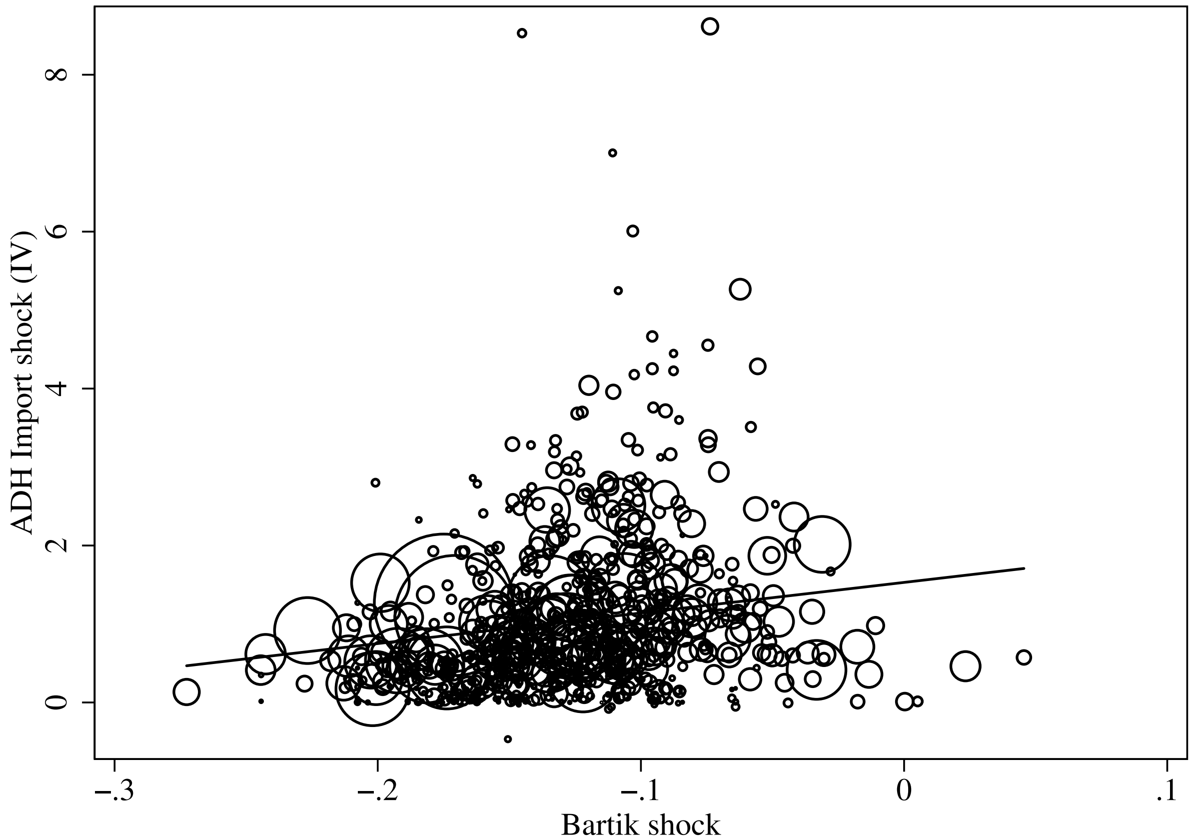

Note: Scatterplots are shown between instrumental variables related to Chinese import competition and Bartik labor demand instruments. The figure describes the correlation between a given CZ's Bartik instrument for 1980 to 1990 against its Chinese import exposure for 1990-2000. I use population weights at the beginning of the period relevant for the Chinese import shock.

APPENDIX B

Parental Proximity and Earnings After Job Displacements

B.1 Appendix: Additional Robustness Exercises (For Online Publication)

In this section we present several robustness checks to our baseline results in Section 2.3. The baseline results are remarkably robust to these different specifications, controls, and samples.

Appendix Figure B5 presents our baseline results together with results where PSID weights are not used. The results are similar and tell the same qualitative story.

Appendix Figure B6 presents results where we include controls for local labor market conditions. In particular we use employment numbers by county from County Business Patterns (CBP) and population information from the National Historical Geographic Information System, where we linearly interpolate between census years. These data are available from 1969 onwards. The figure suggests that controlling for county-level employment-to-population ratios does not affect the results. When we use county-level unemployment rates from the Local Area Unemployment Statistics (LAUS) as controls for local labor market conditions we obtain similar results but those data are only available after 1980, substantially reducing our sample.

Although census tracts are very small areas, sometimes residences that are geographically close might be in different tracts. Using latitudes and longitudes of block groups, we have computed as-the-crow-flies distances between parents and their adult children in the PSID. Appendix Figure B7 shows the earnings results when we group workers based on whether they lived within 3/4 of a mile of their parents prior to the displacement event (roughly a 15 minute walk at average walking speeds) or farther away. We see that this distance based measure of proximity yields virtually identical results.

It is possible that our baseline effect differs by gender. In particular, women might be more likely to benefit from parental proximity since they were largely responsible for childcare and household chores during the bulk of our data. To get at this, we created a dataset that included both heads and wives, since heads in the PSID are very likely to be males. Appendix Figure B8 presents the results from estimating our baseline specification on this pooled sample of wives and heads. We find that the results are very similar.[1]

Similar to the exercise in the main text where we break out proximity by same tract, same commuting zone but not same tract, and farther away, we also look at workers who are actually co-residing with their parents as opposed to living in the same neighborhoods as their parents, in the spirit of Kaplan (2012). Appendix Figure B9 presents the results from this exercise. We cannot reject the null hypothesis that the estimates for the coresiding young adults are different from those sharing the same neighborhood with their parents. The point estimates suggest that the recovery for these two groups of workers is similar and better than for those who live outside of their parents' neighborhoods at the time of displacement. This finding suggests that the effect of parental proximity is not limited to transfers while coresiding.

We have also found that the earnings differences for the two groups persists even if one includes additional interactions of the displacement dummies in equation (2.1) with whether the worker was displaced while living in the county where they grew up.[2] Appendix Figure B10 presents the earnings trajectories for those who are in their parents' neighborhoods and those who are farther away, and not in their home county, after these additional interactions are included in the baseline specification. The results are very similar to the baseline results, and those who are in the county they grew up in at the time of displacement have similar earnings losses to those who are not in their parents' neighborhoods and not in their home county. These findings suggest that parental proximity has an effect on post-displacement earnings that is independent from other local factors.

[1] Estimating the specification separately for men and women suggests that females have a slightly smaller benefit from living at home, but sample sizes are small, and the results are made more difficult to interpret because women, on average, tend to suffer smaller earnings losses following displacement (Ruhm, 1987).

[2] This is based on a retrospective question in the interview.

B.2 Appendix: Reweighting Based on Other Characteristics (For Online Publication)

In this section we present two alternative versions of our reweighting procedure that address two possible limitations inherent in reweighting based on observable characteristics. The first concern is that we may miss unobservable differences between people who live different distances from their parents, and that these differences may be driving our results. The second, essentially, is the opposite – that our reweighting could be overfitting the data. The reweighting could be going too far and selecting a group of workers who have less marketable unobservable skills than workers who are living close to their parents, since we emphasize people who have similar paying jobs and who do not have the benefits of living close to their families and friends. It could be that people who live farther form their parents, with jobs of constant observable quality, may be living farther away because they are unable to find comparable jobs nearby.

We address these concerns by repeating the reweighting exercise using two different sets of characteristics, and testing if our results are robust to alternative specifications. The broad finding is that our results are quite similar, and sometimes more precisely estimated, when we include different sets of characteristics.

We proceed with two strategies. First, we include as many additional characteristics as possible, including characteristics of parents in this case. Second, we pare down our list of characteristics to those that are predetermined when someone decides whether to live near their parents. Most notably, this omits all characteristics of people's jobs since workers can choose jobs based on their preferences about locations.

To implement each strategy, we estimate our baseline specification, equation 2.1, with weights, shown in equation 2.3, using different sets of co-variates.

B.2.1 Including Parents' Characteristics

For the first exercise, we include all of the same characteristics as in Section 2.4, the average employment to population ratio in the census tract that the worker's parents live in, and several characteristics of both the mother and the father, entered separately. These characteristics include a dummy for whether the parent is college educated, the parent's total years of schooling, a dummy for whether the parent is currently employed, the parent's age, and the parent's age squared. We also include a dummy for whether the parent was interviewed in any of the three years before the worker's potential displacement. If the parent was not interviewed we set the dummy for whether the parent was interviewed to one and all of the parent's other characteristics to zero. Otherwise, we follow the methodology in

Section 2.4.[3]

The results, shown in Panel A of Appendix Figure B11, are very similar to the propensity score reweighted results in Figure 2.10. The initial effect of a job displacement is around \$11,000 for each group (around a third of initial earnings). The group of workers who live in their parents' neighborhoods earn about the same amount as if they were not displaced after about six years, however, while the other group permanently earns about \$7,000 less. Compared with Figure 2.10, more of the differences between groups are significant at the five percent level, and confidence intervals are slightly smaller.

B.2.2 Predetermined Characteristics

For the second exercise, we restrict the characteristics to those that are predetermined at the time that the worker decides whether to live in their parents' neighborhood. These characteristics are a dummy for whether the worker is college educated, their total number of years of schooling, their age, a dummy for their gender, a dummy for being African American, and the employment-to-population ratio in the place where they live. Besides the limited set of worker characteristics, we follow the methodology in Section 2.4. For example, our estimates are based on the same sample of workers.

The results, shown in Panel B of Appendix Figure B11 are similar to Panel A, and also to Figure 2.10. The initial effect of a job displacement is also around \$11,000 for each group, though there is some suggestion that it is smaller for the group living closer to their parents. The smaller drop might be because this reweighting results in workers who live closer losing jobs that pay slightly less, on average. Still, the earnings trends are similar to before, with workers who live in their parents' neighborhoods earning about the same amount after six years. The group that lives near their parents earns slightly more in this specification, and the group living further away earns slightly less, so the differences are slightly larger than in the other reweightings. The estimates are somewhat less precise, but the larger differences between the estimates counteract this in terms of statistical significance.

[3]Parents' characteristics are averages for the three years before the displacement, which matches the characteristics for the worker. In cases where the parent was interviewed but we still do not have information about the characteristic (primarily item nonresponse) we include the average value of the characteristic for all PSID respondents. Including separate dummy variables for for each characteristic of each parent resulted in a log likelihood function that was not concave. Including only observations where we had information for all of the characteristics when a parent was surveyed resulted in very similar results, but a much smaller sample. We include the values for either the relevant parent or the relevant in law, based on which is a PSID sample member. In the rare cases where we have information about both, we average the two values.

B.3 Appendix: Reweighted Earnings of Non-Displaced Workers (For Online Publication)

One way to see if the propensity score reweighting is finding workers with similar counterfactual outcomes is to simply plot the earnings trajectories of each control group, suitably reweighted. We do this in Appendix Figure B12 by plotting earnings before and after years where workers were at risk of a displacement, according to our definition, but where they did not actually lose a job. Though there still are some differences after we apply the propensity reweights, the reweighted earnings are very similar before the potential displacement, and continue to behave quite similarly afterwards. The two lines are practically indistinguishable after we include the controls from our baseline regression specifications.

Panel A of Appendix Figure B12 shows that the propensity score reweighting procedure removes almost all of the earnings differences between the two groups, and brings them down to the lower level of earnings among workers who were displaced while living in the same neighborhood as their parents. Before the potential displacement, which is the period we use to match workers, the earnings of the two groups are within \$1,000 of one another and their pre-trends are roughly parallel. Even from period zero (the potential displacement) to period ten, a period that we leave out of our logit specification, the earnings trends still track each other quite well. This implies that matching on initial earnings, education, occupations, gender, and other factors is enough to find workers with similar employment prospects. This is very much in contrast to means using the PSID longitudinal weights.

Panel B of Appendix Figure B12 shows that earnings are even more similar between the two groups after we adjusted for an important control in our regression specification (equation 2.1) – a quartic term in worker age, estimated separately for each group. Taking out the age quartic highlights each worker's good fortune in avoiding a displacement in year zero and makes the lines completely statistically, and economically, indistinguishable. The lines are well within \$1,000 of one another throughout the panel.

B.4 Appendix: Including Additional Interactions in the Baseline Regression (For Online Publication)

To complement our reweighting approach, we also examined the effects of including interactions with other baseline characteristics, in the same way we separate out the effect of being closer to one's parents. We take another characteristic X_{ia}^{C}, like the person's earnings before displacement, and interact it with both the age quartic and the displacement dummies.

To be specific, we estimate the following equation, which includes all of the terms in our baseline specification, equation 2.1, as well as some additional terms:

$$e_{iat} = \alpha_{ia} + \gamma_t + X_{iat}(\beta^A + \beta^H H_{ia} + \beta^C X_{ia}^C) + \sum_{k=-4}^{10+} (D_{it}^k \delta^k + D_{it}^k H_{ia} \zeta^k + D_{it}^k X_{ia}^C \xi^k) + \epsilon_{iat} \quad \text{(B.1)}$$

The fixed effect, α_{ia} already controls for an effect of X_{ia}^C that is constant across time, but the additional interactions also control for time varying effects around the displacement. For example, if earnings losses are bigger after layoffs from jobs that pay more, then this would be reflected in negative values of ξ^k for $k > 0$. If this were driving our result that workers who live closer to their parents suffer smaller earnings losses, then including this term would also move the value of ζ^k closer to zero.

As before, the effects of a displacement for different groups are different linear combinations of $\hat{\delta}^k$ and $\hat{\zeta}^k$ terms, and we plot these as a simple way of understanding the impact of these different specifications. We plot $\hat{\delta}^k + \hat{\zeta}^k$ as the effect for people living near their parents and $\hat{\delta}^k$ for people living farther from their parents. Since we are not including the $\hat{\xi}^k$ terms, the effect is for the omitted group where X_{ia}^C is a dummy variable and the value at the mean of X_{ia}^C (since we de-mean X_{ia}^C) when it is a continuous variable. Note that the difference between the two lines is, due to functional form, unchanged regardless of the value of X_{ia}^C.

Appendix Figure B13 shows the coefficient estimates with several different interactions. The light gray lines, reproduced from Figure 2.2, show the baseline earnings losses for people living in the same neighborhood as parents (dashed line) and people living away from parents (solid line). The darker lines in Panels A and B of Appendix Figure B13, show the same coefficient plot if one also allows effects to vary by how much they earned before the before displacement. In Panel A we include an interaction with a linear earnings control and in Panel B we include an interaction with a high/low earnings dummy. Controlling for initial incomes generally makes the initial earnings losses much more similar between people at different distances from their parents. Controlling for income, however, does little to the finding that the two paths diverge later on. Panel C of Appendix Figure B13 presents the earnings plots when we include an interaction with a dummy for college education. As with the income interactions, this reduces the difference between the two groups but does not remove the long-run divergence.

B.5 Appendix: Search Intensity and Switching Industries (For Online Publication)

In Section 2.5 we outlined two possible mechanisms: housing transfers and parental employment networks. Here we look at two more: job search intensity and industry switching. We find that neither is a likely explanation for our baseline earnings results.

B.5.1 Search Intensity

This section outlines some results about how search intensity for young workers varies with proximity to parents. We are motivated by the idea that parents may provide additional encouragement to their children after the job displacement, which may help with the job search process (Dalton, 2013). Our analysis, however, documents no statistically significant relationship between the search intensity of unemployed young adults and living close to parents.

The search activities data we use here only started in 1988 (as opposed to 1968 for the main analysis) and we stop the analysis in 2013, yielding 18 years of data. This means that the sample used for this exercise will be different from the one in the main text. Nonetheless, unless we think that the relationship between search intensity and parental proximity has changed from the 70s and 80s to the time thereafter, the present analysis should be representative of the entire period. Other than that, we use the same "stacked" version of the data in this analysis as described in Section 2.2. The search activity questions ask how respondents seared for jobs, e.g. they ask if respondents checked with private employment agency, if they checked with friends or relatives, and if they placed or answered ads.

Appendix Table B2 presents summary statistics on search intensity for those living in the same neighborhood as their parents and for those living farther away, by labor force status at the time of the interview, for younger (25-35) and older (36-55) adult children. The last two columns of this table suggest that young workers living close to their parents are perhaps more likely to engage in some form of job activity than those living farther away. A comparison between those two columns should be informed by the fact that those living close to their parents are more likely to be unemployed and unemployed people are more likely to search. The latter can be seen in Appendix Table B2 by comparing the search activities of the employed and the unemployed. On the former, the unemployment rate of those living close to their parents is far higher than for those living farther away. Within labor force status, those at home, if anything, appear to search less than those farther away, however when unemployed, they appear to be more likely to check with friends or relatives.

The bottom panel of the table shows that older workers are, on average, less likely to be searching for a new job than young workers, regardless of labor force status. The patterns of search activities by proximity to parents for older adults are similar to younger workers, although the differences between the searching behavior of those close to their parents and farther away is more similar than for younger workers. In particular, older workers who are unemployed at the time of the survey are no more likely to check with their friends or relatives than those who live farther away.

In order to go beyond these basic comparisons of means, we estimate the following linear probability model:

$$search_{it} = \alpha + \beta Sametract_{it} + \gamma X_{it} + \epsilon_{it} \tag{B.2}$$

where $search_{it}$ is a dummy for whether worker i reported any job search activity in period t, $Sametract_{it}$ is a dummy for whether worker i is living in their parents' neighborhood in time period t, and X_{it} includes a host of controls.

Table B3 presents the results of this analysis, by labor force status. The first column reproduces the average difference in the probability of searching for a job between unemployed young adults living in the same neighborhoods as their parents and those living farther away from Table B2. Columns (2) and (3) add increasing number of controls, including demographic controls, such as age, race, and education, and year fixed effects. These results rule out large negative effects of parental proximity on young adult search activity and suggest a negative relationship between the two that is not statistically significant. Column (3), for example, suggests that living in the same neighborhood as one's parents reduces the probability of unemployed young adults engaging in search activities by 11 pp (on an average of around 85 percent), but the standard errors are large. Column (4) shows that employed young adults who live in their parents' neighborhoods, conditional on the controls in column (3), are slightly less likely to be engaged in search activities than young adults living away from home, but the point estimate is virtually zero and precisely estimated.

Column (5) pools unemployed and employed young adults and includes worker fixed effects in addition to the other (time-varying) controls in columns (3) and (4). This approach uses variation in proximity to parents within an individual's observations to identify the effect of parental proximity on search activity as opposed to variation in proximity to parents across individuals. The results of column (5) are also not statistically significant and small. We take these results as not suggesting large differences in job search activity for those living in their parents' neighborhoods and those living farther away.

We are also not able to say much about how search activity changes around a displacement event for those living close to their parents and those living farther away. In particular, when estimating equations like equation (2.1), but with search activity as an outcome variable, we are unable to reject that the search activity of these two groups move in the same way around a displacement event. Taking these results at face value, we conclude that, although variations in search intensity among young adults living close to and farther away from home could be partially responsible for the markedly differential post-displacement earnings trajectories we observe for these two groups, the effect would likely have to be through something other than higher job finding probabilities for those at home, based on our results on unemployment duration (Section 2.3.2). Ultimately, further research with different data would be needed to make a more definitive statement.

B.5.2 Industry Switching

Previous work, including Jacobson et al. (1993) and Stevens (1997), has documented that industry and occupation switchers experience larger post-displacement earnings losses than workers who retain a job in their former line of work. We document industry switching around the displacement event for workers who were in the same neighborhood as their parents prior to the displacement event and those who lived farther away. We estimate equation (2.1) but use as an outcome variable a dummy, D_switch_{iat}, that equals one if the worker i switches one-digit industry between survey year t and $t+1$ at base age a.

Appendix Figure B15 presents this probability of switching industries by parental proximity. Both groups of young adults switch industries more frequently around a displacement event than in other periods, consistent with previous work (Burda and Mertens, 2001). This switching rate stays elevated for several years after the displacement event. Appendix Figure B15 also suggests that workers living in the same neighborhood as their parents prior to the displacement event experience markedly sharper increases in their probability of switching industries than workers living farther away. Based on the prior work cited above, this would predict larger post-displacement earnings losses for workers living close to parents and would thus work against our baseline findings. As such these results suggest that industry switching cannot explain our baseline findings. In results not shown, occupation switching is similar for the two groups around a displacement event.

B.6 Appendix: Measures of Transfers (For Online Publication)

This section presents the technique that we use to determine how much children were able to save in rent by living with their parents as well as some simple analysis of the PSID's question about cash transfers (help) from friends or relatives.

B.6.1 Calculating Implied Savings on Rent

To calculate the implied amount that a family unit saves on rent, we rely on the OECD equivalence scale and an assumption about the user cost of capital to back out the cost of a dwelling where the family unit could live in the same amount of comfort.

We use an equivalence scale to make comparisons between larger houses that have many people living in them and smaller houses that have fewer people living in them. The OECD equivalence scale is among the most commonly used equivalence scales that accounts for both crowding and also returns to scale in household consumption.

$$E(A, C) = 1 + 0.7(A - 1) + 0.5C$$

Mechanically, each adult additional (A) counts for 70 percent of the initial adult, and each child (C, 14 or younger) counts for 50 percent of the initial adult. A given value of the scale, $E(A, C)$, implies someone living alone in a house that costs, say, a dollars would be indifferent to living in a house costing $E(A, C) \times a$ dollars if they were to live with $A - 1$ other adults and C children.

In a case where a child lives in a house that their parents are renting, it is possible to back out the implied amount the child would have to pay to live alone in a house of a similar quality. Say the child would live in a family unit with A_C adults and C_C children, and the parents in a family unit with A_P adults and C_P children. Then, given that the parent's rent is R, the child would have to pay the following to live separately in a house of a similar quality.

$$\frac{R}{E(A_P + A_C, C_P + C_C)} E(A_C, C_C)$$

Intuitively, the formula first converts the rent into a per person level of consumption within the larger household by dividing by the equivalence scale for the larger household. Then the formula multiplies the per person level of consumption by the value of the equivalence scale for the child's household. This give the amount of rent the child's household would have to pay to enjoy the same standard of living. The difference between this counterfactual rent and the rent that the child actually pays is the implied savings from living with parents.

One complication in practice is that parents oftentimes own their houses, which means there is no direct measure of parents' rents. To compute an annual rent equivalent in these settings, we employ a user cost of capital equal to 0.0785 (following Albouy and Zabek (2016)). The user cost gives, essentially, the implied rental payment that the household pays for using the house for the year, as opposed to renting it out to another family or selling it.

It will depend on the depreciation of the house, the interest rate of the mortgage, property tax rates, and any specific tax incentives for home ownership. For simplicity, we set it to a fixed value and we only use it in situations where we need to convert the value of someone's house into a value on the rental market.

B.6.2 Cash Transfers from Friends and Relatives

In addition to transfers of housing, children can receive cash transfers from their parents. Appendix Figure B16 shows how these change around displacement, using our main event study specification including fixed effects, an age quartic, and other controls. For these plots we use an annual question in the PSID that asks how much money a household received from friends and relatives. Much more detailed transfer data exist in two single year transfer supplements to the PSID. These are of limited use in our context, however, because our estimation strategy is only possible when we have a data across many years.

Appendix Figure B16 shows that workers who are living away from their parents appear to receive larger monetary transfers two years after a displacement. This is apparent both for extensive margins (Panel A) and intensive margins (Panel B). Workers who live closer do not appear to receive any more money around a displacement, though this series is noisy and it has very large standard errors. As with housing transfers, the amounts are fairly small; the increase around a displacement is estimated to be about $150 per year, which is about one percent of the estimated earnings losses after a displacement for this group.

B.7 Appendix Figures (For Online Publication)

Figure B1: Average Log Earnings for Young Displaced Workers by Proximity to Parents

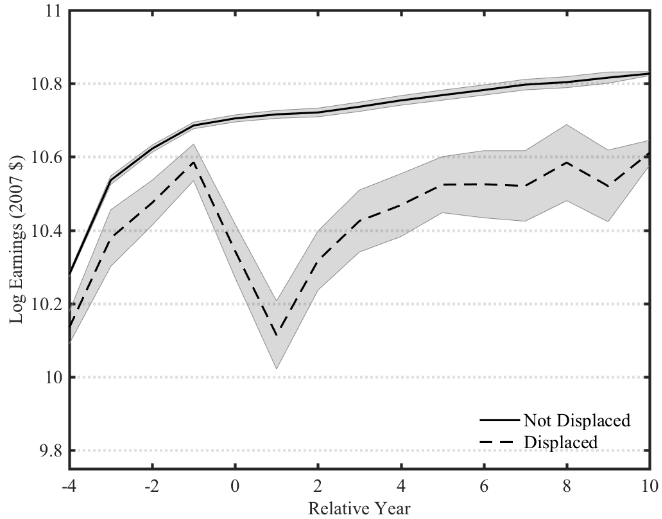

Panel A: Average Log Earnings for Young Displaced and Non-Displaced Workers

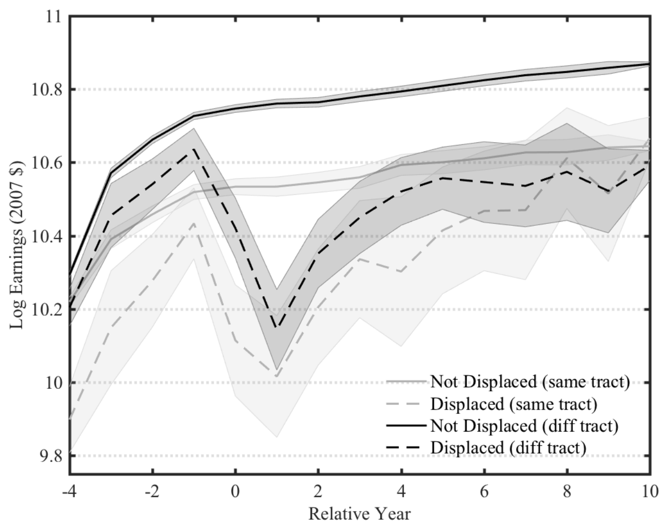

Panel B: Average Log Earnings for Those In Their Parents' Neighborhoods and Not

Note: This figure depicts log earnings and is analogous to Figure 2.1 in the main text that depicts earnings in levels, including its qualitative results. The shading represents 95 percent confidence intervals based on clustered standard errors, at the worker level. More information on the specification, definitions, etc. is in Figure 2.1.

Figure B2: Earnings Losses for Young Displaced Workers (Excl. Zeroes)

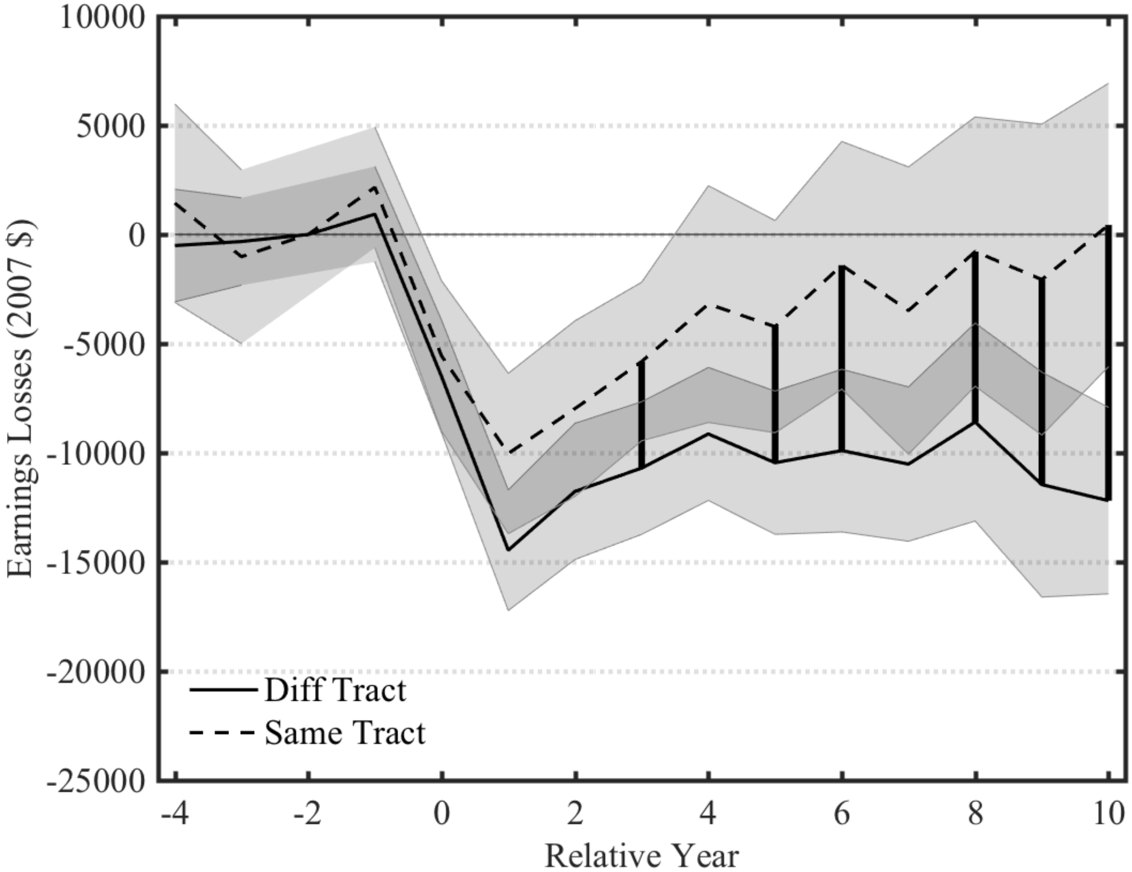

Note: Dropping observations with zero earnings gives the same result as the baseline finding in Figure 2.2: In the medium- and long-run, young workers living in the same neighborhoods with their parents experience a full recovery in earnings after a displacement event. Young workers not living in their parents' neighborhoods experience large and permanent earnings losses, amounting to around 30 percent of their pre-displacement earnings even 10 years after the displacement event. Plotted are regression coefficients from equation (2.1) where we drop years where workers reported zero earnings. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. We cluster the standard errors at the worker level. Figure 2.2 contains more information about the specification, definitions, etc.

Figure B3: Percent Earnings Losses for Young Displaced Workers

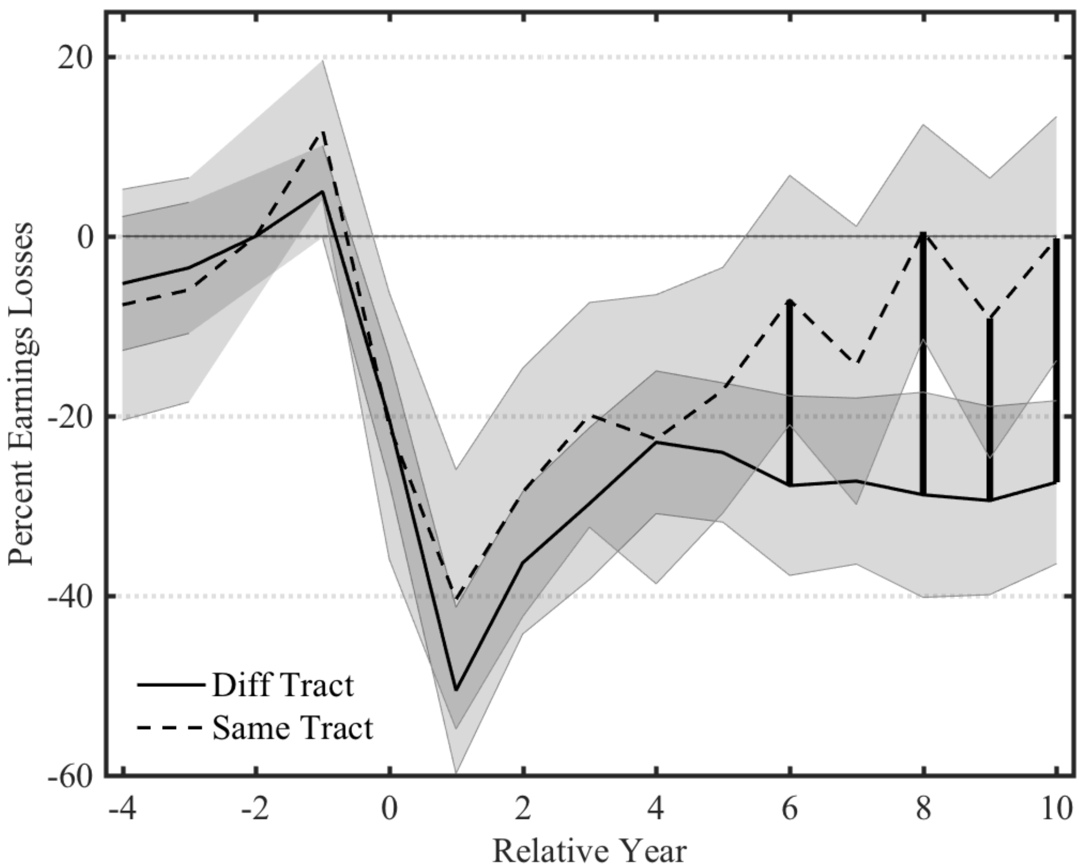

Note: Using log earnings instead of earnings in levels gives the same result as the baseline finding in Figure 2.2: In the medium- and long-run, young workers living in the same neighborhoods with their parents experience a full recovery in earnings after a displacement event. Young workers not living in their parents' neighborhoods experience large and permanent earnings losses, amounting to almost 30 percent of their pre-displacement earnings even 10 years after the displacement event. Plotted are regression coefficients from equation (2.1) with log income on the left hand side. To obtain percentage changes we plot $(e^{\hat{\delta}^k} - 1) \times 100$. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. We cluster the standard errors at the worker level. Figure 2.2 contains more information about the specification, definitions, etc.

Figure B4: Probability of Switching Commuting Zones Around Displacement

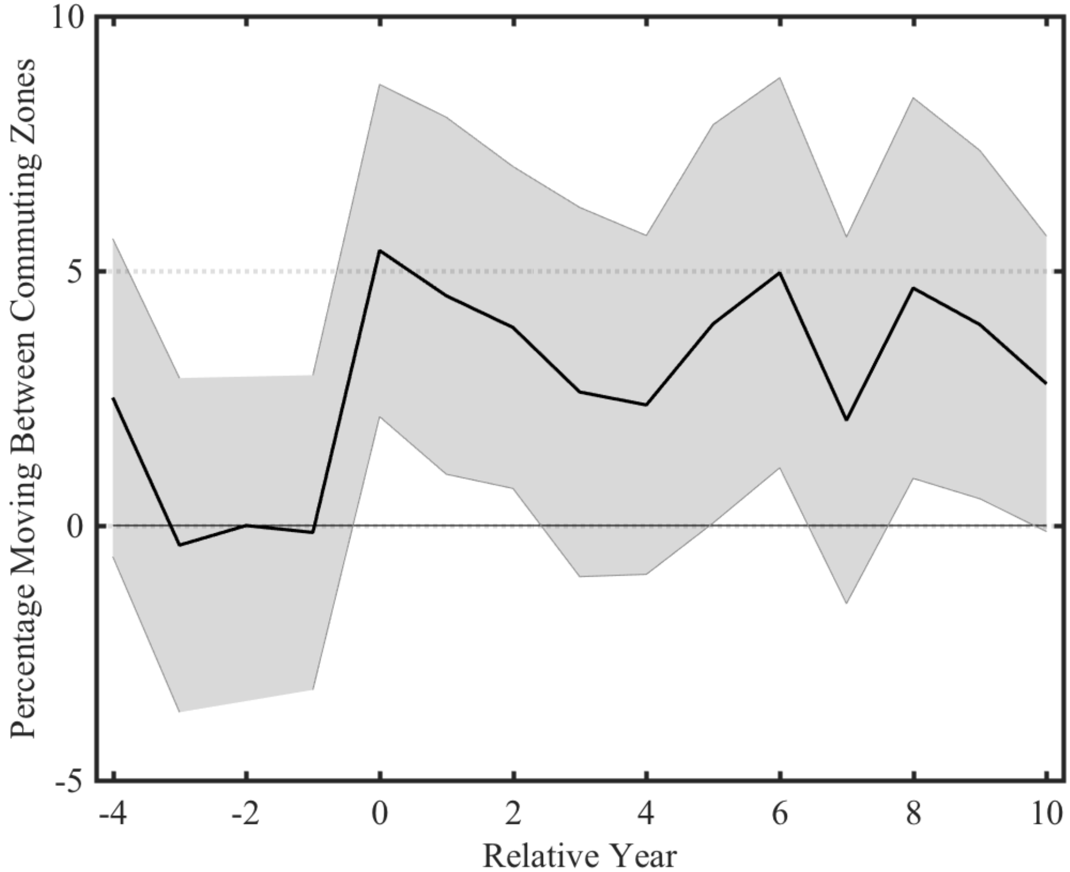

Note: At the time of displacement, the annual commuting zone switching probability rises by around 5pp. In light of an average annual probability of switching commuting zones at around 5pp, this increase represents a sharp increase in geographic mobility. Plotted are regression coefficieints from a linear probability model with a specification very similar to equation (2.1). The main differences are the outcome, moving between commuting zones in the year in question, and that we pool both groups of workers to increase precision. The shading represents 95 percent confidence intervals based on clustered standard errors, at the worker level. We use commuting zones as the relevant measure of geography because they most closely resemble the "regional labor markets" that Huttunen and Salvanes (2015) use with Norwegian data.

Figure B5: Earnings Losses for Young Displaced Workers (No PSID Weights)

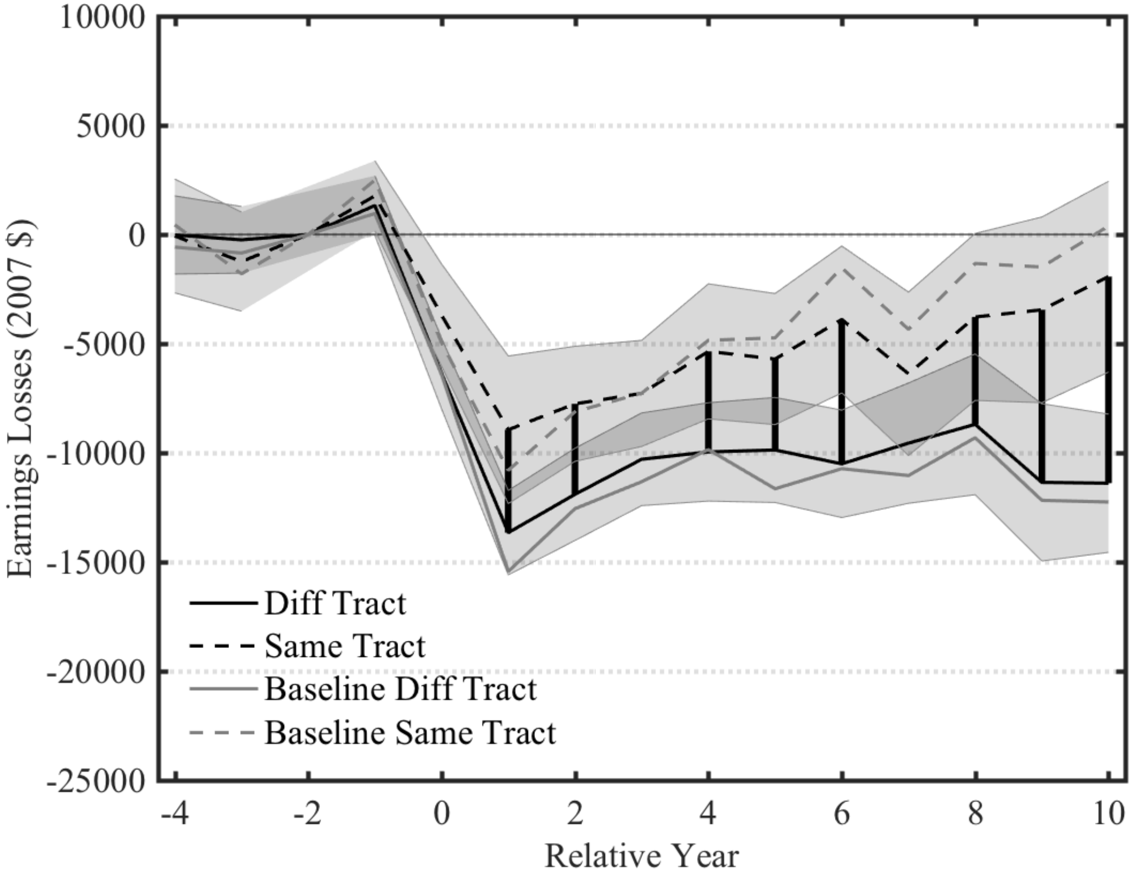

Note: These are the baseline results from equation (2.1) when we do not use the PSID longitudinal weights. They tell the same story as the baseline results in Figure 2.2, although those living near their parents at the time of displacement see slightly less of a benefit. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. We cluster the standard errors at the worker level.

Figure B6: Earnings Losses for Young Displaced Workers (Controlling for Local Labor Market Conditions)

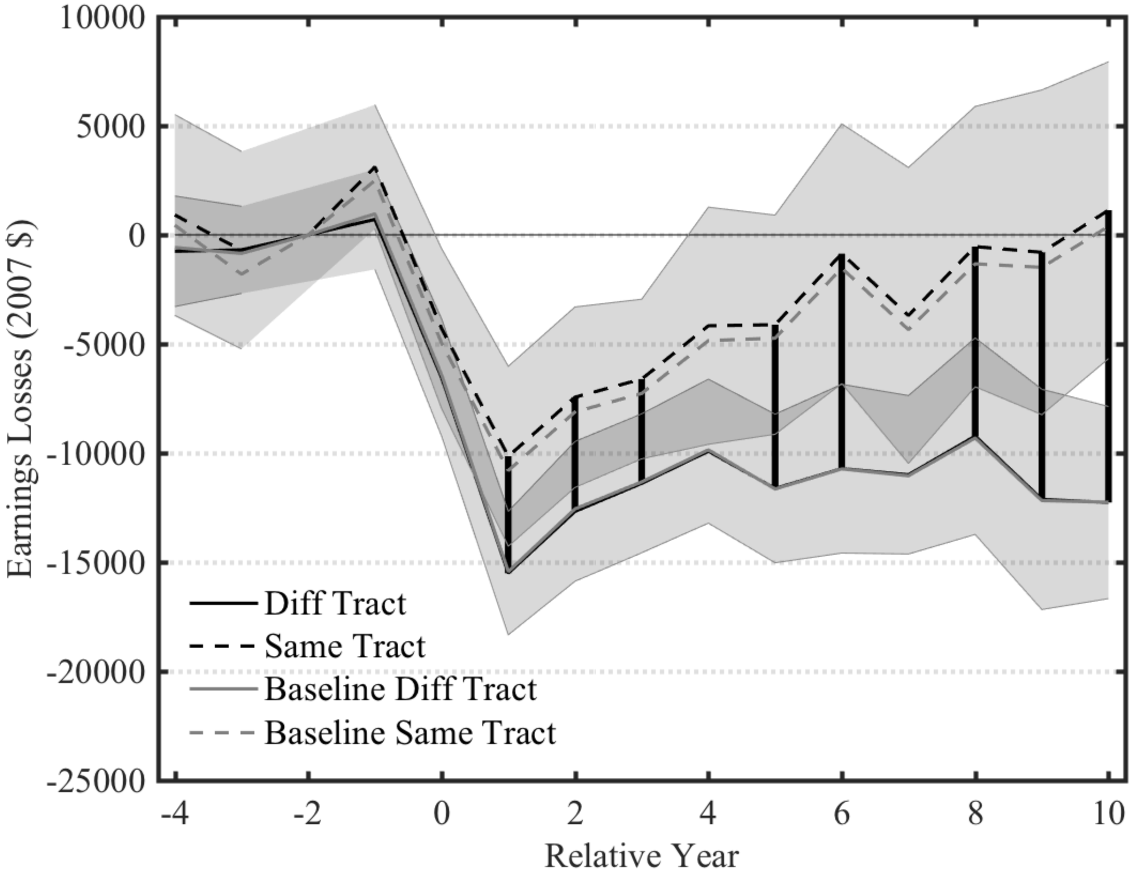

Note: These results control for employment-to-population ratios at the county level in our baseline equation (2.1). The results are virtually the same as in the baseline specification. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. We cluster the standard errors at the worker level. See Appendix B.1 for more details on how we measure local labor market conditions. The specification is very similar to our baseline specification in Figure 2.2.

Figure B7: Earnings Losses for Young Displaced Workers (Using Distance Measures)

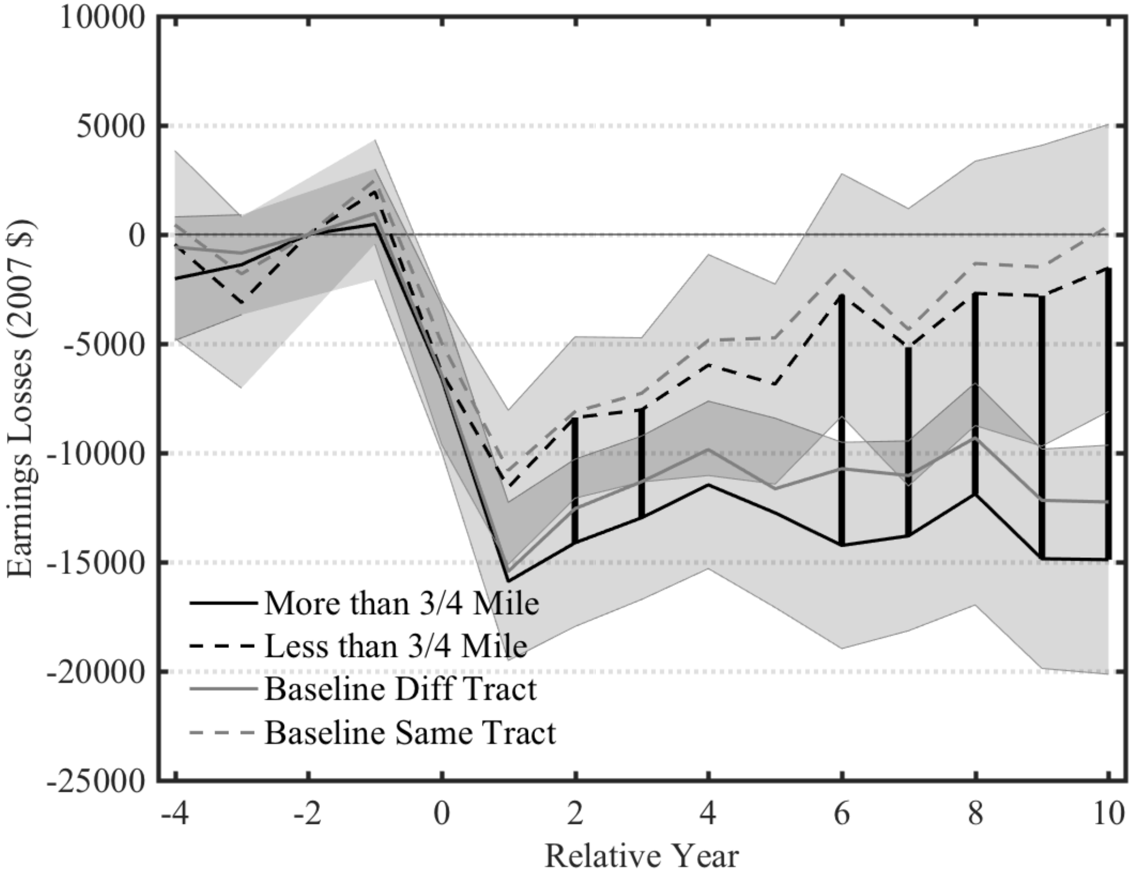

Note: These results estimate equation (2.1) where we define closeness by distance to parents and less than 3/4 miles is close. The results are very similar to the baseline specification in Figure 2.2. If anything, this approach strengthens the findings slightly. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. We cluster the standard errors at the worker level. See Figure 2.2 for full details of the specification.

Figure B8: Earnings Losses for Young Displaced Workers (Heads and Wives)

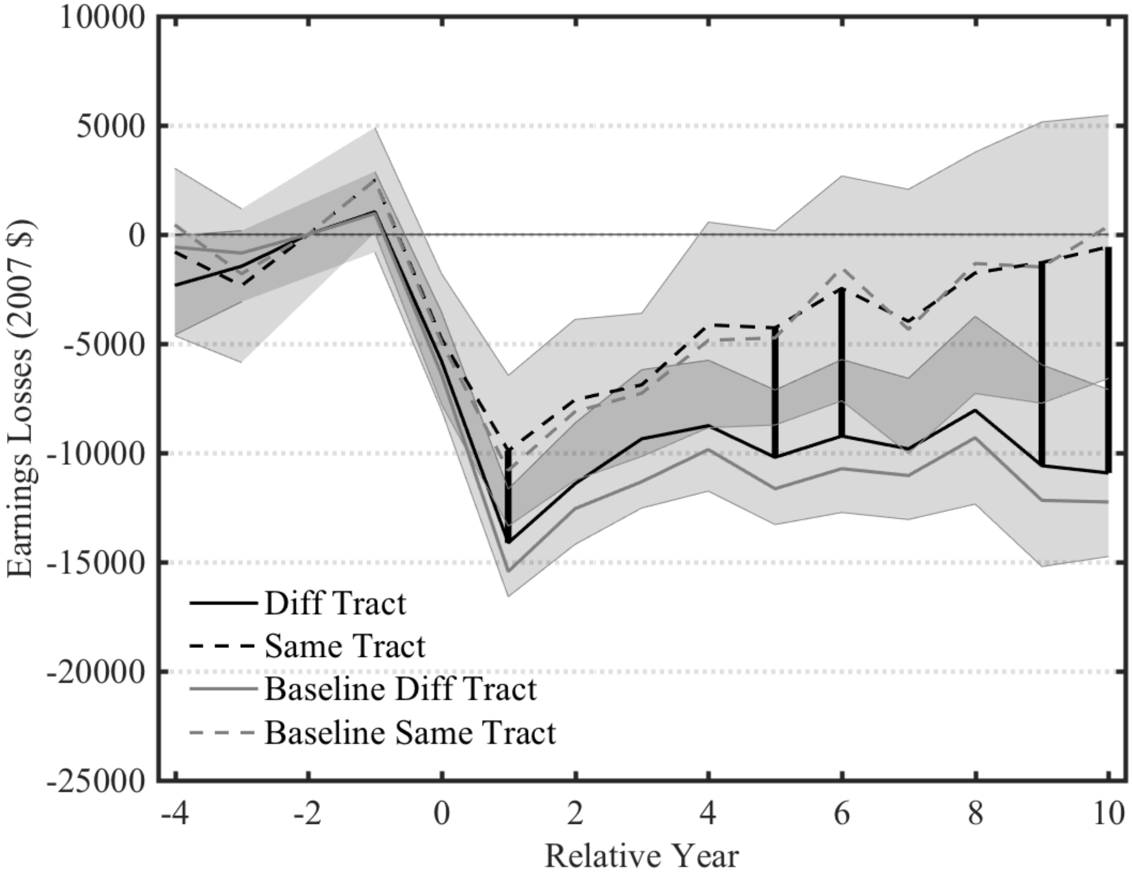

Note: These results present the coefficients from equation (2.1) using both heads and wives as opposed to just heads as in our baseline sample. The results are very similar to the baseline results in Figure 2.2. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. We cluster the standard errors at the worker level. See Figure 2.2 for full details of the specification.

Figure B9: Earnings Losses for Young Displaced Workers (Same Tract vs. Coresiding)

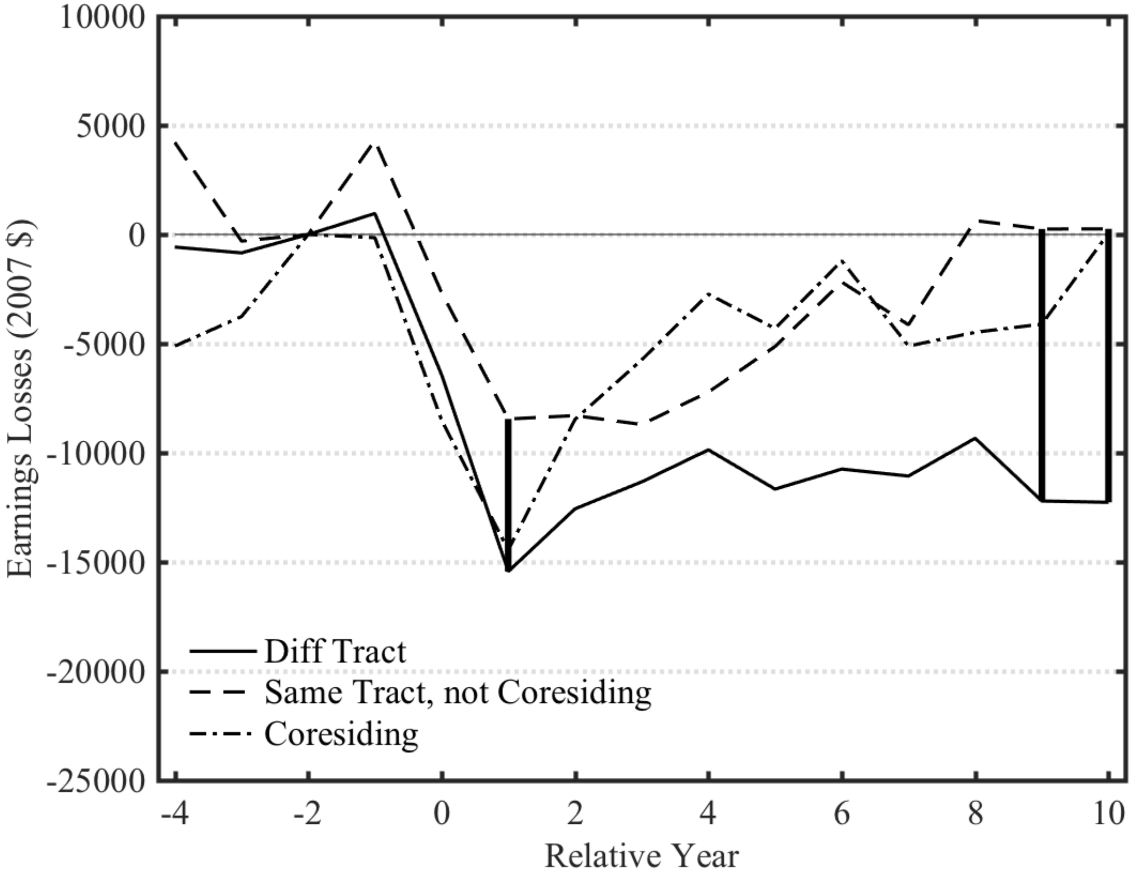

Note: The post-displacement earnings recoveries look similar for those workers who actually live in the same house as their parents (coresiding) and those who live in the same tract as their parents but are not coresiding. Both groups appear to do better than those who live outside their parents' neighborhoods. The figure plots regression coefficients from a specification similar to equation (2.1) describing the impact of a job displacement on the earnings of workers who lived in a different census tract from their parents, those who lived in the same census tract but in a different house, and finally those who lived in the same house as their parents. The groups are mutually exclusive. The figure includes vertical bars that connect the line for workers who live in the same tract (not coresiding) with the line for workers who live in a different tract. We include these when the estimates are statistically significantly different from one another at the five percent level. We cluster the standard errors at the worker level.

Figure B10: Earnings Losses for Young Displaced Workers with Home County Interactions

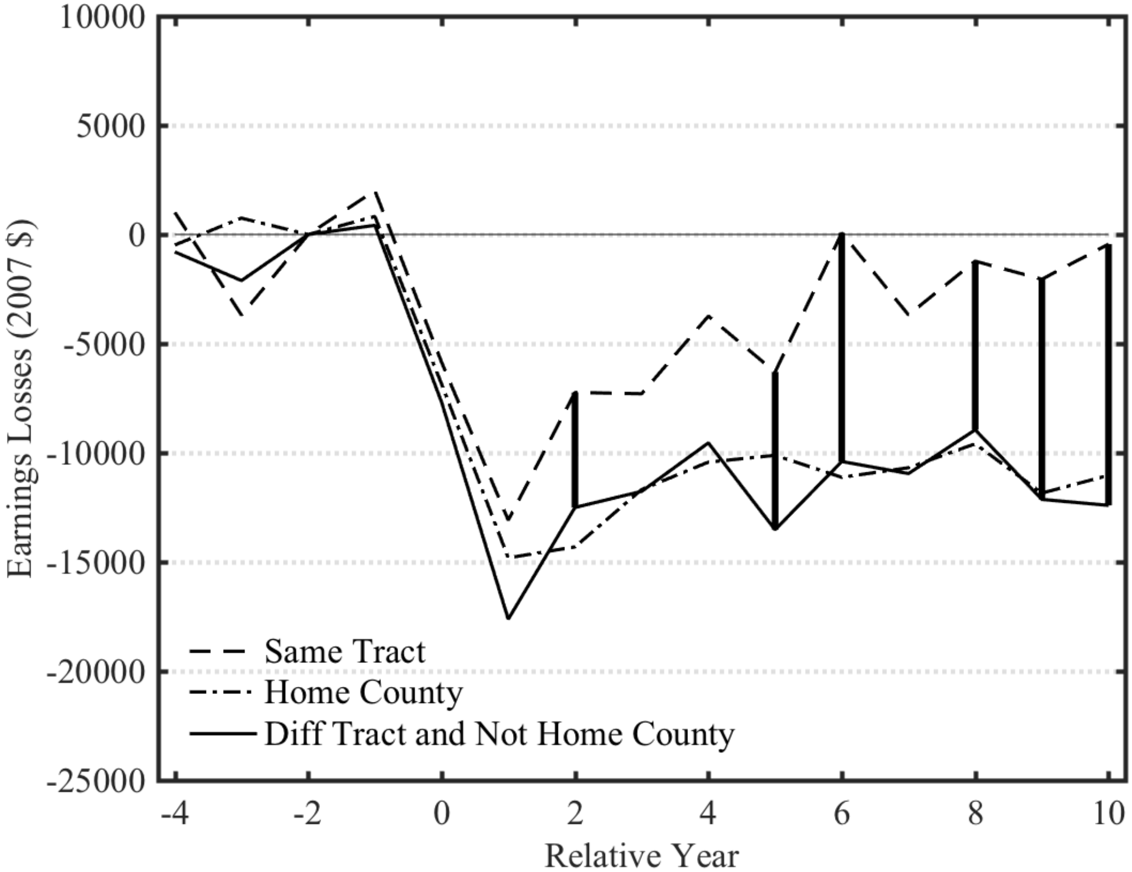

Note: The post-displacement earnings recoveries look similar to our baseline results even after interacting the displacement dummies in equation (2.1) with whether a worker lived in the county where they grew up at the time of displacement. The earnings losses for those in their home county look similar to those who are neither in their parents' neighborhoods or their home county. The figure plots regression coefficients from a specification similar to equation (2.1) describing the impact of a job displacement on the earnings of workers who lived in the same census tract as their parents, those who lived in the county they reported that they grew up in, and those who lived away from both their parents and the place that they grew up. For the former two the lines correspond to people where the specified category is true, but the other is not. The figure includes vertical bars that connect the line for workers who live in the same tract with the line for workers who live in a different tract within a county that they did not grow up in. We include these when the estimates are statistically significantly different from one another at the five percent level. We cluster the standard errors at the worker level. More information on these specifications is in Appendix B.1.

Figure B11: Reweighted Regressions Based on Alternative Reweighting Specifications

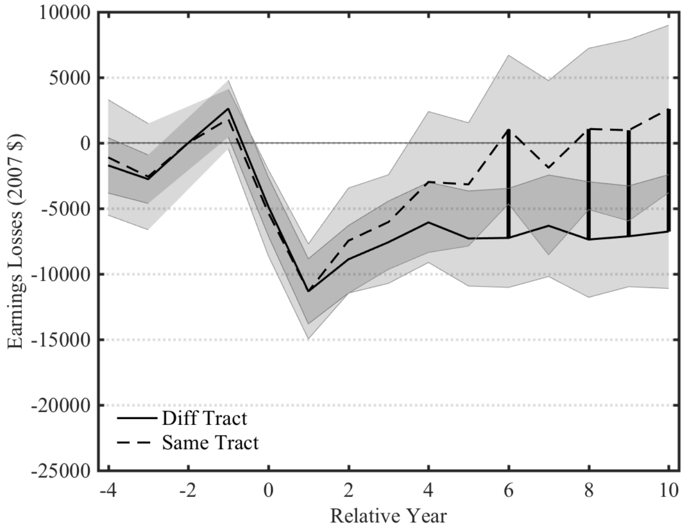

Panel A: Including Parents' Characteristics

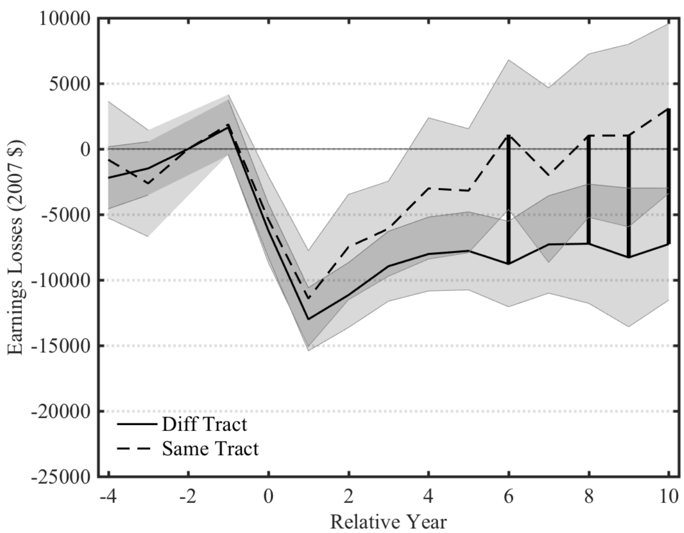

Panel B: Including Only Predetermined Characteristics

Note: The results from the propensity score reweighting do not appear to be sensitive to including either additional controls for parents' characteristics, or to including only predetermined characteristics, like educational levels, age, and race. The figure plots propensity score weighted regression coefficients from equation (2.1) describing the impact of a job displacement on the earnings of young workers. The weights in Panel A are calculated to match parents' characteristics between the different samples, in addition to the main characteristics. The weights in Panel B are calculated to only match predetermined characteristics between the sample. See Appendix B.2 for more details on these two types of weights, see Figure 2.10 for the original specification, and see Section 2.4 for more information on the reweighting scheme.

Figure B12: Mean Earnings For the Reweighted Control Samples

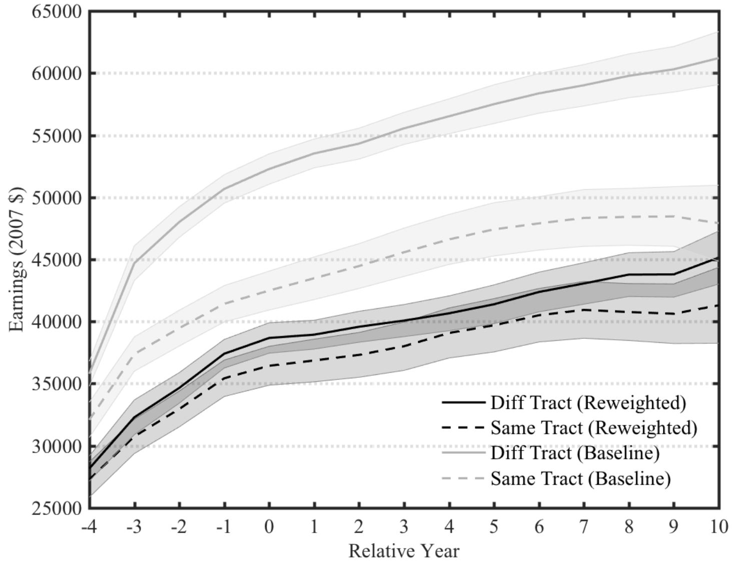

Panel A: Reweighted Means Without Age Trend

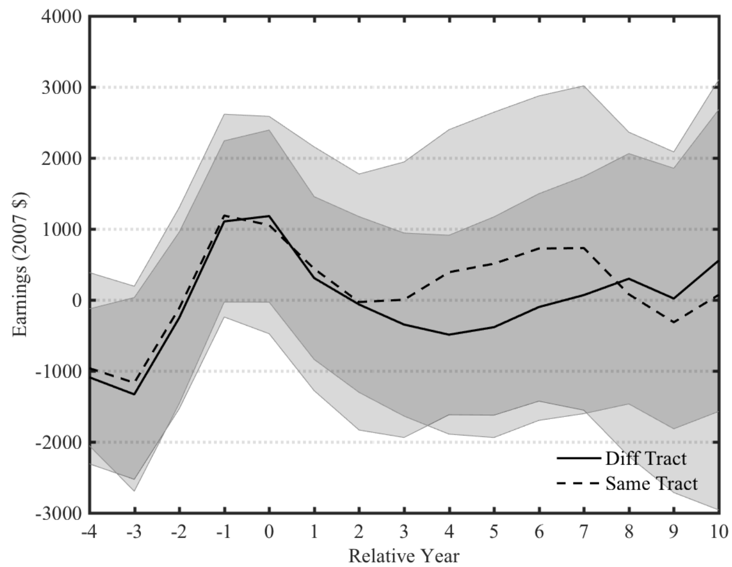

Panel B: Reweighted Means Including Age Trend

Note: The propensity score reweighting appears to be successful in terms of reweighting each control group so that they are mostly similar to each other. Panel A shows that after applying the propensity score weights, but without adjusting for age quartics, non-displaced workers living close to their parents and farther away in relative year '-1' have similar earnings trajectories, except for a small level shift. Panel B shows the average earnings for the non-displaced after removing an age quartic. Not surprisingly, the differences that remain between the two groups are quite small.

Figure B13: Including Additional Interactions in the Baseline Specification

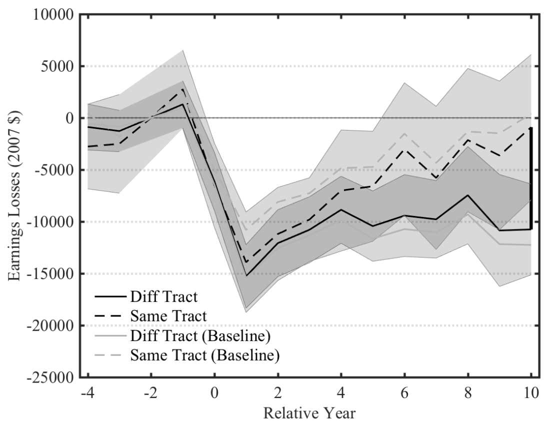

Panel A: Baseline Income as Linear Interaction

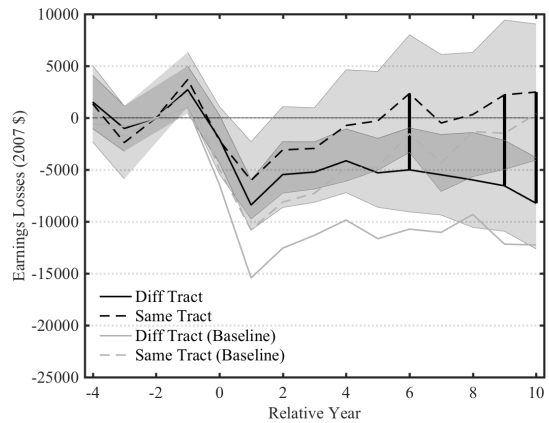

Panel B: Baseline Income as a Dummy

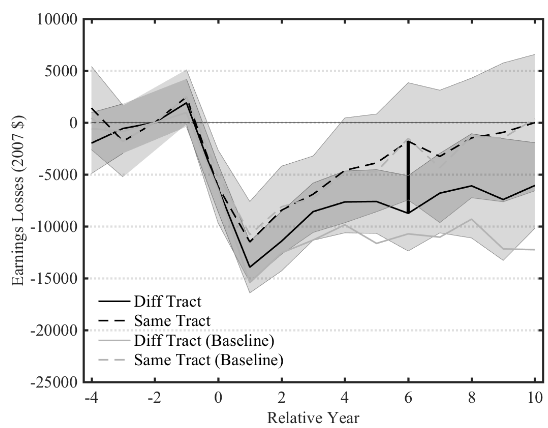

Panel C: College Education as a Dummy

Note: Additional interactions with the displacement dummies (equation B.1) do not change the effect of parental proximity on the post-displacement earnings outcomes. Although interacting with earnings prior to job loss generally makes the initial earnings losses similar for the two groups, the two paths still diverge later on. Plotted are regression coefficients from estimating equation (B.1). Panel A shows the coefficients when one includes an additional interaction with a linear term earnings, Panel B shows the results after including an interaction with a dummy for having above average earnings, and Panel C shoes results after including an interaction with a dummy for being college educated. Lighter lines reproduce the baseline results from Figure 2.2. Appendix B.4 includes more details on the specification.

Figure B14: Reweighting on the Subsample with Common Support

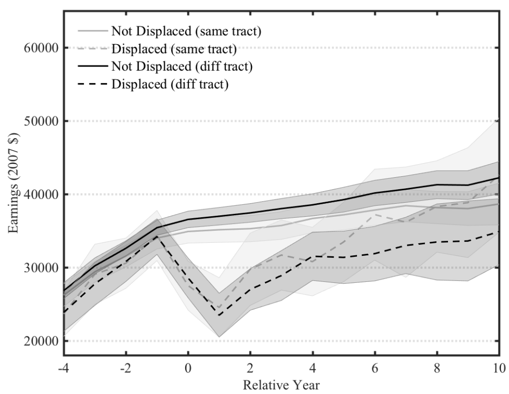

Panel A: Reweighted Means

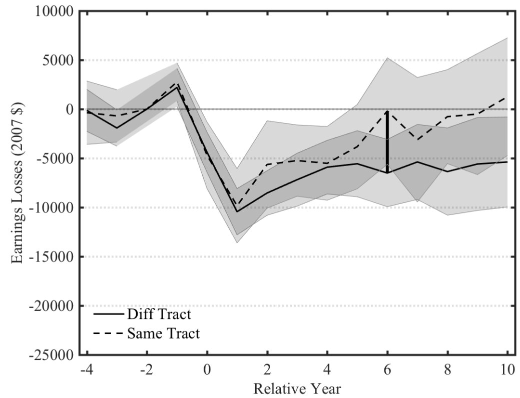

Panel B: Reweighted Regressions

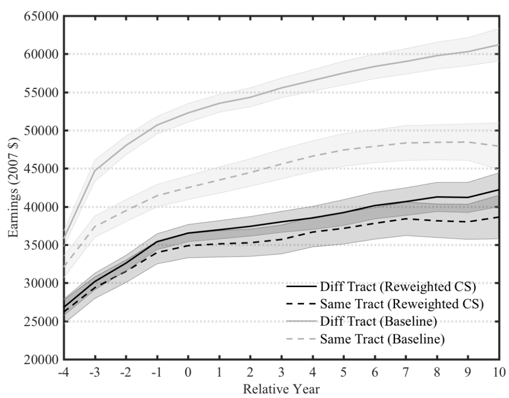

Panel C: Reweighted Means Without Age Trend

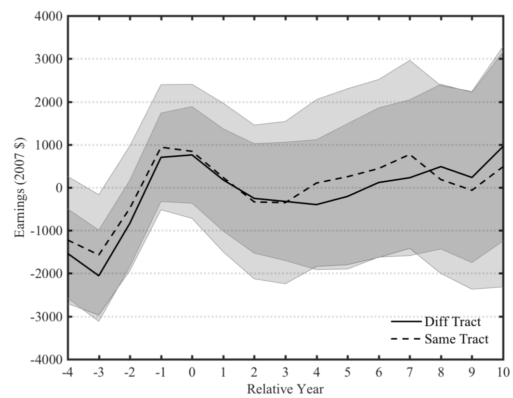

Panel D: Reweighted Means Including Age Trend

Note: Restricting the sample to a subset of observations with common support gives qualitatively similar results to reweighting the whole sample, though with less precision. These panels replicate Figure 2.9, Figure 2.10, and Appendix Figure B12, but with datasets where the sample is restricted to observations where there is common support between the group that was displaced at home and the various other reweighting groups. Panel A shows propensity score reweighted sample means around a potential displacement, Panel B shows propensity score weighted regression coefficieints around a displacement, Panel C shows average earnings among young workers who are not displaced, and Panel D shows average earnings for workers who are not displaced, once a quartic trend in the worker's age is removed. Figure 2.9, Figure 2.10, and Appendix Figure B12 contain more details on the methodology.

Figure B15: Probability of Switching Industries

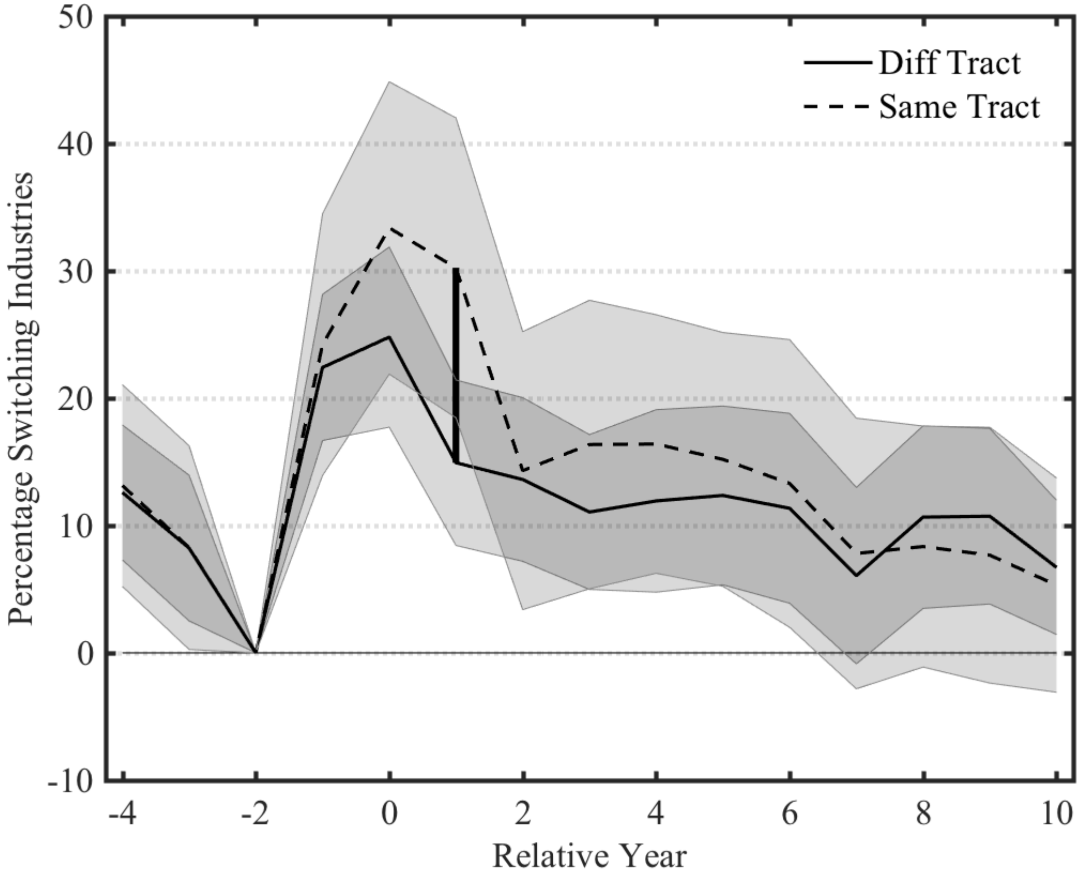

Note: Those who live in the same tract as their parents prior to displacement are more likely to switch industries at the time of job loss than those who live farther away. These figures plot regression coefficients from equation (2.1) describing the impact of a job displacement on switching one-digit PSID industries. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. We cluster the standard errors at the worker level.

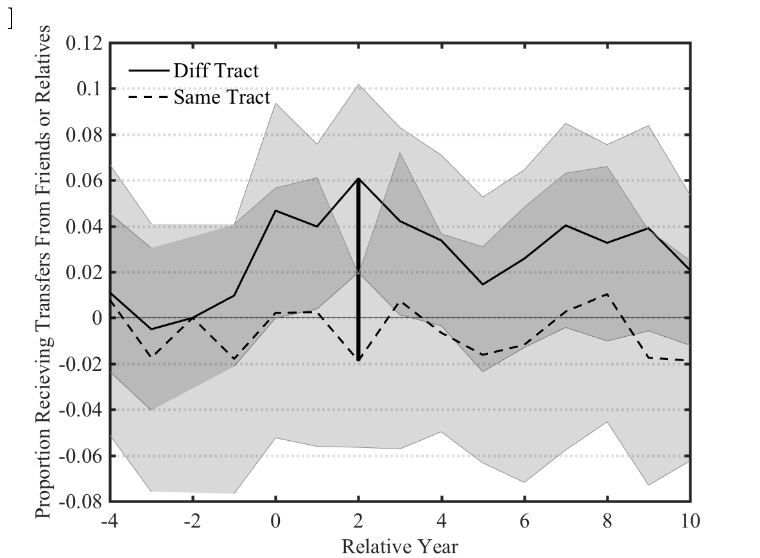

Panel A: Receiving a Monetary Transfer

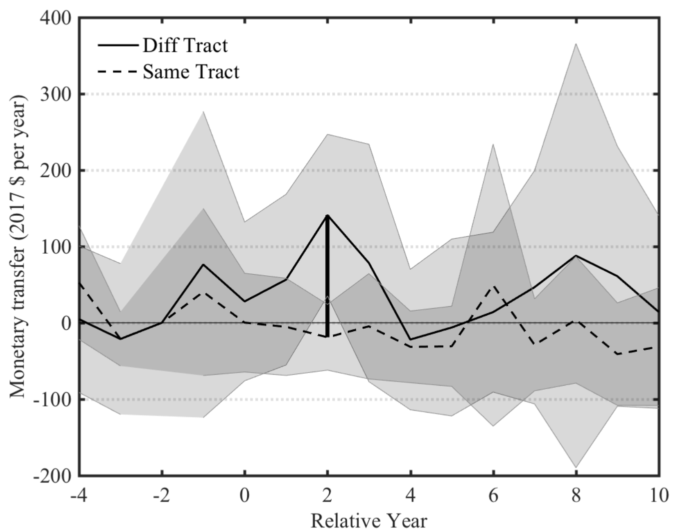

Panel B: Total Monetary Transfers

Note: The annual question in the PSID asking about help from friends or relatives gives relatively noisy results around a displacement, but there is some evidence that workers who live farther from their parents are more likely to receive a small amount of monetary help after a displacement. These figures plot regression coefficients from equation (2.1) describing the impact of a job displacement on monetary transfers. The measure in Panel A is the proportion of workers who report that they received help from friends or family. The measure in Panel B is the reported dollar value per year. The shading represents 95 percent confidence intervals, and any vertical bars connecting the two lines signify that the estimates are statistically significantly different from one another in that year, at the five percent level. Shading and statistical significance are based on standard errors computed by clustering at the worker level. See Figure 2.11 and Appendix B.6 for for more details on the specification and the data.

Figure B17: Working in a Parent's Industry for Older Workers

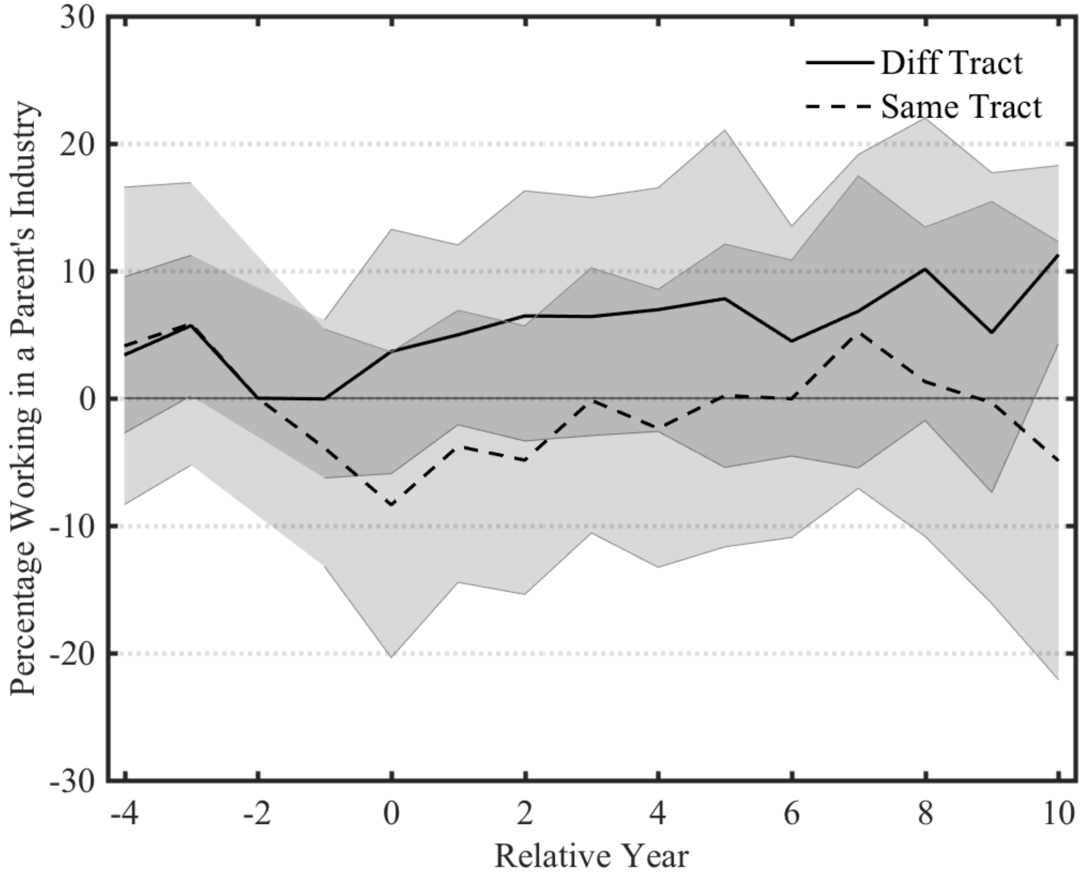

Note: Older workers who do not live near their parents may be more likely to work in their parents' industries after a displacement event, though the estimates are quite imprecise. These figures plot regression coefficients from equation (2.1) describing the impact of a job displacement on the proportion of older (36 to 55 year old) workers who work in the same one-digit PSID coded industry as a parent. The shading represents 95 percent confidence intervals, and the lack of vertical bars connecting the two lines signify that none of the estimates are statistically different from one another at the five percent level in any years. Shading and statistical significance are based on standard errors computed by clustering at the worker level.

Figure B18: Working in a Parent's Industry by Commuting Zone

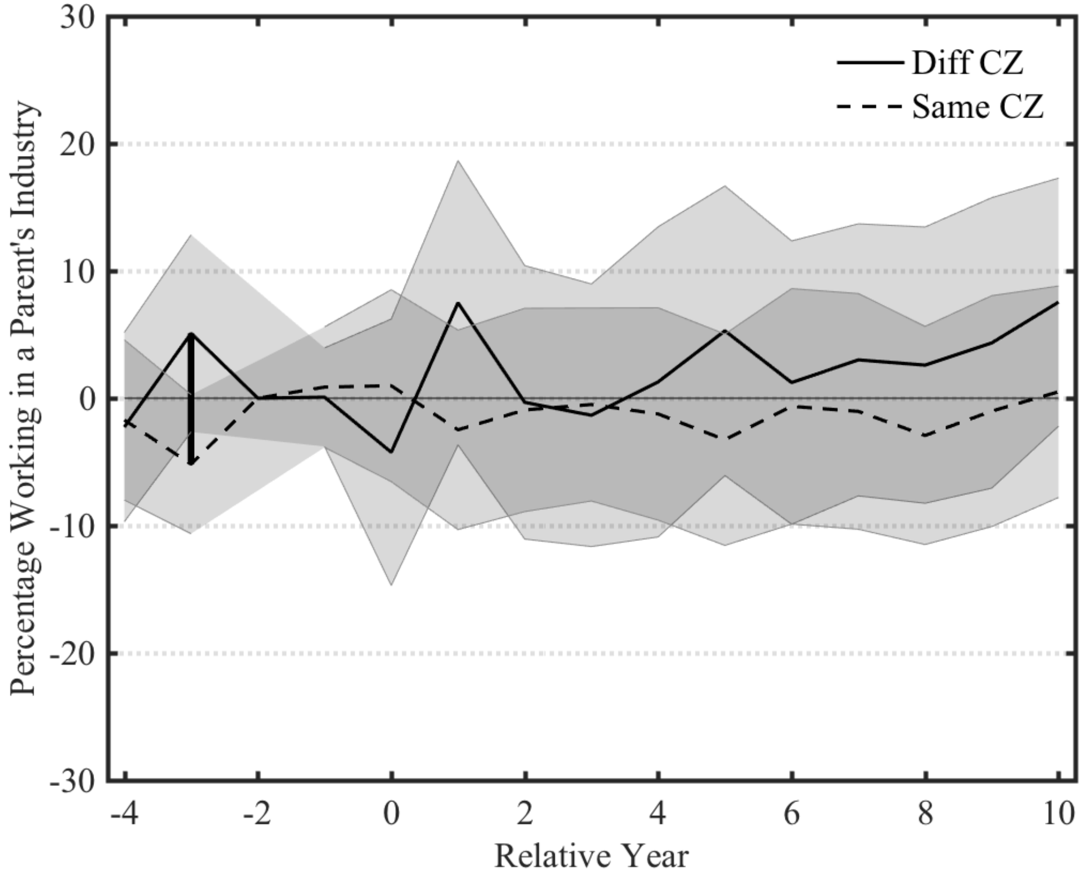

Note: Younger workers who live in the same CZ as their parents are not any more likely to work in their parents' industries after a displacement, though the estimates are imprecise. These figures plot regression coefficients from equation (2.1) describing the impact of a job displacement on the proportion of younger workers who work in the same one-digit PSID coded industry as a parent. The shading represents 95 percent confidence intervals, and vertical bars connecting the two lines signify that the estimates are statistically different from one another at the five percent level. Shading and statistical significance are based on standard errors computed by clustering at the worker level.

B.8 Appendix Tables (For Online Publication)

Table B1: Means Before and After Reweighting the Sample with Common Support

Variable	Panel A: PSID Weights				Panel B: Reweighted			
	Same Tract		Different Tract		Same Tract		Different Tract	
	Displaced	Not Displaced	Displaced	Not Displaced	Displaced	Not Displaced	Displaced	Not Displaced
Earnings	$31,600	$36,500	$38,200	$41,900	$31,600	$31,700	$31,000	$32,800
	[1.00]	[0.03]	[0.00]	[0.00]	[1.00]	[0.36]	[0.43]	[0.93]
Average Change in Earnings	$2,100	$2,100	$3,000	$3,000	$2,100	$2,000	$3,400	$2,400
	[1.00]	[0.36]	[0.97]	[0.39]	[1.00]	[0.25]	[0.64]	[0.58]
Years of Schooling	11.96	12.45	12.29	12.72	11.96	12.12	12.01	12.15
	[1.00]	[0.00]	[0.00]	[0.00]	[1.00]	[0.84]	[0.61]	[0.93]
Share in Goods Industries	0.55	0.46	0.59	0.39	0.55	0.50	0.54	0.42
	[1.00]	[0.06]	[0.81]	[0.00]	[1.00]	[0.33]	[0.86]	[0.02]
Share Manager/Professional	0.13	0.18	0.15	0.21	0.13	0.16	0.11	0.15
	[1.00]	[0.11]	[0.03]	[0.00]	[1.00]	[0.92]	[0.54]	[0.91]
Employer Tenure	5.26	6.54	5.13	6.45	5.26	5.29	5.17	5.30
	[1.00]	[0.00]	[0.65]	[0.00]	[1.00]	[0.81]	[0.76]	[0.98]
Unemp Rate in County	7.66	7.31	7.87	7.19	7.66	7.50	7.94	7.62
	[1.00]	[0.34]	[0.98]	[0.04]	[1.00]	[0.89]	[0.45]	[0.96]
Age	27.97	28.57	28.30	28.53	27.97	27.70	27.60	27.76
	[1.00]	[0.00]	[0.02]	[0.00]	[1.00]	[0.34]	[0.32]	[0.43]
Number of Children	1.37	1.25	1.23	1.22	1.37	1.18	1.27	1.22
	[1.00]	[0.72]	[0.32]	[0.26]	[1.00]	[0.42]	[0.83]	[0.56]
Fraction Male	0.82	0.80	0.83	0.85	0.82	0.81	0.75	0.83
	[1.00]	[0.82]	[0.34]	[0.39]	[1.00]	[0.89]	[0.36]	[0.73]
Number of Records	177	3,021	351	8,394	177	3,021	351	8,394

Note: After applying the propensity score weights on the sample with common support, the sample of workers who live in the same tract as their parents and those living farther away are statistically indistinguishable in terms of many observable characteristics. This table reports means for each group in the sample with common support using PSID weights in the first four columns and the propensity score weights in the last four columns. For each variable, we report the mean and a p-value in brackets of a Wald test that this mean is the same as the value in the first column. See the initial version, Figure 2.2, for more details on the table specification.

Table B2: Summary Statistics of Search Intensity by Proximity to Parents and Labor Force Status

	Unemployed (A)	Unemployed (H)	Working (A)	Working (H)	Pooled (A)	Pooled (H)
Panel A: Young Workers						
$\mathbb{P}[\textit{any search activity}]$	0.86	0.86	0.070	0.068	0.080	0.094
$\mathbb{P}[\textit{checked w/ frnds or rels}]$	0.25	0.32	0.023	0.018	0.028	0.029
$\mathbb{P}[\textit{searched but not w/ frnds or rels}]$	0.62	0.72	0.062	0.061	0.070	0.088
Panel B: Older Workers						
$\mathbb{P}[\textit{any search activity}]$	0.66	0.75	0.039	0.037	0.048	0.058
$\mathbb{P}[\textit{checked w/ frnds or rels}]$	0.22	0.22	0.013	0.015	0.017	0.023
$\mathbb{P}[\textit{searched but not w/ frnds or rels}]$	0.53	0.63	0.035	0.034	0.042	0.054

Note: Young workers living in the same neighborhoods as their parents are more likely to engage in search activities than young workers living farther away (pooled results). A similar pattern holds for older workers. Parenthetical (A) stands for "away," i.e. those not in the same neighborhoods as their parents at the time of the survey, and (H) stands for "home," i.e. those living in their parents' neighborhoods.

Table B3: Any Job Search Activity for Young Workers

	(1)	(2)	(3)	(4)	(5)
	Unemployed	Unemployed	Unemployed	Working	Pooled
Same tract	-0.005	-0.059	-0.082	-0.001	0.003
	(0.077)	(0.100)	(0.115)	(0.009)	(0.013)
Unemployment					0.75***
					(0.050)
Observations	1,741	1,509	1,509	57,460	58,969
R-squared	0.000	0.038	0.085	0.020	0.140
Demographic controls	NO	YES	YES	YES	YES
Year FEs	NO	NO	YES	YES	YES
Worker FEs	NO	NO	NO	NO	YES

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

Note: Young workers (25 to 35 year olds) living in the same neighborhoods as their parents are no more likely to engage in search activities than young workers who live farther away. This does not depend on employment status. Standard errors adjust for clustering at the worker level.

APPENDIX C

Housing Inequality

C.1 Data appendix

C.1.1 Integrated Public Use Microdata Series samples

Our Census data are provided by the Integrated Public Use Microdata Series, from usa.ipums.org detailed in Ruggles et al. (2010). For each year we use the following samples.

- 1930: 5 percent sample
- 1940: 1 percent sample
- 1960: 1 percent
- 1970: 1 percent state and metro (depending on state or commuting zone) and format 1 or 2 (depending on the questions needed).
- 1980: 5 percent state sample
- 1990: 5 percent sample
- 2000: 5 percent sample
- 2009-12 ACS five year 2012 sample with interviews in 2008 dropped

C.1.2 Sample restrictions

As mentioned in the main text, we include households in the United States with three restrictions: We omit houses outside the continental US, we exclude houses used to generate revenue, and we similarly exclude owner occupied multiple family residences. Appendix table C1 shows the percentage of houses we exclude for each year due to being used to generate revenue or encompassing multiple units. We present these separately for both owners and renters.

The first row for each panel shows the number of houses in the continental US that are included in our sample. this is quite high for each year, and in later years it is over 90 percent. The next few rows show reasons why houses are excluded. By far the most important reason is because they are farms. For example, in 1930 25 percent of households lived on a farm. Multiple family houses and houses used for commercial purposes are also excluded in our analysis for the years where the census identifies them. In 1960, the first year we can identify them, these housing categories made up roughly 6 percent of owner occupied housing.

The final few rows show categories of houses that were excluded from questions in some years. From 1960-1980 mobile homes and trailers were excluded and in 1970 condos were as well. In other years they are included. These categories, again, make up a small proportion of housing units in these years, They never reach six percent, and often are well below that mark.

Table C1: Exclusions Imposed in Sample Selection

Panel A: Home Owners

	1930	1940	1960	1970	1980	1990	2000	2012
Included	0.748	0.758	0.843	0.873	0.881	0.910	0.900	0.922
Farm	0.252	0.242	0.093	0.051	0.029	0.023	0.019	0.017
Commercial use			0.013	0.010	0.015	0.020	0.033	0.015
Multiple families			0.051	0.065	0.075	0.047	0.048	0.046
Mobile home or similar			0.015	0.038	0.055	0.082	0.083	0.067
Condominium				0.002	0.007	0.015	0.018	0.032
Sample size	516,941	585,459	275,602	1,388,542	2,292,767	2,826,214	3,244,508	3,128,981

Panel B: Renters

	1930	1940	1960	1970	1980	1990	2000	2012
Included	0.798	0.837	0.977	0.986	0.988	0.989	0.986	0.993
Farm	0.202	0.163	0.023	0.008	0.005	0.004	0.003	0.002
Commercial use				0.007	0.007	0.007	0.011	0.005
Mobile home or similar			0.003	0.011	0.023	0.050	0.046	0.044
Condominium					0.017	0.045		
Sample size	629,593	850,940	175,464	848,785	1,296,401	1,346,603	1,505,211	1,281,471

NOTE: The gives the (unweighted) numbers of houses falling into categories that either merit exclusion or fall out of the universe of houses we have data for in some years. Farms, units used for commercial purposes, and multiple family owner occupied structures are excluded from our statistics. Mobile homes (as well as tents, vans, trailers, and boats) and condominiums are not asked about in certain years, so they are absent in those years but included when they are asked about.

C.1.3 Questionnaires

Relevant enumerator instructions for 1940:

Figure C1: Enumerator Instructions for 1940

431. ***Column 5. Value of Home, if Owned, or Monthly Rental, if Rented.***—If the home is owned, as indicated by the entry "O" in col. 4, enter in col. 5, on the line for the head of the household, the current market value of the home, as nearly as it can be ascertained. Unless the home has been recently purchased, it will be necessary to estimate its value. The estimate should represent the amount for which the home, including (except on a farm) such land as belongs to it, would sell under ordinary conditions—not at forced sale. The assessor's valuation, on which taxation is based, is usually not a safe guide.

432. Where a person owns a house with living accommodations for more than one household and his household occupies only a portion of the house, as where the owner of a two-family house rents part to another household, estimate the value of the portion of the house occupied by the owner's household (which for a two-family house may be about one-half of the total value), and enter this amount in col. 5 for the owner's household. The entry in col. 5 for the household or households renting a portion of the structure will be the amount paid in monthly rental. Where any considerable portion of the house is used for business purposes, such as a store, deduct the value of this portion—except that the value of one or two rooms used as an office by a dentist, lawyer, or contractor, etc., need not be deducted.

433. For the home of a farm operator who owns, and lives on, his farm (or who owns that part of the farm on which the dwelling stands), obtain an estimate of the value of the dwelling in which he lives, *excluding the land on which it is built.* (This figure should represent a reasonable fraction of the value of all farm buildings reported on the Farm schedule.)

434. Make it clear to your informant that the values returned on the census schedule are not to be used in any way in connection with taxation and are not open to public inspection.

435. If the home or dwelling unit is rented, as indicated by "R" in col. 4, enter in col. 5 to the nearest dollar the actual amount paid each month as rent, or enter one-twelfth of the annual rental, in case payment is not made monthly. *Do not enter fractions of a dollar.*

436. If no money rent is paid, as where a workman receives the use of a house as part of his wages, enter in col. 5 the estimated monthly rental value based on the monthly rental paid for similar dwelling units in the neighborhood.

437. In the case of a tenant farm operator, that is, one who pays rent in some form for the farm, including his dwelling (rather than for the dwelling alone), estimate the monthly rental value of the dwelling in which he lives. This estimate should be based, if possible, on the rent actually paid for similar dwellings nearby, making allowance for the fact that rents are usually lower in the open country than in town.

438. If there is no other basis for estimating the rental value of the home of a farm tenant (or in some instances a nonfarm tenant), you may consider that 1 percent of the total value of the dwelling is a fair monthly rental. For example, if $1,000 seems to be a reasonable estimate of the total value of the dwelling, enter $10 as the monthly rental value.

439. Whenever the value reported to you for a dwelling seems a great deal higher or lower than the value for similar structures in the same neighborhood, question your informant further to make sure that he has properly understood the question and that the value is the current market value of the living quarters.

C.2 Pareto Interpolation and Extrapolation

Table C2: House Value Interval Boundaries Over Time

Year	Intervals
1960	10
1970	11
1980	24
1990	25
2000	24

NOTE: The gives the numbers of intervals for each year the census data are intervalled.

We use a pareto distribution to interpolate values both in years where data are intervalled and where they are continuous. Where data are intervalled, we interpolate within the specified intervals. Where data are continuous, we divide the data into equally size intervals. We start with 25, but decrease the number until each interval has a reasonable "width," or distance form its minimum to maximum value.[1]

The pareto distribution's cumulative distribution function (CDF) has the from $F(y) = 1 - (k/y)^{\alpha}$ for $y \geq k$. It follows that the logarithm of the complementary CDF (or "tail" function) has a slope of $-\alpha$ with respect to $\ln(y)$. The greater the slope α, the faster observations "die off," and the less disperse the distribution. For any value y_i, define

$$p_i \equiv ln[1 - F(y)] = \alpha \ln(k) - \alpha \ln(y) \tag{C.1}$$

The interpolation method (Pareto (1896)) estimates α by differencing expression C.1 at the values for the two endpoints of a bracket $[y_i, y_{i+1}]$, where $i = 1, ..., I$ indexes the intervals, namely

$$\hat{\alpha}_i = \frac{\ln(y_{i+1}) - \ln(y_i)}{p_i - p_{i+1}}, i = 1, ..., I - 1, \tag{C.2}$$

and $\hat{k}_i = y_i p_i^{1/\alpha}$. Even when the distribution may not be considered Pareto globally, this procedure fits a simple line segment to the log survivor function in terms of $\ln(y)$.

Since the Pareto distribution is unbounded, determining a value of α beyond the top code, y_I in the data can be difficult. This problem is largely remedied if the mean is available, as it is for housing values in 1930, 1940, and 2012. There, the property that the mean value of $\bar{y}_I \equiv E[y|y \geq y_I] = \alpha y_I/(\alpha - 1)$, imples a value of

$$\hat{\alpha}_I = \frac{\bar{y}_I}{\bar{y}_I - y_I} \tag{C.3}$$

In other situations it is impossible to obtain a mean above a given cutoff. In these cases we infer a value of α by taking the average of values from the linear interpolation method (C.2) for the top 10 percent of the distribution that do not correspond to the top interval.[2]For rent values in 1930 and 1940, we code the top pareto parameter to be 2.5, which roughly corresponds to this.

For the bottom-most interval we employ a simple strategy of imposing a uniform dis-

[1]The precise definition is that the procedure will decrease the number of intervals until each interval has a max value that is 1.01 times its min value or greater. We do this because the procedure fails for intervals that are excessively "narrow" as sometimes occurs where respondents round values.

[2]In the extremely rare case where all values in the top 10 percent are in the top bin, we check the top 11, 12, 13, etc. percent, stopping at the first percentile where an estimate exists.

tribution so our choices impact our results as minimally as possible. To do this we need to establish a minimum possible value for the variable we are concerned with. The bottom value is determined to be, for rents, 0.03 percent of average household incomes per year (corresponding to a household that makes 10 percent of the average household income and spends 1/3 of it on rent). For home values it is the same value, but divided by 0.0785, which is the rent to value ratio we use to convert housing values into user costs of housing. For household incomes it is the 0.1 times average household income in that year. Average (national) household incomes for this purpose across the many years in our sample come from Chao and Utgoff (2006). How we treat bottom values matters little except for in the variance of logarithms, which has undesirable properties at the bottom of the distribution (for instance, it cannot account for zero values).

Figure C2 shows several illustrations of our imputation procedure for different types of situations. We plot both the CDF and the tail function, which is linear in the pareto distribution. The second plot is especially relevant for determining the fit in the very top of the distribution.

Panel A of Figure C2 shows both the empirical cumulative density function for home prices in 2012 and the pareto CDF used in the paper. Note that this is a year where respondents reported a continuous quantity, though the distribution plot shows that many did round. The dashed lines show intervals in each bin used for rounding. As in other years, we divided the data into 25 bins to perform the CDF interpolation. The dashed lines show the min and max proportions for each bin, as well as its placement along the scale.

Figure Panel B, which shows a plot of the tail function, which shows how the Pareto distribution matches the top of the distribution. Overall the fit is quite good, with the line tracking the distribution until it encounters the state by year topcodes for the top 0.5 percent of the distribution.

Similarly Panel C and Panel D show the same for the 1970 distribution where the variable is divided into 11 intervals. Here the distribution meets the interval boundaries exactly and the interpolation simply smooths the area between them. The top code extrapolation has the same slope as the one for 2012, but this slope matches the slope of the previous interval rather well, implying the extrapolation is reasonable for this year as well.

Figure C2: Example CDF Plots for the Pareto Interpolation

Panel A: CDF of House Values for 2012

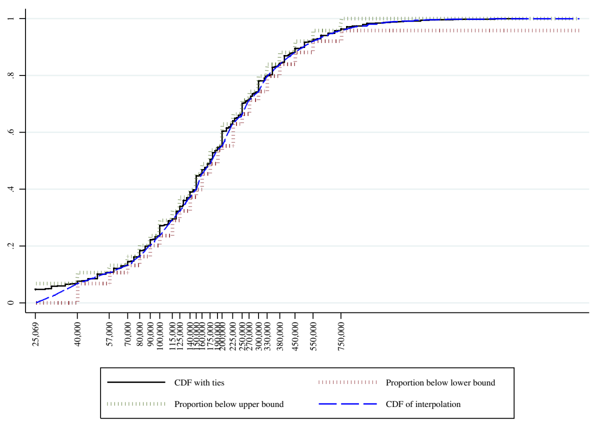

Panel B: Tail Function of House Values for 2012

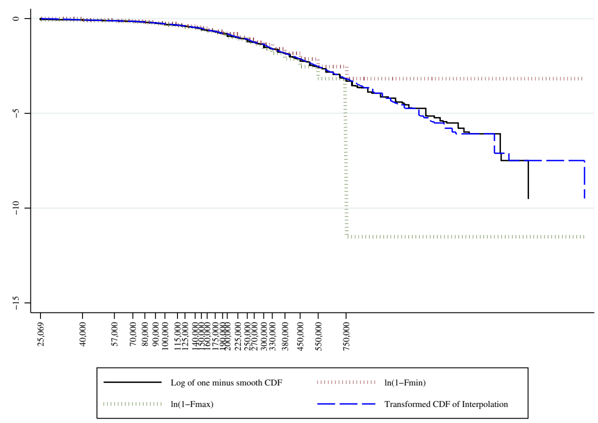

Panel C: CDF of House Values for 1970

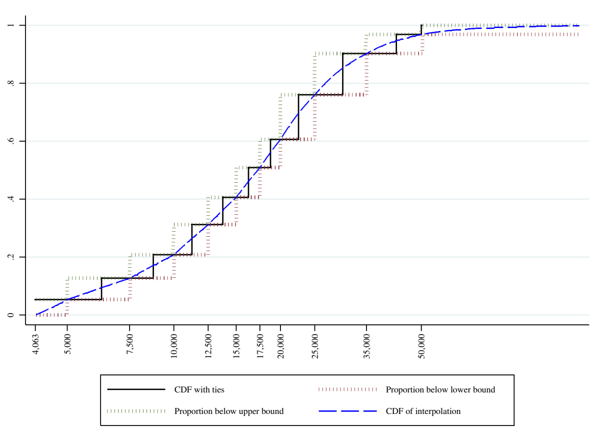

Panel D: Tail Function of House Values for 1970

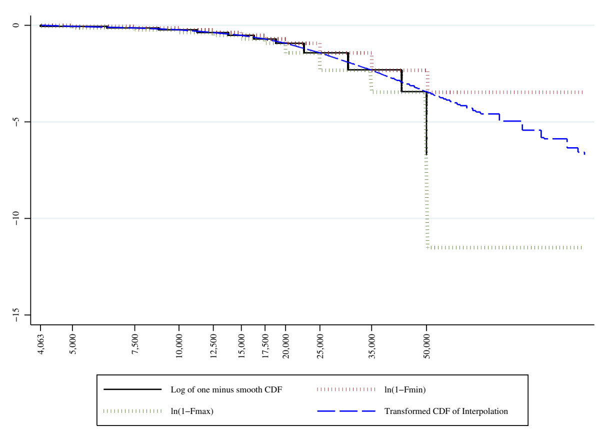

NOTE: The term “tail function” corresponds to the logarithm of one minus the the CDF.

C.3 Geographic Assignment

We keep a consistent geographic sample across all years in the sample by primarily using 1990 local labor markets, or commuting zones, as described by Tolbert and Sizer (1996). Local labor markets are a natural geography for this analysis since housing values are closely related to local income levels. For 1970-2011 we match geographically using the probabilistic mapping made available by David Dorn via his website and used in, for example, Dorn (2009). For 2012 we replicate the probabalistis matching using updated PUMA definitions for the 2012 census and for 1930 and 1940 we use the direct county to commuting zone mapping in Tolbert and Sizer (1996) since counties are available. Unfortunately it is not possible to match to commuing zones in for 1960 based on the publicly available samples, so we instead use states as our geographic entity where we use data from 1960 (this is generally marked in each table).

This mapping uses the most disaggregated geography available in each census to project this geography onto counties and then project these counties onto commuting zones. Unfortunately, the PUMAS and county groups available often contain multiple counties and these counties sometimes fall into separate commuting zones. In these cases we compute a probability that a given person resides in one commuting zone based on the proportion of people in her identified geography that in fact live in the given commuting zone based on non-restricted census summary tables. Where these probabilistic matches occur, we create multiple observations – one for each possible commuting zone – and weight them by the probability the original observation is in that commuting zone, times their initial weight if applicable. Samples for 1930 and 1940 contain county information so the match is exact and we avoid the first projection.

C.4 Deflating

Since the inequality measures we use are scale invariant, deflating is generally unimportant. We deflate values using the Consumer Price Index excluding structures published by the BLS, from the Federal Reserve Bank of St.Louis' FRED system. The series mnemonic is: `CUUR0000SA0L2`. Before 1935, when structures were not identified separately, we combine the data series from the general CPI (`CPIAUCNS`) and the CPI for rent of primary residence (`CUUR0000SEHA`) by using an OLS regression to predict the shelters excluded series for the years it is available prior to 1941 based on the general CPI and the CPI for rent of primary residence. This regression has good out of sample prediction powers for years up to 1960, so we believe it is a reasonable approximation for the initial five years.

C.5 Re-weighting decomposition specifics

Formally, using notation similar to Fortin et al. (2011), we have two groups of houses, one group of houses observed in t and another in t'. We have prices for each group and some houses may appear in both groups. We want to compute a statistic, $\nu(\mathcal{P})$, that is a function of a distribution of house prices, $\mathcal{P} \to \mathbb{R}$. With the above caveats in mind, we denote the distribution $\mathcal{P}_{t,t'}$ where the first subscript refers to the year we would like to apply prices from and the second year is the year we would like to match in terms of observables. So we re-weight observed houses in t so that they corresponded to the distribution of characteristics of houses in t'. Note that in this notation $\mathcal{P}_{t,t}$ is simply the distribution of houses in period t. Putting this together, we are seeking to recover:

$$\nu\left(\mathcal{P}_{t,t'}\right) \tag{C.4}$$

Conditional independence, or ignorability, in this context means that once we condition on all observable characteristics (S and L), the unobservable characteristics (ϵ) are independent of whether we observe the house in t or t'. Here this would imply that once we take into account each observable detail about the house there is no special trait which we omitted that differs systematically between t and t'.

To actually estimate $\nu\left(\mathcal{P}_{t,t'}\right)$, we compute re weighting factors following DiNardo et al. (1996) which Hirano et al. (2003) and Firpo (2007) formally establish are efficient. $\Phi(S, L)$ with the following form:

$$\Phi(S, L) = \frac{\Pr(\tau = t'|S, L)/\Pr(\tau = t')}{\Pr(\tau = t|S, L)/\Pr(\tau = t)} \tag{C.5}$$

Where τ represents the year the house was observed in (either t' or t). The intuition is that we would like to heavily weight the observations from year t that look very much like the observations in year t'. We do this by using the odds that an observation came from t' given its observables.

We compute the probabilities using logit regressions that pool houses observed in both t and t'. By regressing a dummy indicator of a house being observed in t' on a flexible functional form of S and L In this context we mostly use dummy variables for different categories. we are able to compute predicted probabilities $\hat{\Pr}(\tau = t'|S, L)$ and $\hat{\Pr}(\tau = t|S, L)$. To compute $\hat{\Pr}(\tau = t')$ and $\hat{\Pr}(\tau = t)$ we simply use the sample proportions (appropriately weighted). Applying equation C.5 then gives a value of $\hat{\Phi}(S, L)$ that we use to re-weight each statistic.

We then multiply $\hat{\Phi}(S, L)$ by the weights already applied to the houses observed in t to simulate the density that would exist under the counter-factual that houses had identical characteristics to houses in t' but were priced in terms of the prices that prevailed in t.

BIBLIOGRAPHY

Aguiar, Mark and Mark Bils. 2011. Has Consumption Inequality Mirrored Income Inequality? Tech. rep., National Bureau of Economic Research.

Aguiar, Mark, Mark Bils, Kerwin Kofi Charles, and Erik Hurst. 2017. Leisure Luxuries and the Labor Supply of Young Men.

Albouy, David. 2009. The Unequal Geographic Burden of Federal Taxation. *Journal of Political Economy* 117, no. 4:635–667.

———. 2016. What are cities worth? Land rents, local productivity, and the total value of amenities. *The Review of Economics and Statistics* 98, no. 3:477–487.

Albouy, David and Bert Lue. 2015. Driving to opportunity: Local rents, wages, commuting, and sub-metropolitan quality of life. *Journal of Urban Economics* 89:74–92.

Albouy, David and Mike Zabek. 2016. Housing Inequality. *NBER Working Paper Series* 21916.

Alesina, Alberto, Yann Algan, Pierre Cahuc, and Paolo Giuliano. 2015. Family Values and the Regulation of Labor. *Journal of the European Economic Association* 14, no. 4:599–630.

Alesina, Alberto and Robert Barro. 2002. Currency Unions. *The Quarterly Journal of Economics* 117, no. 2:409–436.

Arons, Rachel. 2012. A Mythical Bayou's All-Too-Real Peril.

Arrow, Kenneth J. 1950. A Difficulty in the Concept of Social Welfare. *Journal of Political Economy* 58, no. 4:328–346.

Atkinson, Anthony. 1970. On the measurement of inequality. *Journal of Economic Theory* 2:244–263.

Atkinson, Anthony, Thomas Piketty, and Emmanuel Saez. 2011. Top Incomes in the Long Run of History. *Journal of Economic Literature* 49, no. 1:3–71.

Attanasio, Orazio and Luigi Pistaferri. 2014. Consumption inequality over the last half century: Some evidence using the new PSID consumption measure. *American Economic Review* 104, no. 5:122–126.

Autor, David and David Dorn. 2013. The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review* 103, no. 5:1553–1597.

Autor, David, David Dorn, and Gordon Hanson. 2013. The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review* 103, no. 6:2121–2168.

Autor, David, Christopher J. Palmer, and Parag A. Pathak. 2014. Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts. *Journal of Political Economy* 122, no. 3:661–717.

Barro, Robert, Xavier Sala-i Martin, Olivier Jean Blanchard, and Robert E Hall. 1991. Convergence across states and regions. *Brookings papers on economic activity* :107–182.

Bartik, Timothy. 1991. Who Benefits from State and Local Economic Development Policies? Tech. rep., Kalamazoo, MI.

———. 1993. Who Benefits from Local Job Growth: Migrants or the Original Residents? *Regional Studies* 27, no. 4:297–311.

———. 2009. What proportion of children stay in the same location as adults, and how does this vary across location and groups? Tech. Rep. Upjohn Institute Working Paper 09-145.

Bartik, Timothy and Nathan Sotherland. 2015. Migration and Housing Price Effects of Place-Based College Scholarships. *Upjohn Institute working paper* 15, no. 245.

Baum-Snow, Nathaniel and Ronni Pavan. 2013. Inequality and City Size. *Review of Economics and Statistics* 95, no. 5:1535–1548.

Beaudry, Paul, David A. Green, and Benjamin M. Sand. 2014. Spatial equilibrium with unemployment and wage bargaining: Theory and estimation. *Journal of Urban Economics* 79:2–19.

Becker, Gary. 1960. An economic analysis of fertility. In *Demographic and economic change in developed countries*, ed. George Roberts. National Bureau of Economic Research, 209–240.

Becker, Gary and H Gregg Lewis. 1973. On the Interaction between the Quantity and Quality of Children. *Journal of Political Economy* 81, no. 2:S279–88.

Bentolila, Samuel, Claudio Micehacci, and Javier Suarez. 2010. Social Contact and Occupational Choice. *Economica* 77, no. 305:20–45.

Berry, Brian J L and Donald C Dahmann. 1977. Population Redistribution in the United States in the 1970s. *Population and Development Review* 3, no. 4:443–471.

Bishop, Kelly C. 2008. A Dynamic Model of Location Choice and Hedonic Valuation. *Unpublished, Washington University in St. Louis* .

Black, Dan A, Natalia Kolesnikova, Seth G Sanders, and Lowell J Taylor. 2013. Are children normal? *The Review of Economics and Statistics* 95, no. 1:21–33.

Blackorby, Charles and David Donaldson. 1990. A Review Article: The Case against the Use of the Sum of Compensating Variations in Cost-Benefit Analysis. *The Canadian Journal of Economics / Revue canadienne d'Economique* 23, no. 3:471–494.

Blanchard, Olivier and Lawrence Katz. 1992. Regional evolutions. *Brookings Papers on Economic Activity* 23, no. 1:1–73.

Blatter, Marc, Samuel Muehlemann, and Samuel Schenker. 2012. The costs of hiring skilled workers. *European Economic Review* 56, no. 1:20–35.

Bonnet, Odran, Pierre-Henri Bono, Guillaume Chapelle, and Etienne Wasmer. 2014. Does housing capital contribute to inequality? A comment on Thomas Piketty's Capital in the 21st Century. *Sciences Po Economics Discussion Papers* 7.

Bound, John and Harry J. Holzer. 2000. Demand shifts, population adjustments, and labor market outcomes during the 1980s. *Journal of Labor Economics* 18, no. 1:20–54.

Bricker, Jesse and Brian Bucks. 2013. Household Mobility over the Great Recession: Evidence from the U.S. 2007-09 Survey of Consumer Finances Panel.

Bui, Quoctrung and Claire Cain Miller. 2015. The typical american lives only 18 miles from mom. NY Times, The Upshot.

Burda, Michael C and Antje Mertens. 2001. Estimating wage losses of displaced workers in Germany. *Labour Economics* 8, no. 1:15–41.

Busso, Matias, Jesse Gregory, and Patrick Kline. 2013. Assessing the Incidence and Efficiency of a Prominent Place Based Policy. *American Economic Review* 103, no. 2:897–947.

Cadena, Brian C. and Brian K. Kovak. 2016. Immigrants equilibrate local labor markets: Evidence from the great recession. *American Economic Journal: Applied Economics* 8, no. 1:257–290.

Cao, Yipei and Frank P Stafford. 2017. Relocation, Intergenerational Transfers and Extended Family Insurance.

Carrington, William J, Enrica Detragiache, and Tara Vishwanath. 1996. Migration with Endogenous Moving Costs. *The American Economic Review* 86, no. 4:909–930.

Chao, Elaine L. and Kathleen P. Utgoff. 2006. 100 Years of US Consumer Spending Data for the Nation, New York City, and Boston. Tech. Rep. Report 991, United States Department of Labor.

Chari, Amalavoyal V, John Engberg, Kristin N Ray, and Ateev Mehrotra. 2015. The Opportunity Costs of Informal Elder-Care in the United States: New Estimates from the American Time Use Survey. *Health Services Research* 50, no. 3:871–882.

Chetty, Raj. 2009. Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods. *Annual Review of Economics* 1, no. 1:451–488.

Chetty, Raj and Nathaniel Hendren. 2017. The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effect. Tech. Rep. May.

Chetty, Raj, Nathaniel Hendren, and Lawrence F Katz. 2015. The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. *American Economic Review* , no. August:90.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States. *Quarterly Journal of Economics* 129, no. 4:1553–1623.

Chipman, Js and Jc Moore. 1980. Compensating variation, consumer's surplus, and welfare. *The American Economic Review* 70, no. 5:933–949.

Coate, Patrick. 2017. Parental Influence on Labor Market Outcomes of Young Workers.

Coen-Pirani, Daniele. 2010. Understanding gross worker flows across U.S. states. *Journal of Monetary Economics* 57:769–784.

Compton, Janice and Robert A Pollak. 2015. Proximity and Co-residence of Adult Children and their Parents in the United States: Descriptions and Correlates. *Annals of Economics and Statistics* , no. 117/118:91–114.

Corak, Miles and Patrizio Piraino. 2011. The Intergenerational Transmission of Employers. *Journal of Labor Economics* 29, no. 1:37–68.

Coulson, N. Edward and Lynn M. Fisher. 2009. Housing tenure and labor market impacts: The search goes on. *Journal of Urban Economics* 65, no. 3:252–264.

Cowell, Frank. 2011. *Measuring inequality*. Oxford University Press.

Crump, Richard K., V. Joseph Hotz, Guido W. Imbens, and Oscar A. Mitnik. 2009. Dealing with limited overlap in estimation of average treatment effects. *Biometrika* 96, no. 1:187–199.

Cutler, David M. and L. F. Katz. 1992. Rising inequality? Changes in the distribution of income and consumption in the 1980's. *American Economic Review* 82, no. 2:546–551.

Dalton, Michael. 2013. Family Transfers in Response to Unemployment.

Davis, Morris A., Jonas D. M. Fisher, and Marcelo Veracierto. 2013. Gross Migration, Housing and Urban Population Dynamics. Tech. Rep. FRB of Chicago Working Paper 2013-19.

Davis, Morris A. and Jonathan Heathcote. 2007. The price and quantity of residential land in the United States. *Journal of Monetary Economics* 54, no. 8:2595–2620.

Davis, Steven J. and Till M. von Wachter. 2011. Recessions and the costs of job loss. *Brookings Papers on Economic Activity* Fall, no. 1:1–72.

Diamond, Rebecca. 2016. The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *American Economic Review* 106, no. 3:479–524.

DiNardo, John, Nicole M Fortin, and Thomas Lemieux. 1996. Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach. *Econometrica* 64, no. 5:1001–1044.

Donovan, Colleen and Calvin Schnure. 2011. Locked in the House: Do Underwater Mortgages Reduce Labor Market Mobility?

Dorn, David. 2009. Essays on Inequality, Spatial Interaction, and the Demand for Skills. Ph.D. thesis.

Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker. 2017. Factors Determining Callbacks to Job Applications by the Unemployed: An Audit Study. *RSF: The Russell Sage Foundation Journal of the Social Sciences* 3, no. 3:168–201.

EU. 2012. Consolidated versions of the Treaty on European Union and the Treaty on the Functioning of the European Union. *Official Journal of the European Union* 55, no. C 326:1–407.

Fallick, Bruce C. 1996. A Review of the Recent Empirical Literature on Displaced Workers. *Industrial and Labor Relations Review* 50, no. 1:5–16.

Farber, Henry S. 2011. Job Loss in the Great Recession: Historical Perspective from the Displaced Workers Survey, 1984-2010. *National Bureau of Economic Research Working Paper Series* No. 17040, no. 5696:1984–2010.

Feenstra, Robert C., Maurice Obstfeld, and Katheryn N. Russ. 2014. In search of the Armington elasticity. *NBER Working Paper* , no. 20063:1–44.

Feler, Leo and Mine Z Senses. 2015. Trade Shocks and the Provision of Local Public Goods .

Ferreira, Fernando, Joseph Gyourko, and Joseph Tracy. 2010. Housing busts and household mobility. *Journal of Urban Economics* 68, no. 1:34–45.

Firpo, Sergio. 2007. Efficient semiparametric estimation of quantile treatment effects. *Econometrica* 75, no. 1:259–276.

Fortin, Nicole, Thomas Lemieux, and Sergio Firpo. 2011. Chapter 1 Decomposition Methods in Economics. In *Handbook of labor economics*, vol. 4. 1–102.

Friedman, Milton. 1957. *A theory of the consumption function.* Princeton: Princeton University Press.

Ganong, Peter and Daniel Shoag. 2012. Why Has Regional Convergence in the U.S. Declined? *SSRN Electronic Journal* .

Gemici, Ahu. 2011. Family Migration and Labor Market Outcomes.

Glaeser, Edward and Joshua Gottlieb. 2008. The Economics of Place-Making Policies. *Brookings Papers on Economic Activity* 2008, no. 1:155–253.

Glaeser, Edward and Joseph Gyourko. 2002. The Impact of Zoning on Housing Affordability. Tech. Rep. NBER Working Paper No. 8835.

———. 2005. Urban Decline and Durable Housing. *Journal of Political Economy* 113, no. 2:345–000.

———. 2008. *Rethinking federal housing policy: how to make housing plentiful and affordable*, vol. 76.

Glaeser, Edward, Joseph Gyourko, and Albert Saiz. 2008. Housing supply and housing bubbles. *Journal of Urban Economics* 64:198–217.

Glaeser, Edward and Jacob Vigdor. 2012. The end of the segregated century: racial separation in America's neighborhoods, 1890-2010.

Goldin, Claudia and Robert A. Margo. 1992. The Great Compression: The Wage Structure in the United States at Mid- Century. *The Quarterly Journal of Economics* 107, no. 1:1–34.

Goldschmidt, Deborah and Johannes F Schmieder. 2015. The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2017. Bartik Instruments : What , When , Why , and How .

Granovetter, Mark S. 1995. *Getting a Job: A Study of Contacts and Careers*. Chicago: University of Chicago Press.

Green, Richard K., Stephen Malpezzi, and Stephen K Mayo. 2005a. Metropolitan-Specific Estimates of the Price Elasticity of Supply of Housing, and Their Sources. *The American Economic Review* 95, no. 2:334–339.

Green, Richard K., Stephen Malpezzi, and Stephen K. Mayo. 2005b. Metropolitan-specific estimates of the price elasticity of supply of housing, and their sources.

Gregory, Jesse. 2013. The Impact of Post-Katrina Rebuilding Grants on the Resettlement Choices of New Orleans Homeowners. *Unpublished manuscript* .

Guerrieri, Veronica, Daniel Hartley, and Erik Hurst. 2013. Endogenous gentrification and housing price dynamics. *Journal of Public Economics* 100:45–60.

Gyourko, Joseph, Christopher Mayer, and Todd Sinai. 2013. Superstar Cities. *American Economic Journal: Economic Policy* 5, no. 4:167–199.

Gyourko, Joseph, A. Saiz, and A. Summers. 2008a. A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index.

Gyourko, Joseph, Albert Saiz, and Anita Summers. 2008b. A new measure of the local regulatory environment for housing markets: The Wharton Residential Land Use Regulatory Index. *Urban Studies* 45, no. 3:693–729.

Hamermesh, Daniel S. 1989. What Do We Know About Worker Displacement in the U.S.? *Industrial Relations* 28, no. 1:51–59.

Hank, Karsten. 2007. Proximity and Contacts between Older Parents and Their Children: A European Comparison. *Journal of Marriage and Family* 69, no. 1:157–173.

Hanushek, Eric A and John M Quigley. 1980. What is the Price Elasticity of Housing Demand? *Source: The Review of Economics and Statistics* 62, no. 3:449–454.

Harberger, Arnold C. 1964. The Measurement of Waste. *The American Economic Review* 54, no. 3:58–76.

Harris, John and Michael Todaro. 1970. Migration, Unemployment and Development: A Two-Sector Analysis. *American Economic Review* 60, no. 1:126–142.

Hellerstein, Judith K, Mark J Kutzbach, and David Neumark. 2015. Labor Market Networks and Recovery from Mass Layoffs: Evidence from the Great Recession Period.

Henley, Andrew. 1998. Residential Mobility, Housing Equity and the Labour Market. *The Economic Journal* 108, no. 447:414–427.

Hines, James R. 1999. Three Sides of Harberger Triangles. *Journal of Economic Perspectives* 13, no. 2:167–188.

Hirano, Keisuke, Guido W. Imbens, and Geert Ridder. 2003. Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score.

Howard, Greg. 2017. The Migration Accelerator: Labor Mobility, Housing, and Aggregate Demand.

Hsieh, Chang-Tai and Enrico Moretti. 2015. Housing Constraints and Spatial Misallocation. Working Paper 21154, National Bureau of Economic Research.

Hudomiet, Péter. 2014. The role of occupation specific adaptation costs in explaining the educational gap in unemployment .

Huttunen, Kristiina and Kjell G Salvanes. 2015. Job Loss, Family Ties and Regional Mobility. *IZA Discussion Paper* , no. 8780:0–54.

Imbens, G W. 2000. The role of the propensity score in estimating dose-response functions. *Biometrika* 87, no. 3:706–710.

Ioannides, Yannis M and Linda Datcher Loury. 2004. Job Information Networks, Neighborhood Effects, and Inequality. *Journal of Economic Literature* 42, no. 4:1056–1093.

Jackson, Ted. 2016. Stay or go ? Isle de Jean Charles families wrestle with the sea.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. 1993. Earnings losses of displaced workers. *American Economic Review* 83, no. 4:685–709.

Jarosch, Gregor. 2015. Searching for Job Security and the Consequences of Job Loss.

Kaplan, Greg. 2012. Moving Back Home: Insurance against Labor Market Risk. *Journal of Political Economy* 120, no. 3:446–512.

Kaplan, Greg and Sam Schulhofer-Wohl. 2013. Understanding the long-run decline in interstate migration. Tech. Rep. Federal Reserve Bank of Minneapolis Working Paper 697.

———. 2017. Understanding the Long-Run Decline in Interstate Migration. *International Economic Review* 58, no. 1:57–94.

Kennan, John and James Walker. 2011. The effect of expected income on individual migration decisions. *Econometrica* 79, no. 1:211–251.

Kleemans, Marieke. 2015. Migration Choice under Risk and Liquidity Constraints. *Working Paper* , no. December:1–63.

Kleiner, Morris M. and Alan B. Krueger. 2013. Analyzing the extent and influence of occupational licensing on the labor market. *Journal of Labor Economics* 31, no. 2:S173–202.

Kletzer, Lori G. 1998. Job Displacement. *Journal of Economic Perspectives* 12, no. 1:115–136.

Kletzer, Lori G and Robert W Fairlie. 2003. The Long-Term Costs of Job Displacement for Young Adult Workers. *Indusrial and Labor Relations Review* 54, no. 4:682–698.

Kline, Patrick and Enrico Moretti. 2014a. Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority. *The Quarterly Journal of Economics* 129, no. 1:275–331.

———. 2014b. People, Places and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs. *Annual Review of Economics* 6:629–62.

Konrad, K.A. and H. Kunemund. 2002. Geography of the family. *American Economic Review* 92, no. 4:981–998.

Kopczuk, Wojciech. 2004. Top Wealth Shares in the United States, 1916- 2000: Evidence from Estate Tax Returns. In *National tax journal, june 2004*.

———. 2015. What Do We Know about the Evolution of Top Wealth Shares in the United States? *Journal of Economic Perspectives* 29, no. 1:47–66.

Kramarz, Francis and Oskar Nordström Skans. 2014. When Strong Ties Are Strong: Family Networks and Youth Labor Market Entry. *Review of Economic Studies* 81, no. 3:1164–1200.

Krolikowski, Pawel. 2017a. Choosing a Control Group for Displaced Workers.

———. 2017b. Job Ladders and Earnings of Displaced Workers. *American Economic Journal: Macroeconomics* 9, no. 2:1–31.

Krueger, Dirk and Fabrizio Perri. 2006. Does income inequality lead to consumption inequality? Evidence and theory. *Review of Economic Studies* 73, no. 1:163–193.

Kuhn, Peter J. 2002. Summary and Synthesis. In *Losing work, moving on: International perspectives on worker displacement*, ed. Peter J Kuhn, chap. 1. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research, 1–104.

Kurth, Joel and Christine MacDonald. 2015. Volume of abandoned homes 'absolutely terrifying'.

Lerman, Robert and Shlomo Yitzhaki. 1985. Income inequality effects by income. *The Review of Economics and Statistics* 67, no. 1:151–56.

Lin, I-Fen and Hsueh-Sheng Wu. 2010. Does Social Support Mediate the Association Between Functional Disability and Depression?

Lindo, Jason M. 2010. Are Children Really Inferior Goods?: Evidence from Displacement-Driven Income Shocks. *Journal of Human Resources* 45, no. 2:301–327.

Ludwig, Jens, Greg J. Duncan, Lisa A. Gennetian, Lawrence F. Katz, Ronald C. Kessler, Jeffrey R. Kling, and Lisa Sanbonmatsu. 2013. Long-term neighborhood effects on low-income families: Evidence from moving to opportunity. In *American economic review*, vol. 103. 226–231.

Matlack, Janna L. and Jacob L. Vigdor. 2008. Do rising tides lift all prices? Income inequality and housing affordability. *Journal of Housing Economics* 17, no. 3:212–224.

Mayo, Stephen K. 1981. Theory and estimation in the economics of housing demand. *Journal of Urban Economics* 10, no. 1:95–116.

McKinnish, Terra. 2008. Spousal Mobility and Earnings. *Demography* 45, no. 4:829–849.

Meyer, Bruce and James Sullivan. 2010. Consumption and Income Inequality in the US Since the 1960s. *University of Chicago manuscript* .

Mincer, Jacob. 1978. Family Migration Decisions. *Journal of Political Economy* 86, no. 5:749–773.

Modestino, Alicia Sasser and Julia Dennett. 2012. Are American homeowners locked into their houses? The impact of housing market conditions on state-to-state migration.

Molloy, Raven, Christopher L Smith, and Abigail Wozniak. 2011. Internal Migration in the United States. *Journal of Economic Perspectives* 25, no. 3:173–196.

Monras, Joan. 2015. Economic shocks and internal migration .

Moretti, Enrico. 2013. Real wage inequality. *American Economic Journal: Applied Economics* 5, no. 1:65–103.

Munshi, Kaivan and Mark Rosenzweig. 2016. Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap. *American Economic Review* 106, no. 1:46–98.

National Center for Health Statistics. 2014. Data File Documentations, Natality.

Neumark, David and Helen Simpson. 2015. Place-Based Policies. In *Handbook of regional and urban economics*, vol. 5. 1197–1287.

Notowidigdo, Matthew J. 2013. The Incidence of Local Labor Demand Shocks.

Notowidigdo, MJ. 2011. The incidence of local labor demand shocks. Tech. Rep. NBER Working Paper 17167, NBER.

Oswald, Andrew J. 1996. A conjecture on the explanation for high unemployment in the industrialized nations: Part 1. *Warwick economic research papers* , no. 475.

Oswald, Florian. 2015. Regional Shocks, Migration and Homeownership.

Panel Study of Income Dynamics. 2017. Restricted use dataset. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan.

Peiser, Richard B. and Lawrence B. Smith. 1985. Homeownership Returns, Tenure Choice and Inflation. *Real Estate Economics* 13:343–360.

Piketty, Thomas and Emmanuel Saez. 2003. Income Inequality in the United States, 1913–1998. *The Quarterly Journal of Economics* 118, no. February:1–39.

Piketty, Thomas and Gabriel Zucman. 2014. Capital is back: Wealth-income ratios in rich countries, 1700-2010. *The Quarterly Journal of Economics* .

Polinsky, A Mitchell and David T Ellwood. 1979. An empirical reconciliation of micro and grouped estimates of the demand for housing. *The Review of Economics and Statistics* :199–205.

Poterba, James and Todd Sinai. 2008. Tax Expenditures for Owner-Occupied Housing: Deductions for Property Taxes and Mortgage Interest and the Exclusion of Imputed Rental Income. *American Economic Review* 98, no. 2:84–89.

Rainer, Helmut and Thomas Siedler. 2009. O Brother, Where Art Thou? The Effects of Having a Sibling on Geographic Mobility and Labour Market Outcomes. *Economica* 76, no. 303:528–556.

Ramey, Valerie A and Matthew D. Shapiro. 2001. Displaced Capital : A Study of Aerospace Plant Closings. *Journal of Political Economy* 109, no. 5:958–992.

Rappaport, Jordan. 2004. Why are population flows so persistent? *Journal of Urban Economics* 56, no. 3:554–580.

Reardon, Sean F and Kendra Bischoff. 2011. Income inequality and income segregation. *American Journal of Sociology* 116, no. 4:1092–1153.

Roback, Jennifer. 1982. Wages, rents, and the quality of life. *The Journal of Political Economy* 90, no. 6:1257–1278.

Rognlie, Matthew. 2014. A note on Piketty and diminishing returns to capital.

Rosen, Sherwin. 1979. Wage-based indexes of urban quality of life. In *Current issues in urban economics*, eds. Mahlon Straszheim and Peter Mieszkowski. Baltimore, MD: Johns Hopkins University Press.

Rosenbaum, PR and DB Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70, no. 1:41–55.

Rossi-Hansberg, Esteban, PierreDaniel Sarte, and Raymond Owens. 2010. Housing Externalities. *Journal of Political Economy* 118, no. 3:pp. 485–535.

Rubinowitz, Leonard S and James E Rosenbaum. 2000. *Crossing the class and color lines : from public housing to white suburbia*. Chicago: University of Chicago Press.

Ruggles, Steven, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek. 2010. Integrated Public Use Microdata Series: Version 5.0.

Ruhm, Christopher J. 1987. The Economic Consequences of Labor Mobility. *Industrial and Labor Relations Review* 41, no. 1:30–42.

Saez, Emmanuel and Gabriel Zucman. 2014. Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data. *National Bureau of Economic Research Working Paper Series* No. 20625.

Saiz, Albert. 2010. The geographic determinants of housing supply. *The Quarterly Journal of Economics* 125, no. 3:1253–1296.

Saks, RE and Abigail Wozniak. 2011. Labor reallocation over the business cycle: new evidence from internal migration. *Journal of Labor Economics* 29, no. 4:697–739.

Sandell, Steven H. 1977. Women and the Economics of Family Migration. *The Review of Economics and Statistics* 59, no. 4:406–414.

Schmieder, Johannes F. and Till von Wachter. 2016. The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation. *Annual Review of Economics* 8:547–581.

Schmutz, Benoît and Modibo Sidibe. 2016. Frictional Labor Mobility.

Scott, James C. 1998. *Seeing like a state : how certain schemes to improve the human condition have failed*. New Haven: Yale University Press.

Sjaastad, Larry A. 1962. The Costs and Returns of Human Migration.

Small, Kenneth A and Harvey S Rosen. 1981. Applied Welfare Economics with Discrete Choice Models. *Econometrica* 49, no. 1:105–130.

Solon, Gary, Seven J. Haider, and Jeffrey M. Wooldridge. 2015. What are we weighting for? *Journal of Human Resources* 2, no. 50:301–316. URL https://doi.org/10.3368/jhr.50.2.301.

Stevens, Ann. 1997. Persistent Effects of Job Displacement: The Importance of Multiple Job Losses. *Journal of Labor Economics* 15, no. 1:165–188.

Suárez Serrato, Juan Carlos and Owen Zidar. 2014. Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms. Tech. Rep. NBER Working Paper 20289.

Tolbert, Charles M. and Molly Sizer. 1996. U.S. Commuting Zones and Labor Market Areas: A 1990 Update. Tech. Rep. Rural Economy Division, Economic Research service, U.S. Department of Agriculture Staff Paper AGES-9614, Washington, DC.

Topa, Giorgio. 2011. Labor markets and referrals. *Handbook of Social Economics* 1, no. 1 B:1193–1221.

Topel, Robert. 1986. Local labor markets. *The Journal of Political Economy* 94, no. 3:S111–S143.

Train, Kenneth E. 2009. *Discrete choice methods with simulation*. New York: Cambridge University Press, second ed.

United States Census Bureau. 2016. Historical Living Arrangements of Adults.

Valletta, Robert G. 2013. House lock and structural unemployment. *Labour Economics* 25:86–97.

Van Nieuwerburgh, Stijn and Pierre-Olivier Weill. 2010. Why Has House Price Dispersion Gone Up? *Review of Economic Studies* 77, no. 4:1567–1606.

Varian, Hal R. 1984. *Microeconomic analysis*. New York; London: W.W.Norton & Company, 2nd ed.

Verbrugge, Randal. 2008. The puzzling divergence of rents and user costs, 1980-2004. *Review of Income and Wealth* 54, no. 4:671–699.

von Wachter, Till, Jae Song, and Joyce Manchester. 2009. Long-Term Earnings Losses due to Mass Layoffs During the 1982 Recession: An Analysis Using Longitudinal Administrative Data from 1974 to 2004. Tech. rep., Department of Economics, Columbia University.

Watson, Tara. 2009. Inequality and the measurement of residential segregation by income in american neighborhoods. *Review of Income and Wealth* 55, no. 3:820–844.

Willis, Robert. 1973. A new approach to the economic theory of fertility behavior. *Journal of Political Economy* 81, no. 2.

Wilson, William Julius. 1987. *The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy*. Chicago: University of Chicago Press.

Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data*, *MIT Press Books*, vol. 58. Cambridge, MA: MIT Press.

Yagan, Danny. 2013. Moving to opportunity? Migratory insurance over the Great Recession. Tech. Rep. Unpublished manuscript.

———. 2017. Employment Hysteresis from the Great Recession.

Yannay Spitzer. 2015. The Dynamics of Mass Movements.

Zabek, Mike. 2017. Local Ties in Spatial Equilibrium.

Zhang, Fudong. 2016. Inequality and house prices.