ABSTRACT

Title of Dissertation: ECONOMIC MOBILITY AND FRICTIONAL LABOR MARKETS

Matthew Staiger
Doctor of Philosophy, 2021

Dissertation Directed by: Professor John Haltiwanger
Professor Sebastian Galiani
Department of Economics

My dissertation consists of three chapters related to labor economics. The first chapter investigates how the earnings of young workers are affected by the *intergenerational transmission of employers*–which refers to individuals working for the same employer as a parent. My analysis of U.S. linked survey and administrative data indicates that 7 percent of young workers find their first stable job at the same employer as a parent. Using an instrumental variables strategy that exploits exogenous variation in the availability of jobs at the parent's employer, I estimate that working for the same employer as a parent increases initial earnings by 31 percent. The earnings benefits are attributable to parents providing access to higher-paying employers. Individuals with higher-earning parents are more likely to work for the employer of their parent and experience greater earnings benefits when they do. Thus, the intergenerational transmission of employers amplifies the extent to which earnings persist from one generation to the next. The second chapter uses administrative data on earnings and participation in subsidized housing to study how the demolition of 160 public

housing projects–funded by the HOPE VI program–affected the adult labor market outcomes for 18,500 children. Children from HOPE VI projects earn 14 percent more at age 26 relative to children from comparable non-HOPE VI projects. These gains are not driven by improvements in household or neighborhood environments that promote human capital development in children. Rather, subsequent improvements in job accessibility represent a likely pathway. The third chapter uses U.S. linked employer-employee data to examine cyclical worker flows across firms ranked by productivity. In expansions high-productivity firms grow faster by hiring workers from low-productivity firms. The rate at which these job-to-job flows move workers up the productivity ladder is highly procyclical and productivity growth slows during recessions because this job ladder collapses. In contrast, flows into nonemployment from low-productivity firms disproportionately increase in recessions, which leads to an increase in productivity growth. I thus find evidence of both sullying and cleansing effects of recessions. The cleansing effect dominates early in downturns but the sullying effect lingers well into the recovery.

ECONOMIC MOBILITY AND FRICTIONAL LABOR MARKETS

by

Matthew Staiger

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2021

Advisory Committee:
Professor Sebastian Galiani, Co-Chair/Advisor
Professor John Haltiwanger, Co-Chair/Advisor
Professor Judith Hellerstein, Advisor
Dr. Erika McEntarfer, Advisor
Professor Rajshree Agarwal, Dean's Representative

Acknowledgments

I would like to thank my advisers John Haltiwanger, Sebastian Galiani, Judy Hellerstein, and Erika McEntarfer for their valuable feedback and guidance during graduate school. I would also like to thank Rajshree Agarwal for valuable feedback and agreeing to be the Dean's Representative on my committee. In relation to my first chapter, I would like to thank Katharine Abraham, Sandra Black, Sydnee Caldwell, Raj Chetty, John Coglianese, Melissa Kearney, Nuno Limão, Seth Murray, Giordano Palloni, Nolan Pope, John Sabelhaus, Claudia Sahm, John Shea, Doug Staiger ("uncle Doug"), Bob Staiger ("dad"), Cody Tuttle, Mateo Uribe-Castro, Daniel Vincent, John Wallis, Derek Wu, Moises Yi, Sammy Young, Emily Wiemers, and all the participants of the UMD 708 and U.S. Census Bureau seminar series for their feedback and input. In relation to my second chapter, I would like to thank my co-authors John C. Haltiwanger, Mark J. Kutzbach, Giordano Palloni, Henry O. Pollakowski, and Daniel H. Weinberg. In relation to my third chapter, I would like to thank my co-authors John Haltiwanger, Henry Hyatt, and Erika McEntarfer. Lastly, I would also like to thank the Kauffman Foundation, the Washington Center for Equitable Growth, and the Economic Club of Washington, D.C. for generously providing me with funding during graduate school. Thanks to Anna and the rest of my friends and family for all their support during graduate school as well as Ray for always encouraging me to get out for a walk and clear my head.

Table of Contents

Chapter 1: The Intergenerational Transmission of Employers and the Earnings of Young Workers

Disclaimer

Any analysis, opinions, and conclusions expressed herein are those of the author alone and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. See U.S. Census Bureau Disclosure Review Board bypass numbers CBDRB-: FY20-002, FY20-186, FY20-CED006-0020, FY20-CED0006-0025, and FY21-CES011-001.

1.1 Introduction

In the United States, earnings are highly persistent from one generation to the next.[1] The fact that children born into poverty are likely to remain in poverty as adults runs counter to the ideal of equality of opportunity and may be indicative of untapped human potential. But the justification and design of an effective policy response depends on the mechanisms through which parents shape the earnings of their children. Much of the research on intergenerational mobility attributes differences in earnings by family background to differences in human capital (Black and Devereux, 2011). However, family connections in the labor market may also play a role. Indeed, most jobs are found through a social contact (Ioannides and Loury, 2004), with potentially important implications for earnings given the substantial variation in pay policies across firms (Manning, 2011). But despite their potential importance, it is not well understood how family connections shape the intergenerational persistence in earnings.

I investigate how the earnings of young workers are affected by individuals working for the same employer as a parent, which I refer to as the *intergenerational transmission of employers*. My paper therefore sheds light on one particular type of family connection; connections that operate within the current employer of the parent. Previous research suggests that this is an important way in which parents provide access to jobs. For example, in the context of Sweden, Kramarz and Skans (2014) find that 10 percent of individuals find their first job at the same employer as a parent. Previ-

[1] Intergenerational mobility in the United States is low both relative to the past (Chetty et al., 2017) and relative to other developed countries (Solon, 2002).

ous research on the intergenerational transmission of employers is largely descriptive and I advance the literature by investigating implications for the intergenerational persistence in earnings.[2] Specifically, I ask how the intergenerational persistence in earnings would change if no one worked for the same employer as their parent–i.e., if individuals who do work for their parent's employer instead worked at their next best option. I focus on outcomes at the first stable job, which has important and long-lasting effects on an individual's career.[3]

The intergenerational transmission of employers will increase the intergenerational persistence in earnings if individuals with higher-earning parents benefit more. However, the benefits–which depend on the likelihood and earnings consequences of working for a parent's employer–could be increasing or decreasing in parental earnings. On the one hand, higher-earning parents may be better able to provide access to high-paying jobs. On the other hand, individuals from disadvantaged backgrounds may be more reliant on their parents to find a decent-paying job. Which force dominates is an empirical question, which I answer by estimating descriptive statistics of how common it is to work for a parent's employer and the causal earnings consequences of doing so.

I begin by showing that it is not uncommon for an individual to work for their

[2]Other papers that study the intergenerational transmission of employers include Corak and Piraino (2011), Bingley et al. (2011), Stinson and Wignall (2018), Eliason et al. (2019), and San (2020). Of these papers, Eliason et al. (2019) and San (2020) are most closely related and also find that parents affect the earnings of their children by providing access to higher-paying firms. However, neither of these papers study how parental connections affect the intergenerational persistence in earnings. Eliason et al. (2019) and San (2020) focus on understanding on how parental connections affect overall earnings inequality and the earnings gap between ethnic groups, respectively.

[3]Both theoretical (e.g., Jovanovic and Nyarko, 1997; Gibbons and Waldman, 2006) and empirical (e.g., Von Wachter and Bender, 2006; Khan, 2010; Oreopoulos et al., 2012; Altonji et al., 2016; Arellano-Bover, 2020) evidence suggests that early career experiences can have a large and persistent effect on earnings. See Section 1.3.1 for the definition of the first stable job.

parent's employer. I link survey data from the 2000 Decennial Census to administrative data from the Longitudinal Employer-Household Dynamics (LEHD) program and study 10 recent cohorts. I find that 7 percent of individuals work for a parent's employer at their first stable job, and 29 percent do so at some point between the ages of 18 and 30.[4] Individuals with higher-earning parents are more likely to work for a parent's employer. There are several possible explanations for why someone might work for their parent's employer. For example, children may simply tend to work in the same industry and live in the same geographic region. However, individuals are 70 times more likely to work for their parent's employer relative to another employer in the same industry, commuting zone, and size category. Rather, the evidence is more consistent with parents acting as a social contact to help children who otherwise would have struggled to find a decent-paying job. Indeed, individuals with less education who are searching for a job in periods of high unemployment are more likely to work for a parent's employer.

Next, I find large earnings benefits of working for a parent's employer. Estimating causal effects is difficult because individuals who work for a parent's employer likely differ from those who do not. In an ideal experiment, I would prohibit some employers from hiring the children of current employees and use this random assignment as an instrument. My actual empirical strategy mimics this ideal experiment and exploits exogenous variation in the availability of jobs at the parent's employer. Specifically, I instrument for whether an individual works for their parent's employer with the hiring

[4] My estimates of the rate of transmission are consistent with other estimates from the United States (Stinson and Wignall, 2018).

rate at that employer and include fixed effects for the parent's employer and the local labor market. Intuitively, my empirical strategy compares individuals whose parents work for the same employer but who enter the labor market at different times when there are relatively more or less job opportunities at the parent's employer (measured by the hiring rate). The key assumption is that, conditional on the local labor market fixed effects, differences in earnings between the individuals are attributable to differences in the propensity to work for their parent's employer. I find that individuals earn 31 percent more at their first stable job when working for their parent's employer relative to their next best option. Individuals with higher-earning parents experience larger gains.

These earnings gains appear to be explained by parents providing access to higher-paying employers. Following Abowd et al. (1999), I estimate employer-level pay premiums and find that working for a parent's employer leads individuals to work for employers that pay all workers 30 percent more, which is virtually identical to the effect on individual earnings. A wide class of models (e.g., Postel-Vinay and Robin, 2002) illustrate how search frictions lead to job ladders, whereby more productive firms offer higher wages. Consistent with these models, I find that parents provide access to firms on a higher rung of the job ladder as measured by productivity, average wages, and worker flows. A narrative consistent with my results is that there is a group individuals who, without help from their parents, have limited labor market options and would end up at low-paying firms such as a fast food restaurant. However, their parents provide access to jobs at better-paying firms such as a manufacturing plant. Indeed, access to jobs in higher-paying industries explains 75 percent of the effect on individual earnings.

Lastly, I find that the intergenerational transmission of employers leads to a modest increase in the degree to which earnings persist across generations. I develop a methodology that allows me to quantify the difference between observed measures of the intergenerational persistence in earnings and measures that correspond to a counterfactual world in which no one worked for the employer of a parent as a function of the benefits of working for a parent's employer conditional on parental earnings. These benefits depend on the likelihood of working for a parent's employer and the earnings consequences conditional on doing so, two objects that I estimate in my paper. I find that the elasticity of the initial earnings of an individual with respect to the earnings of their parents would be 10 percent lower if no one worked for the employer of a parent.

Non-Black males with high-earning parents are the largest beneficiaries of the intergenerational transmission of employers. Consistent with Chetty et al. (2020), I find that, conditional on parental earnings, Black males have lower expected earnings than White males. On average, the intergenerational transmission of employers explains 10 percent of this conditional Black-White gap in initial earnings. The intergenerational transmission of employers disproportionately benefits sons of high-earning parents but daughters of low-earning parents. On average, daughters benefit more than sons, and the gender pay gap in initial earnings would be 4 percent larger if no one worked for a parent's employer.

My main contribution is to show that the positive association between the earnings of an individual and the earnings of their parents is attributable, in part, to parents using their connections to provide access to higher-paying employers. For some individuals, a job at their parent's employer offers better pay relative to jobs they could find through alternative search methods. Individuals from high-income backgrounds benefit the most from these connections because their parents are more likely to hold positions of authority at high-paying firms. Most explanations of the intergenerational persistence in earnings focus on the development of human capital during childhood. In contrast, I show that parents continue to affect the labor market outcomes of their adult children by using their connections to provide access to jobs. Given that parents could have contacts at other employers, my results likely understate the importance of parental labor market networks more broadly defined.

My conclusions depend on the estimates of the earnings consequences, whose credibility is supported by a number of supplemental analyses. Existing evidence of the earnings consequences of working for a parent's employer–or, more generally, finding a job through a social contact–is mixed, in part, because it is difficult to fully account for factors that affect both earnings and the method of job finding.[5] While my empirical strategy exploits exogenous variation in the propensity to work for a parent's employer, there are several potential issues. First, employers may offer higher wages when hiring more intensively. However, my estimates are robust to controlling

[5]For example, Kramarz and Skans (2014) control for observable differences between children who do and do not work with their parents and find negligible earnings benefits in Sweden. In contrast, Stinson and Wignall (2018) use data from the United States and find large benefits using an individual fixed effects estimator. More generally, estimating the causal effect of finding a job through a social contact has proven difficult (Topa, 2011). Although, a number of recent papers convincingly establish that social contacts can improve labor market outcomes by reducing the duration of unemployment (Beaman, 2012; Cingano and Rosolia, 2012; Glitz, 2017), helping workers find jobs at high-paying firms (Schmutte, 2015), and strengthening workers' bargaining positions (Caldwell and Harmon, 2019).

for the earnings of other new hires, the earnings growth of existing employees, and the employment growth rate at the parent's employer. Second, the hiring rate at the parent's employer may be correlated with local labor market conditions even after conditioning on labor labor market fixed effects. However, using a placebo exercise, I show that the earnings of the child is related only to the hiring conditions at their parent's employer and is unrelated to hiring conditions at other similar employers in the area. These results–and other results discussed below–support the credibility of my estimates.

My results also provide novel evidence that firm-level pay policies are an important determinant of earnings. A substantial portion of earnings inequality is attributable to differences in average pay across firms. But competing explanations emphasize the role of dispersion of firm-level pay policies versus the sorting of workers into firms (Manning, 2003). Prior research finds that moves to higher-paying firms are associated with earnings growth (e.g., Abowd et al., 1999; Haltiwanger et al., 2018). However, the changes in earnings are not necessarily explained by differences in firm-level pay policies since worker mobility is endogenous, and factors that lead workers to change firms could be correlated with factors that have an independent effect on earnings. A number of recent papers (e.g, Schmeider et al., 2020; Lachowska et al., 2020) study workers who leave their employers for exogenous reasons and find that changes in earnings are predicted by changes in firm pay premiums. I provide complementary evidence of the importance of firm-level pay policies in determining individual earnings since my empirical strategy isolates exogenous variation in the employers that individuals end up joining.

The remainder of the paper is structured as follows. Section 2 presents the conceptual framework. Section 3 discusses the data. Section 4 documents descriptive patterns in the intergenerational transmission of employers. Section 5 estimates the earnings consequences of working for the employer of a parent. Section 6 investigates implications for the intergenerational persistence in earnings. Section 7 concludes.

1.2 Conceptual Framework

This section presents a conceptual framework that relates the intergenerational transmission of employers to the intergenerational persistence in earnings. Let y_{ij} denote the log earnings of individual i at their first stable job, which is at employer j. And let y_p denote the log of the life-time earnings of i's parents. A common measure of the intergenerational persistence in earnings is the intergenerational elasticity of earnings (IGE), which is the coefficient obtained from regressing y_{ij} on y_p and is denoted $\rho(y_{ij}, y_p)$. It is important to note that most estimates of the IGE use a measure of life-time earnings for both the parent and their child. In contrast, I focus on initial labor market outcomes of the child. My objective is to understand how the intergenerational transmission of employers affects the intergenerational persistence in earnings, as measured by $\rho(y_{ij}, y_p)$.

I use the potential outcomes framework to characterize the role of the intergenerational transmission of employers. Let $y_{ij(1)}$ denote the individual's earnings if they work for their parent's employer and let $y_{ij(0)}$ denote their earnings if they work for the employer that is their next best option (i.e., where the individual would work if they did not work for their parent's employer). The treatment effect of working for a parent's employer is the difference between potential outcomes and is denoted $\beta_i = y_{ij(1)} - y_{ij(0)}$. Thus,

$$y_{ij} = D_i\beta_i + y_{ij(0)} \tag{1.1}$$

where D_i is an indicator equal to one if the individual works for their parent's employer. It is possible that working for a parent's employer could affect when and even whether an individual finds their first stable job. This poses potential challenges to estimating the earnings benefits. Section 1.5.3 discusses this point in more detail.

I quantify how the intergenerational transmission of employers affects the intergenerational persistence in earnings by comparing the observed IGE, $\rho(y_{ij}, y_p)$, to the IGE that corresponds to the counterfactual in which no one worked for their parent's employer, $\rho(y_{ij(0)}, y_p)$.[6] Combining equation 1.1 with the identity $\rho(y_{ij}, y_p) \equiv \frac{\text{cov}(y_{ij}, y_p)}{\text{var}(y_p)}$, yields,

$$\rho(y_{ij}, y_p) - \rho(y_{ij(0)}, y_p) = \frac{\text{cov}(D_i\beta_i, y_p)}{\text{var}(y_p)} \tag{1.2}$$

To estimates $\text{cov}(D_i\beta_i, y_p)$ I develop and use the following approximation:

$$\text{cov}(D_i\beta_i, y_p) \approx \mathbb{E}\Big[\mathbb{E}\big[D_i|r_p\big]E\big[\beta_i|r_p, D_i = 1\big]\mathbb{E}\big[y_p|r_p\big]\Big] - \mathbb{E}\big[D_i\big]\mathbb{E}\big[\beta_i|D_i = 1\big]\mathbb{E}\big[y_p\big] \tag{1.3}$$

where r_p is the percentile rank of parental earnings.

The approximation relies on two insights. First, the expected value of the product of two random variables is approximately equal to the product of their expected val-

[6] As discussed in more detail in Section 1.6.1, this is a partial equilibrium analysis, which assumes that $y_{ij(0)}$ does not change if individual i does not have the option to work at their parent's employer.

ues if there is little variation in one of the variables: $\mathbb{E}[D_i \beta_i y_p | r_p] \approx \mathbb{E}[D_i \beta_i | r_p] \mathbb{E}[y_p | r_p]$. Second, by iterated expectations, the average benefit of working for a parent's employer is the product of the proportion of individuals who work for their parent's employer and the earnings benefits conditional on doing so: $\mathbb{E}[D_i \beta_i] = \mathbb{E}[D_i] \mathbb{E}[\beta_i \mid D_i = 1]$. I validate the methodology by showing that estimates of the IGE based on the micro data are virtually identical to estimates derived from the approximation. See Appendix A.3 for details.

Equations 1.2 and 1.3 illustrate that the intergenerational transmission of employers will increase the intergenerational persistence in earnings if the average benefits, $\mathbb{E}[D_i \beta_i \mid r_p]$, are increasing in parental earnings. These benefits are a function of the proportion of individuals who work for a parent's employer, $\mathbb{E}[D_i \mid r_p]$, and the earnings benefits of doing so, $\mathbb{E}[\beta_i \mid r_p, D_i = 1]$. Thus, my goal is to understand how these two objects vary with parental earnings. I estimate the former in Section 1.4 and the latter in Section 1.5.

My methodology is a significant improvement over the descriptive analysis in Corak and Piraino (2011) and Stinson and Wignall (2018). These two papers estimate a standard intergenerational earnings regression as well as a modified specification in which they control for whether an individual works for their parent's employer. They then attempt to determine how the transmission of employers shapes intergenerational mobility by comparing the estimated coefficients on parental earnings between the two specifications and by examining sign of the coefficient on the interaction between parental earnings and employer transmission. As both papers acknowledge, the modified intergenerational earnings regression is likely to deliver biased estimates of the earnings benefits of employer transmission, which makes their estimates difficult to interpret. In contrast, I use causal estimates of the earnings benefits of working for a parent's employer in order to quantify how the IGE would change if no one worked for their parent's employer. More generally, many papers estimate the causal relationship between characteristics of parents–such as income (Shea, 2000), education (Black et al., 2005), or labor market networks (Magruder, 2010)–and outcomes of their children.[7] These causal estimates are informative, but they fall short of quantifying the extent to which different channels shape intergenerational associations. My methodology helps to bridge the gap between the focus on causal identification and the broader research agenda that seeks to understand why economic outcomes persist across generations.

1.3 Data

I rely on two main sources of data: (1) the Hundred Percent Census Edited File (HCEF), which measures the relationship between parents and children who are living together in 2000 and (2) data from the LEHD program to measure labor market outcomes of both parents and their children between 2000 and 2016. The HCEF contains all responses from the 2000 Decennial Census Short Form and, in principle, includes all individuals living in the United States in 2000.[8] The LEHD is an employer-employee linked dataset produced by the U.S. Census Bureau and is constructed from two core

[7] Of these papers, Magruder (2010) is most closely related to my paper. Magrduer (2010) finds that parental labor market networks help young unemployed workers find a job in the context of South Africa.

[8] In practice, some individuals are not surveyed in the 2000 Decennial Census and non-respondents are more likely to be minorities or lower-income households. See Appendix A.2.1 for details.

administrative datasets: (1) unemployment insurance (UI) records, which provide job-level earnings records and (2) the Quarterly Census of Employment and Wages, which provides establishment-level characteristics. The earnings records in the LEHD capture roughly 96 percent of private non-farm wage and salary employment in the United States (Abowd et al. 2009). Employers are identified by a state-level employer identification number (SEIN), which typically captures the activity of a firm within a state and industry.[9] The LEHD covers most jobs, but a notable exception is self-employment. While previous work, such as Dunn and Holtz-Eakin (2000), documents strong patterns of intergenerational persistence in self-employment, I focus on more formal employer-employee relationships.

The sample frame is defined based on the HCEF and includes children who are living with their parents in 2000 and who were born after June 30th of 1982 and before July 1st of 1992.[10] The cohorts were chosen so as to focus on a set of individuals who are young enough to likely have lived with their parents in 2000–the oldest individual in the sample was 17 years old when data collection for the 2000 Decennial Census took place–but old enough to have likely entered the labor market by 2016–the youngest individual in the sample was 24 years old by the end of 2016. There are approximately 37 million individuals in the sample frame. See Appendix A.2.1 for details.

I implement two sets of sample restrictions. First, I require that the individuals

[9] A worker could have positive earnings at multiple employers in a given quarter. In such cases, I measure the characteristics of the employer providing the majority of earnings in that quarter.

[10] Over 90 percent of individuals within this age range live with a parent in 2000. Children are individuals whose relationship to the household head is: son/daughter, adopted son/daughter or step son/daughter. I exclude individuals living in U.S. territories in 2000.

and their parents found in the HCEF can be linked to the LEHD. In order to account for non-random attrition from the sample due to issues associated with linking records across the two data sources, I construct sample weights and use them to produce all descriptive statistics. Second, I drop cases in which the earnings of the children or parents are likely to be affected by coverage issues in the LEHD. Of the 37 million children in the sample frame, approximately 21 million (57 percent) meet the two sets of restrictions. Based on these sample restrictions and the source of earnings data, my analysis should be viewed as representative of working families, a category which excludes very low income households (approximately the bottom 10 percent of households) and very high income households (approximately the top 1 percent of households). See Appendix A.2.2 for details.

1.3.1 Measurement of Key Variables

My paper focuses on initial labor market outcomes and thus I need to define when individuals enter the labor market. Conceptually, I define entry as the first period in which work becomes the primary activity. My empirical definition of entry is the first quarter in which the individual earns at least $3,300 per quarter–which approximately corresponds to working 35 hours per week at the federal minimum wage–in the current and two consecutive quarters, and receives positive earnings from the same employer for those three quarters.[11] I refer to the employment spell at this employer as the first stable job. Approximately 80 percent of individuals (17 million individuals) that

[11] Dollar values are converted to 2016 dollars using the Consumer Price Index for All Urban Consumers.

meet the sample restrictions have entered the labor market by the end of 2016.

There are many possible ways to define entry, but three pieces of evidence suggest that my approach is reasonable.[12] First, individuals experience a dramatic and persistent increase in earnings upon entry. Average quarterly earnings in the three years prior to entry is $1,258 compared to $6,597 in the three years after entry. Figure A.1 provides more detailed evidence by plotting the average quarterly earnings in the three years before and after entry. Second, the age of entry generally lines up with common perceptions of when individuals start their careers. For example, 89 percent of children enter the labor market between ages 18 and 26. Figure A.3 depicts the distribution of the age at which the children enter the labor market and compares this distribution to results based on an analogous measure constructed from the National Longitudinal Survey of Youth 1997 cohort (NLSY97).[13] The timing of entry is quite similar in the two data sources. Furthermore, 83 percent of workers in the NLSY97 data are not enrolled in school at the time of labor market entry, which suggests that my measure is not primarily picking up jobs held by students. Third, the first stable job is indeed stable as the average duration of employment at the first stable job exceeds two years.

I construct a measure of the lifetime earnings of the parents. Without data on the full labor market history, a common approach is to calculate parental earnings as the average earnings over a limited number of years. In addition to the measurement

[12] Kramarz and Skans (2014) use a similar set of criteria to identify the first stable job.

[13] The analogous measure constructed from the NLSY97 is the first time an individual works at least 35 hours for 36 consecutive weeks (or three quarters). An alternative approach is to focus on labor market outcomes after all schooling is completed and I also present results for this definition of entry.

issues raised by Solon (1989) and Zimmerman (1992), this approach problematic when using the LEHD since there is no way to distinguish between zero earnings and missing data.[14] Instead, I construct a measure of lifetime parental earnings by estimating a regression of quarterly earnings on an individual fixed effect and a third degree polynomial in age within cells defined by the interaction between state of residence in 2000, sex, and race.[15] The measure of the lifetime earnings of each parent is the imputed value of earnings between ages 35 and 55. For one-parent households, parental earnings is the lifetime earnings of the parent. For two-parent households, parental earnings is the average of the lifetime earnings of both parents. The parental earnings percentile ranks are calculated within each cohort of children using sample weights.[16] See Appendix A.2.4 for details.

1.4 Intergenerational Transmission of Employers

I begin the empirical analysis by documenting descriptive patterns related to the intergenerational transmission of employers. Table 1.1 presents summary statistics. The first column presents results for the entire sample. The second through fifth columns present results for subsamples defined by whether the first stable job is at the employer of neither parent, the secondary earner, both parents, or the primary

[14] Earnings data could be missing either because a state may not report to the LEHD in a given time period or because the job may not be covered in the LEHD frame.

[15] The data are a panel measured at a quarterly frequency that include all strictly positive earnings records between 2000 and 2016 for the parents in the sample. Quarters with zero earnings are not included in the sample. I further restrict the panel to observations when the individuals are between the ages of 30 and 60 and drop individuals that have fewer than 4 quarters of strictly positive earnings over the entire time period. Parents not included in this sample are assumed to have zero lifetime earnings.

[16] Cohorts consist of individuals born between July 1st of year t and June 30th of year t+1.

earner, respectively.[17] The bottom row indicates that 7 percent of individuals work for the employer of either parent at their first stable job. A comparison across columns indicates that individuals who work for a parent's employer tend to stay at their first stable job longer, are less likely to be employed in the unskilled service sector, are more likely to work in the manufacturing/production sector, and earn slightly less.[18]

One interpretation is that parents are a social contact and influence the hiring or job search process. This would be consistent with Loury (2006), who finds that 10 percent of males found their current job through a parent, as well as with the literature that finds ubiquitous use of informal search methods (Ioannides and Loury 2004; Topa 2011) and that labor market networks influence where individuals work (Bayer et al. 2008; Hellerstein et al. 2011). However, there are other possible interpretations.

Individuals are much more likely to work for their parent's employer relative to other similar employers in the same local labor market, which suggests that the intergenerational transmission of employers is not explained by the tendency for parents and children to work in the same sector and location. Table 1.1 indicates that individuals who work for a parent's employer are no more likely to work for large employers and over 70 percent of these individuals are located in urban areas. This suggests that the tendency to work for a parent's employer is not driven by cases in which a single employer dominates a local labor market. To investigate the issue more rigorously, I calculate the proportion of individuals who work for an employer of the same

[17] The primary earner is defined as the parent with the greatest earnings in the year prior to the quarter in which the child entered the labor market.

[18] I group two-digit North American Industry Classification System (NAICS) industry codes into three sectors: unskilled services, skilled services, and manufacturing/production. See Appendix A.2.5 for details.

Table 1.1: Summary Statistics

	Full Sample	First Job at the Employer of: Neither Parent	Secondary Earner	Both Parents	Primary Earner
A. Individual Characteristics					
male	0.50	0.49	0.46	0.61	0.60
White non-Hispanic	0.79	0.78	0.83	0.84	0.79
Black non-Hispanic	0.08	0.08	0.05	0.02	0.07
Asian non-Hispanic	0.02	0.02	0.02	0.04	0.02
Hispanic	0.09	0.09	0.09	0.08	0.09
born in United States	0.97	0.97	0.97	0.96	0.97
B. Household Characteristics					
parents are married	0.78	0.78	0.96	0.98	0.78
parent has unmarried partner	0.03	0.03	0.04	0.02	0.03
primary earner is male	0.57	0.56	0.77	0.80	0.54
parental earnings / 1,000	51.42	51.26	53.15	67.28	51.28
C. First Stable Job					
age at first job	20.94	21.00	20.10	19.81	20.08
tenure at first job (quarters)	10.07	9.77	13.40	18.03	13.67
log of quarterly earnings	8.74	8.74	8.62	8.70	8.72
skilled services	0.37	0.37	0.45	0.31	0.37
unskilled services	0.46	0.47	0.36	0.31	0.28
manufacturing/production	0.18	0.16	0.19	0.39	0.36
employer size $<$ 50	0.28	0.28	0.27	0.62	0.30
50$\leq$employer size$<$500	0.31	0.32	0.26	0.15	0.29
500$\leq$employer size	0.40	0.40	0.47	0.23	0.40
located in urban area	0.77	0.78	0.71	0.71	0.72
Sample Size					
proportion of full sample		0.93	0.02	0.01	0.04
observations	17,010,000	15,830,000	298,000	137,000	746,000

Notes: The table presents the average value of the variable defined in the row. Column 1 presents results for the full sample and columns 2-5 present results for the sample of children who, at their first stable job, worked for the employer of neither parent, the secondary earner, both parents, or the primary earner, respectively.

Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

size category and located in the same census tract as the employer of the primary earner.[19] I find that individuals are 43 times more likely to work for the employer of their parent compared to another employer in the same census tract. I calculate a similar statistic for employers that are in the same commuting zone, size category, and industry (defined at the three-digit NAICS industry code) and find that individuals are 70 times more likely to work for the employer of the primary earner.[20] These results suggest that geography, industry, and employer size are poor explanations for the intergenerational transmission of employers.

Individuals are also more likely to work for the current employer of their parent relative to past or future employers, which casts doubt on the possibility that the intergenerational transmission of employers is explained by intergenerational transmission of human capital or preferences. I identify past employers (the employer of the primary earner when the child was 10 years old) and future employers (the employer of the primary earner in 2016). Separately for past and future employers, I limit the sample to cases where the past or future employer existed in the quarter in which the child entered the labor market but the current employer of the primary earner differed. Within these two samples, I find individuals are 6 and 4 times more likely to work for their parent's current employer relative to the past and future employers, respectively. The fact that the child is more likely to work for a past or future employer of the parent relative to other employers in the same local labor market could be explained by the presence of other social contacts.

[19] Employer size categories are: small (employees< 50), medium ($50 \leq$employees< 500), and large ($500 \leq$employees).

[20] These statistics are derived from the estimates in Panel A of Table 1.3.

Thus, it is the presence of a parent at an employer, not the characteristics of the employer, that leads individuals to work for their parent's employer. This suggests that the intergenerational transmission of employers occurs primarily because parents influence the hiring or job search process.[21] The extent to which the employer benefits depends largely on how parents provide access to jobs–be it through nepotism or reducing information asymmetries between the child and employer. While distinguishing between the alternative explanations is difficult, descriptive evidence suggests that employer transmission tends to benefit children with more limited labor market opportunities. Table A.1 links responses to the American Community Survey to a subset of records and shows that, conditional on parental earnings, individuals with lower levels of educational attainment are more likely to work for a parent's employer. Table A.2 shows that, conditional on the age of entry, the transmission of employers is more likely to occur when unemployment is high.[22] Figures A.4 and A.5 illustrate that the industries in which employer transmission is more common tend to offer higher wages (conditional on observable worker characteristics) and exhibit higher rates of unionization. These results provide suggestive evidence that the intergenerational transmission of employers benefits the child in the form of higher wages but does not necessarily benefit the employer by providing access to more productive workers.

Individuals with higher-earning parents are more likely to work for their parent's

[21] It is also possible that non-monetary benefits could make it more likely for individuals to want to work for their parent's employer. However, this explanation seems less likely, in light of the large earnings benefits found in the next section.

[22] I condition on the age of entry because older individuals are less likely to work for the employer of a parent and average age of entry is older later in the sample period (when unemployment is higher).

employer, although the relationship is nonlinear.[23] Figure 1.1 presents the proportion of individuals who work for the employer of either parent at their first stable job by parental earnings, sex, and race/ethnicity. There is a strong positive association between the likelihood of working for a parent's employer and parental earnings in the bottom quintile and top decile of the parental earnings distribution and a weak (slightly negative for sons) association elsewhere. For daughters, the patterns are similar across the race/ethnicity categories. In contrast, Black sons are substantially less likely to work for the employer of a parent relative to other groups throughout the parental earnings distribution.

A plausible explanation for the relationship between parental earnings and the intergenerational transmission of employers is that higher-earning parents are more likely to be employed and hold a position of authority within their employer. The percent of primary earners that are employed when their child enters the labor market rises steeply from 55 percent to 84 percent between the 1st and 20th percentiles of the parental earnings distribution and eventually plateaus at 94 percent. The percent of primary earners whose earnings are in the top percentile within their employer when the child enters the labor market rises gradually from 4 percent to 14 percent between the 1st and 90th percentiles of the parental earnings distribution and then rises steeply to 41 percent in the top percentile. Thus, the nonlinear relationship between the probability of working for a parent's employer and parental earnings

[23]Sons are more likely to work for the employer of a parent at their first stable job relative to daughters, with 7.8 percent of sons doing so compared to 6.0 percent of daughters. Sons and daughters are more likely to work for the employer of the primary earner relative to the secondary earner, but the difference is larger for sons. Individuals are at least twice as likely to work with the parent of the same sex (See Table A.3).

Figure 1.1: Intergenerational Transmission of Employers

(A) Daughters

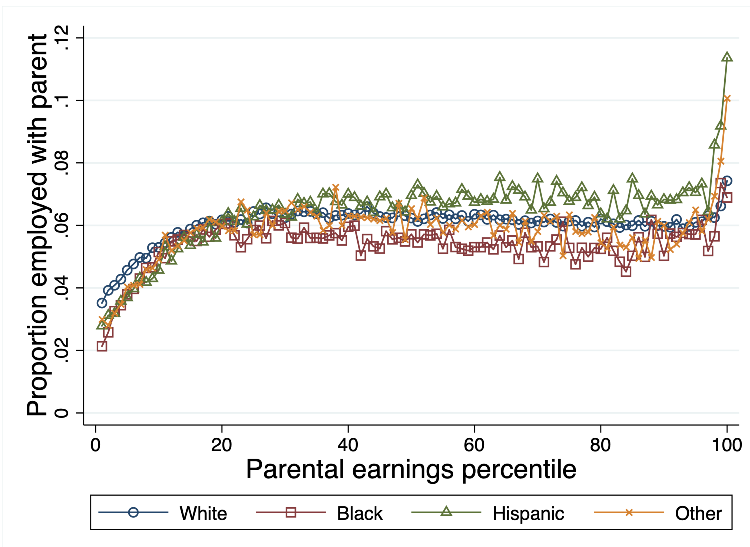

(B) Sons

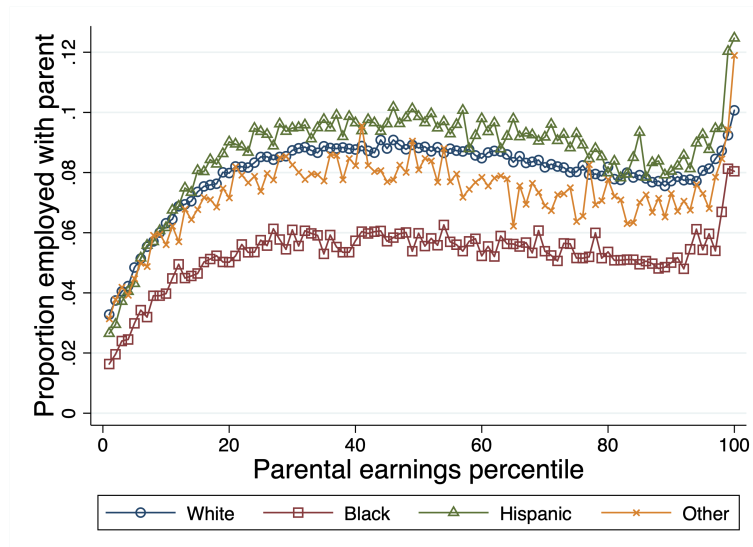

Notes: The figures plot the proportion of individuals whose first stable job is at the employer of either parent. Each statistic is reported separately by sex (daughters in Panel A and sons in Panel B), the percentile of parental earnings (defined by the x-axis), and race/ethnicity. All statistics are calculated using sample weights.

Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

closely tracks the probability that the parent is employed or is a top earner within their employer.[24]

The nonlinear relationship between the intergenerational transmission employers and parental earnings is also present in longer-run measures. Within the sample of individuals who turn 30 by the end of 2016, 28 percent of daughters and 29 percent of sons work for the employer of a parent between the ages of 18 and 30. These estimates are consistent with Stinson and Wignall (2018), who find that 22 percent of sons have shared an employer with their father by the time they are 30 years old.[25] Figure A.7 presents how these estimates vary across the parental earnings distribution and illustrates that the nonlinear patterns observed at the first stable job are replicated in these longer-run measures.

1.5 Earnings Consequences

This section estimates the earnings consequences of working for a parent's employer. I begin by considering a structural earnings equation in order to highlight potential mechanisms and illustrate the challenges associated with estimating causal parameters. Building on the notation from Section 1.2, let the log earnings at the first stable job (y_{ij}) be additive in an individual component (α_i), an employer component (ψ_j), and an individual-employer component (ζ_{ij}), where i denotes the individual and j denotes the employer. Working at a parent's employer affects where an individual

[24]Figure A.6 presents these results in detail by plotting the proportion of parents that are employed and that are top earners within their employer against the percentile of parental earnings.

[25]Similar estimates for other countries include 40 percent in Canada (Corak and Piraino 2011) and 28 percent in Denmark (Bingley et al. 2011).

works and thus may affect earnings through ψ_j or ζ_{ij}. Equation 1.1 can be rewritten as,

$$y_{ij} = D_i\big(\beta_i^{\psi} + \beta_i^{\zeta}\big) + \big(\alpha_i + \psi_{j(0)} + \zeta_{ij(0)}\big) \tag{1.4}$$

where D_i is an indicator equal to one if the first stable job is at the parent's employer. The treatment effect of working for a parent's employer consists of an employer component ($\beta_i^{\psi} = \psi_{j(1)} - \psi_{j(0)}$) and an individual-employer component ($\beta_i^{\zeta} = \zeta_{ij(1)} - \zeta_{ij(0)}$), where $j(1)$ denotes the parent's employer and $j(0)$ denotes the next best option.

Equation 1.4 highlights two mechanisms through which working for the parent's employer could affect earnings. First, the term β_i^{ψ} indicates that parents may provide access to employers that pay all workers higher (or lower) wages. This mechanism is consistent with the model of labor market networks developed in Mortensen and Vishwanath (1995) as well as models which show how imperfect competition in the labor market leads to dispersion in employer-level pay policies. Second, the term β_i^{ζ} illustrates that employers might offer different wages to children of current employees relative to otherwise similar workers. This could happen if parents reduce information asymmetries between workers and employers (e.g., Montgomery, 1991; Dustmann et al., 2016) or if working with a parent affects worker productivity (e.g., Heath, 2018).

Equation 1.4 also highlights the empirical challenges associated with estimating causal parameters. In the previous section I found that individuals were more likely to work for a parent's employer if they were less educated–this could be modeled as a negative correlation between α_i and D_i–and if they were searching for a job in labor markets with higher levels of unemployment–this could be modeled as a negative correlation between $\psi_{j(0)}$ or $\zeta_{ij(0)}$ and D_i. These patterns suggest that a naive comparison between individuals who do and do not work for their parent's employer would understate the earnings benefits. More generally, an empirical strategy that identifies causal parameters must account for the possibility that the characteristics and outside options of individuals are related to the probability that they take a job at their parent's employer.

1.5.1 Instrumental Variables Strategy

I use an instrumental variables strategy that exploits exogenous variation in the availability of jobs at the parent's employer. In order to explain the empirical strategy, consider estimating the following equation via two-stage least squares,

$$\begin{aligned} D_i &= \tilde{\pi}^1 + \gamma Z_{j(1)t-1} + \tilde{u}_i \\ y_{ij} &= \tilde{\pi}^2 + \beta_i D_i + \tilde{v}_i \end{aligned} \tag{1.5}$$

where t is the quarter in which the individual starts their first stable job and $Z_{j(1)t-1}$ is the average quarterly hiring rate at the parent's employer in the four quarters prior to the quarter in which the child begins their fist stable job.[26] Intuitively, the parent's employer will be more likely to make a job offer to the child of a current employee when they are hiring more intensively. By taking the average hiring rate over the preceding four quarters, I avoid measuring the hiring rate in the quarter in which the

[26] I follow the methodology used to produce the Quarterly Workforce Indicators and calculate the End-of-Quarter Hiring Rate, which is the number of new hires that remain with the employer for at least one additional quarter divided by the average of the total employment at the employer at the beginning and end of the quarter.

child starts their first job and ensure that the hiring rate is not affected by seasonal variation.

The stylized model highlights two main reasons why the independence assumption–which is a key assumption needed to interpret the estimates as causal–is unlikely to hold.[27] First, the hiring rate at the parent's employer could be correlated with local labor market conditions that directly affect the earnings of the child–this could be modeled as a positive correlation between $\psi_{j(0)}$ or $\zeta_{ij(0)}$ and $Z_{j(1)t-1}$. Second, employers that hire more intensively may tend to employ more highly educated workers who have more highly educated children–this could be modeled as a positive correlation between α_i and $Z_{j(1)t-1}$.

I include two-way fixed effects in the empirical model to address the concern that the hiring rate at the parent's employer could be related to time-varying local labor market conditions or time-invariant characteristics of the parent's employer. Specifically, I estimate the following equation via two-stage least squares,[28]

$$\begin{aligned} D_i &= \pi^1 + \gamma Z_{j(1)t-1} + X_i\Gamma^1 + \delta^1_{j(1)} + \lambda^1_{l(j(1),t)} + u_i \\ y_{ij} &= \pi^2 + \beta_i D_i + X_i\Gamma^2 + \delta^2_{j(1)} + \lambda^2_{l(j(1),t)} + v_i \end{aligned} \tag{1.6}$$

where $\delta_{j(1)}$ is a fixed effect for the parent's employer; $\lambda_{l(j(1),t)}$ is a fixed effect for the local labor market in which the parent's employer is located, which is defined by the interaction between the state, industry (two-digit NAICS code), and calendar year;

[27]Note that $\beta_i = \beta_i^{\psi} + \beta_i^{\zeta}$, $\tilde{\pi}^2 = \alpha_i + \psi_{j(0)} + \zeta_{ij(0)}$, and $\tilde{v}_i = \mathbb{E}[\alpha_i + \psi_{j(0)} + \zeta_{ij(0)}] - \tilde{\pi}^2$. The independence assumption is $\{\alpha_i, \psi_{j(1)}, \psi_{j(0)}, \zeta_{ij(1)}, \zeta_{ij(0)}\} \perp\!\!\!\perp Z_{j(1)t-1}$.

[28]I estimate all regressions without sample weights since the empirical strategy explicitly accounts for the reasons weights should be used when estimating causal effects (Solon et al. 2015). In practice, I find that the using sample weights makes little difference for the results.

X_i is a vector of demographic characteristics; and u_i and v_i are regression residuals, which are clustered at the level of the parent's employer.[29]

I implement two sample selection criteria when estimating the specification. First, since I exploit variation in the hiring rate at the parents' employer, I require that the parent is employed at the time the child enters the labor market. For much of the analysis I focus on estimating the effect of working for the employer of the parent who is the primary earner and require that the primary earner has at least one year of tenure in the quarter in which the child enters the labor market. The tenure restriction helps address concerns that children and parents might be responding to common economic shocks affecting firms in the local labor market. Second, I drop all singleton observations because these observations do not contribute to the identification of any parameters in the model and retaining them would bias estimates of the standard errors.[30]

The estimates from equation 1.6 have a causal interpretation under three assumptions. First, the hiring rate must affect the probability of working for a parent's employer. This assumption is testable and I present the relevant empirical evidence in Section 1.5.2. Second, the hiring rate must have a monotonic affect on the probability of working for a parent's employer. With the two sets of fixed effects in the

[29]The vector of demographic characteristics includes: the log of the annual earnings of the parent in the year prior to entry; a fixed effect for the cohort of the child; and an interaction between the sex of the child and their race, ethnicity, and an indicator equal to one if born in the United States. The race categories include White, Black, Native American, Asian, Pacific Islander, and other. Ethnicity is defined as Hispanic and non-Hispanic.

[30]A singleton refers to an observation which has a unique value of a fixed effect. For example, if there only existed one observation for a given parent's employer, then the outcome would be perfectly predicted by the employer fixed effect and this observation would not contribute to the identification of any other coefficients.

model, this assumption implies that for any two employers and any two periods, the employer that experiences a larger increase in the hiring rate also experiences a larger increase in the propensity to hire a child of a current employee.[31] While not directly testable, Section 1.5.3 presents some empirical evidence to support the plausibility of this assumption.

Third, the independence assumption requires that the hiring rate is only related to the earnings of the individual through the effect on working for the parent's employer.[32] The covariates directly address two main concerns. First, the state-by-industry-by-year fixed effects address the possibility that the hiring rate at the parent's employer might be correlated with local labor market conditions. Second, the fixed effects for the parent's employer address the concern that the hiring rate may be correlated with time-invariant characteristics of the employer that are correlated with the characteristics of the parents and their children. The vector of demographic variables accounts for additional individual-level heterogeneity not captured by the employer fixed effect; although, the demographic controls do not play a major role in identification.[33] In Section 1.5.3 I present evidence to suggest that the covariates achieve their stated objective and I also explore other possible violations of the

[31] The hiring rate may be correlated with the composition of new hires if some types of workers are relatively more likely to be hired than others when the employer is hiring more intensively. However, this is not a violation of the monotonicity assumption as long as the absolute probability–as opposed to the probability relative to other workers–of a given worker being hired is weakly increasing in the hiring rate. To see why, consider the following example. The parent's employer only makes job offers to the high ability individuals when hiring is relatively low and makes job offers to both high and low ability individuals when hiring is relatively high. While this affects the interpretation of the estimates (the estimates identify the average effect for low ability individuals in this case), it does not necessarily affect the validity of the instrument. I make this point formally in the context of the stylized model presented in Appendix A.4.

[32] Independence requires that $\{\alpha_i, \psi_{j(1)}, \psi_{j(0)}, \zeta_{ij(1)}, \zeta_{ij(0)}\} \perp\!\!\!\perp Z_{j(1)t-1} \mid \{X_i, \delta_{j(1)}, \lambda_{l(j(1),t)}\}$.

[33] The main estimates are qualitatively similar when including no demographic controls.

independence assumption.

With two-way fixed effects, the identifying variation comes from the difference across employers in the differences in the hiring rate over time. Intuitively, the first-stage compares individuals whose parents work for the same employer but who enter the labor market at different times. I ask if the individual is more likely to work with their parent if they enter the labor market when their parent's employer is hiring more intensively, and whether this difference is larger relative to individuals who enter the same local labor market in the same periods but whose parent's employer experiences a relatively smaller growth in the hiring rate. In this way, the empirical strategy exploits variation in the hiring rate that is orthogonal to both time-invariant characteristics of the parent's employers and time-varying conditions of the local labor market.

If the three identifying assumptions are met, the two-stage least squares estimator identifies a *local average treatment effect* (LATE), which is the average effect for the *compliers*–the population whose treatment status depends on the value of the instrument (Imbens and Angrist 1994). I first focus on understanding the consequences of working for a parent's employer for this population. After presenting the main results, Section 1.5.5 explores the relationship between the LATE and other causal parameters of interest.

1.5.2 Effect on Initial Earnings

Table 1.2 presents estimates from equation 1.6 of the earnings consequences of working for the employer of a parent (the primary earner) at the first stable job. Column 1 presents the estimates from the first-stage and demonstrates that the hiring rate at the parent's employer is highly predictive of whether or not the child works there, with an associated F-statistic of 1,434.[34] Column 2 presents the reduced form estimates, illustrating that there is a positive and statistically significant relationship between the hiring rate and initial earnings, which are measured during the first full-quarter of employment at the first stable job.[35] Column 4 presents the second stage estimates, which indicate that working for a parent's employer leads to a 31 percent increase in initial earnings. Column 3 presents Ordinary Least Squares (OLS) estimates for comparison, which are positive but significantly smaller than the two-stage least squares estimates. The OLS estimates could be negatively biased if, for example, low-ability children with limited labor market opportunities are most likely to accept job offers from their parents' employers. It is plausible that the OLS estimates would suffer severely from bias since the data lack meaningful measures of human capital.

The estimated earnings benefits of working for the employer of a parent are large

[34] To assess magnitude of the first-stage, note that a one standard deviation increase in the residualized hiring rate leads to an 8% increase in the probability of working for the parent's employer. I also estimate placebo regressions in which I replace all variables related to the employer of the parent with variables that correspond to the placebo employers considered in Section 1.4, including employers in the same census tract or local labor market and past or future employers. Both the point estimates and F-statistics associated with the true employers are an order of magnitude larger (see Panel B of Table 1.3).

[35] A full-quarter employment spell occurs when a worker receives strictly positive earnings from the same employer in the current, previous and subsequent quarter and variation in earnings is less likely to be driven by differences in the duration of an employment spell within a quarter. The definition of the first stable job implies that every worker experiences a full-quarter employment spell in the second quarter at their first stable job.

Table 1.2: Effect on Initial Earnings

	works for parent's employer	log of quarterly earnings		
	(1)	(2)	(3)	(4)
hiring rate	0.119***	0.036***		
	(0.003)	(0.003)		
works for parent's employer			0.032***	0.307***
			(0.002)	(0.029)
estimator	OLS	OLS	OLS	2SLS
F-statistic	1,434			
mean	0.056			
control mean		8.737	8.737	8.737
control s.d.		0.427	0.427	0.427
observations	11,460,000	11,460,000	11,460,000	11,460,000

Notes: Each column presents results from a separate regression. The outcome variable in column 1 is an indicator equal to one if the individual works for their parent's employer (primary earner) at the first stable job. The outcome variable in columns 2-4 is the log of the first full-quarter earnings at the first stable job. The main independent variable in column 1 is the average quarterly hiring rate at the parent's employer and the main independent variable in columns 2-4 is an indicator equal to one if the individual works for their parent's employer. The results in columns 1-3 are estimated by Ordinary Least Squares (OLS) and the results in column 4 are estimated by two-stage least squares (2SLS), where the instrument is the average quarterly hiring rate at the parent's employer in the four quarters prior to entry. All specifications include a fixed effect for the parent's employer; a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and the standard vector of demographic characteristics. Standard errors are clustered at the level of parent's employer and are presented in parentheses.

Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

but not inconsistent with other evidence of the importance of place of work in determining earnings. For example, the estimated effect is about twice as large as the union wage premium (Farber et al., 2018) and about two standard deviations of the inter-industry wage premium (Katz and Summers, 1989). Another way to assess the magnitude of my estimates is to compare them to the college premium–the relative wage of college versus high school educated workers–which is about 68 log points (Acemoglu and Autor, 2011). In the context of the United States, Stinson and Wignall (2018) estimate specifications with individual fixed effects and find that sons and daughters who work for the employer of their father experience an increase in earnings by 22% and 8%, respectively. My results differ more dramatically relative to Kramarz and Skans (2014), who study the school-to-work transition in Sweden and find small wage losses in the short run, which appear to be offset by stronger wage growth in the medium run; this finding is supported by Eliason et al. (2019), who use more recent data from Sweden.

1.5.3 Validity of the Empirical Strategy

One potential issue is that employers might offer higher wages when hiring more intensively. I assess this concern by controlling for the log of average earnings of all new hires at the parent's employer in the preceding year. This only reduces the main estimates from 0.307 to 0.299 (see column 2 of Table A.4). However, changes in the earnings of new hires might partially reflect a change in the composition of workers being hired. Columns 3 and 4 of Table A.4 take an alternative approach and control for the earnings growth of the parents and all workers at the employer, respectively, in the year prior to entry. The idea is that changes in offer wages are likely to be correlated with earnings growth for current workers. Again, the estimated earnings benefits are largely unaffected. Lastly, column 5 of Table A.4 shows that the results are also robust to controlling for the growth in employment in the year prior to entry (point estimate is 0.307). In general hiring and employment growth are positively correlated, but, conditional on the covariates in the model, the hiring rate captures variation in job opportunities that is orthogonal to more general measures of firm health. Thus, the results do not appear to be affected by a correlation between the hiring rate and time-varying wage setting policies.

The empirical specification might not adequately control for local labor market conditions. I investigate this by estimating placebo specifications in which I replace all variables related to the employer of the parent with variables related to the placebo employer. Panel C of Table 1.3 presents the estimates from the reduced form. In column 3 the placebo employer is an employer in the same commuting zone, three-digit industry, and size class as the parent's employers. Note that the local labor market fixed effects in the empirical model are defined at a higher level of geography (state versus commuting zone) and higher level of industry (two-digit versus three-digit) relative to the placebo employers. The estimates illustrate that the hiring rate at the placebo employers is unrelated to the earnings of the child. Thus, the positive relationship between initial earnings and the instrument is unlikely to be driven by local labor market conditions. The estimates in column 2 indicate earnings are also unrelated to the hiring rate at placebo employers in the same census tract and size

Table 1.3: Placebo Analysis

	Same Labor Market			Past Employer		Future Employer	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Employer Transmission							
proportion at same employer	0.0561	0.0013	0.0008	0.0444	0.0078	0.0472	0.0107
B. First Stage							
hiring rate	0.1187***			0.0781***		0.0893***	
	(0.0031)			(0.0073)		(0.0058)	
placebo hiring rate		0.0003**	0.0002*		0.0030**		0.0027***
		(0.0001)	(0.0001)		(0.0010)		(0.0006)
F-statistic	1,434	10	6	113	9	235	20
C. Reduced Form							
hiring rate	0.0364***			0.0204		0.0183*	
	(0.0033)			(0.0106)		(0.0086)	
placebo hiring rate		-0.0016	0.0018		0.0092		0.0057*
		(0.0015)	(0.0015)		(0.0060)		(0.0025)
employer	parent's employer	census tract	CZ and industry	parent's employer	past employer	parent's employer	future employer
observations	11,460,000	10,470,000	10,530,000	1,114,000	1,209,000	1,434,000	1,554,000

Notes: The estimates in: columns 1-3 are based on the full sample, columns 4-5 are based on the sample of primary earners who have a different employer in the quarter in which their child finds their first stable job relative to their employer when the child was ten years old, and columns 6-7 are based on the sample of primary earners who have a different employer in the quater in which their child finds their first stable job relative to their employer in 2016. The samples used in columns 4-5 and 6-7 also require that the past and future employers, respectively have strictly positive employment in the quarter in which the child finds their first stable job. The results in columns 1, 4, and 6 correspond to the current employer of the parent who is the primary earner. The results in columns 2, 3, 5, and 7 correspond to the placebo employers. The plabebo employers are defined as: (column 2) an employer in the same census tract and employer size category as the employer of the primary earner; (column 3) an employer in the same employer size cateory, commuting zone, and three-digit industry code as the employer of the primary earner; (column 5) the past employer; and (column 7) the future employer. Panel A presents the proportion of children who worked for the employer defined by the column. Panel B presents estimates from the first-stage specification. Panel C presents estimates from the reduce form specifcation. Standard errors are clustered at the level of the employer defined by the column and are presented in parentheses. All statistics are calculated on a sample without singleton observations.

Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

category as the parent's employer. Columns 5 and 7 indicate that the reduced form is relatively more positive for past and future employers. However, the magnitudes are substantially smaller compared to the first stage when using the parent's employer (see columns 4 and 6). I do not view these results as problematic since it is possible that young workers have access to these employers through other connections. This interpretation is consistent with the results in Panels A and B.

Local labor market conditions could also lead to a violation of the monotonicity assumption. If there are more job opportunities at all firms when the parent's employer is hiring, an increase in the hiring rate at the parent's employer could actually accompany a reduction in the probability that the individual works there. I measure the aggregate hiring rate in the three sectors–unskilled services, skilled services, and manufacturing/production–in the commuting zone in which the parent's employer is located and include a vector of controls that interacts these aggregate hiring rates with the sector of the parent's employer. This modified specification directly controls for hiring conditions at all employers in the local labor market. The point estimate (standard error) from the first and second stage are 0.118 (0.003) and 0.297 (0.029), respectively. These controls have little impact on the results, which provides additional evidence that local labor market conditions are not biasing the estimates.

Where the parent works is not randomly assigned, which raises two concerns. First, the sample excludes parents that are not employed and some parents may lose their jobs when the hiring rate is lower and the employer is not doing well. It is likely that this would produce negative bias, since lower-earning parents would be more likely to appear at employers with higher hiring rates. Second, parents may anticipate that their child will struggle to find a job and move to employers that have more job opportunities when their child is starting their career. If parents are more likely to do this for children with lower earnings potential, then this would also lead to negative bias. The sample selection criteria requiring parents to have at least one year of tenure likely helps to address these concerns, as the concerns are more applicable to parents that are less attached to their employer. As an initial check, I estimate the main specification on a sample of parents with at least five years of tenure and continue to find large positive earnings benefits for this sample with a point estimate (standard error) of 0.23 (0.048).

I use comparisons between siblings to further investigate potential issues that could arise from parents sorting into employers. Specifically, I estimate one specification that includes a fixed effect for the parent's employer and another that includes a fixed effect for the parent's employer by household, which limits the identifying variation to comparisons between siblings. Both regressions are estimated on the same subsample, which retains cases for which at least two siblings entered the labor market when the primary earner was at the same employer. The estimates (standard errors) from the specification with the employer fixed effect and the household by employer fixed effect are 0.199 (0.040) and 0.155 (0.045), respectively (see Table A.5). The two estimates are qualitatively and quantitatively similar, which suggests that the results are not driven by unobserved differences across households. These estimates are smaller than the main estimates but this does not necessarily indicate any issues related to the validity of the empirical strategy.

The hiring rate at the parent's employer could be related to earnings through some other channel. First, the option to work for the parent's employer might raise an individual's reservation wage, leading them to match with better employers even if they do not end up working with their parent. Second, if the hiring rate is correlated with other measures of parental financial well-being, individuals might stay in school longer absent financial constraints. Both mechanisms ought to delay entry into the labor market. However, the estimates in Table A.6 illustrate that working for a parent's employer leads individuals to find their first stable job almost a year earlier and makes them slightly less likely to be employed in the three years prior to entry, which might indicate a smoother transition between school and work. Thus, there is no evidence that the earnings gains are driven by an increase in educational attainment or in the time spent searching for a job. This is not surprising in light of evidence from Hilger (2016) and Fradkin et al. (2018) who find that parental job loss during adolescence does not meaningfully impact educational attainment or job quality through extended search.

It is potentially problematic that working for a parent's employer affects the timing of entry. There are two stories for why the hiring rate at the parent's employer could affect the timing of entry. First, if there are job opportunities in the current period, the individual may start their career earlier if they anticipate not being able to find a better option in future periods. Second, if the parent's employer is not hiring when the individual decides to start looking for work, they may not find their first stable job until the parent's employer is hiring at later date. Both stories are more relevant for individuals who have more limited labor market options. This would then likely bias the estimates downward because individuals with low-earnings potential would be disproportionately likely to work for a parent's employer when the hiring rate is high. My main empirical specification measures the hiring rate at the parent's employer in the four quarters prior to when the individual enters the labor market. I assess the sensitivity of the estimates to shifting this window of measurement four quarters earlier and four quarters later. Shifting this window of measurement backwards one quarter or forwards three quarters yields qualitatively similar results with point estimates (standard errors) that range from 0.25 (0.013) to 0.42 (0.054). Outside of this range the point estimates grow larger (point estimates between 0.49 and 0.72), but first stage grows weaker.[36] See Table A.7 for all estimates. Taken together, these results suggest that is unlikely that issues related to timing of entry are driving the positive earnings benefits.

Lastly, it is possible that working for a parent's employer could affect whether an individual finds a first stable job. This would likely produce negative bias, since individuals with the lowest earnings potential would be disproportionately likely to appear at their parent's employers. Figure A.2 presents age-earnings profiles between the ages of 17 and 30 for different groups of workers defined by when they enter the labor market. For workers that ever enter the labor market, annual earnings rise dramatically and persistently at the time of entry. For workers that never enter the labor market, earnings remain persistently low (average annual earnings is only $1,814 at age 30). Workers who never enter the labor market simply never participate in work in a meaningful way. Based on this observation, it seems unlikely that an

[36]The F-statistic from the first stage falls to 177 when the hiring is measured between eight and four quarters prior to the quarter of entry.

individual would satisfy the earnings restriction for labor market entry only if they had the option to work for their parent's employer.

1.5.4 Mechanisms and Other Results

One possible channel through which working for a parent's employer could affect earnings is by matching individuals to firms that offer higher pay to all workers. I investigate this in column 1 of Table 1.4, where the outcome is the employer-level pay premium estimated from the full sample of workers in the Unites States via a model with worker and employer fixed effects as in Abowd et al. (1999), hereafter referred to as AKM (see Appendix A.2.6 for details). Working for the parent's employer leads individuals to work for employers that pay all workers 30.4 percent more. A comparison to the main results in Table 1.2 reveals that virtually the entire impact on individual earnings is explained by an improvement in the employer pay premium.[37]

A wide class of models illustrate how search and matching frictions lead to dispersion in firm-level pay policies.[38] In these models more productive firms pay more and poach workers from less productivity firms. Consistent with this class of models, columns 2-4 of Table 1.4 illustrate that working for the employer of a parent leads individuals to work at employers that are more productive (measured by revenue per

[37]I estimate a specification where the outcome variable is individual log earnings minus the employer pay premium and the estimated effect falls to 0.004 with a standard error of 0.03. This provides additional evidence that the earnings benefits are driven by access to higher-paying employers.

[38]Dispersion in firm-level pay policies also arise out of static models in which heterogeneous preferences over a firm's non-wage characteristics lead to imperfect competition (Card et al. 2018). While these models could also be used to interpret my results, dynamic models that emphasize the role of frictions (e.g., Burdett and Mortensen 1998; Postel-Vinay and Robin 2002) offer a more explicit explanation for the outcomes related to poaching hires and subsequent job mobility.

worker), that are more likely poach workers from other employers when hiring, and whose employees are paid more on average.[39] These results suggest that working for a parent's employer increases earnings by allowing individuals to start their careers on a higher rung of the job ladder. Column 4 suggests that individuals who work for their parent's employer end up at smaller firms. While job ladder models typically predict that larger firms will occupy higher rungs of the job ladder, Haltiwanger et al. (2018) find that firm age complicates this prediction because there are productive young firms that have not had ample time to grow into large firms. Consistent with this explanation, column 6 indicates that working for a parent's employer leads individuals to work for younger firms.

A number of papers find a systematic relationship between individual earnings and the identity of the employer. For example, workers tend to experience earnings growth when they move up the firm job ladder defined by productivity (Haltiwanger et al., 2017), poaching flows (Bagger and Lentz, 2019), and average pay (Haltiwanger et al., 2018). Furthermore, evidence from the AKM empirical model suggests that different workers who move between the same employers experience similar changes in earnings. One interpretation of this evidence is that some firms pay higher wages than others. However, this interpretation is complicated by the fact that worker mobility is endogenous: workers on an upwards (or downwards) career trajectory, might tend to move to certain firms.

My empirical strategy isolates exogenous variation in where individuals are em-

[39]The outcomes in columns 2-4 correspond to the rank of time-invariant characteristics of the first stable employer relative to the national distribution of employers. See Appendix A.2.7 for a description of how the employer- and firm-level variables are constructed.

Table 1.4: Effect on Employer Characteristics

		National Rank					Sector		
	employer pay premium (1)	revenue per worker (2)	poaching hires (3)	average earnings (4)	log firm size (5)	firm age in years (6)	unskilled services (7)	skilled services (8)	manufacturing/ production (9)
works for parent's employer	0.304***	4.767*	9.217***	25.060***	-1.979***	-3.040***	-0.433***	0.062*	0.372***
	(0.024)	(2.118)	(1.638)	(1.741)	(0.244)	(0.866)	(0.036)	(0.031)	(0.028)
control mean	0.359	51.75	55.33	42.47	6.72	22.78	0.375	0.460	0.165
control s.d.	0.366	27.52	23.38	26.53	3.40	12.28	0.484	0.498	0.371
observations	11,460,000	9,391,000	11,460,000	11,460,000	11,460,000	11,460,000	11,460,000	11,460,000	11,460,000

Notes: Each column presents estimates from a separate regression estimated by two-stage least squares. All outcomes measure characteristics of the employer of the individual at the first stable job. The outcome in column 9 is the employer-level pay premium estimated via a model with worker and employer fixed effects. The outcomes in columns 2-4 correspond to the rank of the employer based on time-invariant measures of log revenue per worker, the proportion of hires made through poaching workers from other employers and the average log earnings of all workers at the employer, respectively. The outcomes in column 5 and 6 are log firm size and firm age, respectively. The outcome in columns 7-9 are indicator variables equal to one if the employer of the child is in the unskilled service sector, skilled service sector, or the manufacturing/production sector, respectively. The endogenous variable is an indicator equal to one if the individual works for their parent's employer (primary earner) at the first stable job, which is instrumented for using the average quarterly hiring rate at the parent's employer in the four quarters prior to entry. All specifications include a fixed effect for the parent's employer, a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer, and the standard vector of demographic characteristics. Missing data in the outcome varaibles in columns 3 and 4 are imputed with the mean of the control group and the specifications include an indicator equal to one if the outcome is imputed. Standard errors are clustered at the level of parent's employer and are presented in parentheses.

Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

ployed and thus I provide more direct evidence that the firm-level pay policies are an important determinant of earnings. My results indicate that, for the complier population, the parent's employer occupies a higher rung of the job ladder than the employer that is the next best option and individuals earn more at their parent's employer. The fact that the estimated effect on individual earnings is virtually identical to the estimated effect on the employer pay premium is most easily explained by the following three statements being true: (1) my instrumental variables estimator identifies a causal parameter, (2) the AKM empirical model identifies an employer pay premium, and (3) the earnings benefits of working for a parent's employer are driven by parents providing access to higher-paying employers. In this way, my results offer novel support of the plausibility of the assumptions underlying the AKM empirical model. Importantly, the identifying assumptions imposed by AKM are entirely distinct from the assumptions required to interpret my two-stage least squares estimates as causal. Identification of the AKM model relies on assumptions on the relationship between unobserved error term and the individual- and employer-level components of earnings, whereas my empirical strategy makes no assumptions about the relationship between these variables.

Part of the effect on the employer pay premium is explained by parents providing access to employers in higher-paying industries. Columns 7-9 of Table 1.4 present estimates in which the outcome is an indicator equal to one if the child works in one of three broad sectors. Working for a parent's employer reduces the probability of working in the unskilled service sector by 43 percentage points and increases the probability of working in the manufacturing/production sector by 37 percentage points.

The effect on the industry of employment has large predicted earnings consequences. Table A.8 presents estimates in which the outcome variable is the industry-level earnings premium (estimated analogously to the employer-level pay premium). Working for a parent's employer increases the two- and six-digit industry pay premium by 0.167 and 0.230, respectively. Thus, 75 percent of the effect on individual earnings is explained by individuals working in different six-digit industries. To the extent that young workers are aware of pay differences across industries, these results cast doubt on the possibility that parents simply provide information to their children about where to look for high-paying jobs.

Working for a parent's employer leads individuals to stay at their first employer longer. Column 1 of Table 1.5 indicates that working for a parent's employer increases the probability of remaining at the first employer for at least three years by 17.4 percentage points. Columns 2 and 3 illustrate that this effect is entirely driven by a reduction in the probability of making a job-to-job transition. If the outcomes in columns 2 and 3 are viewed as proxies for quits and fires, respectively, then these results suggest that working for a parent's employer allows individuals to gain access to employers that pay more and are generally more desirable than their outside option and so they choose to remain at those employers, whereas the employers are not gaining access to better workers and so they are no less likely to fire these workers.[40] However, employers may benefit from the lower quit rates if hiring and retention are costly.

[40] Fallick et al. (2019) find a strong association between transitions into nonemployment and earnings losses, which lends credibility to this interpretation of job-to-job and job-to-nonemployment transitions.

Table 1.5: Effect on Earnings and Job Mobility Three Years After Entry

	First Move in Three Years			Annual Earnings After Entry		
	stay	j2j	j2n	year one	year two	year three
	(1)	(2)	(3)	(4)	(5)	(6)
works for parent's employer	0.174***	-0.182***	0.008	7,363***	7,226***	4,790***
	(0.034)	(0.033)	(0.038)	(1,002)	(1,284)	(1,440)
control mean	0.286	0.279	0.435	26,460	25,660	26,740
control s.d.	0.452	0.449	0.496	15,660	19,290	21,130
observations	10,200,000	10,200,000	10,200,000	10,200,000	10,200,000	10,200,000

Notes: Each column presents estimates from a separate regression estimated by two-stage least squares. The outcome in column 1 is an indicator variable equal to one if the child stayed at their first employer for three years. The outcomes in columns 2 and 3 are indicator variables equal to one if the child left their first employer within three years and made a job-to-job (j2j) or a job-to-nonemployment (j2n) transition for their first move, respectively. Note that the outcome variables in columns 1-3 are mutually exclusive and exhaustive. The outcome varaibles in columns 4-6 are the annual earnings one, two and three years after the quarter of entry, respectively. The endogenous variable is an indicator equal to one if the individual works for their parent's employer (primary earner) at the first stable job, which is instrumented for using the average quarterly hiring rate at the parent's employer in the four quarters prior to entry. The sample includes individuals who found their first stable job prior in 2014:Q1 or earlier. The F-statistic from the first stage for this sample is 1,532. All specifications include a fixed effect for the parent's employer, a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer, and the standard vector of demographic characteristics. Standard errors are clustered at the level of parent's employer and are presented in parentheses.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Columns 4-6 of Table 1.5 illustrate that the earnings benefits of working for a parent's employer are quite persistent. Working for the parent's employer leads to an increase of \$7,363 in the first year after entry into the labor market. The effects are persistent but by the third year the magnitude of the effect falls to \$4,790. Figure A.8 presents estimates of the effect on annual earnings one to six years after entry for a group of individuals for whom I am able to observe these outcomes. There is less statistical precision in the later years but the point estimates suggest that the earnings benefits are quite persistent, with annual earnings benefits that exceed \$5,000 even six years after entry.

Table 1.6 investigates heterogeneous effects by estimating the main specification on subgroups of workers defined by sex and the quintile of parental earnings. On average, the benefits for daughters (0.424) is larger than for sons (0.312). While earnings benefits for daughters are also larger within each parental earnings quintile, large standard errors prevent me from concluding whether or not there are meaningful differences between the earnings effects by sex. A comparison of estimates across columns 1-5 in Panel A indicates that children with parents higher up in the parental earnings distribution experience greater earnings benefits from working for a parent's employer. For example, the estimated effect for individuals whose parents are in the fifth quintile (highest earnings) is 0.328 compared to 0.189 for individuals whose parents are in the first quintile (lowest earnings). The estimates in Panels B and C illustrate that the positive association between the effects on earnings and parental earnings is entirely driven by sons.

I use the same empirical strategy to investigate the earnings consequences of

Table 1.6: Heterogeneous Effects by Sex and Parental Earnings

	log of quarterly earnings					
	(1)	(2)	(3)	(4)	(5)	(6)
A. All						
works for parent's employer	0.189*	0.273***	0.234***	0.362***	0.328***	0.307***
	(0.082)	(0.063)	(0.062)	(0.077)	(0.085)	(0.029)
F-statistic	238.9	358.0	446.8	330.4	316.2	1,434
observations	1,350,000	1,987,000	2,297,000	2,462,000	2,487,000	11,460,000
B. Daughters						
works for parent's employer	0.384*	0.413***	0.356*	0.443*	0.417*	0.424***
	(0.176)	(0.129)	(0.164)	(0.176)	(0.188)	(0.057)
F-statistic	64.3	131.2	100.0	106.1	130.9	679.8
observations	586,000	876,000	1,029,000	1,128,000	1,152,000	5,387,000
C. Sons						
works for parent's employer	0.075	0.222**	0.316***	0.405***	0.398***	0.312***
	(0.115)	(0.083)	(0.077)	(0.099)	(0.116)	(0.036)
F-statistic	97.7	198.3	245.7	176.9	161.3	854.2
observations	600,000	909,000	1,067,000	1,149,000	1,148,000	5,501,000
Sample Description						
parental earnings quintile	first	second	third	fourth	fifth	all

Notes: This table presents estimates on subsamples defined by the interaction between the quintile of parental earnings (defined by the column) and sex (defined by the panel). The outcome in all columns is the log of the first full quarter of earnings at the first stable job. The endogenous variable is an indicator equal to one if the individual works for their parent's employer (primary earner) at the first stable job, which is instrumented for using the average quarterly hiring rate at the parent's employer in the four quarters prior to entry. All specifications include a fixed effect for the parent's employer, a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer, and the standard vector of demographic characteristics. Standard errors are clustered at the level of parent's employer and are presented in parentheses.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

working for the employer of the parent who is the secondary earner.[41] The results, presented in Table A.9, indicate that working for the employer of the secondary earner leads to an increase in initial earnings by 29 percent. Thus, there is no evidence that working for the employer of the secondary earner produces different earnings benefits compared to working for the employer of the primary earner. In addition, Table A.10 presents the estimated effect of working for the father's and mother's employer. Working for the father's employer leads to a 52 percent and 33 percent increase in initial earnings for daughters and sons, respectively. Working for the mother's employer leads to a 28 percent and 34 percent increase in initial earnings for daughters and sons, respectively. Thus, there are substantial earnings benefits associated with working for the employer of either, particularly for daughters who are able to get a job at their father's employer.

1.5.5 Interpreting the Local Average Treatment Effect

If the three identifying assumptions are satisfied, the two-stage least squares estimator identifies a LATE, which is the average treatment effect for the compliers. This section explores the relationship between the LATE and other causal parameters of interest. First, I provide a theoretical argument for why in my context, in which working for a parent's employer is determined by the decisions of multiple agents, the LATE may be a reasonable approximation of the *average treatment effect on the treated* (ATT)–which is the average treatment effect for individuals who work for their

[41] In order to avoid estimating effects of working with the primary earner, I limit the sample to cases in which the secondary earner is employed with a year of tenure in the quarter of entry and does work work at the same employer as the primary earner.

parent's employer. Second, I present empirical evidence to assess the plausibility of this interpretation.

Let $Y_i(d, z)$ denote the potential outcome of individual i who has the treatment status $D_i = d \in \{0, 1\}$ and instrument value $Z_i = z \in \{\underline{z}, \bar{z}\}$ where $\underline{z} < \bar{z}$. Let D_{zi} denote the treatment status of i when $Z_i = z$. Furthermore, assume the following: (Independence) $\{Y_i(D_{\bar{z}i}, \bar{z}), Y_i(D_{\underline{z}i}, \underline{z}), D_{\bar{z}i}, D_{\underline{z}i}\} \perp\!\!\!\perp Z_i$, (Exclusion) $Y_i(d, \underline{z}) = Y_i(d, \bar{z}) \equiv Y_{di}$ for $d = \{0, 1\}$, (First Stage) $\mathbb{E}[D_{\bar{z}i} - D_{\underline{z}i}] \neq 0$, and (Monotonicity) $D_{\bar{z}i} \leq D_{\underline{z}i}$ $\forall$ i. Under these assumptions, the instrumental variables estimator identifies a LATE, which is the average treatment effect for the compliers (i.e., the population for which $D_{\bar{z}i} < D_{\underline{z}i}$).

In the standard selection framework of Roy (1951), the LATE will likely depend on the specific values of the instruments, since selection into treatment is determined by a single agent who weighs the benefits (treatment effects) against the costs (instruments). To see this more formally, consider the selection model in which $D_{zi} = \mathbb{1}\{\beta_i > z\}$, where $\beta_i = Y_{1i} - Y_{0i}$ is the individual-level treatment effect. It immediately follows that the LATE, which is $\mathbb{E}[\beta_i | \underline{z} < \beta_i < \bar{z}]$, will generally depend on the values of the instruments.

In my context, selection is determined by the choices of more than one agent–the young worker and their parent's employer–and this potentially breaks the link between the instruments and the treatment effects. To see why, consider an alternative selection model in which the individual works for their parent's employer if and only if the employer makes them a job offer and they choose to accept the offer. The employer's decision to make an offer depends on the instruments and is defined as,

$O_{zi} = \mathbb{1}\{\eta_i^O > z\}$. The child's decision to accept the offer depends on the benefits and is defined as, $A_{zi} = \mathbb{1}\{\beta_i > \eta_i^A\}$. Where η_i^O and η_i^A are unobserved error terms whose values are defined independent of D_i and Z_i.[42] Treatment status is then defined as, $D_{zi} = O_{zi} \times A_{zi}$.

The LATE and ATT are equal if the employer's decision to make an offer is unrelated to the child's decision to accept. Formally, if $\{\eta_i^O, \eta_i^A\} \perp\!\!\!\perp Z_i$ and $\{\beta_i, \eta_i^A\} \perp\!\!\!\perp \eta_i^O$, then

$$\underbrace{\mathbb{E}\big[\beta_i | \{\eta_i^A < \beta_i\}, \{\underline{z} < \eta_i^O < \bar{z}\}\big]}_{\text{LATE}} = \underbrace{\mathbb{E}\big[\beta_i | \{\eta_i^A < \beta_i\}, \{Z_i < \eta_i^O\}\big]}_{\text{ATT}} \tag{1.7}$$

Under these conditions, both the compliers and the individuals working for their parent's employer are a random sample of individuals who would accept an offer from their parent's employer if made one. Importantly, because of the multi-agent nature of the selection problem, the LATE and ATT may be equivalent even in the presence of selection on gains and selection bias. Appendix A.4 develops a stylized behavioral model and provides a more detailed discussion of the intuition by focusing on a specific case of equation 1.7.

The assumptions that imply the equality of the LATE and ATT also imply that the estimated treatment effects should not be sensitive to the variation exploited in the instrument; I test that implication here. To do so, I regress the instrument on the covariates from equation 1.6 and compute the residualized value, which is the

[42] More formally, let $\eta_i^x(d, z)$ denote the potential outcome with treatment status $D_i = d$ and instrument value $Z_i = z$. Then I assume that $\eta_i^x = \eta_i^x(d, z)$ for $x \in \{O, A\}$.

source of identifying variation.[43] I then compute terciles based on the residualized instrument, partitioning the sample into periods in which employers have a relatively low, medium and high rate of hiring. I estimate equation 1.6 on samples defined by different combinations of the three terciles. The point estimate (standard error) is 0.44 (0.05), 0.31 (0.029), and 0.23 (0.11) when excluding observations from third, second and first terciles, respectively (see Table A.11). While there is some variation across the samples, the two-stage least square estimates are not excessively sensitive to range of variation exploited in the instrument.

An alternative approach to assessing the representativeness of the two-stage least squares estimates is to characterize the compliers. My data lack variables that strongly predict individual earnings benefits, but I can estimate the size of the complier population. The methodology developed by Abadie (2003) applies to binary instruments, so I construct three binary instruments which are equal to one when the residualized hiring rate exceeds the 25th, 50th, and 75th percentiles. The estimated effect on log earnings when using these three binary instruments is 0.44, 0.42, and 0.25, respectively. The fact that these estimates are qualitatively similar to those obtained using the continuous instrument provides some evidence that the complier population for these instruments is not fundamentally different. For the three instruments, I find that 3.6, 2.8, and 16 percent of the population is in the complier population, respectively.[44] Given that 5.6 percent of individuals in the estimation sample work for their parent's employer, the compliers represent a meaningful per-

[43]The distribution of the residualized hiring rate is both symmetric and smooth (see Figure A.9).
[44]Table A.12 presents estimates of the size and characteristics of the complier population.

centage the treated population.[45] Thus, the results provide additional evidence that the instrumental variables estimates are informative of the ATT.

1.6 Intergenerational Persistence in Earnings

The results from Sections 1.4 and 1.5 show that individuals with higher-earning parents are both more likely to work for the employer of a parent and benefit more when the do. This suggests that the intergenerational transmission of employers increases the intergenerational persistence in earnings. This section implements the methodology described in Section 1.2 in order to quantify the difference between the observed measure of the intergenerational persistence in earnings and a measure that correspond to a counterfactual world in which no one works for the employer of a parent.

Panel A of Table 1.7 presents estimates of the IGE. Columns 1-3 present estimates for daughters, sons, and the full sample. The elasticities, which range from 0.13 to 0.16, are substantially lower than typical estimates of IGE from the literature (Black and Devereux, 2011). To investigate this discrepancy, I produce alternative estimates of the IGE, which measure the earnings of the children in 2016 (when the children are between the ages of 24 and 35). In Table A.13, columns 1-3 of Panel A present estimates based on samples that include children with zero earnings in 2016 (by taking the hyperbolic sine of earnings) and the estimates of the IGE are closer to 0.4, which is comparable to other estimates from the United States. Thus, the low estimates of

[45] These estimates suggest that 49, 25, and 69 percent of the individuals who work for their parent's employer are in the complier population, respectively.

the IGE appear to be an artifact of focusing on labor market outcomes at the time of entry. Panel B of Table A.13 presents estimates of the IGE for a subsample of the children with strictly positive earnings in 2016. Within this sample, the estimated IGE for the full sample is 0.235, which is much closer to the estimates based on initial labor market outcomes. These results highlight the fact that the IGE is sensitive to how observations with zero earnings are dealt with and suggests that my estimates of the IGE based on initial labor market outcomes are lower primarily because they condition on positive earnings.[46]

The intergenerational transmission of employer leads to a modest reduction in the IGE. Panel B of Table 1.7 indicates that the IGE would be 5 percent lower if no one worked for the employer of the parent who is the primary earner.[47] The intergenerational transmission of employers has a larger effect on the IGE for sons because, relative to daughters, both the probability of working for a parent's employer and the earnings consequences are more strongly related to parental earnings. Columns 1 and 2 indicate that the counterfactual IGE for daughters and sons would be 2 percent and 11 percent lower, respectively. Panel C of Table 1.7 indicates that the IGE would be about 10 percent lower if no one worked for the employer of either parent. For the case of working for the employer of the secondary earner, I am unable to estimate

[46]To further investigate these patterns, Figure A.10 plots the avergae log earnings of children against parental earnings and illustrates that the strength of the intergenerational relationship in earnings is dampened in the lower parts of the distribution. This may be explained by the fact that I focus on the earnings at the first stable job, when many workers are earning the minimum wage.

[47]Estimates of the ATT can be found in Table 1.6. I allow all estimates of the ATT to vary by parental earnings quintile. For the counterfactual estimates presented in column 3 of Table 1.7 I use the pooled estimates of the ATT presented in Panel A of Table 1.6. For the counterfactual estimates presented in columns 1 and 2, I use the appropriate sex-specific estimates. Estimates of the proportion of individuals who work for the employer of the primary earner, secondary earner, or both parents are presented in Figure A.11.

Table 1.7: Intergenerational Elasticity of Earnings

	sample		
	daughters (1)	sons (2)	all (3)
A. Observed			
IGE	0.1565	0.1298	0.1430
	(0.0002)	(0.0003)	(0.0002)
B. No Transmission with Primary Earner			
percent change in IGE	-2.04%	-10.79%	-4.73%
	(6.52)	(5.02)	(3.30)
C. No Transmission with Either Parent			
percent change in IGE	-3.87%	-23.09%	-9.68%
	(12.25)	(9.39)	(6.16)
observations	8,416,000	8,591,000	17,010,000

Notes: The results in columns 1-3 corresesopnd to daughters, sons, and all children, respectively. Panel A presents the observed intergenerational elasticity of earnings (IGE), which is denoted $\rho(y_{ijt}, y_p)$ and is estimated with sample weights via weighted least squares. Panels B and C present the percent by which the IGE estimates in Panel A would change if no children were to work for the employer of the parent who is the primary earner or either parent, respectively. The percent change is defined as, $\frac{\rho(y_{ijt}, y_p) - \rho(y_{i(j0)t}, y_p)}{\rho(y_{ijt}, y_p)} \times 100$. The treatment effects used to construct the counterfactual estimates are estimated via two-stage least squares and are estimated separately by the quintile of the parental earnings distrbution for the results in column 3 and are estimated separately by quintile of the parental eanrings distribution and the sex of the child for columns 1 and 2. Standard errors are presented in parentheses and are calculated using the delta method and take into account the uncertainty in the estimated earnings consequences.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

heterogeneous effects by both parental earnings and sex. Thus, I assume that the effect of working for the employer of the secondary earner is the same as working for the employer of the primary earner. As previously discussed, this appears to be true, at least in the full sample.

The standard errors, presented below in parentheses, indicate that uncertainty in the estimates of the ATT create some uncertainty regarding the magnitude of these effects. Table A.14 replicates the analysis but uses the point estimates and standard errors from Table 1.2, which assumes no heterogeneity in effects by parental earnings or sex. Here the magnitudes are smaller, suggesting a 2 percent decline in IGE if no one worked for a the employer of either parent, but they are also much more precisely estimated (standard error is 0.20). While there is some uncertainty around the exact magnitude, both sets of results suggest that the transmission of employers leads to a modest decrease in the IGE.

In addition to the IGE, I consider an alternative measure of the intergenerational persistence in earnings: the conditional expected rank (CER). The CER is defined as, $E[r_{ij}|r_p]$, where r_{ij} is the percentile rank of the earnings of the child, calculated within cohorts and using sample weights. Figure 1.2 presents the average earnings benefits–defined as $\mathbb{E}[D_i\beta_i]$–by sex, race/ethnicity and parental earnings.[48] The benefits are largest for non-Black males whose parents are in the top two quintiles of the earnings distribution. The results by race/ethnicity should be interpreted with some caution since I do not have sufficient power to estimate earnings consequences by parental

[48]The two-stage least squares estimates of the effect of working for the employer of the primary on the earnings rank of the children are presented in table A.15.

Figure 1.2: Average Benefits of Working for Parent's Employer

(A) Daughters

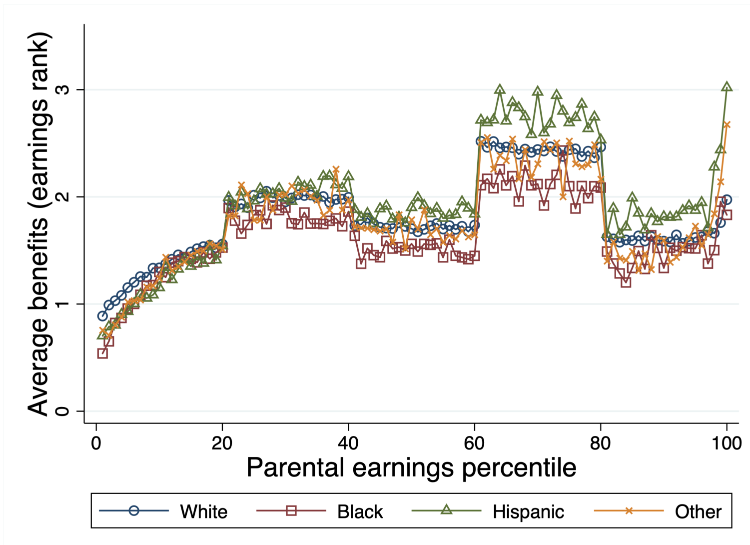

(B) Sons

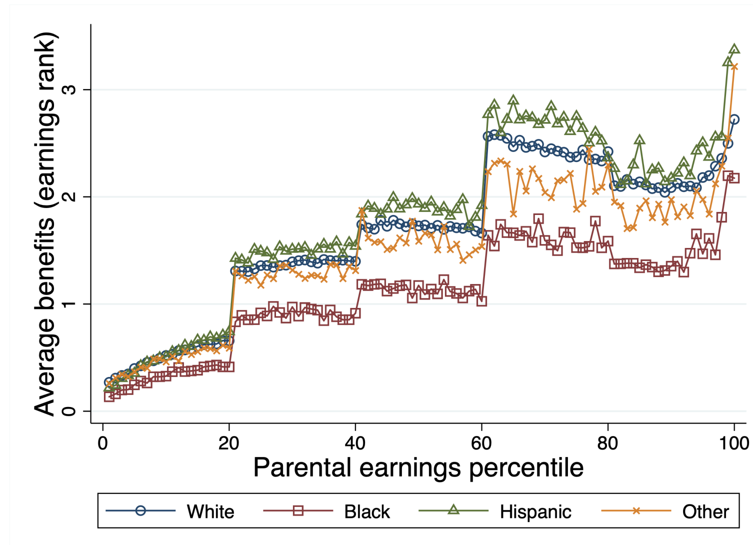

Notes: Each point presents the average earnings benefits from working for the employer of either parent for the sub group defined by sex (daughters in Panel A and sons in Panel B), the percentile of parental earnings (defined by the x-axis), and race/ethnicity. The earnings benefits are measured as the percentile rank. The treatment effects used to construct the counterfactual estimates are estimated via two-stage least squares and are estimated separately by the quintile of the parental earnings distribution and the sex of the child. All statistics, aside from the two-stage least squares estimates, are calculated using sample weights.

Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

earnings, sex, and race/ethnicity and instead assume that, within groups defined by sex and the parental earnings quintile, average treatment effects do not differ by race/ethnicity.

Section 1.4 found that Black sons are less likely to work for the employer of a parent relative to other sons in the same parental earnings percentile. This result is interesting in light of recent work by Chetty et al. (2020), who find that, conditional on parental income, Black males have lower expected income compared to White males. The solid lines in Figure 1.3 present the CER measures for Black and White sons and replicate the finding of Chetty et al. (2020): Black sons earn less on average relative to White sons with parents in the same earnings percentile. The dashed lines below represent the counterfactual CER. Black sons benefits less from working for a parent's employer. Panel B of Figure 1.3 presents the proportion of the Black-White gap (vertical distance between the solid lines in Panel A) that is explained by the intergenerational transmission of employers at each percentile of the parental earnings distribution. On average, the transmission of employers explains about 10 percent of the conditional Black-White earnings gap. While other factors clearly play an important role in determining this earnings gap, my results suggest that young Black males are at a relative disadvantage in part because they are less likely to have an employed father who can help them find work.

The difference in the average benefits of the intergenerational transmission of employers between sons and daughters varies across the parental earnings distribution. Panel A of Figure 1.4 plots the observed and counterfactual CER.[49] Estimating a

[49] The counterfactual is defined as $\mathbb{E}\big[r_{ij(0)}|r_p\big] = \mathbb{E}\big[r_{ij}|r_p\big] - \mathbb{E}\big[D_i|r_p\big] \times \mathbb{E}\big[\beta_i|D_i = 1, r_p\big]$.

Figure 1.3: Black-White Earnings Gap for Sons

(A) Observed and Counterfactural

(B) Explained by Employer Transmission

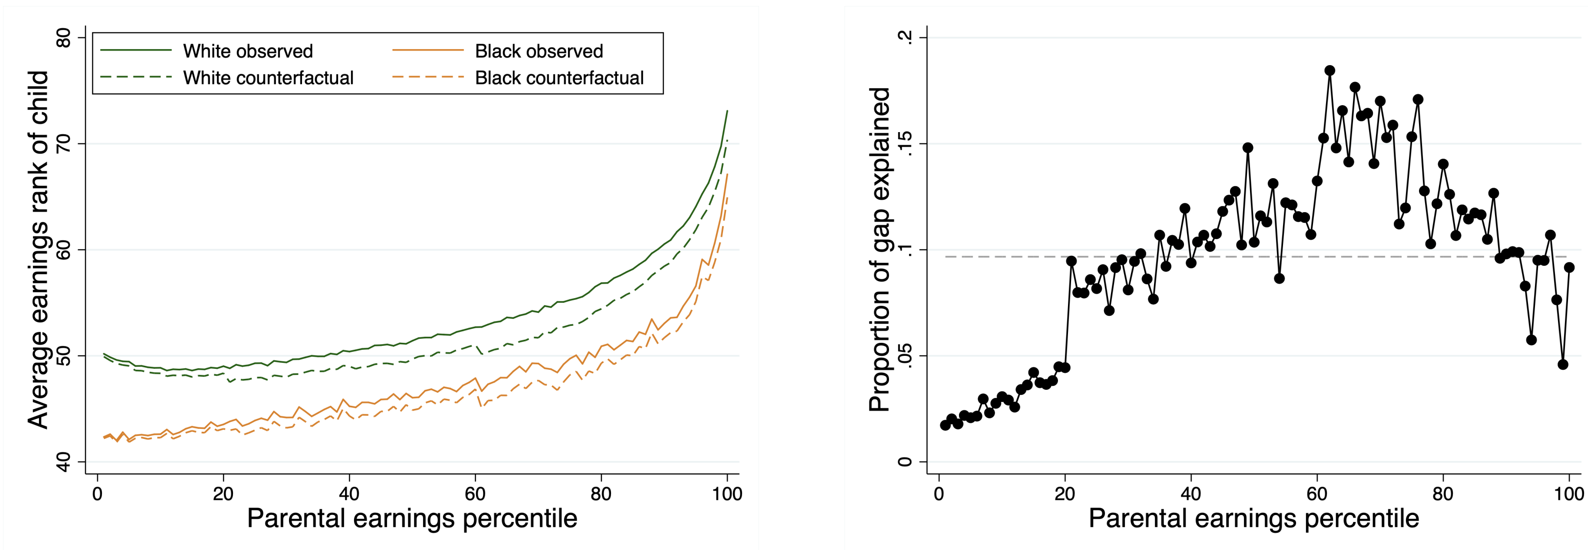

Notes: The solid lines in Panel A present the conditional expected rank (CER) separately for White and Black sons. The dashed lines represent the counterfactual CER that correspond to the counterfactual in which no individual works for the employer of either parent. Panel B plots the proportion of the Black-White earnings gap that is explained by the transmission of employers for each percentile of the parental earnings distribution. The dashed line represents the average across all percentiles. The treatment effects used to construct the counterfactual estimates are estimated via two-stage least squares and are estimated separately by the quintile of the parental earnings distribution and sex of the child. All statistics, aside from the two-stage least squares estimates, are calculated using sample weights.

Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Figure 1.4: Gender Earnings Gap

(A) Observed and Counterfactural

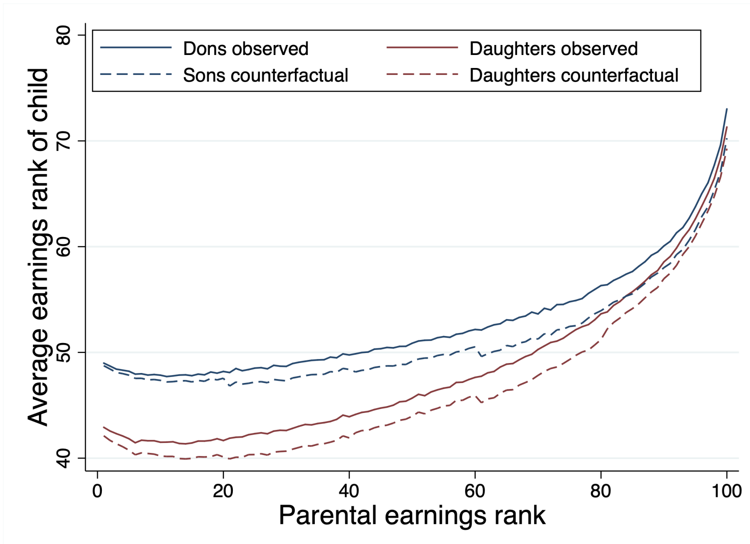

(B) Expected Benefits

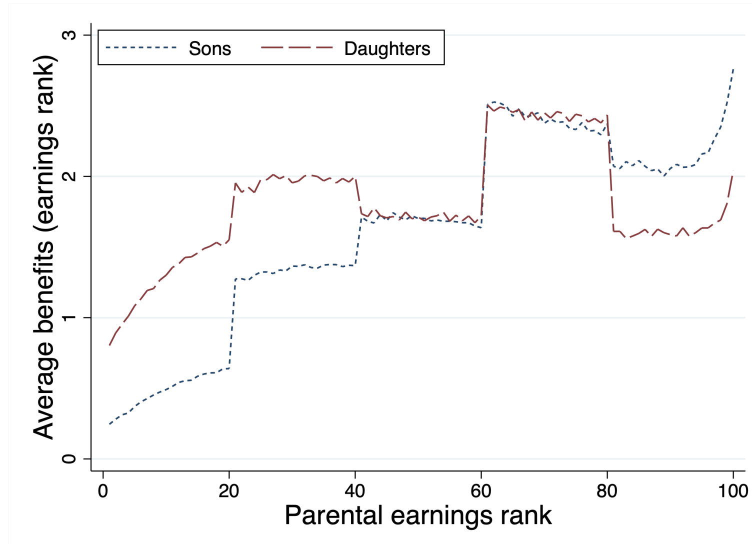

Notes: The solid lines in Panel A present the conditional expected rank (CER) separately for sons and daughters. The dashed lines below them represent the counterfactual CER that correspond to the counterfactual in which no individual works for the employer of either parent. Panel B plots the expected earnings benefit, which is the difference between the observed and counterfactual CER. The treatment effects used to construct the counterfactual estimates are estimated via two-stage least squares and are estimated separately by the quintile of the parental earnings distribution and the sex of the child. All statistics, aside from the two-stage least squares estimates, are calculated using sample weights.
Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

linear regression using the results in Panel A reveals that, if no one worked for the employer of either parent, the slope of the rank-rank relationship would be 13 percent and 3 percent lower for sons and daughters, respectively.[50] While sons are more likely to work for the employer of either parent, daughters experience larger earnings benefits conditional on doing so. Panel B illustrates how this plays out across the parental earnings distribution and plots the expected benefits of working for a parent's employer–which corresponds to the difference between the observed CER and the counterfactual CER–against the parental earnings percentile. Daughters benefit more from working for a parent's employer in the bottom two quintiles of the parental earnings distribution while sons benefits more in the top quintile and both benefit equally elsewhere. Averaging these effects across the parental earnings distribution indicates that the earnings gap between sons and daughters would be 4 percent larger if no one worked for the employer of either parent.

1.6.1 Key Insights from Stylized Model

I synthesize my findings by developing a stylized model of intergenerational mobility in which parents affect the earnings of their children by shaping the development of their human capital and by providing access to higher-paying employers. I summarize the key insights from the model in this section and refer the reader to Appendix A.4 for the details. Relative to other models of intergenerational mobility, the novel features of my model are that I (1) incorporate an employer-specific component into

[50]Panel A of Figure A.12 plots analogous results for the pooled sample of sons and daughters. The counterfactual slope of the rank-rank relationship is 2 percent and 5 percent lower for the scenario in which no individual worker for the employer of the primary earner and secondary earner, respectively.

individual earnings and (2) explicitly model the choices that lead individuals to work for a parent's employer.[51]

The key aspects of the model setup are as follows. Motivated by my finding that parents provide access to higher-paying employers as well as the literature on imperfect competition in the labor market, I depart from existing models of intergenerational by allowing earnings to depend on not only the human capital of the child, but also the pay premium associated with the employer to which they match. I assume that: there is a positive correlation between parental earnings and human capital, children with higher levels of human capital tend to match to higher-paying employers absent parental contacts, and working at the parent's employer affects earnings solely through the effect on the employer pay premium. Individuals will work for the employer of their parent if and only if the employer makes an offer and the child accepts the offer. The employer's decision to make a job offer depends on the human capital of the child and the parent, whereas the child's decision to accept the offer depends on the earnings benefits.

There are two insights from the model. First, the effect of the intergenerational transmission of employers on the intergenerational persistence in earnings is theoretically ambiguous in sign. On the one hand, higher-earning parents are better able to produce high-paying job offers. On the other hand, children of lower-earning par-

[51] Magruder (2010) and Corak and Piraino (2012) and the only two papers that have developed models of intergenerational mobility that incorporate parental contacts. Neither paper considers the role of employer pay premiums and neither paper considers the endogenous use of social contacts. Magruder (2010) assumes that parental contacts produce a positive correlation between the employment status of parents and children. Corak and Piraino (2012) assume that the earnings of the parent have a direct positive effect on the earnings of the child. In both papers the effect of parental contacts on intergenerational mobility is determined by the sign of a single parameter.

ents have lower levels of human capital and may be more reliant on their parents to find a decent-paying job. Thus, while I find that the intergenerational persistence in earnings would be lower if no one worked for their parent's employer, this conclusion might differ in other contexts.

Gaining access to higher-paying employers is the *direct effect* of working for a parent's employer. However, parents might account for this direct effect when investing in the human capital of their children; I refer to this as the *indirect effect*. The second insight of the model is that the indirect effect has a theoretically ambiguous effect on the intergenerational persistence in earnings. On the one hand, working for a parent's employer increases the marginal returns to human capital investments by providing access to higher-paying employers. On the other hand, the marginal returns decline because higher-ability individuals are less likely to work for their parent's employer and benefit less when they do. Thus, parental investment decisions could either amplify or dampen the direct effect of the intergenerational transmission of employers on the intergenerational persistence in earnings. The counterfactual exercise analysis should therefore be viewed as a partial equilibrium analysis, which does not account for the possibility that individuals might adjust their investments in human capital if there was no option to work for their parent's employer. While it would be interesting to explore the implications of the indirect effect, quantifying the importance of the direct effect is the obvious starting point, as the indirect effect is unlikely to be important if direct effect is negligible.

1.7 Conclusion

My paper combines survey and administrative data to investigate how the earnings of young workers are affected by the intergenerational transmission of employers. I start with a descriptive analysis, and find that 7 percent of individuals work for the employer of a parent at their first stable job and 29 percent do so at some point between the ages of 18 and 30. This tendency is best explained by parents influencing the hiring or job search process to help children who have limited options in the labor market. I then use an instrumental variables strategy, which exploits exogenous variation in the availability of jobs at the parent's employer, and find that working for the employer of a parent increases earnings by 31 percent. These large earnings benefits are explained by parents providing access to higher-paying employers: Young workers who find their first stable job at the employer of a parent start their careers on a higher rung of the job ladder.

Individuals with higher-earning parents are more likely to work for the employer of a parent, and benefit more when they do, and thus the intergenerational transmission of employers increases the intergenerational persistence in earnings. I develop a new methodology that allows me to quantify this effect using descriptive statistics and causal estimates. I find that the elasticity of the initial earnings of an individual with respect to the earnings of their parents would be 10 percent lower if no one worked for the employer of a parent. Examining patterns by family background, sex, and race/ethnicity reveals that non-Black males with high-earning parents benefit the most from the intergenerational transmission of employers. My results likely understate the importance of parental labor market networks more broadly defined since parents may also provide access to jobs at other employers through social contacts, such as friends, former co-workers, or classmates. This is especially true for the implications for intergenerational mobility if higher-income parents are more likely to have contacts outside of their current employer.

My results relate to the normative assessment of whether rates of intergenerational mobility are too low in the United States, an assessment which depends on whether the economic system that produces the intergenerational persistence in earnings is equitable and efficient. While equity depends on subjective moral values, a core ideal in the United State is that of equality of opportunity, which requires that an individual's success be a function of their hard work and ability rather than the circumstances into which they were born.[52] Thus, from an equity standpoint, my finding that individuals from high-income families disproportionately benefit from their parents' connections should raise concerns about the relatively low levels of intergenerational mobility in the United States. My results do not speak directly to the implications for efficiency and future research should aim to understand whether the use of parental labor market networks leads to gains or losses in productivity.

My results are also informative of the positive assessment of what would be required to achieve equality of opportunity. One view is that the United States is a meritocracy, where economic rewards are determined by hard work and ability. Ac-

[52]According to Roemer (1998), equality of opportunity requires that the outcomes of individuals are not systematically determined by factors for which they are not responsible. Defining what to hold someone responsible for is a subjective judgment. But most people in the United States would likely agree that individuals should not be responsible for their parents' lack of connections in the labor market.

cording to this view, efforts to expand economic opportunity should aim to equip everyone with the skills they need to succeed in the labor market. Government programs such as Head Start, which provides access to early childhood education, and the Pell Grant program, which helps students pay for college, are both examples of programs that promote the development of skills for individuals from disadvantaged backgrounds. However, my results challenge a purely meritocratic view of the labor market, as individuals from high-income families are likely to earn more not only because they are more skilled, but also, because their parents are able to provide access to high-paying firms. If the labor market plays a direct role in propagating intergenerational disadvantage, then achieving equality of opportunity in terms of education will not necessarily produce equality of opportunity in the labor market. Rather, individuals from disadvantaged backgrounds may require additional support throughout their early careers. Gaining a better understanding of the mechanisms through which parents help their children find high-paying jobs may offer ideas for how to help young workers who cannot rely on the connections of their parents to more successfully navigate the labor market.

Chapter 2: The Children of HOPE VI Demolitions: National Evidence on Labor Market Outcomes

Disclaimer

This chapter is joint work with John C. Haltiwanger, Mark J. Kutzbach, Giordano Palloni, Henry O. Pollakowski, and Daniel H. Weinberg. Any opinions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau or the Federal Deposit Insurance Corporation. All results have been reviewed to ensure that no confidential information is disclosed (q.v. U.S. Census Bureau Disclosure Review Board numbers: DRB-B0038-CED-20190405, DRB-B0071-CED-20190829, and CBDRB-FY20-CED0060029). Much of the work for this analysis was done while Mark Kutzbach was an employee of the Census Bureau. John Haltiwanger was a Schedule A (part-time) employee and Matthew Staiger a Pathways Intern of the U.S. Census Bureau at the time of the writing of this paper. This research has been supported by grant number 98082 from the "How Housing Matters" research program of the John D. and Catherine T. MacArthur Foundation, by NSF grant number 1730108, by a Research Partnership grant from the U.S. Department of Housing and Urban Development, and by a grant from the Russell Sage Foundation. This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants: SES-9978093, SES-0339191, and ITR-0427889; a National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation.

2.1 Introduction

The concern that placement in subsidized housing, especially large public housing projects in high-poverty neighborhoods, could negatively affect children has been the focus of a substantial literature.[1] Based partly on this rationale, the last 30 years of federal assisted housing policy has sought to deconcentrate subsidized housing participants, mainly through the provision of Housing Choice Vouchers (hereafter, vouchers) that subsidize low-income families to live in market-supplied housing. A significant effort to spur the dispersion of these households has focused on the demolition of public housing projects paired with support for existing residents to find alternative housing, most notably under the U.S. Department of Housing and Urban Development's (HUD's) HOPE VI program.[2] Despite the growing availability of vouchers and the continued funding of programs intended to reduce the population living in low-quality public housing projects, there is little representative evidence about how these demolitions affect short- or long-term outcomes for exposed children and adults; much of the existing research on subsidized housing is conducted in a limited number of large metropolitan areas and it is not clear how these findings apply to other contexts in the U.S.

This paper explores how the HOPE VI Demolition program affected the adult labor market outcomes of children who resided in demolished projects. Using a unique

[1]For example, Oreopoulos (2003), Jacob (2004), Chetty et al. (2016), Andersson et al. (2018a), and Chyn (2018).

[2]The HOPE VI program, originally known as the Urban Revitalization Demonstration, was the sixth of the Housing Opportunities for People Everywhere Grants, funded by P.L. 102-389 (HUD 2007).

dataset available at the U.S. Census Bureau–which links administrative data on earnings and participation in subsidized housing–we identify approximately 18,500 children exposed to 160 HOPE VI demolitions in diverse environments across the U.S. Our empirical strategy for estimating the causal impacts of the program is based on the observation that even though the HOPE VI program systematically targeted the "worst" public housing projects, there were many similarly distressed projects in equally disadvantaged neighborhoods that were not demolished. We leverage the richness and size of the data by using a stratification with regression estimator (Imbens and Rubin 2015), which combines features of both matching and regression in order to flexibly account for observable differences between the 160 HOPE VI projects and 8,800 public housing projects unaffected by the program.

Our main finding is that exposure to the HOPE VI Demolition program between the ages of 10 and 18 produced substantial long-run labor market benefits, increasing age 26 earnings by 14.2 percent relative to comparable children from non-HOPE VI projects. Interestingly, we find that the positive impacts are driven by children from projects in neighborhoods served by the larger housing authorities typically located in large metro areas. For example, we estimate that HOPE VI increased earnings by 19.5 percent in large (greater than 2,500 units) Public Housing Authorities (PHAs), compared to a statistically insignificant 4.5 percent increase in smaller PHAs.

We start our investigation of mechanisms by studying the short- and medium-term impacts of the program. The demolitions led to large changes in housing circumstances, forcing most HOPE VI households out of their initial projects and into other public housing projects or the voucher program. While households exited subsidized housing at a higher rate in the year after the demolition, there is no evidence that the program displaced households from subsidized housing entirely in later years.[3] Despite the changes in project and subsidy type, the HOPE VI-induced moves did not produce measurable changes in school quality, neighborhood economic and demographic characteristics, or the labor market outcomes of parents. Furthermore, we find no evidence of larger impacts for children who were younger at the time of a demolition. Together, these results suggest that the long-term labor market benefits of HOPE VI demolitions are not driven by increased exposure to better neighborhoods during childhood, at least as such neighborhood effects are characterized elsewhere in the literature (Chetty et al. 2016).

Rather, the strongest evidence suggests that HOPE VI improved long-run labor market outcomes by affecting the characteristics of the neighborhoods where the children moved to and ended up living as adults. Specifically, we find that HOPE VI led to a significant improvement in measures of the geographic proximity of job opportunities–jobs per person, average commute time, and a job proximity index constructed by HUD–in the neighborhoods that the children were living in 2010, 7-13 years after the demolitions.[4] Improved job proximity can reduce job search and commuting costs and therefore reduce job search duration (Andersson et al. 2018b) and encourage individuals on the margin between working and not working to increase

[3]While it is quite common for residents of HOPE VI projects to exit subsidize housing, they are no more likely to do so relative to residents of other similarly distressed projects that were not exposed to the HOPE VI program.

[4]Our results point to a spatial concept of job accessibility which we refer to as job proximity. While it is also possible that the program could have moved children into neighborhoods that provided better access to jobs through labor market networks, we find no evidence that this is the case.

search effort and participate in the labor market (Smith and Zenou 2003). Consistent with job accessibility being an important mechanism, we find that an important part of the observed earnings gains are driven by an extensive margin labor supply response.

The data suggest that improvements in job accessibility occurred through two distinct channels. First, the demolitions transformed the neighborhoods in which the HOPE VI projects were originally located. Public housing projects, particularly those served by large PHAs, often provide housing to many individuals in geographically concentrated areas. This results in neighborhoods that provide limited access to jobs, with many job searchers competing for nearby jobs. The demolition of public housing projects drastically reduced population density with no corresponding decrease in the number of jobs in the neighborhood. Thus, on average, children that remained near the location of their original project, experienced an improvement in job proximity.

Second, HOPE VI increased the likelihood that households moved, and destination neighborhoods provided better access to jobs, even though these neighborhoods were typically geographically close (within the same county) and similarly poor. This pattern of moves is explained by two features in the data: 1) HOPE VI projects in large PHAs were located in neighborhoods that were especially disadvantaged in terms of both poverty rates and job accessibility relative to surrounding areas; and 2) housing prices increase sharply with reductions in neighborhood poverty but there is no similar price gradient with respect to job accessibility. Together, this suggests that while local moves induced by HOPE VI were likely to generate improvements in job accessibility, financial constraints may have prevented reductions in neighborhood poverty. The latter observation is consistent with existing research on the Voucher program (Patterson et al. 2004; Eriksen and Ross 2013; Jacob, Ludwig, and Miller 2013; Collinson and Ganong 2018; Andersson et al. 2018a).

In relation to previous work, an important contribution of this paper is to obtain estimates of the long-term impact of a large assisted housing program that are more representative of the full population of affected projects. Much of the relevant prior empirical research relies on data from a limited set of large metropolitan areas. Figure B.1 plots the distribution of the size of PHAs that participated in three important randomized controls trials including the HUD Moving to Opportunity (MTO) experiment (Ludwig et al. 2013), the Gautreaux program (Rosenbaum 1995), and the Effects of Housing Choice Voucher on Welfare Families project (Mills et al. 2006).[5] Approximately half of all public housing units are located in small PHAs but only two of the ten PHAs in this previous research are located in small PHAs. In contrast, over two thirds of the PHAs that received HOPE VI funding are located in small PHAs. Thus, our results are likely to be more representative of the effects for the broader population in public housing. Chicago is the third largest PHA and is shown separately in Figure B.1 as it is the setting for Chyn (2018), the closest existing paper to our work. Chyn (2018) studies the long-term earnings impacts of public housing project demolitions in Chicago and also finds substantial long-term benefits;

[5]Similarly, non-experimental research on the consequences of changes in access to vouchers or increases in voucher generosity, identifies treatment effects for a small set of non-representative cities, as in Collinson and Ganong (2016). An exception is a companion paper–Anderson et al. (2018a)– which uses a household fixed-effects identification strategy and finds long-term benefits of time spent in public and voucher housing between the ages of 13 and 18. Anderson et al. (2018a) use data from nearly the universe of assisted housing participants so that the results capture the typical effect of participating in the public housing or voucher program. In contrast, the current paper focuses on a population that is more disadvantaged relative to the subsidized housing population as a whole.

estimating that demolitions increased earnings for children in affected buildings by 16 percent relative to unaffected children who resided in the same projects.[6] Our results provide additional evidence on the long-term benefits of the demolitions of distressed public housing projects in contexts beyond Chicago and provide more insight into the mechanisms through which these demolitions affected long-term labor market outcomes.

Our results also shed light on an open puzzle in the existing literature: Does inducing households to move to new neighborhoods have to occur while children are still young in order to have long-run benefits? Chyn (2018) and the results in our paper suggest that demolitions do produce long-run benefits for older children (older than 13 at the time of the demolition). Conversely, in their analysis of the MTO experiment, Chetty et al. (2016) find no evidence of long-run gains for older children who transitioned from public to voucher housing.[7] One explanation for this discrepancy suggested by Chyn (2018) is that the projects in his study were in much more disadvantaged neighborhoods relative to those in MTO. If older children only benefit when the origin neighborhood is especially distressed, this could reconcile the findings from MTO, Chyn (2018), and this paper. We exploit the variation in pre-demolition neighborhood characteristics and find that HOPE VI had the largest

[6] Chyn (2018) measures earnings between the ages of 19 and 32 whereas we focus on labor market outcomes measured at age 26. Some of the projects studied in Chyn (2018) and Jacob (2004), who studied the short-run impacts of the same demolitions, were demolished under the HOPE VI program.

[7] The MTO study randomly assigned 4,600 households living in public housing projects to a control group, a "Section 8" group which was offered standard vouchers, or an experimental group which was offered vouchers that could only be used in census tracts with a 1990 poverty rate below 10 percent. The primary comparison made by Chetty, et al. (2016) is between this experimental group and the control group. Their results thus rely on moves to lower poverty neighborhoods, a case in which it makes sense that younger children should benefit more. Survey and administrative data have provided means of evaluating the impact of the two treatments (Ludwig et al. 2013).

impact on age 26 earnings for projects located in neighborhoods that had higher poverty rates, were more densely populated and had lower measures of job proximity. Intuitively, large distressed public housing projects create an environment in which there are many more people looking for work relative to the jobs available nearby, and this creates barriers to employment. The children located in these neighborhoods–even if they were exposed to the program only later in adolescence–still benefited from the HOPE VI intervention.

The paper proceeds as follows. Section 2.2 provides background on the HOPE VI program and related research and discusses the potential mechanisms through which public housing demolitions could affect the long-term well-being of children in displaced households. Section 2.3 describes the data sources and sample construction. Section 2.4 highlights challenges for the identification of unbiased treatment effects and discusses the stratification with regression estimator. Section 2.5 presents the empirical results, and Section 2.6 concludes.

2.2 Background and Anticipated Impacts of the Program

HUD launched the HOPE VI initiative in response to the report by the National Commission on Severely Distressed Public Housing (NCSDPH), which, in 1992, found that 86,000 of the 1.4 million public housing units nationwide qualified as "severely distressed" (NCSDPH 1992, HUD 2007). HOPE VI consisted of two main programs designed to address this issue: (1) the Demolition program, which provided funding for the demolition of public housing projects and the relocation of affected residents, and

(2) the Revitalization program, which provided funding to redevelop neighborhoods with public housing into low-density, mixed-income communities. The focus of our paper is strictly on the Demolition program and unless otherwise noted, any mention of HOPE VI refers solely to this program.[8] Between 1996 and 2003, HUD awarded $392 million through 285 HOPE VI grants for the demolition of more than 57,000 public housing units. Displaced households were typically either offered an apartment in another public housing project, a voucher, or they were forced out of subsidized housing altogether (Popkin et al. 2004).[9] Research tracking the former residents of a limited set of demolished public housing projects estimates that about half of displaced households moved to a new public housing project, a third were provided with a voucher and the remainder exited subsidized housing altogether (Kingsley et al. 2003; Popkin et al. 2009).

HOPE VI Demolition grants were awarded based on a competitive process in which HUD posted a notice of funding availability, PHAs submitted applications and HUD selected a limited set of awardees (Murphy 2012). Any PHA was eligible to submit an application for the demolition of severely distressed public housing developments (using the NCSDPH criteria). However, at least in the earliest year, HUD explicitly differentiated between PHAs of various sizes in their call for funding (2,500 units or less, between 2,501 and 10,000 units, and over 10,000 units); applicants were

[8]There is some overlap between the Revitalization and Demolition programs so that some recipients of a Demolition grant later received a Revitalization grant. However, the Revitalization intervention typically began years after the demolition occurred. As we discuss in Appendix B.2, we find no evidence that our estimated impact of the Demolition program is affected by the Revitalization program.

[9]Displaced households could also be offered a unit in a revitalized HOPE VI site, but substantial lags were involved.

evaluated within these groups and group size determined the amount of funding for which PHAs were eligible. Our analysis often differentiates between large (more than 2,500 units) and small (2,500 or fewer units) PHAs based on these cutoffs.[10] Each year, HUD classified applicants into one of four priority groups, and grants were awarded (conditional on eligibility and approval) on a first-come, first-served basis by priority group until funds were exhausted.[11] Given limited funding, both the number of applicants and eligible projects exceeded the number of awards.[12] Furthermore, many eligible projects never applied for funding while some non-distressed projects received funding, leaving many distressed-projects unaffected by HOPE VI. Indeed, Turner et al. (2007) estimate that there were between 47,000 to 82,000 severely distressed units that remained in public housing inventory as of 2007 (four years after the last demolition grant award). We return to these points later in our discussion of the empirical strategy.

It is not obvious how we should expect HOPE VI to affect the long-term labor market outcomes of displaced children. A primary goal of the program was to move families out of environments characterized by a "high incidence of crime," physical deterioration "that renders the housing dangerous to the health and safety of its res-

[10]We do not further differentiate the large PHA sample because there are too few HOPE VI projects in PHAs that exceed 10,000 units in our sample to analyze separately.

[11]Different sources give slightly different accounts of the award process. However, the Congressional Research Service Report RL32236, describes the first-come, first-served process and notes that the "priority groups are, in order of priority, (1) approved for a 202 conversion, (2) applied for a 202 conversion, (3) approved for a Section 18 demolition, or (4) approved for a HOPE VI revitalization grant. Section 202 Mandatory Conversion is the conversion of public housing developments to Section 8. If it costs less to give the residents a Section 8 voucher, rather than maintain the low rent public housing building, the building is shut down and the residents are given Section 8 vouchers."

[12]On average only 53 percent of applicants were funded each year. The percentage is based on the authors' calculation using publicly available data (HUD 2007) and the statistic excludes data from 1996, for which we do not know the number of applicants.

idents" and "limited opportunities for meaningful employment of residents."[13] Based on these stated objectives, demolitions could have shaped the development of children by improving the home and neighborhood environments they were exposed to while young. This would be consistent with recent empirical evidence suggesting that neighborhood conditions in childhood can affect the development of human capital, which in turn affect long-term labor market outcomes (Chetty et al. 2014, 2016; Chetty and Hendren 2018). Alternatively, the program could have affected adult labor market outcomes by changing access to jobs in the neighborhoods where children end up living as young adults. Theory highlighting the potential importance of job accessibility dates back to Kain (1968), arguing that the geographic location of jobs and job seekers can have important implications for labor market outcomes; recent empirical evidence in Andersson et al. (2018b) supports this hypothesis.

The program also could have had an adverse effect. Home and neighborhood environment could have worsened if the program forced people from their homes without providing proper relocation support. In addition, by dispersing residents that previously lived close to one another, the program could have disrupted social networks. Indeed, this was a major concern for residents like George Moses, a former long-time resident of public housing and the Chair of the Board of Directors of the National Low Income Housing Coalition, who spoke in objection to the HOPE VI program at a congressional hearing in 2007: "in my neighborhood, people would gather to talk, watch one another's children, and form strong bonds. When we tear these neighborhoods apart, [...] the impact is both immediate and long-lasting."

[13]Quotes are from NCSDPH (1992).

The existing empirical research on HOPE VI is largely descriptive but it suggests that the program had limited success in achieving its short-term goals. Popkin et al. (2004; 2009) find that households affected by HOPE VI experienced large changes in housing and most households moved to neighborhoods with lower poverty rates and less crime, and reported being more satisfied with their new neighborhoods, particularly if they received vouchers. However, most research finds little evidence that HOPE VI affected the short-term labor market outcomes of adults (Goetz 2010; Jones and Paulsen 2011; Popkin et al. 2009) or the health, education or behavioral outcomes of the children (Gallagher and Bajaj 2007). A limitation of this research is that it primarily documents how outcomes changed over time for households exposed to the program. This is particularly problematic in the HOPE VI setting because, even in absence of demolitions, households in public housing exhibit a high degree of residential mobility (McClure 2018). Jacob (2004) is an exception to this descriptive work, obtaining credible causal estimates of the demolition of public housing projects by comparing outcomes for children who resided in buildings that were demolished to children who resided in buildings that were not demolished but were located within the same project. Jacob (2004) finds no evidence of short-term gains in educational outcomes. In the only research on the long-term outcomes of demolitions for children, Chyn (2018) uses a similar empirical strategy and finds positive impacts on adult labor market outcomes. These results suggest that a lack of short-term impacts does not preclude the possibility of longer-term effects on labor market outcomes. However, the results from Jacob and Chyn may not be representative of the HOPE VI program as a whole since their sample is limited to public housing residents in Chicago. An important contribution of our paper is to obtain more representative estimates of the impact of the HOPE VI program by studying 160 demolitions that occurred in diverse environments across the U.S. In contrast to Jacob (2004) and Chyn (2018), we observe a great deal of variation in project and neighborhood characteristics within our empirical sample. This enables us to empirically assess how the impact of the HOPE VI program differed across projects located in heterogeneous pre-program contexts.

2.3 Description of the Data

The data requirements for this project are substantial. We need to be able to identify children and parents affected by public housing project demolitions, track exposed and non-exposed residents as they move across subsidized housing programs and neighborhoods, and match the children's housing and residential experiences to their labor market outcomes as adults. We overcome these challenges by combining two key data sources: (1) HUD-PIC (Public and Indian Housing Information Center) administrative records of participation in subsidized housing, and (2) the Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) Infrastructure Files, an administrative records system for employer-employee matched data. Below, we describe these sources and discuss how we integrate them to construct our sample.

2.3.1 Data Sources

HUD-PIC tracks public housing and voucher recipients during our study period. As part of their housing occupancy verification process, PHAs provide HUD with the identities of residents, which HUD then compiles into an annual relational database. Absent the coverage limitations we discuss below, these files record every individual participating in public or voucher housing in each year between 1997 and 2010. Our analysis makes use of the individual- and household-level files, which include indicators of housing type (public or voucher), identifiers for housing authorities and projects, as well as some individual- and household-level demographic information. HUD provides a public use summary of these data through the HUDUSER web tool, which we use to calculate PHA-level characteristics.

Data from the LEHD program are based on two sources provided by states on a quarterly basis: (1) unemployment insurance (UI) wage records, providing the earnings of each worker at each employer, and (2) employer account reports providing establishment-level data, also known as the Quarterly Census of Employment and Wages (and formerly as the ES-202 program).[14] The state-provided data cover more than 95 percent of wage and salary civilian jobs, including both private sector and state and local government workers. Some omissions remain, including the armed forces, earnings through self-employment, the postal service, family workers, federal

[14]The LEHD program, a partnership that has been established between the Census Bureau, all 50 states, the District of Columbia, and the U.S. Office of Personnel Management, produces public use data tabulations that are widely used by state and local governments such as: Quarterly Workforce Indicators, LEHD Origin-Destination Employment Statistics (LODES), and Job-to-Job Flows. For a description of the LEHD Infrastructure Files, see Abowd et al. (2004). For a description of files available in the Federal Statistical Research Data Centers, see Vilhuber (2018).

workers, and some non-profit and agricultural workers (U.S. Bureau of Labor Statistics 1997, 2017). Nevertheless, the LEHD earnings data enable us to track a large set of children into adulthood and measure their earnings and employment outcomes as well as these outcomes for the parents of these children.[15] The coverage extends from the beginning of state reporting through the last quarter of 2016.[16]

Another strength of our data is our ability to track the residential location of households who leave subsidized housing. We do this using two sources of data. First, we use a measure of annual residential location from the Composite Person Record (CPR), a Census Bureau file created from several federal administrative datasets, which begins in 1999.[17] We identify a residence census tract for each child and adult from 1999-2010 where available (approximately 10 percent of children are missing a CPR residence in each year). Second, we use responses from the 2010 Decennial Census to identify where individuals lived in April 2010. These responses provide an additional data source covering geographic residence of each individual, and also allow us to determine whether that individual is incarcerated in 2010.

We also draw on a number of different publicly available data sources. Most importantly, we characterize the neighborhoods in which individuals live and projects are

[15]Specifically, we measure outcomes for the head of households as identified in the HUD-PIC data. Most children (92 percent) grow up in single-parent households in the HOPE VI sample.

[16]We code earnings as missing if the state in which their project is located was not yet reporting in the LEHD. However, the vast majority of states are reporting to the LEHD by 2005, which is the earliest year in which we measure age 26 earnings for the children. For the small fraction of children who have missing age 26 earnings, we impute these values using earnings from later years. Specifically, we use a panel of non-missing earnings data for all children between ages 18 and 30 to estimate a regression of annual earnings on an individual fixed effect and a third order polynomial in age interacted with gender. We use the estimates to impute missing earnings data at ages 18-26.

[17]The LEHD uses the CPR for imputation models and for the residence component of the LEHD Origin-Destination Employment Statistics (LODES) data (for more information on the sources contributing to the CPR, see Graham et al. 2017).

located using a number of different files including: census tract-level characteristics drawn from the 1990, 2000, and 2010 Decennial Censuses and five-year-average data from the American Community Survey collected between 2008 and 2012;[18] county-level unemployment rates from the U.S. Bureau of Labor Statistics' Local Area Unemployment Statistics program; area median income and characteristics of PHAs in 1997 from HUD USER; the number of jobs per census tract in 2010, by workplace and residence, from LODES; school proficiency and jobs proximity indexes constructed using data from 2013-2014 and provided through HUD Open Data (the job proximity index is based on LODES); land areas as well as crosswalks between various measures of geographies from the U.S. Census Bureau's Geography Relationship Files; and the Census Bureau Gazeteer files to measure the latitude and longitude of the centroid of census tracts.[19] We use the Consumer Price Index-Urban to convert all dollar amounts into 2000 dollars.

2.3.2 Integration and Sample Selection

Our sample construction begins by using the HUD-PIC records to identify children between the ages of 10 and 18 who lived in public housing between 1997 and 2001. The range of years is selected because 1997 is the earliest year when reliable HUD microdata are available and 2001 is the date of the last HOPE VI demolition.[20] The

[18]Obtained from the National Historical Geographic Information System (NHGIS) from IPUMS; see Ruggles et al. (2019).

[19]In a small number of cases, neighborhood-level data are missing for certain variables. In order to avoid changes in the sample composition based on the variables used in the analysis, we impute using higher levels of geography. For example, if a variable is missing for a given census tract, we impute the value with county-level value.

[20]As discussed later in this section, we set the "demolition date" two years prior to the award date.

age range is chosen to allow us to observe earnings up through age 26 for all children in the sample.[21] We choose to focus on age 26 earnings in our main results since most children will have completed their education by this date and work by Chetty et al. (2014) finds that outcomes measured at this age are strongly predictive of later-life measures of labor market success. We then attach data from the LEHD, CPR and 2010 Decennial Census to each record from the HUD-PIC data.[22] An analogous dataset is constructed using the household heads of the children in the sample.

We construct a dataset of public housing projects that describes characteristics of the residents and the neighborhoods in which they are located. To identify the set of projects that received a HOPE VI demolition grant, we start from publicly available data that lists all 285 HOPE VI demolition grant awards.[23] We make several sample restrictions to the full list of projects to exclude those that are not well-suited for our study design (such as excluding senior housing). These sample restrictions, described in Table B.1, reduce the analysis sample to about 160 projects that received HOPE VI demolitions awards.[24] Implementing a similar set of restrictions produces a sample of about 8,800 non-HOPE VI projects.[25]

[21]There is one cohort of children, 10-year-olds who appear in public housing in 2001, for whom we do not observe age 26 earnings because our earning data are only available through 2016. For this cohort, we use observed earnings up through age 25 to impute their earnings at age 26. Specifically, we use a panel of non-missing earnings data for all children between ages 18 and 30 to estimate a regression of annual earnings on an individual fixed effect and a third order polynomial in age interacted with gender. We use the estimates to impute missing earnings data at age 26.

[22]Individuals are identified by a "Protected Identification Key" (PIK) generated by the Census personally identified information, allowing us to attach LEHD data to other data sources. PIKs are linked to approximately 98 percent of person records in the HUD-PIC member file for our study period and we drop the 2 percent of individuals that are not assigned a unique PIK.

[23]For the HOPE VI demolition grant list, see: HOPE VI DEMOLITION GRANTS: FY 1996 - 2003 (available at https://www.hud.gov/sites/documents/DOC_9890.PDF, dated October 2004).

[24]Throughout the paper we often report rounded numbers to limit risk of disclosure.

[25]Specifically, based on the restrictions defined in Table B.1, we apply the following sample restrictions to the non-HOPE VI projects: 1, 5, 6, and 7. Data from HUDUSER indicate that in 1997 there were about 13,400 projects in the U.S. (excluding territories) and about 10,100 projects that were

Our primary analysis dataset combines the project- and individual-level data to create a file in which the unit of observation is at the individual-year level, where an individual will appear in the sample for every year that they appear in public housing. We define the "reference year" as the year in which the individual appears in public housing. We drop individual-year observations that appear in the HOPE VI projects in years in which the demolition did not occur whereas we retain all observations from non-HOPE VI projects.[26] Thus, the reference year for the HOPE VI sample is simply the year of the demolition. This produces a sample with 1,682,000 child-year observations and 1,023,000 household head-year observations.

To identify treated individuals, we need to determine who was living in the project at the time of the demolition. However, identifying the timing of the demolition is complicated by the fact that the PHA may have started to move households out of the project prior to the physical demolition of the building. To address this possibility, we classify households as treated if they resided in a HOPE VI project two years prior to the award date.[27] To simplify language, we refer to the two years prior to the award date as the year in which the demolition occurred. We view this definition of timing as conservative as it minimizes the chances that our estimated treatment

within our size range (between 15 and 3,000 occupied units) that were not senior citizen housing. Thus, even though we lack data on some PHAs that participated in the Moving To Work (MTW) demonstration, our sample appears to cover most of the comparable public housing projects.

[26]We drop projects that received a HOPE VI Revitalization grant but did not receive a HOPE VI Demolition grant, as households in these projects were treated by a different, but closely related program. For the non-HOPE VI sample, we drop individual-year observations who previously appeared in a HOPE VI project. This restriction prevents us from using individuals who moved out of HOPE VI projects and into other public housing projects as control observations.

[27]Because the HUD-PIC data start in 1997, any HOPE VI projects that have an award date prior to 1999 are assigned a demolition year of 1997. The decision to retain the early awardees is in part motivated by reports that there were longer delays between grant awards and demolitions for these projects (GAO 2003). In Appendix B.2 we show that our results are robust to how we treat projects that received HOPE VI grants prior to 1999.

effects are contaminated by selection out of the project prior to the demolition while potentially underestimating the effect if the demolition does not occur until a later time. To evaluate this definition, Figure 2.1 presents changes in project size relative to this demolition date.[28] The figure shows similar trends in project size for HOPE VI and non-HOPE VI projects prior to demolition, with HOPE VI projects declining in size for several years thereafter.

Another related issue apparent from Figure 2.1 is that some of the projects were only partially demolished. While a substantial portion of the households in HOPE VI projects were forced out within five years of the demolition, our sample does include some households who resided in undemolished units and remained in their original housing units. We include these households in the sample as our view is that they are still "treated" by the program since the demolition could have affected the people or characteristics of the neighborhood in which the HOPE VI project was located.[29]

2.4 Empirical Strategy

Our primary goal is to estimate the average effect of HOPE VI demolitions on young adult (age 26) labor market outcomes for children affected by the program–the average treatment effect on the treated. The challenge is that, by design, the projects demolished under HOPE VI were systematically different from those that were not. This is readily apparent from Table 2.1, which presents the mean and standard deviations

[28] We measure changes in project size using the number of occupied units in the HUD-PIC household file.

[29] Indeed, in Section 2.5 we find that the neighborhood in which the project was located is affected in important ways by the demolitions. Furthermore, we find no evidence that the impacts on adult earnings are different for complete versus partial demolitions (see Appendix B.2).

Figure 2.1: Changes in Project Size Relative to Year of Demolition

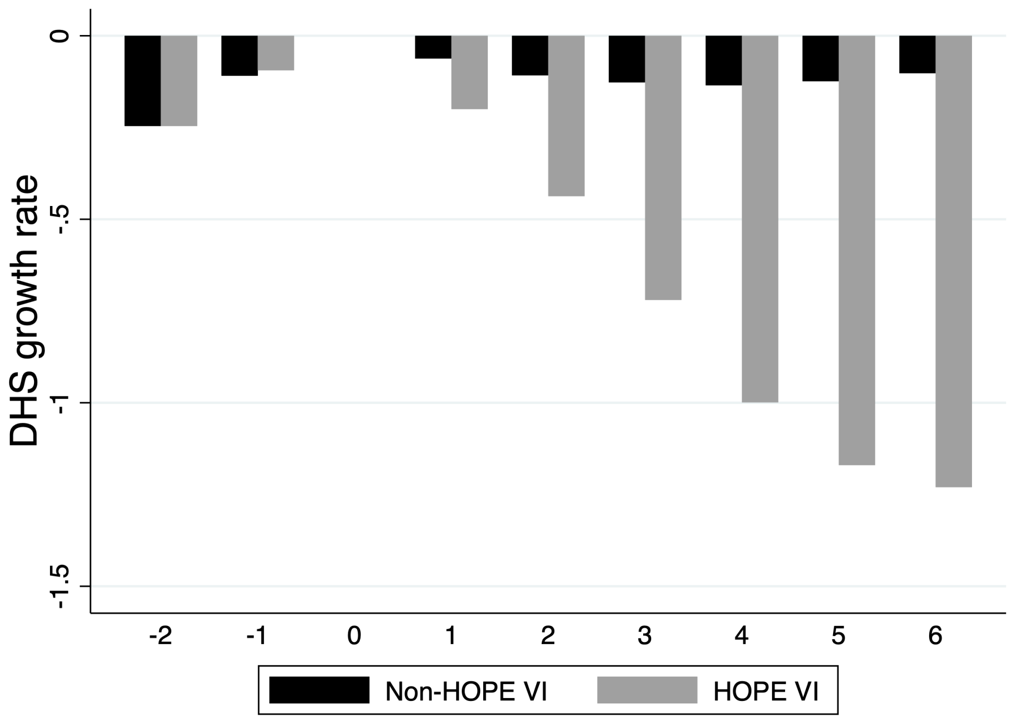

Notes: The figure plots the average DHS growth rate (see Davis, Haltiwanger and Schuh 1996) in project size between the reference year and x years after the reference year, where x corresponds to the value on the horizontal axis. The growth rate in project size between year t (y_t) and year s (y_s) is defined as: $\frac{y_t - y_s}{\frac{1}{2}(y_t + y_s)}$. For HOPE VI projects, the reference year is the year of the demolition, which is defined as the greater of two years prior to the award year and 1997. Averages are calculated using the child-year dataset, implying that the averages are weighted by project size.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

of baseline characteristics for projects and residents of HOPE VI and non-HOPE VI projects as well as the differences between two samples. Along almost every observable dimension, children growing up in HOPE VI projects are more disadvantaged. For example, HOPE VI projects are in census tracts with 52 percent higher poverty rates, the residents have 20 percent lower total annual household income and are almost 50 percent less likely to have a married head of household.

Given these pronounced observable differences and the lack of experimental variation, our empirical strategy aims to estimate causal impacts by accounting for observable baseline differences between HOPE VI and non-HOPE VI projects. We argue that this is a reasonable approach in our context because the number of distressed, eligible projects greatly exceeded the number of HOPE VI awardees and our data infrastructure enables us to observe and characterize the conditions in nearly all public housing projects in the U.S. Thus, there is a large sample of non-HOPE VI projects that are informative of what would have happened to the residents of HOPE VI projects had there been no demolitions. In order to estimate the causal impacts of the Demolition program, we employ the stratification with regression estimator proposed by Rosenbaum and Rubin (1983, 1984) and discussed at length in Imbens and Rubin (2015) and Imbens (2015). The method combines features of both matching and regression in the following steps: (1) nearest-neighbor matching to trim the sample, (2) groups similar observations into distinct strata based on an estimated propensity score, (3) estimates strata-level treatment effects using Ordinary Least Squares (OLS) regressions with controls within strata, and (4) calculates aggregate treatment effects as a weighted average of the stratum-level estimates.

Table 2.1: Summary Statistics of Baseline Characteristics

	HOPE VI	All Non-HOPE VI			Control		
	mean	mean	t-stat	Δ	mean	t-stat	Δ
Panel A. Neighborhood							
median household income/1,000	22.0	27.9	-6.68	-0.520	22.9	-0.936	-0.085
	[11.1]	[11.5]			[10.3]		
poverty rate	0.374	0.247	7.84	0.724	0.346	1.60	0.146
	[.206]	[.14]			[.181]		
log(population density)	0.033	-1.24	11.2	0.703	-0.153	1.43	0.122
	[1.42]	[2.13]			[1.63]		
Panel B. Household Head							
household income/1,000	9.00	11.3	-6.93	-0.515	9.65	-1.77	-0.174
	[6.63]	[8.63]			[7.31]		
age	38.6	39.3	-2.36	-0.167	38.6	-0.096	-0.054
	[10.1]	[10.1]			[10]		
female	0.904	0.869	2.46	0.409	0.899	0.338	0.044
	[.294]	[.338]			[.301]		
married	0.078	0.133	-4.52	-0.436	0.081	-0.241	-0.014
	[.268]	[.34]			[.273]		
has disability	0.113	0.121	-1.64	-0.063	0.111	0.284	-0.003
	[.316]	[.326]			[.314]		
number of dependents	2.76	2.54	4.24	0.454	2.63	2.27	0.273
	[1.56]	[1.4]			[1.46]		
white non-Hispanic	0.064	0.207	-12.0	-0.709	0.079	-1.16	-0.034
	[.244]	[.405]			[.27]		
black non-Hispanic	0.684	0.522	3.67	0.660	0.692	-0.151	0.041
	[.465]	[.5]			[.462]		
Hispanic	0.161	0.184	-0.575	-0.083	0.152	0.206	-0.056
	[.368]	[.387]			[.359]		
Panel C. Children							
age	13.6	13.6	-2.33	-0.288	13.5	0.300	-0.232
	[2.58]	[2.57]			[2.58]		
female	0.509	0.507	0.431	-0.137	0.512	-0.631	-0.124
	[.5]	[.5]			[.5]		
has disability	0.020	0.027	-2.79	-0.077	0.020	-0.011	-0.020
	[.14]	[.161]			[.14]		

Notes: This table presents summary statistics for the baseline variables listed in the rows. The variables in Panel A, B and C are characteristics of: (A) the census tract in which the projects were located measured in 1990, (B) the households or head of households and (C) the children. Column 1 presents the mean for the HOPE VI sample. Columns 2-4 (5-7) present statistics calculated from a sample that include all non-HOPE VI (control) projects. Columns 2 and 5 present the mean of the non-HOPE VI projects. Columns 3 and 6 present the t-statistic from a regression of the baseline variable in the row on an indicator for HOPE VI. Standard errors are clustered at the project level. Columns 4 and 7 present the normalized difference of the row variable between the HOPE VI and non-HOPE VI observation. Normalized differences are calcualted from data collapsed to the project level and are defined as $\Delta = (\bar{x}_1 - \bar{x}_0)/(\sqrt{(s_1^2 + s_0^2)/2})$, where $\bar{x}_d$ and s_d is the sample average and variance for HOPE VI (d=1) and non-HOPE VI (d=0) observations, respectively. The standard deviation for each sample is presented in brackets below the mean.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

There are three principal advantages of the stratification with regression estimator over the more traditional OLS estimator. First, trimming the sample and using the stratification structure helps us relax the linear functional form assumptions implicit in OLS. As a rule of thumb, linear regression techniques will tend to be sensitive to the specification when the value of normalized differences between the treatment and control groups exceed one-quarter (Imbens and Woolridge 2009).[30] Table 2.1 demonstrates that many important baseline variables have normalized differences that exceed this threshold.[31] Second, many choices on how to adjust for observable differences between HOPE VI and non-HOPE VI projects are governed by the data, which helps mitigate concerns that the choice of specification is influenced by ex-post analysis of results. Third, the stratification with regression methodology presents a number of ways in which we can evaluate the plausibility of the identifying assumptions, some of which are specific to the method and have no clear analogue under OLS. These are discussed in Section 2.5.3.

Construction of the strata is implemented in three steps. First, we trim the sample of non-HOPE VI projects to reduce the observable differences between the HOPE VI and non-HOPE VI samples. To do so, we start with a project-level dataset that includes all projects after imposing the restrictions mentioned in Section 2.3.[32] We use a project-level, as opposed to an individual-level dataset because the treatment

[30] Let $\bar{x}_d$ and s_d be the mean and standard deviation of the variable x for the HOPE VI (d=1) and non-HOPE VI (d=0) samples, respectively. Then the normalized difference is defined as $(\bar{x}_1 - \bar{x}_0)/\sqrt{(s_1^2 + s_0^2)/2}$.

[31] The solid line in Figure B.2 makes a similar point by presenting the distribution of the normalized differences for all baseline variables calculated on the full sample.

[32] Project-level characteristics are measured in the year of the demolition for HOPE VI projects, whereas for non-HOPE VI projects they are equal to the average of observed values between 1997 and 2001.

is assigned at the project-level. For each HOPE VI project, we use nearest neighborhood matching to identify and retain the five nearest neighbors among the non-HOPE VI projects. Matching is conducted with replacement; distance is measured using the Euclidean distance metric based on observable project and neighborhood characteristics (see Appendix B.4 for list of variables used in matching); and we require exact matching on the size (large or small) of the PHA. The resulting dataset, which we refer to as the matched sample, contains all 160 HOPE VI projects and a subsample of 570 matched non-HOPE VI projects, which we refer to as control projects.

We thus drop non-HOPE VI projects that are fundamentally different and unlikely to be informative of counterfactual outcomes for HOPE VI residents. The dashed line in Figure B.2 illustrates the success of this trimming by presenting the distribution of the normalized differences of all baseline variables in the matched sample. The differences are much smaller relative to those calculated in the full sample with nearly all smaller than one-quarter. The final three columns of Table 2.1 make a similar point by presenting summary statistics and difference measures for HOPE VI and matched controls for a subset of important baseline variables. This step does not reduce the external validity of the estimates since we retain all HOPE VI projects and our goal is to estimate the average treatment effect on the treated.

In the second step, we estimate a project-level propensity score defined as the probability that a project receives a HOPE VI Demolition grant, conditional on observable characteristics. To determine the covariates included in the propensity score model, we use a data-driven method described by Imbens and Rubin (2015). Specifically, we start by estimating a logistic regression of receipt of HOPE VI on a set of covariates that we think are important for predicting treatment (average household income and the proportion of household heads who are black non-Hispanic). Next, we estimate a separate logistic regression for each baseline variable that we consider adding to the model and calculate the log likelihood for each logistic regression. If the value of the log likelihood ratio test statistic for a given set of covariates is larger than it is for the models with the other potential covariates and sufficiently greater than the initial log likelihood, then we include the covariate in the model.[33] We iteratively apply this procedure until no more covariates are selected. We then create interaction terms between all the selected covariates and repeat this process to determine which second-order terms to include in the model. Figure 2.2 plots the distribution of the linearized estimated propensity score for HOPE VI and control projects.[34] The figures indicate that there is good overlap between the estimated propensity scores of the treated and control projects.

In the third step we use a data-driven method to group projects into distinct strata based on the estimated propensity score. We start by separating the projects into two strata based only on PHA size (small and large). This distinction is motivated by the fact that HUD differentiated between these PHAs in the application process. However, it also has the added benefit of avoiding comparisons between individuals who reside in fundamentally different economic environments (e.g., a comparison of someone living in a rural county to an individual living in major metropolitan area). We then expand the number of strata for each initial large- and small-stratum.

[33]We include additional first-order (second-order) terms only if the likelihood ratio statistic for the test of the null hypothesis that the additional covariate is equal to zero exceeds 2.5 (4.21).

[34]As a confidentiality protection measure, we Winsorize each distribution at the 5th and 95th percentiles, which overstates the lack of overlap at the tails of the distribution.

Figure 2.2: Distribution of Estimated Propensity Score in Matched Sample

(A) Large PHAs (B) Small PHAs

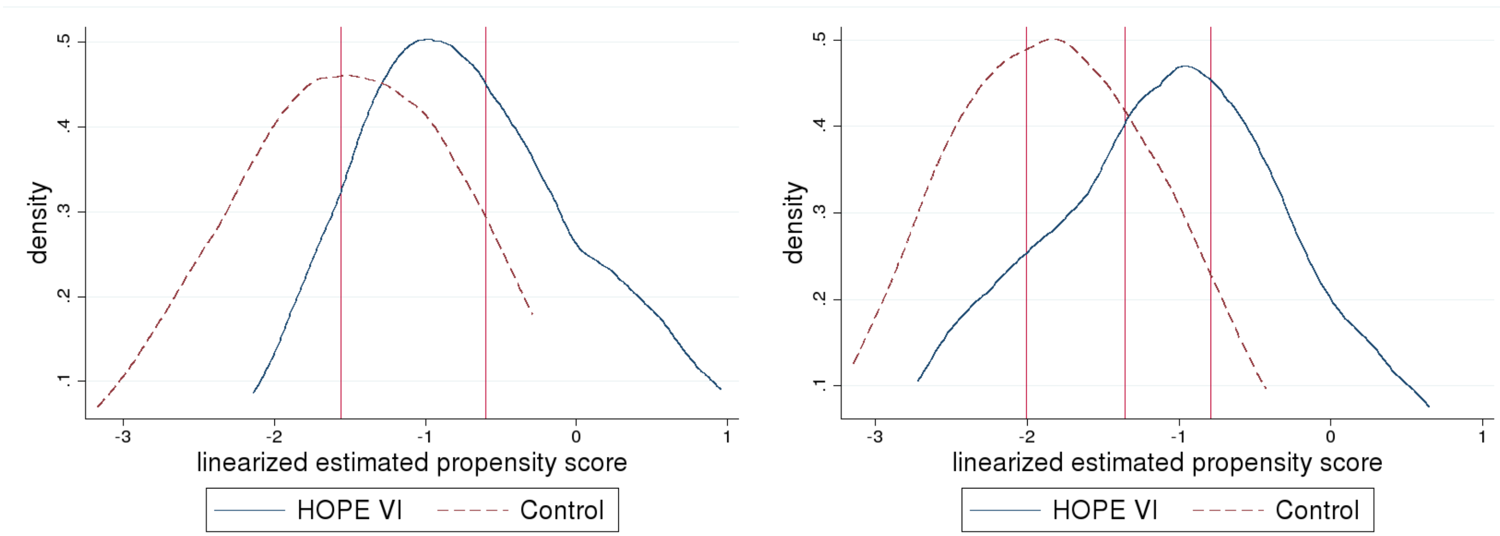

Notes: The figures present the kernel densities of the linearized estimated propensity score for large (panel A) and small (panel B) Public Housing Authorities (PHAs) for the matched sample. The estimated propensity score is the predicted value from a logistic regression of an indicator for HOPE VI on a vector of observable characteristics. The propensity score is estimated on project-level data and the covariates in the model include: average total household income of project residents; natural log of the number of occupied units in the project, the proportion of households in the PHA with a majority of income from wages or business income, average gross monthly rent in the PHA, and the first and second order terms (interactions terms too) of the proportion of household heads in the project that are black non-Hispanic and married. The vertical lines indicate the boundaries of the strata. To pass disclosure review requirements, each of the four distributions are Winsorized at the 5th and 95th percentiles. Let $\hat{\rho}$ denote the estimated propensity score, then the linearzied estimated propensity score is, $ln(\frac{\hat{\rho}}{1-\hat{\rho}})$.

Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

The adequacy of the existing strata is assessed by calculating a t-statistic for each stratum where the null hypothesis is that the average value of the estimated linearized propensity score is the same for the treated and control projects in that stratum. If the null hypothesis is rejected (i.e. the absolute value of the t-statistic exceeds 1.645), then the stratum is split into two new strata by grouping projects above and below the median linearized propensity score.[35] The newly generated strata are required to have at least 3 HOPE VI and control projects and 50 total projects in order to prevent issues related to small sample sizes in the analysis.[36] The process is then repeated until either the null hypothesis of no difference between treatment and control projects in the linearized propensity score is not rejected for any stratum, or splitting the stratum at the median treatment project's linearized propensity score would result in too few projects in one of the newly generated strata. This process divides the sample into seven distinct strata. On average, each stratum contains about 18,000 unique children from 100 different projects, 15 percent of whom reside in HOPE VI projects. The boundary points of the strata are depicted by the vertical lines in Figure 2.2 and Table B.2 presents the sample size within each stratum.

This procedure does an excellent job of eliminating differences in observable characteristics between control and treatment groups within each stratum. To demonstrate this point, we regress 92 different baseline variables on an indicator for HOPE VI within each of the seven strata and calculate a t-statistic to summarize the differences between control and treatment observations (standard errors are clustered

[35] Let p denote the propensity score, then the linearized propensity score is defined as ln(p/(1-p)).

[36] We require 50 total projects in each stratum because we cluster standard errors at the project level.

at the project level). We plot the distribution of the absolute value of the resulting 644 t-statistics in Figure B.3 and compare it to the distribution one would expect from the absolute value of t-statistics from a standard normal distribution. The figure illustrates that, if anything, there is more balance within stratum than would be expected from random assignment. Table B.3 provides a more detailed view by presenting the proportion of test statistics that have a p-value of less than 0.10 for neighborhood-, project- and individual-level characteristics. If balance were good, we would expect that the share of significant test statistics would be approximately 10 percent. For the most part, we find that this pattern applies. For example, column 6 of Table B.3 suggests that only 12 percent of the 276 p-values calculated within the large PHA sample had a p-value of less than 0.10 and only 6.2 percent of 368 p-values calculated within the small PHA sample had a p-value of less than 0.10.[37]

An advantage of this methodology is that many of the choices about how to adjust for observable differences between HOPE VI and non-HOPE VI projects are determined by the data. However, the method does depend on six tuning parameters, which must be defined by the researcher.[38] We chose a set of tuning parameters

[37]Without making any adjustments for multiple hypothesis testing, we should observe slightly more than 10 percent of tests rejected at the 10 percent level.

[38]These parameters are: (1) the number of matches to use when trimming the sample, (2) the threshold for the likelihood ratio test to include first-order terms for the estimation of the propensity score, (3) the threshold for the likelihood ratio test to include second-order terms for the estimation of the propensity score, (4) a threshold for the test statistic used to determine whether the estimated propensity scores of control and treated projects are sufficiently similar within strata, (5) the minimum number of control projects that must be included in each stratum and (6) the minimum number of treated projects that must be included in each stratum. We view the first three tuning parameters as both the most consequential, since they determine which projects serve as controls for each HOPE VI project, and the most likely to require values specific to applications that differ in number of observations and heterogeneity within the sample. Thus, we use standard values for the fourth through sixth tuning parameters but select "optimal" values for the first through third parameters.

that robustly eliminates baseline differences between HOPE VI and control projects within strata.[39] There are two important considerations to note here. First, the criteria used for selecting tuning parameters are only based on how well the method eliminates observable differences between HOPE VI and control projects and do not use the outcome variables. Thus, we avoid concerns of specification search. Second, in practice our main findings are robust to alternative choices of tuning parameters (see Appendix B.2).

Using the stratification structure, we implement our estimator in two steps. First, we separately estimate the following OLS specifications within each of the strata:

$$y_{itps} = \alpha_s + D_p\delta_s + X_{itps}\beta_s + \epsilon_{itps} \tag{2.1}$$

where y is a labor market, neighborhood, or household outcome; i is the individual; t is the year in which that individual appears in public housing; p is the project; s is the stratum the project was assigned to in the first stage; D is an indicator equal to one if the project received a HOPE VI demolition award; X is vector of observable individual-, household-, project-, and neighborhood-level characteristics; and ϵ is an error term which we cluster at the project level.[40] Because the specifications are run

[39] To do this, we implement the stratification 33 different times using different values of the number of matches (3, 5 or 7) and different values of the second and third tuning parameter. (As a rule of thumb, Imbens and Rubin (2015) find 1.00 and 2.71 work well for the values of the second and third tuning parameters. We vary the value of the second tuning parameter from 1.0 to 6.0 and set the value of the third tuning parameter to 1.71 higher than the second.) We then create a score for each iteration based on the resulting balance of all baseline covariates across HOPE VI and control observations. We find balance is achieved most robustly when using five matches. Thus, we opt to use the specification that delivers the best balance of baseline covariates (lowest-ranked score) when using five matches.

[40] There are a small number of cases in which the outcome variable is missing. To avoid disclosure issues related to releasing results from multiple sample, we impute these missing values with the mean value in the control group and then include an interaction between an indicator for this

within each stratum, all of the estimated coefficients are stratum-specific.

All specifications include controls for the year in which the individual appears in public housing (with the HOPE VI individuals only appearing in one year), and a standard set of project-level controls that include characteristics of the project (average total income of resident households, proportion black non-Hispanic, and proportion of household heads that are female); area median income in 1990; characteristics of census tract in 1990 (proportion on public assistance, median income, and poverty rate); and the county-level unemployment rate in 1996.[41] The standard vector of individual-level covariates included in the specifications estimated on the child-level dataset includes the interaction between sex and mutually exclusive race/ethnicity categories (black non-Hispanic, white non-Hispanic, Hispanic, and other race or race not specified non-Hispanic); the number of dependents in the household; household size; an indicator for disability; a fixed effect for age at the time of appearing in public housing; an indicator for whether the head of household has a disability; an indicator for whether the household head is female; the marital status of head of household; the age of the head of household, and total household income.[42] While individuals from HOPE VI projects only appear once in the sample, individuals from control projects may appear multiple times in the sample with an observation for each year they ap-

imputation and treatment status in the regression. In this way imputed values do not contribute to the identification of the treatment effect. In unreported results we estimate all specifications with missing data without this imputation and confirm that the results are not materially different.

[41] The large number of individuals within each stratum allows us to include a large set of individual-level controls in our stratum-level regressions. Since the number of projects per stratum is more limited, we are careful to include a smaller number of project-level controls in the regression analysis.

[42] The standard vector of individual-level covariates included in the specifications estimated on the household head-level dataset includes age, race, sex, number of dependents, household size, disability status, marital status, and total household income.

pear in public housing between 1997 and 2001. Nearly all of these individuals appear in the same project and thus clustering standard errors at the project level allows us to take these "duplicate" observation into account when calculating standard errors with each stratum.[43]

The stratum-specific treatment effects are then aggregated across strata, using the stratum's share of the total of treated individuals as weights. Let N_t^s be the number of treated individuals in stratum s and N_t be the total of treated individuals across all strata including both the large and small PHA groups. The weight for each stratum is given by $w_s = N_t^s / N_t$, and the estimate of the average treatment effect on the treated, $\hat{\delta}^{att}$, and the corresponding standard error, $se(\hat{\delta}^{att})$, are given as:

$$\hat{\delta}^{att} = \sum_{s=1}^{S} \hat{\delta}_s w_s \tag{2.2}$$

$$se(\hat{\delta}^{att}) = \sqrt{\sum_{s=1}^{S} [se(\hat{\delta}_s) w_s]^2} \tag{2.3}$$

where the weighted averages are taken across all S strata (S=7 for the main specification).[44]

Our methodology will produce unbiased estimates of the average treatment effect on the treated under the Conditional Independence Assumption; conditional on the covariates and stratification in the model, assignment of a HOPE VI demolition is

[43] Appendix B.2 shows that our main results are robust to dropping all observations that appear in more than one project and shows that the standard errors are not significantly affected by the presence of these individuals.

[44] The implicit assumption needed to construct the standard errors is that observations across strata are independent. We argue that this is reasonable based on the fact that no project appears in more than one stratum and standard errors are clustered at the project level.

as good as random. While this assumption is not empirically testable, we conduct a number of analyses to assess its plausibility. Our method successfully eliminates observable differences between HOPE VI and control projects, which provides some initial support for the Conditional Independence Assumption. After presenting the main results we discuss other checks intended to assess the validity of the empirical approach.

2.5 Results

2.5.1 Long-Run Effects on Children

Table 2.2 presents the main finding of the paper. On average, exposure to a HOPE VI demolition led to substantial improvements in the long-run labor market outcomes of the children who resided in those projects. Panel A presents the results pooling across large and small PHAs. Columns 1-4 correspond to the estimates for four different labor market outcomes measured in the year that the child turns 26: the number of quarters worked, an indicator equal to one if earnings are strictly positive in all four quarters, total earnings divided by 1,000, and the inverse hyperbolic sine (IHS) of annual earnings.[45] All coefficients are estimated using the stratification with regression methodology, and the standard set of covariates used in the stratum-level regressions. We find that, on average, the HOPE VI program increased age 26 earnings by 14.2 percent, annual earnings by $622, the number of quarters worked

[45]We use the IHS of earnings rather than the more traditional log of earnings because estimated coefficients can be interpreted in the same way as with a log-transformed dependent variable but, unlike with the log of earnings, IHS is defined for zero earnings. The IHS is defined as $log[y_i + (1 + y_i^2)^0.5)]$ where y_i is total earnings for individual i (see Burbidge et al. 1988).

by 0.057, and the probability that an individual worked all four quarters by 1.6 percentage points. HOPE VI clearly had important positive impacts on adult labor market outcomes.

While the overall impact of the program was positive, there is heterogeneity across different housing environments. Panels B and C of Table 2.2 present results separately for large and small PHAs. The positive impacts are generally stronger in large PHAs, with differences that are often economically important in size. For example, the IHS earnings specification suggests a 19.5 percent increase in age 26 earnings for children in large PHAs and only a 4.5 percent increase for those in small PHAs.[46] We provide additional evidence below that there is meaningful heterogeneity by PHA size in the effect of the program.

We explore heterogenous effects by child age at the time of the demolition, race, and sex by estimating a model in which the indicator for HOPE VI is interacted with these characteristics. Table 2.3 presents the resulting estimates for large PHAs.[47] Column 1 indicates that the impacts of the program are no different for older and younger children.[48] Specifically, children exposed to HOPE VI when they were 10 years old experienced an earnings gain of 20.5 percent while this gain is 18.9 percent for 18-year-olds; a difference that is neither economically or statistically significant.

[46] The long-run benefits found in large PHAs are robust to measuring earnings at alternative times. Figure B.4 in Appendix B.1 presents estimates of the effect of HOPE VI on the IHS of earnings measured between ages 18 and 26. The effect of the program grows over time, starting around zero at age 18 and rising to about 0.2 by age 23, after which point the effects stabilize through age 26.

[47] Not surprisingly, we also find little evidence of heterogeneous effects in small PHAs. The one exception is that there is some evidence that white children may have benefited more than non-white children in small PHAs. See Appendix B.2 for details.

[48] In unreported results we also find that the lack of heterogeneous effects by age is robust to estimating alternative specifications that employ project or household fixed effects.

Table 2.2: Earnings Outcomes

	qrtrs worked (1)	worked 4 qrtrs (2)	earnings / 1,000 (3)	IHS earnings (4)
Panel A. All PHAs				
HOPE VI	0.057***	0.016***	0.622**	0.142**
	(0.021)	(0.006)	(0.282)	(0.056)
control mean	2.16	0.404	8.33	6.3
	[1.73]	[0.482]	[34]	[4.53]
observations	258,000	258,000	258,000	258,000
Panel B. Large PHAs				
HOPE VI	0.076***	0.019***	0.529*	0.195***
	(0.027)	(0.007)	(0.287)	(0.073)
control mean	2.14	0.4	8.44	6.24
	[1.73]	[0.481]	[40.5]	[4.55]
observations	149,000	149,000	149,000	149,000
Panel C. Small PHAs				
HOPE VI	0.022	0.009	0.794	0.045
	(0.035)	(0.009)	(0.601)	(0.087)
control mean	2.2	0.41	8.12	6.4
	[1.72]	[0.483]	[16.4]	[4.48]
observations	109,000	109,000	109,000	109,000

Notes: Panels A, B, and C present estimates from the stratification with regression estimator for all, large, and small Public Housing Authorities (PHAs), respectively. All outcome variables are annual labor market outcomes measured in the year in which the child turns 26. In columns 1-4 the outcome variables are: the number of quarters worked, an indicator equal to one if the child had positive earnings for all four quarters, earnings/1,000 winsorized at the 99^{th} percentile, and the inverse hyperbolic sine (IHS) of earnings. All stratum-level regressions control for the base year in which the child appears in public housing, the year in which the child turns 26, and the standard vector of individual- and project-level covariates. Standard errors are clustered at the project level and are presented in parentheses. The mean and standard deviation, presented in brackets, of the outcome for the control group are a weighted aggregate of stratum-level statistics, where the weights are proportional to the number of treated individuals in a strata.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** p≤0.01, ** p≤0.05, * p≤0.10.

This offers some initial evidence that the impacts of the program are not driven by differences in human capital accumulation from exposure to neighborhoods of varying quality, at least through the exposure model typically considered in this literature (as in Chetty et al. 2016; Chetty and Hendren 2018). Column 2 indicates that males experience significantly larger earnings benefits while column 3 suggests that non-white children also benefit more. While we do not have enough power to estimate a model with the full set of interactions between race and sex, column 4 presents estimates from a specification in which we compare the effects for non-white males to all other children. We find that non-white males appear to be the primary beneficiaries of the program.

2.5.2 Short- and Medium-Term Effects for Head of Households

To better understand the mechanisms through which HOPE VI demolitions affected long-term labor market outcomes, we explore the short- and medium-term effects of the program, starting with housing outcomes for households one, three, and five years after the demolition. In Table 2.4, column 1 shows that HOPE VI led to a 15 and 18 percentage point reduction in the probability that the household head lives in the same housing project five years after the demolition in large and small PHAs, respectively (relative to 33 percent and 28 percent of control households remaining in their original project). Column 2 and 3 indicate that HOPE VI pushed households into both voucher and other public housing with a slightly larger shift into voucher housing. Five years after the demolition, HOPE VI households in large housing authorities

Table 2.3: Heterogeneous Effects by Demographics, for Large PHAs

	IHS of Earnings at Age 26			
	(1)	(2)	(3)	(4)
HOPE VI	0.189	0.071	-0.194	0.059
	(0.119)	(0.085)	(0.236)	(0.085)
HOPE VI×(18-age at demolition)	0.002			
	(0.020)			
HOPE VI×male		0.254**		
		(0.122)		
HOPE VI×black			0.425*	
			(0.252)	
HOPE VI×Hispanic			0.422	
			(0.326)	
HOPE VI×other			0.354	
			(0.295)	
HOPE VI×male×non-white				0.287**
				(0.126)
observations	149,000	149,000	149,000	149,000

Notes: The table presents estimates from the stratification with regression estimator for large Public Housing Authorities (PHAs) only. The outcome variable is the inverse hyperbolic sine (IHS) of annual earnings measured in the year the child turns 26. Each column presents results from a seperate regression in which the inidcator for HOPE VI is interacted with a different individual-level variable. Note that there are four mutually exclusive race/ethnicity categories, including: white (non-Hispanic), black (non-Hispanic), Hispanic, and other non-Hispanic. All stratum-level regressions control for the base year in which the child appears in public housing, the year in which the child turns 26 and the standard vector of individual- and project-level covariates. Standard errors are clustered at the project-level and are presented in parentheses.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** p≤0.01, ** p≤0.05, * p≤0.10.

are 9.8 percentage points (98 percent) more likely to be in voucher housing and 5.9 percentage points (70 percent) more likely to be in a new public housing project; the analogous figures in small housing authorities are 10.7 percentage points (106 percent) and 9 percentage points (134 percent) for voucher housing and new public housing projects, respectively.

Column 4 of Table 2.4 illustrates that while there is evidence that households were displaced from assisted housing one year after the demolition in large PHAs, HOPE VI did not push households out of subsidized housing in the longer-run.[49] We emphasize that many households in HOPE VI projects did end up leaving subsidized housing within a five-year period, but that the rate at which they did so was similar in the control group–48.5 percent and 54.9 percent of control household heads departed assisted housing within five years in large and small PHAs, respectively. This finding is consistent with other work that finds high rates of turnover in low-quality public housing projects (McClure 2018).

In addition to altering the type of housing, HOPE VI also increased the likelihood of migration to new neighborhoods. Column 6 of Table 2.4 indicates that HOPE VI increased the probability of moving to a new census tract five years after the demolition by 13.0 and 17.2 percentage points in large and small PHAs, respectively.

[49]The category "other public" refers to individuals who appear in the HUD-PIC files but are not in the same project or in voucher housing. The vast majority of these individuals are actually in public housing but there may be a small percentage who participate in the Section 8 Moderate Rehabilitation Program, which is the other assisted housing program covered by the HUD-PIC files. In addition, the category "non-subsidized" refers to individuals who do not appear in the HUD-PIC files. The HUD-PIC files cover both the public housing and voucher programs, which are by far the largest programs subsidizing housing costs for renters. Thus, while there may some households in this group that participate in other subsidized housing programs not covered in the HUD-PIC data, the numbers are likely to be very small.

Table 2.4: Household Head Housing Outcomes

	Housing Type				Moved to New	
	same project (1)	voucher (2)	other public (3)	non-subsidized (4)	county (5)	tract (6)
Panel A. Large PHAs						
A1. 1 year after						
HOPE VI	-0.115***	0.014***	0.023*	0.077***		
	(0.028)	(0.005)	(0.013)	(0.024)		
control mean	0.754	0.026	0.035	0.185		
A2. 3 years after						
HOPE VI	-0.135***	0.085***	0.062***	-0.011	-0.018**	0.085***
	(0.024)	(0.017)	(0.015)	(0.015)	(0.008)	(0.033)
control mean	0.479	0.073	0.068	0.379	0.100	0.535
A3. 5 years after						
HOPE VI	-0.150***	0.098***	0.059***	-0.007	-0.029***	0.130***
	(0.027)	(0.019)	(0.018)	(0.018)	(0.010)	(0.028)
control mean	0.332	0.099	0.084	0.485	0.142	0.646
observations	87,000	87,000	87,000	87,000	87,000	87,000
Panel B. Small PHAs						
B1. 1 year after						
HOPE VI	-0.019	0.016***	0.017**	-0.014		
	(0.022)	(0.006)	(0.008)	(0.020)		
control mean	0.697	0.027	0.027	0.248		
B2. 3 years after						
HOPE VI	-0.207***	0.092***	0.106***	0.009	0.016	0.125***
	(0.023)	(0.018)	(0.021)	(0.019)	(0.011)	(0.030)
control mean	0.416	0.071	0.056	0.458	0.134	0.523
B2. 5 years after						
HOPE VI	-0.184***	0.107***	0.090***	-0.013	0.021	0.172***
	(0.017)	(0.018)	(0.014)	(0.014)	(0.014)	(0.025)
control mean	0.283	0.101	0.067	0.549	0.171	0.621
observations	66,000	66,000	66,000	66,000	66,000	66,000

Notes: Panels A and B present estimates from the stratification with regression estimator for large and small Public Housing Authorities (PHAs), respectively. Outcomes are measured: one year after the refernce year in panels A1 and B1, three years after the reference year in panels A2 and B2, and five years after the reference year in panels A3 and B3. The outcomes in columns 1-4 are indicator variables with a value equal to one if the head of household appears in the same project, other public housing, voucher housing, or other housing after the reference year (categories are mutually exclusive). In columns 5-7 the outcomes are indicators equal to one if the head of household moved to a new state, county, and census tract, respectively. Each stratum-level regression contains a fixed effect for the base year in which the household appears in public housing as well as the standard set of project- and individual-level covariates. Standard errors are clustered at the project-level and are presented in parentheses. The mean and standard deviation, presented in brackets, of the outcome for the control group are a weighted aggregate of strata-level statistics, where the weights are proportional to the number of treated individuals in a strata.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** p≤0.01, ** p≤0.05, * p≤0.10.

Column 5 indicates that these moves to new neighborhoods were typically occurring without moves across county boundaries. The HOPE VI-induced residential mobility is therefore extremely local.

Given this increased mobility of HOPE VI households, we examine the average characteristics of the neighborhoods in which they resided between one and five years after the demolition.[50] Table 2.5 illustrates that, in both large and small PHAs, HOPE VI did not lead households to move to higher quality neighborhoods as measured by census tract school quality and poverty rate, or to demographically distinct neighborhoods as measured by the share of residents that are White non-Hispanic. The estimated impacts on the school proficiency index are not statistically distinguishable from zero and less than 1 percent of the magnitude of a control group standard deviation; similarly, while the point estimates for census tract poverty rates are negative in both small and large PHAs, we are unable to reject the null of no effect and they are only around 4 percent of the control group mean poverty rate in both PHA size groups. This is partially consistent with existing evidence from Chicago: Chyn (2018) finds evidence of short-term moves to more advantaged neighborhoods, but he also finds that these effects fade quickly over time.

Finally, we estimate the effect of HOPE VI on labor market outcomes for the head of household. Table 2.6 presents estimates of the impact of the program on the number of quarters worked and the IHS of annual earnings measured five and ten years after the demolition for the heads of household. We find no evidence that

[50]We find similar patterns if we instead use the timing as in Table 2.5 and measure characteristics of neighborhoods 1, 3, and 5 years after the demolition.

Table 2.5: Household Head Neighborhood Outcomes

	school proficiency index (1)	poverty rate (2)	share white non-Hispanic (3)
Panel A. Large PHAs			
HOPE VI	0.084	-0.017	0.022
	(1.660)	(0.012)	(0.017)
control mean	23.900	0.416	0.192
	[16.900]	[0.181]	[0.235]
observations	87,000	87,000	87,000
Panel B. Small PHAs			
HOPE VI	-0.497	-0.012	0.020
	(2.130)	(0.011)	(0.022)
control mean	30.500	0.289	0.367
	[20.200]	[0.116]	[0.303]
observations	66,000	66,000	66,000

Notes: Panels A and B present estimates from the stratification with regression estimator for large and small Public Housing Authorities (PHAs), respectively. All outcomes are an average characteristic of the census tracts in which the head of household resided in 1-5 years after the reference year. The characteristic in column 1 is the school proficiency index, which measures of the quality of the public schools in that area. In columns 2 and 3 the characteristics are the share of residents who are below the poverty line and white non-Hispanic, respectively. Each stratum-level regression contains a fixed effect for the base year in which the household appears in public housing as well as the standard set of project- and individual-level covariates. Standard errors are clustered at the project-level and are presented in parentheses. The mean and standard deviation, presented in brackets, of the outcome for the control group are a weighted aggregate of strata-level statistics, where the weights are proportional to the number of treated individuals in a strata.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** p$\leq$0.01, ** p$\leq$0.05, * p$\leq$0.10.

HOPE VI improved or depressed labor market outcomes for the parents. Together, none of the results in this section suggest that the long-run benefits for children are driven by measurable improvements in the home or neighborhood environment that would be likely to affect human capital accumulation while young.

2.5.3 Assessing the Validity of the Empirical Strategy

As stated above, our ability to interpret the estimates as causal relies on the Conditional Independence Assumption. While we have previously shown that the methodology does a good job eliminating observable differences between HOPE VI and control projects, it is still possible that the results are biased by unobserved differences or functional form assumptions implicit in the stratum-level regressions. In this section we implement three types of analyses to address these concerns: "pseudo treatment," "pseudo outcome," and "sensitivity/robustness" analyses.

First, we implement a pseudo treatment analysis in which we define a group of projects that were not affected by HOPE VI as pseudo treatment projects. We then estimate pseudo treatment effects by re-implementing the full trimming and stratification with regression method with the pseudo treatment group in place of the true treatment group and omitting the true treatment group from the sample. Estimating null effects for projects that, a priori, should not have systematically different potential outcomes for resident children from comparable projects provides evidence that the methodology is able to adequately correct for baseline differences. This analysis is most convincing if the pseudo treatment projects are, absent exposure

Table 2.6: Household Head Earnings Outcomes

	5 Years After		10 Years After	
	qrtrs worked (1)	IHS earnings (2)	qrtrs worked (3)	IHS earnings (4)
Panel A. Large PHAs				
HOPE VI	-0.001	-0.065	-0.043	-0.134
	(0.036)	(0.092)	(0.032)	(0.082)
control mean	1.960	5.650	1.700	4.840
	[1.810]	[4.800]	[1.840]	[4.950]
observations	87,000	87,000	87,000	87,000
Panel B. Small PHAs				
HOPE VI	-0.005	0.011	0.001	0.004
	(0.039)	(0.099)	(0.041)	(0.109)
control mean	2.070	5.880	1.790	5.050
	[1.810]	[4.750]	[1.860]	[4.940]
observations	66,000	66,000	66,000	66,000

Notes: Panels A and B present estimates from the stratification with regression estimator for large and small Public Housing Authorities (PHAs), respectively. All outcomes are annual labor market outcomes of the head of household measured 5 and 10 years after the reference year for columns 1-2 and 3-4, respectively. In odd and even numbered columns the outcome variables are the number of quarters worked and the inverse hyperbolic since of annual earnings, respectively. Each stratum-level regression contains a fixed effect for the base year in which the household appears in public housing as well as the standard vector of project- and individual-level covariates. Standard errors are clustered at the project-level and are presented in parentheses. The mean and standard deviation, presented in brackets, of the outcome for the control group are a weighted aggregate of strata-level statistics, where the weights are proportional to the number of treated individuals in a strata.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** p≤0.01, ** p≤0.05, * p≤0.10.

to the HOPE VI program, more similar to the true HOPE VI treatment projects than the full set of control projects. Thus, we implement the pseudo treatment analysis using the set of projects that applied for but never received funding for the HOPE VI Demolition or Revitalization programs.[51] Table B.4 presents the estimated effects of the pseudo treatment, which are never statistically different from zero and standard errors are similar in size to those from our main results in Table 2.2. The results indicate that after the matching methodology is applied to the group of failed applicants, there is no evidence of positive bias.[52] Thus, the pseudo treatment analysis bolsters confidence in the validity of our methodology.

Second, we implement pseudo outcomes analyses. Here we select a variable measured prior to the demolition, designate it as a pseudo outcome, and re-implement the trimming and stratification process after excluding any variable that is derived from the pseudo outcome from being included in any other part of the matching or regression analysis. For example, if household income were the pseudo outcome, we would implement the matching and estimation of the propensity score without using

[51]There were too few failed applicants identified in the public data for only the demolitions program, so we pooled applicants across the two programs. However, given that the two programs targeted a similar group of projects and that the projects look similar along observable characteristics at baseline, we argue that this is an informative exercise. Figure B.5 provides evidence to show that failed applicants had similar observable characteristics to the HOPE VI demolition awardees at baseline. Note that failed applicants were subject to the same set of restrictions as all other non-HOPE VI projects.

[52]If anything, there appears to be a negative pseudo treatment effect, which could suggest that HOPE VI projects are negatively selected relative to counterfactual projects and our main estimates may provide lower bounds on the true effect of HOPE VI. Alternatively, these negative (statistically insignificant) associations could be explained if the applicant projects were exposed to alternative, less effective programs in place of HOPE VI. The fact that they might have been exposed to other programs complicates the interpretation of the estimated effect of HOPE VI when the failed applicants are included in the set of controls. While we include the failed applicants in our set of potential controls, in practice they make up only small portion of the matched sample used to estimate the main results. Indeed, our results are robust to excluding failed applicants from the set of matched controls.

the average income at the project level. We then use the stratification with regression estimator to estimate a pseudo outcome effect in which the pseudo outcome is the outcome variable and we include the full set of controls (excluding the pseudo outcome). The results from these analyses are displayed in Table B.5. Each row presents the results for one of the 18 pseudo outcomes, with columns 1-3 presenting estimates for the large, small, and pooled samples, respectively. Overall, the results confirm the ability of the methodology to remove differences between HOPE VI and control projects. Column 3 indicates that only 2 of the 18 pseudo outcome estimates are statistically significant when pooling across housing authority sizes. We do, however, reject the null of no pseudo outcome effect for household income. This likely indicates that household income is a critical variable in the matching process for which there is not a close substitute.

Third, we assess the robustness of the estimates to alternative variables used in the regression adjustment. Table B.6 presents estimates of the effect of HOPE VI for four different specifications that either (1) use the baseline stratification structure or simply define two strata by large and small PHAs and (2) do or do not include covariates in the model. Column 3 and 4 use the baseline stratification structure but do and do not include covariates in the model, respectively. For large PHAs, the estimated effect of HOPE VI on the IHS of earnings at age 26 is 0.157 without controls compared to 0.195 with controls.[53] For small PHAs, estimates with and without controls are similarly small across the two specifications (0.005 and 0.045).

[53]While the point estimate is smaller in column 3 (the specification that uses the stratification structure without covariate adjustment) relative to column 4 (the baseline specification), the estimate would be statistically significant if the standard error from the main specification were used to conduct the hypothesis test.

Thus, once the stratification structure is implemented, the main role of the covariates in the model is to increase precision. This finding suggests that the choice of which covariates are included in the stratum-level regressions and how they are included (functional form) are not driving the results. In addition, the similarity between the standard errors in column 2 and 4 mitigates concerns related to inadequate sample sizes for clustering standard errors at the project level within strata and to individuals in control projects appearing in multiple projects across distinct strata.[54]

2.5.4 Mechanisms

What are the mechanisms through which HOPE VI affected long-run labor market outcomes? In other research that finds long-term labor market benefits of exiting public housing when young, Chyn (2018) and Chetty et al. (2016) find evidence of an exposure model: environment shapes the development of human capital with an influence that is increasing in the duration of exposure and particularly important for young children. However, the evidence presented in Sections 5.1 and 5.2 is inconsistent with the exposure model in our context. Specifically, we find no direct evidence that HOPE VI improved childhood environment by increasing the earnings of parents or improving neighborhood quality along the dimensions typically considered by the literature. Furthermore, we do not find larger impacts for children that were younger at the time of the demolition, a finding that is central to the exposure model. While the evidence suggests a different mechanism than is highlighted in the existing lit-

[54]Appendix Table B.7 shows that the main results are also robust to using OLS and restricting the sample of control projects to: 1) projects in the same PHA as a HOPE VI project, or 2) projects that applied for but did not receive HOPE VI funding. See Appendix B.2 for details.

erature, our results do not conflict with the exposure model findings. Our analysis focuses on older children–between the ages of 10-18 at the time of the demolition–for which prior research has found limited potential for exposure effect-type mechanisms. Additionally, MTO provided assistance for households to facilitate moves to lower poverty neighborhoods and explicitly required moves to lower poverty neighborhoods in the experimental treatment arm. No similar incentives existed for the households affected by the HOPE VI program. It is possible that we would see exposure effects for younger cohorts of children, or if the program at study had included more encouragement for beneficiaries to move to higher quality neighborhoods, as was the case with MTO.

Changes in the exposure to, or involvement in, criminal activity is another mechanism that both motivated the creation of the HOPE VI program and has been explored in the literature on neighborhood effects. While our measures related to crime are admittedly limited, we do not find any evidence that HOPE VI affected the likelihood of incarceration. Following the methodology of Andersson et al. (2018a), we link individuals to the 2010 Decennial Census File to determine whether they reside in an adult correctional facility at the time of the survey. Table B.8 indicates that the effect of HOPE VI increased the probability of being incarcerated in 2010 by 0.001 and 0.005 in large and small PHAs, respectively. These effects are both economically and statistically insignificant.[55]

Rather than affecting the environment in which the children grew up, HOPE VI

[55]In unreported results, we show that this null result also holds when limiting the sample to males, who are at higher risk of being incarcerated.

could have affected labor market outcomes by instead influencing where they live as young adults. We investigate this possibility by studying residential outcomes of the children measured in 2010.[56] As a starting point we estimate a number of specifications in which the outcome variable is an indicator equal to one if the distance between the project and the location of residence in 2010 exceeds some threshold. The results are presented in Figure 2.3. In both large and small PHAs, HOPE VI pushed children away from the neighborhoods in which their projects were located, but the resulting moves were quite local. About one-half of all children lived within five miles of their project in 2010, and HOPE VI increased the likelihood of moving to a new neighborhood within a 5-mile radius of the project but had no discernible effect on moving farther away. Thus, while HOPE VI induced households to move, it did not increase the likelihood that they moved far from their original locations.

It is possible that HOPE VI could have affected labor market outcomes by dispersing residents and breaking apart peer groups. Such disruptions could be either beneficial or detrimental, depending on the characteristics of the network. To investigate this, we use residential location in 2010 to measure the distance between adult children and each of their former public housing co-residents. We create four variables to characterize network dispersion: the average log distance to all former

[56]We focus on 2010 because we are best able to measure residential location by combining data from both the 2010 Decennial Census and the CPR. The children are between the ages of 19 and 31 in 2010 and thus these measures of residential location may not correspond exactly to where children are living when we measure their earnings at age 26. However, we do not think this is a major concern because most children will be in their mid-twenties at this time and, as shown in Figure B.4, the effect of the program in Large PHAs starts at around zero at age 18 but increases to about 0.2 by age 23, after which point the effects stabilize through age 26. The longitude and latitude from the internal points of the census tract (the centroid) are from Census Bureau Gazetteer Files for 2010 geography.

Figure 2.3: Effect of HOPE VI on Distance between Project and Residence in 2010

(A) Large PHAs

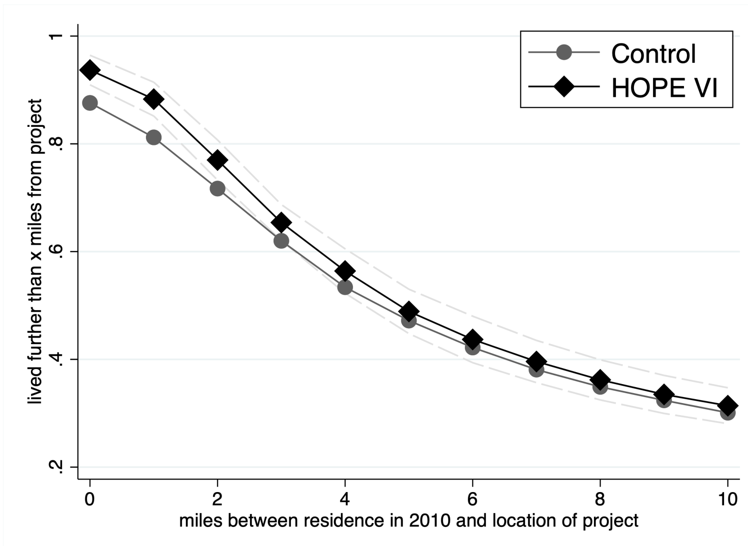

(B) Small PHAs

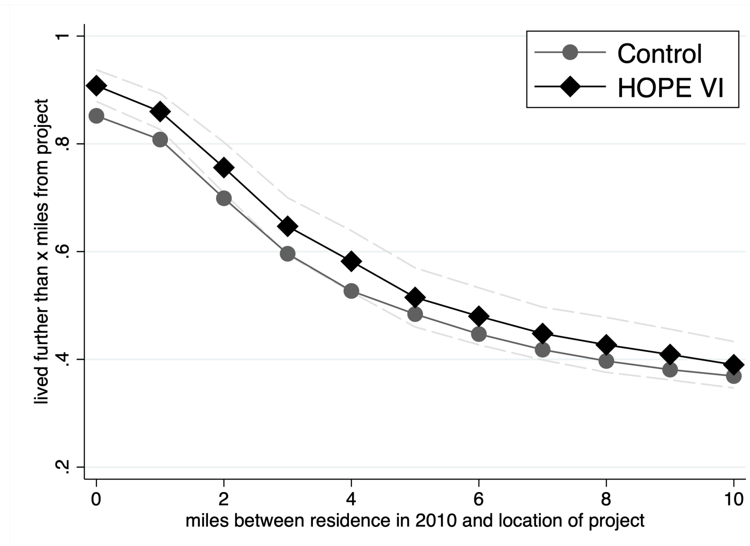

Notes: Panel A and B present results for large and small PHAs, respectively. The grey series with the circle markers plot the proportion of children from control projects for whom the distance between their place of residence in 2010 and the location of their project is strictly greater than the number of miles denoted on the horizontal axis (distance is calculated between the centroids of the census tracts). We use the stratification with regression estimator to estimate the effect of HOPE VI on living further than a given distance. The black series with the diamond markers depicts these results by plotting the control mean plus the estimated effect of HOPE VI. The grey lines denote the 95% confidence interval, where standard errors are clustered at the project-level.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

co-residents and the share of former residents who live within a 1-, 3-, and 5-mile radius. The results, presented in Table B.9, suggest that HOPE VI did not disperse residents geographically in large PHAs. While these are coarse measures, the results provide no evidence that HOPE VI improved labor market outcomes by disrupting peer groups formed in public housing. Note that these results are not inconsistent with those presented in Figure 2.3 since affected households could have remained spatially close if they moved to a nearby neighborhood after the demolition.

HOPE VI could also have influenced children's subsequent labor supply decisions by affecting the probability that children live in subsidized housing or with their parents as young adults. However, the results in column 3 of Table B.8 indicate that HOPE VI had no detectable effect on the probability of being in subsidized housing in 2010; similarly, columns 4 and 5 show that the program had no detectable impact on living near or with parents.

We do find, however, that the program led to a meaningful change in some of the characteristics of the neighborhoods where the children lived as young adults. We estimate the effect of HOPE VI on six characteristics of the census tract in which the individual resided in 2010 including poverty rate, employment rate, a measure of labor market networks (observed network isolation), and three measures of the geographic proximity to jobs (the log of the ratio of jobs to people, the average commute time and a job proximity index that captures the "the accessibility of a given neighborhood as a function of its distance to all job locations within a [Core-Based Statistical Area]").[57]

[57]For a description of the job proximity index see: http://hudgis-hud.opendata.arcgis.com/datasets/jobs-proximity-index. The underlying measure is the same as Shen (1998) and Wang (2007) and is similar to that in Andersson et al. (2018b), though it uses distance for the impedance function rather than travel time. The values of this underlying measure

The results, presented in Table 2.7, illustrate that, within large PHAs, HOPE VI lead to an improvement in the geographic proximity to jobs along all three measures considered. In contrast, there is no evidence that HOPE VI moved children to better neighborhoods in terms of poverty, employment rate, or network isolation. In small PHAs, there is no evidence that HOPE VI improved geographic proximity to jobs, and even some evidence that it led individuals to live in areas with lower job proximity.

Were HOPE VI-induced moves substantial enough to plausibly generate the improvements in job proximity? While these moves tended to be to nearby neighborhoods, moving short distances could still lead to large improvements in job proximity; the housing projects in the sample were often located in neighborhoods that were especially geographically isolated from jobs, even relative to nearby neighborhoods. This can be seen in Figure 2.4, which presents the average commute time, poverty rate and population density in 1990 (before all demolitions) for housing projects by treatment status (HOPE VI and control), PHA size (large and small), and distance to a sample project (whether HOPE VI or control). In large PHAs, the public housing residents (of both HOPE VI and control projects) had substantially higher commute times, poverty rates, and population densities relative to residents of surrounding neighborhoods.[58] Thus, it is plausible that even the local moves induced by HOPE

are percentile ranked with values ranging from 0 to 100 and higher values indicates neighborhoods with better access to jobs. The job proximity index is constructed by HUD using data from LODES (based on LEHD) for 2014. The observed network isolation index measures, for employed residents of a tract, the share of their co-workers who are also neighbors, where high values of this variable could arise if information on job opportunities disseminate through local networks (see Hellerstein et al. 2011 and Hellerstein et al. 2019).

[58]The average commute time is the best available measure of job proximity prior to the demolitions. The job proximity index is not available during this time period since the LEHD data used to construct the measure have limited coverage years prior to 2000.

Table 2.7: Neighborhood Outcomes

	poverty (1)	employment (2)	social isolation index (3)	log(jobs/pop) (4)	avg. commute (5)	jobs proximity index (6)
Panel A. Large PHAs						
HOPE VI	0.003	-0.011**	-0.013	0.019**	-0.883**	2.110**
	(0.007)	(0.006)	(0.03)	(0.009)	(0.388)	(0.859)
control mean	0.296	0.516	1.03	0.285	27	48.1
control sd	[0.179]	[0.122]	[0.951]	[0.347]	[6.51]	[21.2]
observations	149,000	149,000	149,000	149,000	149,000	149,000
Panel B. Small PHAs						
HOPE VI	-0.021***	0.017***	-0.108	0.01	0.424	-1.180*
	(0.007)	(0.006)	(0.084)	(0.011)	(0.351)	(0.703)
control mean	0.273	0.509	1.7	0.318	24	51.7
control sd	[0.148]	[0.112]	[1.35]	[0.317]	[5.55]	[20.2]
observations	109,000	109,000	109,000	109,000	109,000	109,000

Notes: Panels A and B present estimates from the stratification with regression estimator for large and small Public Housing Authorities (PHAs), respectively. All outcomes are characteristics of the census tracts in which children resided in 2010. In columns 1-7 the outcome variable is: poverty rate, employment rate, social isolation index, log of the ratio of jobs to people, average commute time in minutes and the jobs proximity index created by HUD. All stratum-level regressions control for the base year in which the child appears in public housing, the year in which the child turns 26, and the standard vector of individual- and project-level characteristics Standard errors are clustered at the project level and are presented in parentheses. The mean and standard deviation, presented in brackets, of the outcome for the control group are a weighted aggregate of stratum-level statistics, where the weights are proportional to the number of treated individuals in a strata.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$.

VI could have shifted children into neighborhoods with better access to jobs.

HOPE VI neighborhoods were outliers in terms of both job proximity and poverty. An important question then is how the program could have induced moves to new neighborhoods that were better in terms of job accessibility but not poverty. One potential explanation is that neighborhood poverty is more strongly (and negatively) associated with housing prices than job accessibility. For the households participating in public housing, meaningful housing price increases are likely to preclude them from moving to a neighborhood. To assess this possibility, we use publicly available data to identify all counties that contained a HOPE VI project. Within each of these counties we construct population-weighted percentile ranks of neighborhoods based on median rent, average commute time, and poverty rate as measured in 1990. Columns 1-2 of Table B.10 present estimates from bivariate OLS regressions of average commute time on median rent and neighborhood poverty on median rent. There are two key findings from the estimates. First, poverty and median rent are strongly negatively correlated; within a city, neighborhoods with a one percentile higher rank in terms of median rent have a 0.53 and 0.41 percentile lower poverty rate rank in large and small PHAs, respectively. The R-squared values from these regressions are 0.18 (in large PHAs) and 0.11 (in small PHAs). In contrast, there is no evidence that neighborhoods with higher levels of job proximity are more expensive. In large PHAs the R-squared from a regression of the average commute time rank on the median rent rank is 0.003, and the point estimate is small, negative (-0.0189), and not statistically distinguishable from zero.[59] While the average commute time is an imperfect measure

[59]In small PHAs there appears to be evidence that job proximity is negatively related to housing

Figure 2.4: Characteristics of Surrounding Neighborhoods in 1990

(A) Average Commute Time

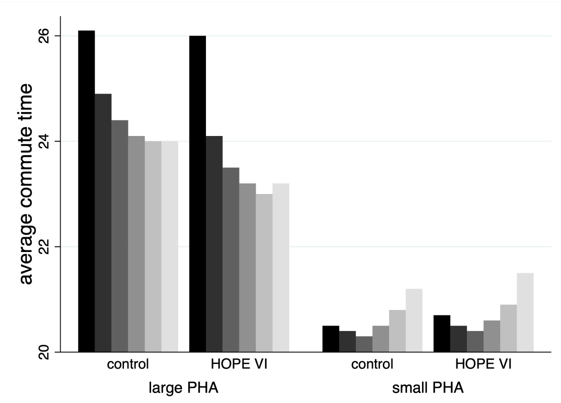

(B) Poverty Rate

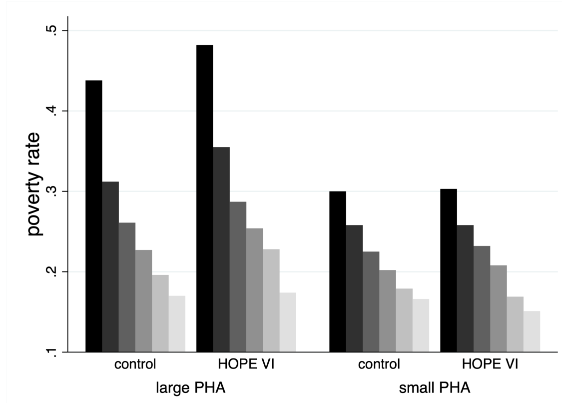

(C) Population Density

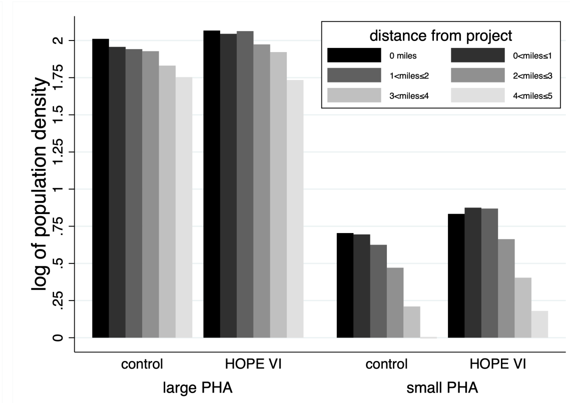

Notes: Panel A, B and C present the average commute time, poverty rate and log of the population density in the census tracts in 1990, respectively. We present characteristics of the neighborhoods surrounding the groups of projects separately by PHA size (large and small) and treatment status (HOPE VI and control). The bars present the average value of census tract characteristic within that group. The six shaded bars within each group summarize this information separately for the census tract in which the project was located as well as surrounding census tracts located within 5 miles of the project.

Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

of job proximity, the lack of relationship between housing costs and job proximity is robust to using the job proximity index constructed by HUD, which is available only in 2010 (see columns 3-4 of Table B.10). The cross-sectional associations therefore support the idea that while moves to neighborhoods with better job accessibility were likely financially feasible for HOPE VI-affected households, higher housing prices may have made moves to lower poverty neighborhoods more difficult. This is consistent with evidence from of quasi-experimental (Collinson and Ganong 2018; Andersson et al. 2018a), experimental (Patterson et al. 2004; Eriksen and Ross 2013; Jacob, Ludwig, and Miller 2013), and observational (Susin 2002; Carlson et al. 2012) studies on the Voucher program. An important caveat is that homes with similar rents in more job-accessible neighborhoods may be lower quality (e.g. smaller), but absolute housing cost may still be the most relevant decision factor for our sample.

In addition to forcing people to move to new neighborhoods, HOPE VI could have also improved job accessibility for households that remained in their original neighborhood or moved extremely short distances by altering the characteristics of the original neighborhoods themselves. To explore this possibility, we measure the average job proximity index of census tracts within half-mile radius bands from zero to five miles around the project. We then attach these neighborhood-level measures to the child-level dataset and implement the stratification with regression methodology as before to estimate the effect of HOPE VI on the characteristics of these neighborhoods. The results for large and small PHAs are presented in Figure 2.5. We see no significant impacts on job proximity in small PHAs at any distance. For large

costs.

PHAs, HOPE VI produced substantial improvements in the job proximity index for the census tract in which the project was located, but these effects dissipate quickly and there appears to be no impact on neighborhoods located farther than half a mile away.[60] That the effects dissipate quickly with distance is reassuring since we would not expect the demolition of a public housing project to drastically transform the population or job density in more distant neighborhoods.[61]

To investigate the origins of the effect on these neighborhood-level measures of job proximity, we estimate the effect of HOPE VI on three characteristics of the census tract in which the project was located: the log of the ratio of jobs to people, the log of the density of jobs, and log of population density.[62] The results, presented in Table B.11, imply that, HOPE VI increased the ratio of jobs to people in large PHAs by 22 percent, and that this impact was driven primarily by a reduction in population density: HOPE VI reduced in population density by 62 percent, a finding that is statistically significant at the 1 percent level. HOPE VI is also associated with a statistically insignificant 4.5 percent increase in job density. A reduction in population density will increase job accessibility by reducing the number of competing searchers in the local labor market (more competing searchers lower the job proximity

[60] The finding that the neighborhood in which the project was located underwent large changes supports our choice to include all, and not just partial, demolitions in the analysis. Household in units that were not demolished were still treated by the program by changes in neighbors and changes in the existing neighborhood.

[61] The fact that HOPE VI affected both the census tract in which the project was located and census tracts within a half mile radius could reflect the fact that projects may have been located in multiple census tracts though we assign each project to a unique census tract. Other research on HOPE VI has generally found that spillover effects of the demolitions dissipate within a mile (e.g. Sandler 2017).

[62] Density is calculated by dividing the number of jobs (or population) by the land area of the census tract, so both measures use the same land area for normalization. Land area cancels out in the job/population ratio.

Figure 2.5: Effect of HOPE VI on Surrounding Neighborhoods

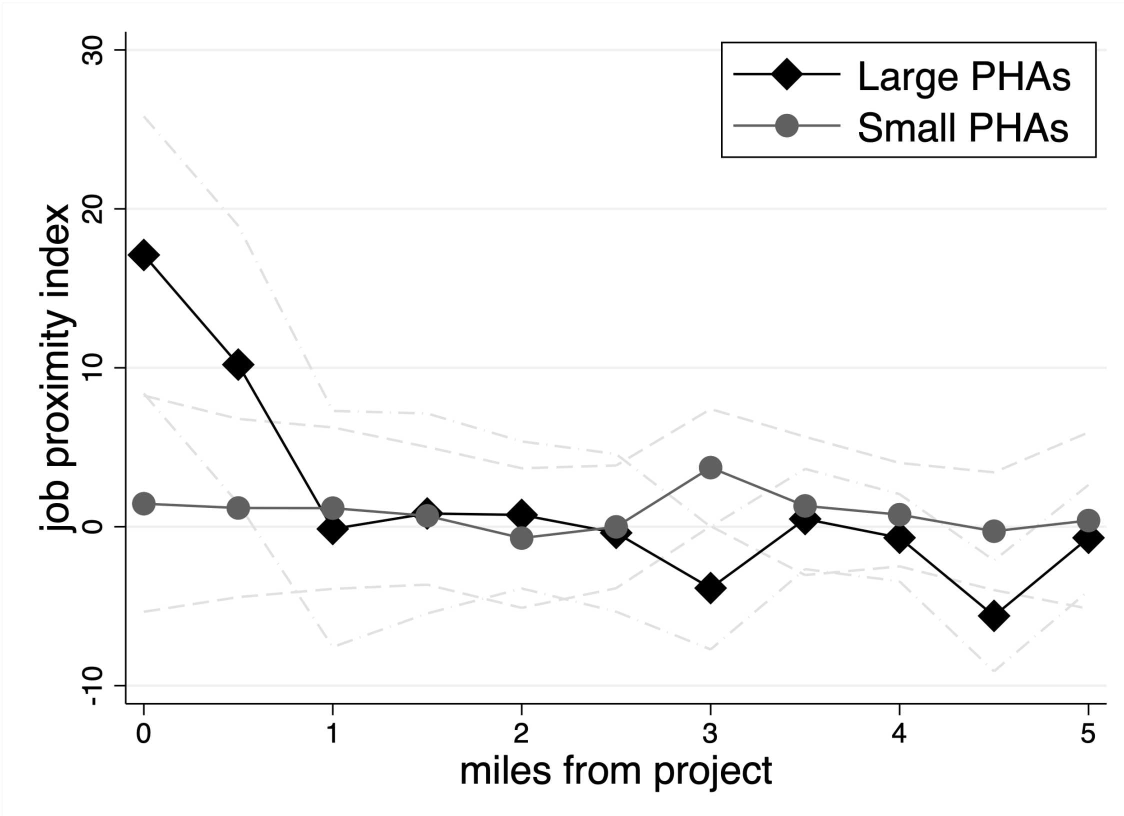

Notes: The black line with diamond markers and the grey line with circle markers plot the estimated effect of HOPE VI for large and small Public Housing Authorities (PHAs), respectively. Each point corresponds to results from a separate specification estimated via the stratification with regression methodology. The outcome for the points at the value of zero on the horizontal axis is the job proximity index (measured in 2010) for the census tract in which the project is located. The outcome for the remaining points correspond to the average job proximity index for other census tracts that are within the number of miles denoted on the horizontal axis (exclusive) and half a mile less than this value (inclusive). All stratum-level regressions are estimated on the child-year dataset and control for the base year in which the child appears in public housing, the year in which the child turns 26, and the standard vector of individual- and project- level covariates. Standard errors are clustered at the project level and are 95% confidence interval is depicted by the dashed light grey lines.

Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

index) as long as the number of jobs in the neighborhood does not also decline. In the case of a public housing demolition, the reduction is for a population likely to compete for a similar set of jobs (Lens 2014; Lens et al. 2019). In small PHAs, we find no effect of HOPE VI on job or population density. This is likely due to the fact that public housing projects typically had far fewer residents in small PHAs, and the demolitions therefore did not displace as many households or lead to meaningful reductions in population density.

The preceding analyses suggest that HOPE VI improved geographic proximity to jobs in large PHAs both by transforming the neighborhood in which the project was located and by moving former residents to new neighborhoods with better accessibility. To investigate the quantitative importance of each channel, we estimate specifications that replace the true measure of job proximity with a counterfactual measure that discards all variation due to changes in the HOPE VI neighborhoods. In order to calculate this counterfactual measure, we use the stratification with regression method to estimate the effect of HOPE VI on the job proximity index, limiting the sample to census tracts within a half-mile radius of the original project; note that these are the areas where HOPE VI directly impacted job proximity, as shown in Figure 2.5. We obtain a predicted value of the job proximity index for HOPE VI neighborhoods in the absence of changes to the original neighborhood by setting all covariates to their true value except for the HOPE VI indicator, which is set to zero instead of one. The counterfactual measure of the job proximity index is equal to this predicted value for all children who resided in HOPE VI projects and still lived within a half-mile of their project in 2010–i.e. children whose neighborhood job proximity was directly affected by the demolitions-induced changes–and is set to the true value of the job proximity index for all other children. Intuitively, we impute the job proximity for individuals from HOPE VI projects who remained within a half-mile of their original project (and therefore benefited from changes in the neighborhood of origin) using the job proximity for individuals from observably similar control projects. Any estimated improvements using this counterfactual measure of job proximity will thus be entirely driven by HOPE VI-induced moves to new neighborhoods. We then estimate the impact of HOPE VI on this counterfactual job proximity measure for large PHAs. The original estimates, presented in Table 2.7, indicate that HOPE VI increased the job proximity index by 2.11. When the counterfactual value of the job proximity index is used as the outcome variable, this estimated impact falls to 1.16, suggesting that improvements in the neighborhood in which HOPE VI projects were located explain about 45 percent of the total impact on the job proximity index in large PHAs. This back-of-the-envelope calculation therefore suggests that, within large PHAs, HOPE VI improved access to jobs by moving children to new neighborhoods and by improving the original neighborhoods that contained the HOPE VI projects, with both channels being quantitatively important.

Improvements in job proximity could affect earnings by reducing job search and/or commuting costs and encouraging individuals on the margin between working and not working to participate in the labor market. Consistent with this hypothesis, we find that an important part of the earnings gains occurs through an extensive margin labor supply response. Using the estimates from Table 2.2, the control means from columns 1 and 4 indicate that the average working child from the control group earns

$3,944 per quarter whereas column 1 indicates that HOPE VI increased quarters worked by 0.076. Using the effect on quarters worked and average earnings per quarter in control projects we calculate that the effect on annual earnings would be $300 (3,944*0.076=$300) if the entire effect were driven by an increase in labor force participation. This is about 57 percent of the estimated effect in column 3, suggesting that extensive margin labor supply responses are the main avenue through which the earnings impacts occur.

As discussed earlier, we find no effect of HOPE VI demolitions on earnings for the heads of household. Given that many of these heads of household are single mothers who qualify for public support and have especially high opportunity costs for time supplied in the labor market, a likely explanation for this discrepancy is that the heads of household have higher reservation wages. Figure B.6 presents the distribution of earnings for household heads and the adult children. Consistent with the theory that household heads have a higher reservation wage, there is a hollowing out of the distribution of labor market earnings for household heads relative to the adult earnings of the children in our main sample; household heads are more likely to have zero earnings (48 percent compared to 35 percent) and less likely to have low levels of strictly positive earnings (10 percent of households heads have earnings in the bottom quartile compared to 18 percent of the adult children).

In sum, there are three reasons that support the job accessibility mechanism as an important driver of our main results. First, we find systematic evidence of improvements in measures of job proximity within large PHAs, where differences in job proximity should be larger and more meaningful. Second, the effect on earnings appears to have a substantial extensive margin component, which is consistent with the hypothesis that HOPE VI primarily affected the costs associated with finding a job and not the rewards from work. Third, the difference in the effect of HOPE VI on earnings in large PHAs versus small PHAs is precisely mirrored by the differences of the effects on the various measures of job accessibility. In the next section we expand upon this last point and show that even within the large PHAs the impacts on earnings are largest in places where we would expect job accessibility to be particularly low absent the intervention.

2.5.5 Reconciling Different Effects in Different Environments

Why does HOPE VI produce substantial long-run labor market gains for children living in large but not small PHAs? One possible explanation is that the program interacted in important ways with local environments. In particular, poor geographic access to jobs might affect labor market outcomes more in the worst neighborhoods. Figure 2.6 presents kernel density plots of the average commute time, poverty rate, and population density in 1990 in the census tracts containing projects in the sample, separately by PHA size (large or small) and HOPE VI treatment status. The figure illustrates that prior to the demolitions, projects in large PHAs, regardless of whether they subsequently received a HOPE VI grant, had significantly higher average commute times, poverty rates, and population densities.[63]

Figure 2.6 also illustrates that there is substantial variation even within the large

[63]Komogorov-Smirnov equality-of-distribution tests confirm that the differences between HOPE VI projects in the large and small PHAs are statistically significant while the differences between the control and HOPE VI projects within large and small PHAs are not statistically different from one another.

Figure 2.6: Characteristics of Census Tracts in 1990

(A) Average Commute Time

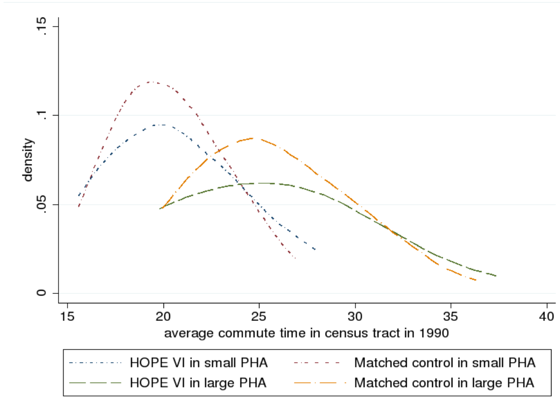

(B) Poverty Rate

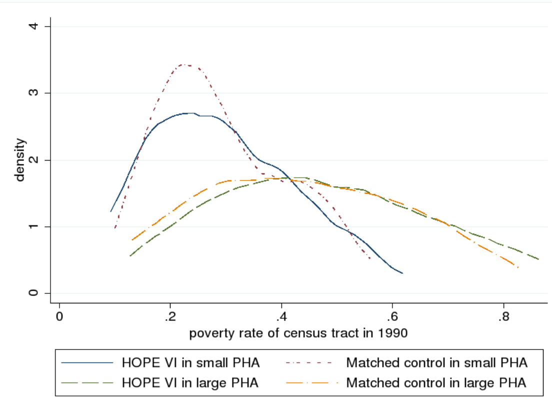

(C) Population Density

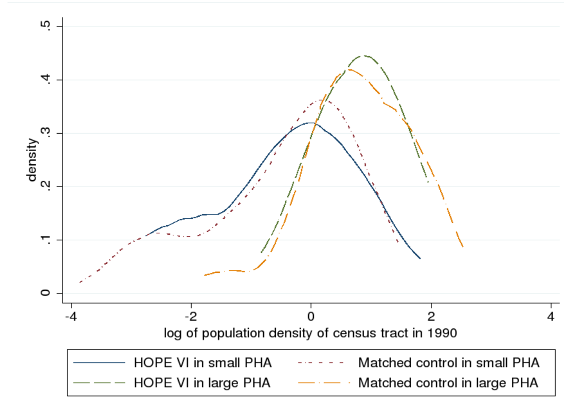

Notes: Panel A, B and C present kernel density plots of the average commute time, poverty rate and log of the population density in census tracts in 1990, respectively. Results are presented separately by groups defined by the interaction between PHA size (large and small) and treatment status (HOPE VI and control). All results are produced from a project-level dataset. To pass disclosure review requirements, each of the distributions are Winsorized at the 5th and 95th percentiles.

Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

PHAs in terms of these baseline characteristics of neighborhoods. We make use of this variation by estimating three specifications in which we interact the indicator for HOPE VI with pre-demolition measures of neighborhood average commute time, poverty, and population density. The results for large PHAs, presented in Table 2.8, suggest that demolitions had stronger effects for projects in neighborhoods that were more densely populated, where commutes were longer, and where the poverty rate was higher in 1990.[64] The heterogeneity is economically meaningful. For example, the results suggest that HOPE VI increased age 26 earnings by 37 percent for children in neighborhoods that had baseline poverty rates one standard deviation above the mean poverty rate among HOPE VI projects. In comparison, children in neighborhoods with poverty rates one standard deviation below the mean only experienced a 10 percent increase in earnings.

Together, the heterogeneity in the effect of HOPE VI both across and within large and small PHAs suggests that the program produced larger labor market gains for children originally residing in high-density, high-poverty neighborhoods, with limited job opportunities nearby. Within these communities, HOPE VI improved labor market outcomes both by shifting children into neighborhoods with better job accessibility and by improving the job accessibility of the original neighborhoods. In contrast, the program offered much smaller (or no) benefits to individuals residing in neighborhoods with better job accessibility prior to the demolition.

The treatment effect heterogeneity is also informative for interpreting findings

[64]Table B.11 presents the results for small PHAs. We find no evidence of meaningful interaction effects here, which is not surprising given that we find no significant effect of HOPE VI in this sample in general.

Table 2.8: Heterogeneous Effects by Neighborhood, for Large PHAs

	IHS of Earnings at Age 26		
	(1)	(2)	(3)
HOPE VI	0.180**	0.189***	0.233***
	(0.072)	(0.065)	(0.071)
log population density	-0.110**		
	(0.052)		
HOPE VI × log population density	0.192**		
	(0.082)		
average commute time		0.080	
		(0.079)	
HOPE VI × average commute time		0.146**	
		(0.073)	
poverty rate			-0.011
			(0.079)
HOPE VI × poverty rate			0.132**
			(0.065)
observations	149,000	149,000	149,000

Notes: The table presents estimates from the stratification with regression estimator for large Public Housing Authorities (PHAs) only. The outcome variable in all specifications is the inverse hyperbolic sine (IHS) of annual earnings measured at age 26. Columns 1-3 presents esitmates from models in which the indicator for HOPE VI is interacted with a characteristic of the census tract in which the project is located measured in 1990. For columns 1-3 these characteristics include the log of the popoulation density, the average commute time in minutes and the poverty rate, all three of which are normalized by substracting by the mean of the control group and dividing by the standard deviation of the control group. All stratum-level regressions control for the base year in which the child appears in public housing, the year in which the child turns 26, and the standard vector of individual- and project-level covariates. Standard errors are clustered at the project-level and are presented in parentheses. The mean and standard deviation, presented in brackets, of the outcome for the control group are a weighted aggregate of stratum-level statistics, where the weights are proportional to the number of treated individuals in a strata.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** p≤0.01, ** p≤0.05, * p≤0.10.

from existing research. As previously discussed, Chyn (2018) and Chetty et al. (2016) both find long-term labor market benefits from exiting public housing when young. However, only Chyn (2018) finds that these benefits extend to older children (older than 13). Our results suggest an explanation for this discrepancy: Chyn (2018) notes that the projects in his study were in much higher poverty neighborhoods than those in MTO. His sample thus included public housing projects that were much more disadvantaged, located in neighborhoods with limited job accessibility. Thus, moving older children out of these projects produced more immediate labor market gains, whereas no such gains occurred for older children in the context of Chetty et al. (2016). Relatedly, while Anderson et al. (2018a) find that time spent in public and voucher housing when young produces long-term labor market benefits of similar magnitudes, our paper highlights the fact that these average effects mask substantial heterogeneity, and that children in the lowest quality public housing projects may benefit from changes in housing. More broadly, the results from our paper highlight how housing and neighborhood can affect long-term outcomes through a multitude of channels that vary in importance with local context.

2.6 Conclusion

This paper uses administrative data on earnings and participation in subsidized housing to study how the demolition of public housing projects—funded by the HOPE VI demolitions program—affected the long-run earnings of resident children. We find that, on average, exposure to a demolition increased earnings at age 26 by 14 percent. However, the benefits appear to be driven by children who lived in neighborhoods that were denser, poorer, and farther from jobs prior to the demolition. In terms of potential mechanisms, we find no evidence that HOPE VI improved the home or neighborhood environment that children were exposed to while young. We do, however, find evidence consistent with HOPE VI improving labor market outcomes by increasing the proximity of job opportunities in the neighborhoods in which the children lived as young adults.

Over the past thirty years, federal housing policy has sought to move families living in subsidized housing out of especially disadvantaged neighborhoods. The results in this paper offer evidence that these moves can generate long-term labor market benefits for children. Interestingly, we find that these moves need not occur in early childhood to produce improvements in adult labor market outcomes.[65] Instead, our findings highlight the important and immediate impact of reducing barriers to young adult employment through increasing the accessibility of formal market jobs. Neighborhoods can affect labor market outcomes through multiple channels, and severely distressed public housing projects can, in some cases, limit job accessibility and discourage labor force participation by creating densely populated neighborhoods with high rates of poverty and a limited number of nearby jobs.

Our results highlight the importance of accounting for the interaction between subsidized housing policies and local context. Much of the research on assisted housing has taken place in a limited set of large metropolitan areas. In the case of public

[65]It is important to emphasize that we are not suggesting that these moves are more beneficial than earlier moves to higher quality neighborhoods. We are not able to investigate this in our study as the youngest children exposed to the demolitions are not old enough to measure adult labor market outcomes.

housing demolitions, our results indicate that the long-run labor market benefits found in other work are specific to this setting (at least for older children), which highlights the possibility that resources may be better spent on alternative interventions in less urban and disadvantaged environments. Research has convincingly documented that housing can have important long-run labor market implications but anticipating the effects of potential interventions requires a more complete understanding of the mechanisms. Future research should continue to focus on better understanding how the impacts of housing policies interact with the characteristics of local environments to produce changes in welfare and contribute to the intergenerational transmission of economic outcomes.

Chapter 3: Cyclical Worker Flows: Cleansing vs. Sullying

Disclaimer

This chapter is joint work with John Haltiwanger, Henry Hyatt, and Erika McEntarfer. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed (CBDRB-FY21-CED006-0002).

3.1 Introduction

Economists have long sought to understand how business cycles affect the reallocation of resources. Do recessions promote economic efficiency by "cleansing" out less productive firms and redirecting labor to more productive uses? Or, does the decline in job mobility in recessions "sully" productivity-enhancing worker reallocation, leaving workers matched to mediocre firms? In this paper we use U.S. linked employer-employee data to decompose the employment growth of high- and low-productivity firms into two components: growth accounted for by job-to-job moves and growth accounted for by flows through nonemployment. We find that job-to-job flows move workers from less productive to more productive firms and the rate at which workers move up this job ladder is highly procyclical. In contrast, less productive firms rely heavily on hiring jobless individuals in expansions and are disproportionately more likely to displace workers back to nonemployment in contractions. In this way, worker flows through nonemployment shift workers away from low-productivity firms in contractions.

We thus find empirical evidence of both cleansing and sullying effects of recessions, which feature in many models of the labor market. Much of the theoretical literature focuses on either cleansing or sullying effects. Schumpeter (1939) originally proposed that recessions may be productivity enhancing, driving out less productive uses of capital and labor and freeing these resources for more efficient use. The notion of cleansing effects of recessions was later revived in Davis and Haltiwanger (1990), Caballero and Hammour (1994), and Mortensen and Pissarides (1994).[1] Barlevy (2002) notes that cleansing effects of recessions appear at odds with observed procyclical job quality. He proposes a model whereby declines in job-to-job moves cause a drag on productivity in recessions, a sullying effect. Sullying effects of recessions feature more recently in a set of papers by Moscarini and Postel-Vinay (2013, 2016) (MPV). The MPV framework in particular yields a rich set of predictions for the cyclical reallocation of workers across the firm productivity distribution that we largely confirm in our empirical analysis.[2] More recent research by Lise and Robin (2017) and Baley, Figueiredo, and Ulbricht (2020) use models with heterogeneity of both workers and firms that induce cyclical variation in sorting over the cycle. Within their frameworks, both of these papers suggest that the cleansing effect dominates the sullying effect.[3] Our empirical results suggest that the cleansing effects dominate at the start of a recession but that the sullying effects continue long into the economic recovery.

In our empirical analysis, we classify firms into high- and low-productivity firms

[1]Unlike Schumpeter, the more recent literature does not argue that recessions are desireable but rather that, conditional on the adverse shock occurring, there may be an acceleration of the ongoing reallocation of resources to more productive uses in recessions.

[2]The MPV model builds on the job ladder framework of Burdett and Mortensen (1998) but importantly for our purposes incorporates business cycle dynamics. In their framework, search frictions prevent workers from immediately moving into desirable job matches and workers move to better firm matches via job-to-job flows. High-productivity firms are able to offer higher wages and thus are able to grow faster in expansions than less-productive firms, who must rely on the pool of unemployed workers in filling vacancies. In recessions, this cyclical job ladder collapses, yielding a sullying effect. Our empirical results are largely consistent with these predicted dynamics. Our results build on the findings in Haltiwanger, Hyatt, Kahn, and McEntarfer (2018) and Haltiwanger, Hyatt, and McEntarfer (2018). This earlier work provides support for a procyclical job ladder in terms of firm earnings and productivity. The current paper is distinguished by explicitly considering the sullying and cleansing contributions of flows through the lens of the impact on the share of employment at high and low productivity firms. This earlier work does not explore decompositions of productivity growth into cleansing and sullying components. Bertheau, Bunzel, and Vejlin (2020) report broadly similar patterns on poaching by high-paying and low-paying firms using Danish data.

[3]However, worker flows across firms ranked by firm productivity are not targeted directly when they estimate their models. While we do not explicitly consider worker heterogeneity in our empirical analysis, we do provide novel evidence on the patterns of worker flows across firms ranked by productivity that are relevant for these models of sorting.

based upon the relative ranking of firms in the measured distribution of firm-level productivity. We find that high-productivity firms grow faster: the differential growth rate averages 0.20 percent of employment on a quarterly basis. Decomposing the growth rate differential into the components due to job-to-job moves versus worker flows through nonemployment, we find that the propensity of high-productivity firms to grow faster is driven by their advantage in poaching workers from less productive firms. This advantage is large: the difference in the growth rates due to job-to-job moves is 0.61 percent per quarter. Low-productivity firms actually lose workers on net through poaching, and so hire relatively more jobless workers to sustain growth: the differential employment growth rate from worker flows through nonemployment (relative to high-productivity firms) is 0.41 percent per quarter.

We find that these patterns of worker reallocation between low- and high-productivity firms differ dramatically over the course of the business cycle. In expansions, high-productivity firms actively poach workers from less productive employers. During recessions, this job ladder collapses yielding a sullying effect. The nonemployment margin also changes dramatically across the cycle. In expansions, low-productivity firms grow primarily through hiring jobless workers. In contractions, worker separations to nonemployment from low-productivity firms spike disproportionately and hires from nonemployment into low-productivity firms decline disproportionately. This cleansing effect peaks earlier in a downturn compared to the collapse of the job ladder that lingers into the early stages of a recovery. Cleansing effects are stronger than sullying effects when the unemployment rate surges during recessions, but they are almost completely absent in the times of high unemployment that follow recessions.

What are the implications of these business cycle dynamics on aggregate productivity growth? To answer this question we use an accounting decomposition of an index of aggregate productivity growth that is the employment-weighted average of a firm-level measure of productivity, which is measured in logs. On average, worker reallocation through job-to-job flows contributes 0.1 log points to the index of overall productivity growth each quarter. This is a substantial contribution to the average quarterly rate of aggregate productivity growth, which is 0.33 based on our data. However, during recessions there is clear evidence of a sullying effect. This is particularly true during the Great Recession. Prior to the recession, in 2006:1, worker reallocation via job-to-job flows contributed 0.13 log points to quarterly aggregate productivity growth but this contribution declined to 0.02 by 2009:2. Acting against this is a cleansing effect that operates via worker flows through nonemployment. In 2006:1, worker reallocation via nonemployment contributed -0.1 log points to quarterly aggregate productivity growth but this increased to 0.08 in 2009:1. We show that the during periods of rising unemployment the cleansing impact on productivity outweighs the sullying impact. However, during periods when unemployment is above trend the sullying impact on productivity outweighs the cleansing impact.

Our primary analysis uses revenue per worker to measure relative productivity within the firm's industry. This has the advantage of allowing us to measure productivity within most industry sectors, but has two key disadvantages.[4] First, we cannot easily compare productivity across sectors and so our main analysis abstracts

[4]Total factor productivity measures are available for only handful of sectors that are a shrinking share of the U.S. economy.

from productivity-enhancing reallocation from less to more productive sectors. Second, revenue per worker reflects both "innate" firm productivity as well as sorting of workers across firms. We assess these issues using the AKM decomposition of earnings.[5] The AKM firm fixed effect represents the pay-premium workers receive independent of worker quality and while not a direct measure of productivity, the relative ranking should reflect productivity differences within and across industries. We find that the cyclical patterns of hires and separations via poaching and nonemployment are very similar whether we rank firms based on the direct measure of labor productivity (revenue per worker) or the indirect measure (AKM firm fixed effect). Thus, our results are do not appear to be driven by cyclical sorting of worker types across firms.[6]

Finally, we consider the implications of the worker flows across firms ranked by productivity for worker earnings. We find that worker movements from low- to high-productivity firms move workers into higher-paying firms, both measured by firm average earnings as well as the AKM firm fixed effect. These earnings changes move strongly with their productivity analogues but are roughly half of the magnitude. These results imply that workers obtain a substantial fraction of the gains from worker movements onto and up the firm productivity job ladder, but that there is suggestive evidence of an incomplete pass-through of gains to workers.

The paper proceeds as follows. Section 3.2 describes the data. Section 3.3 presents

[5] AKM refers to the decomposition developed by Abowd, Kramarz and Margolis (1999).

[6] Haltiwanger, Hyatt and McEntarfer (2018) and Crane, Hyatt, and Murray (2020) do examine cyclical sorting of heterogeneous workers across heterogeneous firms. Both papers find evidence that the assortative matching of workers and firms is countercyclical: a greater share of low-productivity workers are able to match to better firms in expansions.

evidence on worker movements onto, up, and off of the firm productivity job ladder over the cycle. Section 3.4 quantifies the implications of the cyclical worker flows for cyclical variation in productivity. Section 3.5 discusses implications for earnings. Section 3.6 concludes.

3.2 Data

A key contribution of our paper is the matching of U.S. Census Bureau linked employer-employee data to new productivity measures also developed at Census. We will first describe the linked employer-employee data, and how we use it to decompose firm growth via job-to-job moves versus flows through nonemployment. The Longitudinal Employer-Household Dynamics (LEHD) data contain quarterly earnings records collected by state unemployment insurance (UI) programs, linked to establishment-level data from the Quarterly Census of Employment and Wages (QCEW). LEHD employment coverage is quite broad, covering over 95 percent of private sector workers and almost all state and local government employment.[7] State-level data availability varies by year, as states began sharing UI and QCEW data with the Census Bureau at different times. In this paper we use LEHD data for private-sector employers in 28 states from 1998-2015.[8] Our 28 states include many of the largest states so that our sample accounts for 65 percent of U.S. private sector employment.

[7] For a full description of the LEHD data, see Abowd et al. (2009).

[8] Our 28 states are CA, FL, GA ,HI, ID, IL, IN, KS, ME, MD, MN, MO, MT, NC, NJ, ND, NM, NV, PA, OR, RI, SC, SD, TN, VA, WA, and WV. While we restrict our analysis to employers located in our 28-state sample, we use the complete set of available states to construct worker job histories. As described later in this section, our productivity measures reflect the labor productivity of the national firm.

The LEHD data allow us to decompose firm employment growth by worker hires and separations. We use the decomposition developed in Haltiwanger, Hyatt, and McEntarfer (2018) (HHM) and Haltiwanger, Hyatt, Kahn, and McEntarfer (2018) (HHKM) that yields an exact decomposition of firm employment growth due to workers switching jobs (what we call net job-to-job or net poaching flows) and growth due to flows between employment and nonemployment (what we call net nonemployment flows). A challenge for the identification of job-to-job flows in the LEHD data is that the data do not provide information on why a worker left one job and began another. We only have quarterly earnings, from which we infer approximately when workers left and began jobs. HHM and HHKM develop three alternative measures of job-to-job flows, and demonstrate that key findings on the nature of job ladders are robust to different approaches for identifying job-to-job moves in the LEHD data. We use the within/adjacent approach from HHM in this paper. This approach defines job-to-job transitions as those where the new job begins in the same or following quarter as the job separation. Based upon the robustness analysis in HHM, we are confident our main results are not sensitive to the specific rules we use amongst the set of rules they considered.[9]

To measure firm productivity, we use a relatively new firm-level database on productivity from Haltiwanger et al. (2017) based on the revenue and employment data from the Census Business Register and the Longitudinal Business Database (LBD).

[9]They also consider job-to-job flows restricted to those where the transition occurs within the same quarter and those with minimum disruptions in earnings. They find results that are very robust across these alternatives. Each of the different measures is highly correlated with the alternatives (pairwise correlations of about 0.98) and each of the LEHD based job-to-job flow series has a correlation of about 0.96 with CPS based job-to-job flows.

Since the underlying revenue and employment data are from the Census Business Register, this database offers much wider coverage of labor productivity at the firm level than earlier studies that focused on sectors like manufacturing or retail trade. These data allow us to measure the log of real revenue per employee on an annual basis for a wide coverage of the private, non-farm, for-profit firms. Revenue is deflated with the Gross Domestic Product price deflator. This measure of productivity is a standard gross output per worker measure of productivity that is commonly used to measure productivity at the micro and macro level but is a relatively crude measure compared to using total factor productivity (TFP). However, in the empirical literature, this revenue labor productivity measure has been shown to be highly correlated with TFP based measures of productivity across businesses within industries. That is, within detailed industry year cells, Foster, Haltiwanger, and Krizan (2001) and Foster, Haltiwanger, and Syverson (2008) find that the correlation between TFP and gross output (revenue) per worker across businesses is about 0.6 within industries in the manufacturing sector. This finding is consistent with the implications of models with labor market adjustment frictions which motivate our analysis.[10] In our analysis below, we use this revenue labor productivity measure deviated from industry by year means.

The gross output per worker data while offering much wider coverage than earlier studies has some limitations. The data only cover about 80 percent of firms in the Census LBD. The latter cover all firms with at least one paid employee in the private,

[10]See for example Decker et al. (2020). In their calibrated model of labor adjustment frictions, they obtain a correlation of TFP and revenue labor productivity of 0.90.

non-farm sector. One reason is that the revenue data are not available for non-profits. For another, the revenue data derive from different administrative sources than the payroll tax data. Most of the matches between the payroll tax and revenue data are via Employer Identification Numbers (EINs) but firms can use different EINs for filing income taxes and filing quarterly payroll taxes.[11] For such firms, name and address matching is required. Haltiwanger et al. (2017) also show that the missingness of revenue is only weakly related to industry, firm size, or firm age characteristics.[12] We are able to construct measures of labor productivity at the firm (operational control) level given that the Census Business Register has a complete mapping of all EINs owned by any given parent firm. Even with these limitations, we have revenue per worker for more than 4 million firms in each calendar year which we integrate with the LEHD data infrastructure via EINs. For the remaining private-sector employers in the LEHD data for which we cannot match to our productivity data, we impute labor productivity using the size, age, and relative wages paid by the employer within their industry.[13]

[11]Another source of mismatch is sole proprietors file income taxes on their individual income tax returns while payroll taxes are filed via their EIN. Administrative data are available that links the EINs to the filers via the SS-4 form (application for EINs). While this information is incorporated in the Census Business Register, it is imperfect.

[12]The productivity data explicitly excludes North American Industry Classification System (NAICS) 81 which is Other Services. This industry is very heterogeneous, including non-profits such as religious organizations where productivity is not well defined.

[13]The latest year for which we have firm productivity data is 2015, so we end our time series there although the LEHD data are more current. We investigated imputing post-2015 productivity using lagged productivity and other covariates but were not satisfied this 100 percent imputation was of sufficiently high quality. In unreported results, we have found that the patterns of worker flows are robust to excluding the imputed cases. Including the imputed cases facilitates our quantification of shares of employment at high and low productivity firms and in turn the productivity decomposition we use in the analysis.

3.2.1 Productivity, Growth, and Survival

Our measure of firm productivity–based on revenue per worker–exhibits a number of the key features that Syverson (2011) emphasized are common in the literature on firm productivity and dynamics. First, we find tremendous dispersion of revenue labor productivity within narrowly defined sectors. The within industry/year standard deviation of log real revenue per worker is about 0.80. This is in the range of labor productivity dispersion indices reported by Syverson (2004). Second, we find that log real revenue per worker is highly predictive of firm growth and survival, as shown in Table 3.1.[14] We consider two dependent variables for all incumbents in period t-1. The first dependent variable is the Davis, Haltiwanger, and Schuh (1996) firm level growth rate of employment that is inclusive of firm exit from t-1 to t.[15] The second dependent variable is an exit indicator that takes on the value of one if the firm exits between t-1 and t and is zero otherwise. We use a linear probability model for this second specification. Firm exit and growth is organic growth and exit in the manner defined by Haltiwanger, Jarmin, and Miranda (2013) (i.e., it abstracts from changes in ownership or M&A activity). We regress these two outcomes on log productivity in t-1 and on log size in t-1 (log of firm employment in t-1). While these are simple reduced form specifications, these specifications are consistent with standard models of firm growth and survival since these are proxies for the two key state variables for the firm in making growth and survival decisions. The canonical model implies that

[14] For this analysis, we do not restrict the sample to those firms in our LEHD data sample. These regressions use all firm-year observations from the revenue-enhanced Census Business Register.

[15] This measure is given by $g_{it} = (E_{it} - E_{it-1})/(0.5 * (E_{it} + E_{it-1}))$. It is a second order approximation to a log first difference that accommodates entry and exit.

holding initial size constant a firm with higher productivity is more likely to grow and less likely to exit. We find overwhelming evidence in support of these predictions in Table 3.1. A one standard deviation increase in within-industry productivity yields a 20 percentage point increase in net employment growth and 5 percentage point decrease in the likelihood of exit.[16]

These descriptive results give us confidence to proceed with our measure of revenue labor productivity since we produce patterns that others have found using TFP measures in sectors such as manufacturing. In line with the existing literature, our findings of a tight relationship between firm productivity, growth, and survival are consistent with the hypothesis that there are intrinsic differences in productivity across firms that help account for the high rate of reallocation of jobs across firms. In addition, such intrinsic differences in productivity have implications for worker reallocation including the potential role of a productivity job ladder.

3.2.2 Defining High- and Low-Productivity Firms

To help mitigate remaining concerns about measurement error, we use robust measures of the ranking of firms by productivity. We construct time-invariant measures of productivity, defined as the employment-weighted average of firm productivity over the life of the SEIN (the state tax identifier number, the key employer identifier in LEHD data). This approach is broadly consistent with the rank preserving equilibria assumption in the Moscarini and Postel-Vinay (2013) framework. We then compute

[16]Decker et al. (2020) develop a simple model of firm dynamics with adjustment frictions that shows that the relationship between growth and survival from t-1 to t with realizations of labor productivity in period t is very similar as with TFP, holding firm size constant in t-1.

Table 3.1: Productivity, Employment Growth, and Firm Death

	Employment Growth Rate (1)	Firm Death (2)
Productivity	0.216	-0.066
	(0.00011)	(0.00005)
Log of firm size	0.056	-0.045
	(0.00006)	(0.00002)

Notes: Each column presents estimates from a separate regression. The dependent variable in column 1 is the employment growth rate between the current and subsequent year and the dependent variable in column 2 is an indicator equal to one of the firm dies in the subsequent year. The independent variables in each regression are the log of firm size and productivity, which is defined as the log revenue per worked deviated from the industry (defined by the 4-digit NAICS code) average. Standard errors are presented in parentheses.

the employment-weighted and within-industry quintiles of the productivity distribution. Using these quintiles, we define high-productivity firms as those in the top two quintiles and low-productivity firms as those in the bottom three quintiles. In unreported analysis, we have found that results are robust to permitting firms to change ranks over time. This robustness is not surprising given the large differences between high- and low-productivity firms. For example, the within-industry differences in average gross output per worker between high- and low-productivity firms are typically in excess of 85 log points.

As a robustness check on our productivity measure, we also rank firms by the AKM firm fixed effect and average earnings. To construct this measure we estimate an AKM specification by regressing log earnings on person fixed effects, firm fixed effects, and controls for time and worker age.[17] To solve this model, we implement the iterative method proposed by Guimaraes and Portugal (2010). The AKM firm fixed effect abstracts from observable and unobservable individual characteristics and in canonical models, the firm specific pay premia should be closely related to productivity differences across firms. As an additional robustness check, we consider simple non-parametric measures of relative earnings by ranking firms based on the average earnings of full-quarter workers within their industry. We classify firms into two groups, high- and low-ranked, based on the AKM firm effect and average earnings. Just as we do for gross output per worker, we construct employment-weighted within-industry quintiles based on these measures and define high-ranked firms as

[17]To control for time we include a set of year dummies that capture calendar year effects on earnings. To control for worker age, we follow the specification of Card, Cardoso, and Kline (2016). We center age around 40, include a quadratic and cubic transformation of worker age, but omit the linear term.

those in the top two quintiles and low-ranked firms as those in the bottom three quintiles. Particularly for the AKM firm premia rankings, we interpret the results using these rankings as providing an alternative indirect method of ranking firms by their productivity.

3.3 Worker Flows Over the Business Cylcle

We begin by examining how job-to-job moves and worker flows through nonemployment reallocate workers across high- and low-productivity firms. To understand how worker reallocation moves workers from one group of firms to another, we use the following identity:

$$\text{Net Job Flows} = H_t - S_t = \sum_{i\in\{p,n\}} (H_t^i - S_t^i) = \sum_{i\in\{p,n\}} \sum_{j\in\{l,h\}} (H_t^{ij} - S_t^{ij}) \tag{3.1}$$

where H_t is the number of hires and S_t is the number of separations in quarter t. The superscripts denote subsamples defined by the type of worker flow, where $i = p$ denotes poaching (job-to-job) flows and $i = n$ denotes flows through nonemployment, and the type of firm, where $j = h$ denotes high-productivity firms and $j = l$ denotes low-productivity firms.[18] For example, H_t^p denotes poaching hires and H_t^{ph} denotes poaching hires at high productivity firms. We convert all flows to rates by dividing through by employment in time $t-1$.[19] All of the aggregate series we use in this

[18] A given type of firm (e.g., high-productivity) may have workers that are hired by that firm via a job-to-job flow and separate from that firm via a job-to-job flow. We refer to the former as a poaching hire and the latter as a poaching separation.

[19] Hires and separations characterize worker mobility between time $t-1$ and t. For hires we count all worker flows into a firm in our sample at time t. For separations, we count all worker flows out of a firm in our sample at time $t-1$.

section have been seasonally adjusted using the X-12 procedure.

The identity in equation 3.1 decomposes employment growth at high-productivity firms (H_t^h-S_t^h) into net growth due to two components: job-to-job moves of workers or poaching flows (H_t^{ph}-S_t^{ph}) and flows of workers through nonemployment (H_t^{nh}-S_t^{nh}).[20] In the aggregate economy, employment growth is entirely attributable to net worker flows through nonemployment since poaching hires and poaching separations aggregated over both high- and low-productivity firms are equal. However, for any subset of firms in the economy, net poaching need not be zero, as some firms will be more successful poaching workers away from other employers. This "net poaching flows" component of growth captures the comparative growth advantage one group of firms has over another in their ability to attract workers away from other firms.

Figure 3.1 shows our decomposition of net job flows for high- and low-productivity firms. As discussed previously, a key prediction of job ladder models is that job-to-job moves should reallocate workers away from less productive to more productive firms. Figure 3.1(A) shows that this prediction from the theory holds true in the data. The most productive firms have overall positive net employment growth on average and net poaching (H_t^{ph} - S_t^{ph}) is strongly positive. The average net employment growth of high-productivity firms is 0.33 percent per quarter with net poaching (the rate at which job-to-job moves reallocate workers to high-productivity firms) averaging 0.27 percent per quarter. In other words, during the 1998-2015 period, job-to-job moves of workers from less-productive employers account for most (80 percent) of the net employment growth of high-productivity firms.

[20] Correspondingly for low-productivity firms, H_t^l-S_t^l = (H_t^{pl}-S_t^{pl}) + (H_t^{nl}-S_t^{nl}).

Figure 3.1: Poaching and Nonemployment Flows by Firm Productivity

(A) High-Productivity Firms

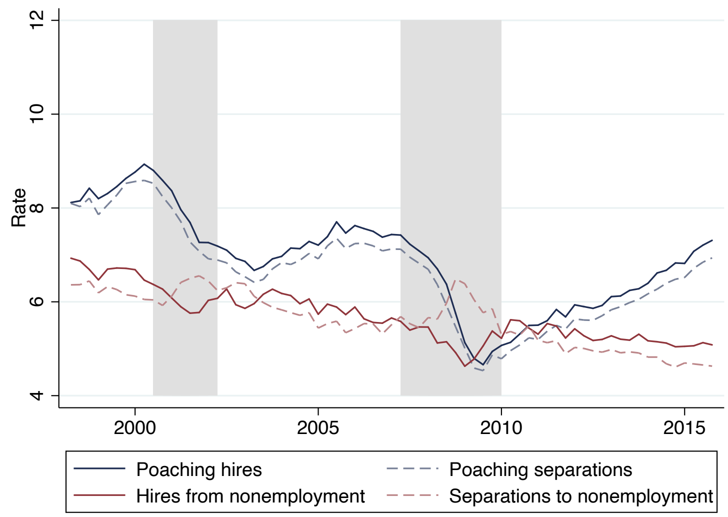

(B) Low-Productivity Firms

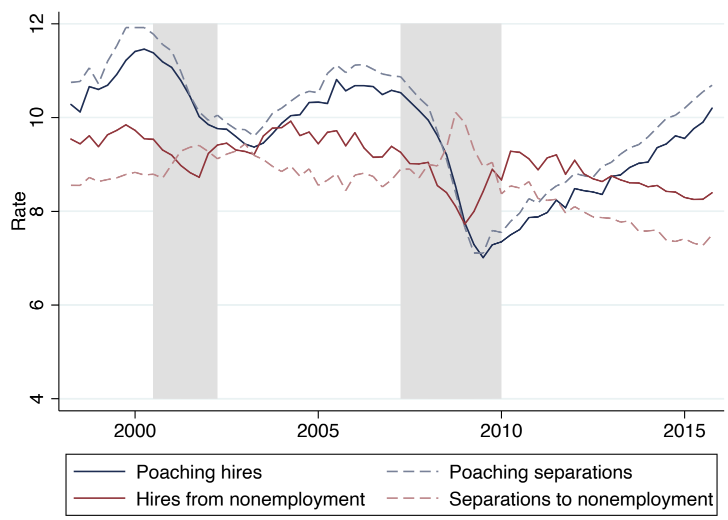

Notes: High-productivity indicates that the firm is in the top two quintiles of the within-industry productivity distribution. Low-productivity indicates the firm is the bottom three quintiles of the within-industry productivity distribution. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

The results of the decomposition are also striking for the less productive firms in the industry. In Figure 3.1(B), low-productivity firms grew at a rate of 0.14 percent per quarter on average from 1998-2015, which is slower than the high-productivity firms. Low-productivity firms lose -0.34 percent employment per quarter from workers "voting with their feet" and moving to firms ranked higher in firm productivity distribution. The positive growth rate for less productive firms is entirely due to strong hiring from nonemployment. In other words, in a typical quarter less productive firms recruit from the pool of unemployed individuals to replace workers moving to better firms. This is also consistent with job ladder models of the labor market. In job ladder models, it is the search and matching frictions that support the presence of low-productivity firms that primarily hire from nonemployment.

The patterns of hires and separations in Figure 1 are instructive for understanding the differences in the cyclical dynamics of job-to-job and nonemployment worker flows. Poaching hires and separations both decline for high-productivity firms in contractions with the decline in poaching hires larger so that net poaching declines significantly. Hires from nonemployment decline sharply for low-productivity firms in contractions accompanied by a surge in separations so that net employment growth declines sharply for low-productivity firms. There are similar qualitative patterns for hires from and separations to nonemployment for high-productivity firms but the magnitudes are smaller.

To more clearly see how worker flows reallocate workers across the productivity ladder, we decompose the average overall net job flow differential between high- and low-productivity groups into the net poaching differential and the net flows from nonemployment differential. Let E^j_{t-1} denote employment at firm type j at time $t-1$. Then $\lambda^i_t = (H^{pj}_t - S^{pj}_t)/E^j_{t-1}$ and $\delta^j_t = (H^{nj}_t - S^{nj}_t)/E^j_{t-1}$ are the employment growth rates at firm type $i \in \{l, h\}$ through net poaching flows and net nonemployment flows, respectively. Figure 3.2 plots the differential rates between high- and low-productivity firms–i.e., $\lambda^h_t - \lambda^l_t$ and $\delta^h_t - \delta^l_t$. The average net poaching differential between high- and low-productivity firms is 0.61 percent per quarter. It is also quite cyclical: a minimum of 0.012 in 2009Q1, and a maximum of 1.01 in 2015Q3. The average net nonemployment differential is -0.41 percent, but this increase at the onset of economic downturns.[21] The negative contribution of flows through nonemployment to the differential growth rates of high- and low-productivity firms implies that search frictions are a drag on productivity-enhancing reallocation in expansions, allowing mediocre firms to attract workers flowing through nonemployment who cannot immediately find better jobs.

For comparison purposes, we also show in Figure 3.2 the differential growth rates for firms ranked by AKM firm fixed effects and average earnings. To make results comparable, the rankings are calculated within industries. These patterns are very similar to those from the job ladder decomposition of growth rates when firms are ranked by productivity, both in levels and business cycle dynamics. The striking similarity between these three measures is again consistent with job ladder models of the labor market, where high-productivity firms offer higher wages, and grow faster than less-productive firms in expansions by poaching workers away from less-productive,

[21]Not shown in Figure 2 is the overall net job flow differential between high and low productivity firms, which is the sum of the poaching and nonemployment margins, and averages 0.20 percent per quarter. It reaches a maximum of 0.58 in 2008Q4 and a minimum of -0.13 in 2004, an expansion year. In the set of regressions in Table 2 we will test the cyclicality of this overall redistribution.

Figure 3.2: Differential Flows between High- and Low-Ranked Firms

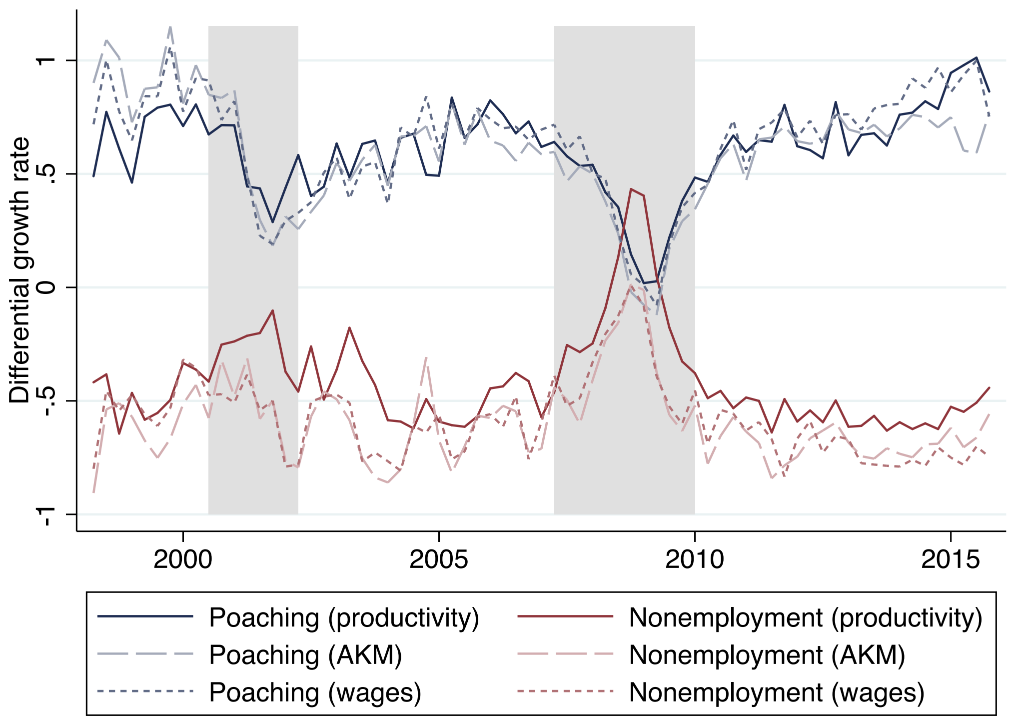

Notes: Differential growth rates are the difference in quarterly employment growth rates between firms in the high category (top two quintiles) and those in the low category (bottom three quintiles). The different series present results in which the high and low categories are defined by productivity, wages (earnings per worker), and the AKM firm effect. For all measures, the quintiles are calculated within NAICS 4-digit industry codes. Results are presented separately for poaching and nonemployment flows. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

lower-paying firms. The strong similarity between the productivity and AKM fixed effects results also provides reassurance about measurement error concerns about the revenue productivity measure. The latter is a relatively crude measure of productivity so the finding that using AKM firm fixed effects yields the same patterns reduces concerns about the use of this measure.

Figure 3.2 also shows pronounced cyclical patterns that differ across the components of net job flows. We quantify the nature of that variation in Table 3.2. Table 3.2 presents the results from regressions where each component of the differential growth rate (net job flows, net poaching flows, and net nonemployment flows) is the regressed on a cyclical indicator and a time-trend. Each column represents the results of these regressions for firms when they are ranked by productivity, AKM firm fixed effects, and average firm earnings per worker. Because the results for all three columns are very similar (as suggested by Figure 3.2) we will focus our discussion on the business cycle dynamics of worker reallocation across high- and low-productivity firms (column 1). Two cyclical indicators are used: the change in unemployment and unemployment deviated from a Hodrick-Prescott (HP) trend. Periods of rising unemployment correspond closely to NBER defined recessions. In contrast, especially over our sample period, unemployment remains substantially above trend well into NBER defined recoveries.

We start by focusing on the net poaching differentials (row B) which shows that net poaching from low- to high-productivity firms decreases in cyclical downturns. This occurs in recessions: for every one percentage point increase in the change in the unemployment rate, high vs. low differential net poaching declines by 0.440

Table 3.2: Differential Net Job Flows Over the Cycle

	High-Low Differential		
	(1)	(2)	(3)
A. Net Job Flows			
Change in unemployment rate	0.136	-0.179	-0.170
	(0.060)	(0.084)	(0.084)
Deviated unemployment rate	-0.086	-0.161	-0.144
	(0.023)	(0.029)	(0.030)
B. Poaching Job Flows			
Change in unemployment rate	-0.440	-0.544	-0.560
	(0.049)	(0.063)	(0.061)
Deviated unemployment rate	-0.097	-0.128	-0.103
	(0.027)	(0.033)	(0.034)
C. Nonemployment Job Flows			
Change in unemployment rate	0.576	0.365	0.390
	(0.043)	(0.047)	(0.049)
Deviated unemployment rate	0.011	-0.033	-0.041
	(0.033)	(0.026)	(0.027)
Definition of Job Ladder			
Firms ranked by	productivity	wages	AKM firm effect
Firms ranked within	industry	industry	industry

Notes: Each cell presents results from a separate regression estimated on national quarterly data. The dependent variable is the differential worker flow rate between firms on the high and long rung of the job ladder, where the type of flow is indicated by the panel label. The job ladder is defined by productivity, wages, and the AKM firm effect in columns 1, 2, and 3, respectively. The independent variables in each regression include a cyclical indicator as well as a linear time-trend and a constant, which are not reported. The cyclical indicators considered include the change in the unemployment rate and the deviations of unemployment from the Hodrick-Prescott Trend. Aside from the linear trend, all dependent and independent variables are measured in percentage point units. Standard errors are presented in parentheses.

percentage points. Poaching flows are also low in the times of high unemployment that follow recessions. For every percentage point that the unemployment rate is above its HP trend, the net poaching differential is lower by 0.097 percentage points. These findings are consistent with a sullying effect of recessions.

There is also evidence of a cleansing effect that works through flows to and from nonemployment (row C). The net flows from nonemployment differentials provide an indication of the reallocation of employment from low- to high-productivity firms that involves transitions to and from nonemployment. These transitions inherently involve intervening spells of nonemployment. In that respect, this type of reallocation is more costly than job-to-job flows since it involves the time and resource costs of nonemployment. Row C indicates that the reallocation that works through the nonemployment margin is countercyclical. When the change in the unemployment rate increases by one percentage point, differential net nonemployment hiring for high vs. low increases by 0.576 percentage points. This is a cleansing effect of recessions that is working in the opposite direction of job-to-job flows and primarily occurs because there is a disproportionately large spike in separations into nonemployment from low-productivity firms.[22] This cleansing effect, however, is almost negligible in the times of high unemployment that follow recessions. An additional percentage

[22] The differential net flows from nonemployment ($\delta_t^h - \delta_t^l$) can be decomposed into the hires component ($H_t^{nh}/E_{t-1}^h - H_t^{nl}/E_{t-1}^l$) minus the separations component ($S_t^{nh}/E_{t-1}^h - S_t^{nl}/E_{t-1}^l$). Regressing these components against the change in the unemployment rate and controlling for the time-trend produces a point estimate of 0.165 for hires and -0.411 for separations (the difference between these coefficients is 0.576, which is the coefficient found in Table 3.2). Thus, when the unemployment rate rises, worker reallocation from low- to high-productivity firms through nonemployment increases primarily because of a disproportionate spike in separations to nonemployment from low-productivity firms but also because a disproportionate decline in the rate at which low-productivity firms hire jobless workers.

point of the unemployment rate above its HP trend is associated with only an increase of 0.011 in the net nonemployment differential.

The coefficient on overall net job flow differentials (row A) is determined by these cleansing and sullying effects. During recessions, cleansing effects are stronger than sullying effects. A one percentage point increase in the change in the unemployment rate is associated with an increase in the relative employment growth of high-productivity firms of 0.136 percentage points. In the times of high unemployment that follow recessions, cleansing effects are small and so sullying effects dominate. An additional percentage point of the unemployment rate above its HP trend is associated with an increase in the relative employment growth of low-productivity firms of 0.086 percentage points. Thus, Table 3.2 illustrates that the relative importance of cleansing and sullying effects varies at different phases of the business cycle.

3.4 Implications for Aggregate Outcomes

What are the implications of the business cycle dynamics described in the previous section on aggregate productivity growth? The decline in productivity-enhancing reallocation through job-to-job moves in slack labor markets should be a drag in productivity growth, while higher rates of job destruction at less-productive firms in downturns ought to free resources for more productive use. The magnitude of the effect on aggregate productivity growth will depend on both the size of the differential employment flows as well as the productivity differential between high- and low-productivity firms. In this section we formalize this intuition and implement a decomposition exercise to quantify how worker reallocation through poaching and nonemployment flows contributes to aggregate productivity growth, and how these components of productivity growth vary over the business cycle.

3.4.1 Worker Reallocation and Employment Shares

We begin by focusing on how worker reallocation affects the share of workers at high- and low-productivity firms. We focus initially on the findings using our gross output per worker measure of productivity.[23] We use the following identity to write changes in the share of employment at high-productivity firms as a function of the differential net poaching rates ($\lambda_t^h - \lambda_t^l$) and differential net nonemployment rates ($\delta_t^h - \delta_t^l$),

$$\Delta\theta_t^h = \tilde{\lambda}_t^h + \tilde{\delta}_t^h + \tilde{\epsilon}_t^h \tag{3.2}$$

where $\tilde{x}^h = (x^h - x^l)\theta_{t-1}^h\theta_{t-1}^l(E_{t-1}/E_t)$ for $x \in \lambda, \delta$ and $\Delta\theta_t^h$ is the change in the share of employment at high-productivity firms between quarter t and $t-1$. This expression shows that the sign of differential net poaching rate, $\lambda_t^h - \lambda_t^l$, determines whether poaching rates will increase or decrease the share of employment at high-productivity firms. The magnitude of this effect also depends on the share of workers at high productivity firms as well as the growth in overall employment. See Appendix C.2 for details.

[23] As we have discussed, this is a relative measure for firms within industries. Using a relative measure within industries somewhat complicates the interpretation of the accounting decomposition we develop and analyze below. Later in the analysis we consider the AKM firm premia measure which overcomes these limitations.

There are two reasons for the existence of the residual term, $\tilde{\epsilon}^h$.[24] First, some workers may move to or from an employer located in a state outside of our 28-state sample. In contrast to the results from the previous section and because we aim to implement an exact decomposition, the counts of hires and separations in this section only include worker flows where both the origin and destination employers are in one of the 28 states in our sample. Second, the administrative code that identifies the employer, the SEIN, can change over time and create a spurious flow of workers between the old and new SEIN. We are able to flag when these changes occur and omit these flows from the poaching and nonemployment flows. However, there is no straightforward way to account for this issue when measuring productivity. Thus, a change in an SEIN could lead to a change in the share of workers at high productivity firms but have no corresponding flow of workers. In unreported results, we directly measure flows of workers in and out of the states in our sample and show that the residual term is primarily attributable to changes in the SEIN over time, not migration in and out of the sample.

Regardless of the source, these residuals flows are both small in magnitude and do not exhibit a clear pattern across the business cycle. Specifically, the average size of the differential net poaching and nonemployment growth rates are three and five times as large as the differential residual flows, respectively. Figure C.1 in Appendix C.1 presents a version of Figure 3.2 that contains the residual flows as well as the poaching and nonemployment flows constructed with and without the restriction that both

[24]Empirically, we measure the residual term as the difference between the observed changes in employment at high- and low-productivity firms and the changes predicted by the poaching and nonemployment flows.

origin and destination employers are in the 28-state sample. The results indicate that the residual flows do not exhibit any notable cyclicality across the business cycle and the differential net poaching and nonemployment growth rates that exclude workers moving in and out of our 28-state sample are very similar (in levels and movements across time) to the results from Figure 3.2. We infer that this residual term is not important for the main results discussed in this section.

Figure 3.3 shows the time series of the main components of the decomposition in equation 3.2 and illustrates how worker reallocation affects the percent of employment at high-productivity firms.[25] Figure 3.3(A) presents the observed changes in the percent of employment at high-productivity firms between the current and previous quarters ($\Delta\theta_t^h$). Figure 3.3(B) presents the components of these changes that are attributable to worker reallocation through poaching flows ($\tilde{\lambda}^h$) and nonemployment flows ($\tilde{\delta}^h$). During expansions, worker flows through nonemployment lead to a reduction in the share of workers at high-productivity firms whereas poaching flows lead to an increase in the share of workers at high-productivity firms. At the onset of a recession, the rate at which nonemployment flows contributes to the growth of employment share of low-productivity firms slows. Indeed, in the Great Recession this change was large enough such that nonemployment flows briefly contributed positively to the growth of the employment share at high-productivity firms. Throughout our sample period, poaching flows always contribute positively to the growth of employment at high productivity firms, but this largely collapses during recessions, particularly in the Great Recession. Consistent with our earlier analysis, the results highlight the

[25] To make the figure more readable, we present the results in percentage point terms.

staggered nature of the timing of these two effects. The cleansing effect begins at the onset of a recession where as the sullying effect starts and peaks later on.

3.4.2 Worker Reallocation and Productivity

We now quantify how changes in the share of workers at high-productivity firms affects an index of aggregate productivity growth using the productivity differential between the two groups of firms. Our analysis focuses on a measure of productivity that is the employment-weighted average of a firm-level measure of productivity, which is measured in logs. Such accounting indices of productivity have been widely used in the literature to quantify the contribution of reallocation effects to productivity (e.g., Olley and Pakes, 1996; Foster, Haltiwanger, and Krizan, 2001; and Melitz and Polanec, 2015). As we discuss in Appendix C.1, an aggregate index based on employment-weighted firm-level labor productivity indices track official statistics from the U.S. Bureau of Labor Statistics (BLS) quite well.[26]

Let, R_t denote the employment-weighted average of the firm-level measure of log revenue per worker in quarter t. Then, $R_t = \theta_t^l R_t^l + \theta_t^h R_t^h$, where R_t^i denotes the

[26]See also Figure A.1 in Decker et al. (2017). Conceptually, this aggregate index is consistent with aggregate productivity in a structural model with a single input (labor), constant returns to scale, and perfect competition in product markets. While these are strong assumptions (although not inconsistent with job ladder models), as noted such indices track official statistics closely. Much of the literature that focuses on misallocation specifies curvature in the revenue function so that there is a well-defined size distribution even in the absence of distortions. Part of the reason for this is that this enables a measure of allocative efficiency that is relative to a frictionless/distortionless benchmark (e.g., Hsieh and Klenow, 2009; Bartelsman et al., 2013; and Blackwood et al., forthcoming). In principle, such curvature is not necessary in models with adjustment frictions such as search and matching frictions. Models with curvature in the revenue function have the property that it is not optimal to allocate all resources to the most productive firm. While this implies caution in using weighted average measures of firm-level productivity in quantitative analysis of models with such curvature, Decker et al. (2020) show that in models with adjustment frictions that this type of aggregate productivity index tracks structural measures of true productivity well even if there is curvature in the revenue function.

Figure 3.3: Changes in Percent of Employment at High-Productivity Firms

(A) Observed Changes

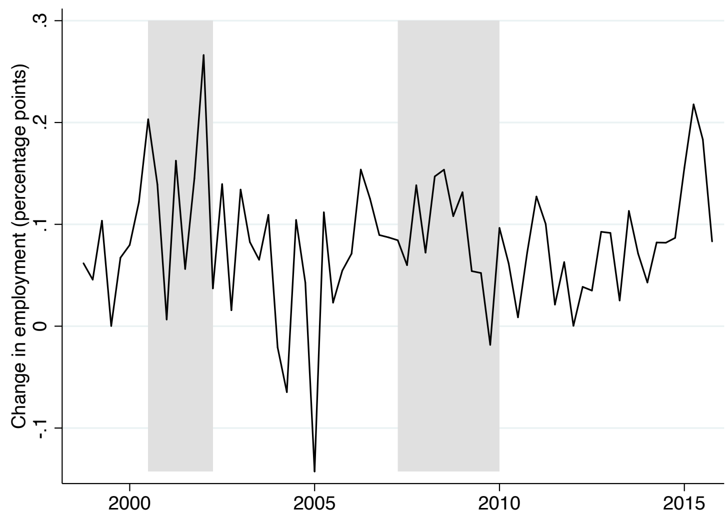

(B) Contribution of Worker Reallocation

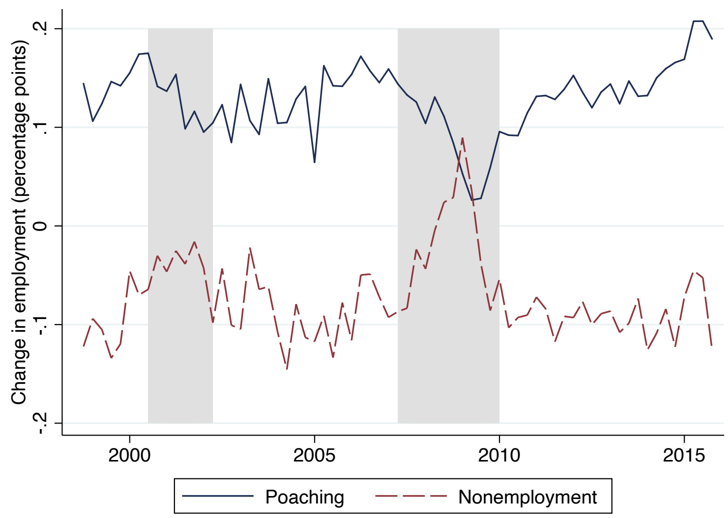

Notes: Panel (A) presents the first difference of the percent of employment at high-productivity firms. Panel (B) presents the change in the percent of employment at high-productivity firms that is attributable to worker reallocation through poaching and nonemployment flows. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

employment-weighted average of the firm-level measure of log revenue per worker within firm type i in quarter t and θ_t^i is the share of employment at firm type i in quarter t.[27] An increase in θ_t^h leads to an increase in R_t but the interpretation of a change in R_t is complicated by the fact that log revenue per worker is a measure of productivity that is not readily comparable across industries. Let P_t denote an unobserved measure of productivity that is the weighted average of a firm-level measure of productivity, which is comparable across industries.

Our goal is to use the observed measure of log revenue per worker in order to quantify how changes in θ_t^h affect P_t. We achieve this goal by making the following two assumptions:

A1. $P_t^i(k) = R_t^i(k) + U_t(k)$, where $R_t^i(k)$ denotes the employment-weighted average of the firm-level measure of log revenue per worker within industry k, firm type i, and quarter t; $U_t(k)$ is an unobserved term that is constant within industry and quarter; and $P_t^i(k)$ is the employment weighted average of the unobserved measure of productivity within industry k, firm type i, and quarter t.[28]

A2. $\Delta\big(cov\big(\theta_t^i(k), \tilde{R}_t^i(k)\big)\big) = 0$, where $\theta_t^i(k)$ is the share of employment at firm type i within industry k and quarter t, $\tilde{R}_t^i(k) \equiv R_t^i(k) - (R_t^l(k) + R_t^h(k))/2$, and Δ denotes the difference between quarter t and $t-1$.

Assumption A1 states that, up to an additive term that is constant within industry and quarter, log revenue per worker is a measure of productivity that is comparable

[27] While the firm type (high- or low-productivity) is time invariant, the firm-level measures of productivity used to construct R_t^i are measured in time t.

[28] Recall that $P_t^i(k)$ is a measure of productivity that is directly comparable across industries.

across industries. Assumption A2 states that the covariance between the share of employment at high-productivity firms and the dispersion of log revenue per worker does not change over time. Appendix C.2 presents empirical evidence that supports the plausibility of assumptions A1 and A2.

Assumptions A1 and A2 allow us to isolate productivity growth that arises from worker reallocation between high- and low-productivity firms. By an accounting identity we can rewrite the unobserved measure of aggregate productivity as an employment-weighted average of firm type by industry averages, $P_t = \sum_k \theta_t(k)[\theta_t^l(k)P_t^l(k) + \theta_t^h(k)P_t^h(k)]$. Combining this accounting identity with assumptions A1 and A2 implies the following expression for this index of aggregate productivity growth,

$$\Delta P_t = \underbrace{(\tilde{R}_{t-1}^h - \tilde{R}_{t-1}^l)\Delta\theta_t^h}_{\text{Worker Reallocation}} + \theta_t^l \Delta \tilde{R}_t^l + \theta_t^h \Delta \tilde{R}_t^h + \Delta\left(\sum_k [\theta_t(k)\bar{P}_t(k)]\right) \tag{3.3}$$

where $\bar{P}_t(k) \equiv [P_t^l(k) + P_t^h(k)]/2$. The first term, $(\tilde{R}_{t-1}^h - \tilde{R}_{t-1}^l)\Delta\theta_t^h$, is the component of productivity growth attributable to worker reallocation between high- and low-productivity firms.[29] The remaining terms capture other aspects of productivity growth, which include productivity growth driven by firm-level innovations unrelated to worker reallocation. See Appendix C.2 for a details.

We can further decompose the component of productivity growth attributable to worker reallocation between high- and low-productivity firms into the components attributable to poaching, nonemployment, and residuals flows. Specifically, combining

[29]While our classification of high- and low-productivity firms is based on a within-industry ranking of log revenue per worker, our worker flows data are not restricted to within-industry job flows.

equation 3.2 with the first term in equation 3.3, yields,

$$\underbrace{(\tilde{R}^h_{t-1} - \tilde{R}^l_{t-1})\Delta\theta^h_t}_{\text{Worker Reallocation}} = \underbrace{(\tilde{R}^h_{t-1} - \tilde{R}^l_{t-1})\tilde{\lambda}^h_t}_{\text{Poaching}} + \underbrace{(\tilde{R}^h_{t-1} - \tilde{R}^l_{t-1})\tilde{\delta}^h_t}_{\text{Nonemployment}} + \underbrace{(\tilde{R}^h_{t-1} - \tilde{R}^l_{t-1})\tilde{\epsilon}^h_t}_{\text{Residual}} \tag{3.4}$$

An increase in the share of workers at high-productivity firms increases the index of aggregate productivity, and the magnitude of the effect is determined by the productivity differential between high- and low-productivity firms. The analysis in this section uses equation 3.4 to decompose productivity growth attributable to worker reallocation (between high- and low-productivity firms) through poaching and nonemployment flows.

Figure 3.4(A) presents the decomposition of productivity growth into components attributable to poaching and nonemployment flows and shows clear evidence of the cleansing and sullying effects of recessions. On average, worker reallocation through poaching flows contributes 0.1 log points to overall productivity growth each quarter (all statistics on productivity changes are quarterly and have not been annualized). This is a substantial contribution to the overall quarterly average rate of productivity growth of 0.33 when aggregating our micro data.[30] However, during recessions there is clear evidence of a sullying effect. In 2006:1 the poaching contribution is 0.13 log points but this declines to 0.02 by 2009:2. In contrast, worker reallocation through nonemployment tends to be a drag on productivity growth, on average, decreasing productivity by 0.67 log points each quarter.[31] However, during recessions there is

[30]As discussed in Appendix C.1, the comparable average quarterly rate using published BLS statistics is 0.43. The aggregate index from our data and from published BLS statistics are highly correlated (0.85).

[31]This drag on productivity is consistent with job ladder models with search and matching fric-

evidence of a cleansing effect since during those times nonemployment flows yield declines in the employment share of low-productivity firms. In 2006:1, the nonemployment component is -0.1 log points but increases to 0.08 in 2009:1. The figure illustrates the staggered nature of these effects in which the cleansing occurs at the outset of the recession–when unemployment rate is rising most rapidly–and the sullying effect peaks relatively further on into the downturn–which the unemployment rate is highest. In addition, the sullying effect lingers well into the recovery.

Figure 3.4(B) presents the combined effect of worker reallocation through poaching and nonemployment flows. During expansions, the total effect of worker reallocation through poaching and nonemployment flows contributes 0.04 log points to overall productivity growth each quarter. In recessions, this increases to 0.08. While this might suggest that cleansing effects dominate, these calculations neglect the fact that the sullying effect lingers well into the recovery. Using an alternative indicator of the cycle which is whether the unemployment rate is above or below trend (using the Hodrick-Prescott filter), we find that the combined effect is 0.07 on average for quarters where the unemployment rate is below trend and 0.03 on average for quarters where unemployment is above trend. This reversal is driven by much larger contributions of poaching flows during periods of low unemployment (0.13) compared to high unemployment (0.09).

Table 3.3 summarizes these patterns. Here we show results from a set of regressions where different components of productivity growth are regressed on a cyclical indicator and a time-trend. We use two alternative cyclical indicators: changes in the

tions. Dispersion in productivity is supported by such frictions in equilibrium.

Figure 3.4: Decomposition of Growth in Productivity Over the Cycle

(A) Poaching and Nonemployment Flows

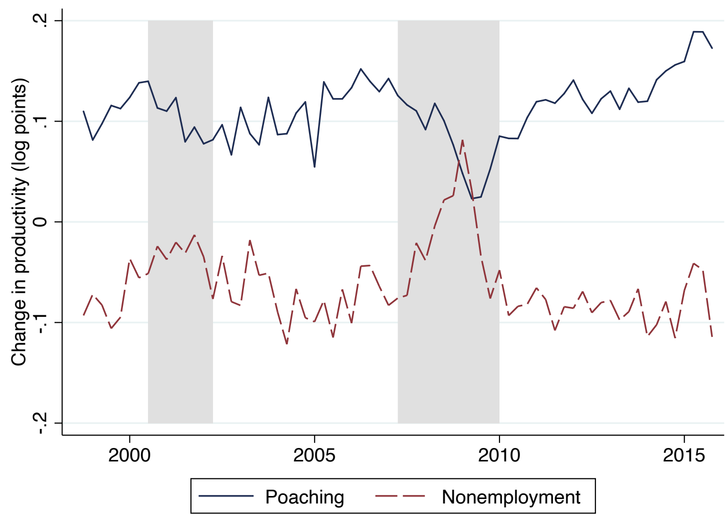

(B) Combined Effect

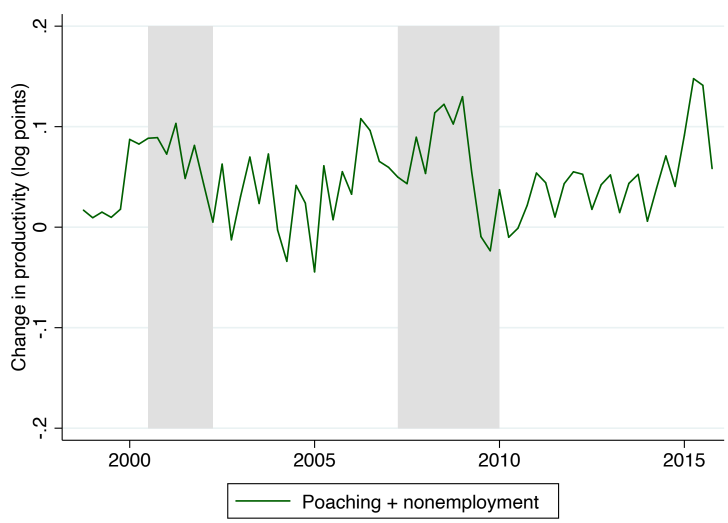

Notes: Panel (A) presents the components of quarterly productivity growth that are attributable to worker reallocation between high- and low-productivity firms through poaching and nonemployment flows and Panel (B) presents the sum of these two components. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

unemployment rate and deviations of the unemployment rate from trend. Column 1(B) shows that productivity-enhancing reallocation through job-to-job moves is procyclical using both measures.[32] Column 1(C) shows that productivity-enhancing reallocation through worker flows through nonemployment is counter-cyclical using both measures. The net effect in Column 1(A), however, does depends on the cyclical indicator. Sullying effects dominate when cyclicality is measured using deviations from the unemployment rate: this is because it takes a while for productivity-enhancing reallocation from job-to-job moves to recover in expansions. In contrast, cleansing effects dominate when cyclicality is measured using changes in the unemployment rate. Taken together, these results imply that the cleansing effect peaks earlier in a downturn compared to the collapse of the job ladder that lingers into the early stages of a recovery. These results also suggest that slow labor market recoveries will be generally more damaging to productivity growth than V-shaped recoveries as slow recoveries exhibit an accompanying slow recovery of job-to-job flows.

3.4.3 Robustness to Using the AKM Firm Premium

We assess the robustness of our results to using an alternative measure of firm performance: the AKM firm premium. While the revenue per worker measure has many strengths, it is not easily comparable across sectors and may partially reflect the sorting of workers across firms in addition to innate differences in firm productivity. The AKM firm premium, which is an alternative proxy for firm performance, is not subject to these same limitations and therefore provides a useful comparison.

[32] We discuss columns 2 and 3 of this table in the next section.

Table 3.3: Productivity Growth from Job Flows Over the Cycle

	Productivity Growth		
	(1)	(2)	(3)
A. Net Job Flows			
Change in unemployment rate	0.040	0.010	0.005
	(0.015)	(0.016)	(0.043)
Deviated unemployment rate	-0.024	-0.026	-0.084
	(0.005)	(0.006)	(0.014)
B. Poaching Job Flows			
Change in unemployment rate	-0.052	-0.062	-0.225
	(0.010)	(0.009)	(0.029)
Deviated unemployment rate	-0.024	-0.021	-0.073
	(0.004)	(0.004)	(0.013)
C. Nonemployment Job Flows			
Change in unemployment rate	0.092	0.072	0.230
	(0.009)	(0.010)	(0.026)
Deviated unemployment rate	0.000	-0.005	-0.012
	(0.006)	(0.005)	(0.015)
Definition of Job Ladder			
Firms ranked by	productivity	AKM firm effect	AKM firm effect
Firms ranked within	industry	industry	full sample

Notes: Each cell presents results from a separate regression estimated on national quarterly data. The dependent variable is the growth in productivity or AKM firm fixed effect as noted by the second to last row. The independent variables in each regression include a cyclical indicator as well as a linear time-trend and a constant, which are not reported. The cyclical indicators considered include the change in the unemployment rate and the deviations of unemployment from the Hodrick-Prescott Trend. Aside from the linear trend, all dependent and independent variables are measured in percentage point units. Standard errors are presented in parentheses.

Figure 3.5 decomposes growth in the aggregate firm premium attributable to worker reallocation across the firm premium ladder through poaching and nonemployment flows. Note that firms are re-ranked into high- and low-premium firms and the worker flows (i.e., $\tilde{\lambda}_t^h$ and $\tilde{\delta}_t^h$) and the differentials in the firm premium (i.e., $\tilde{R}_{t-1}^h - \tilde{R}_{t-1}^l$) are re-calculated to reflect these new rankings and the alternative measure of firm performance. Figure 3.5(A) follows the methodology described in equation 3.4 but uses the AKM firm effect as the measure of productivity instead of revenue per worker. Importantly, firms are ranked within 4-digit NAICS industry codes. Unlike the gross output per worker measure, the AKM firm effect is directly comparable across industries. Thus, Figure 3.5(B) also uses the methodology described in equation 3.4 but, in addition to using the AKM firm effect as the measure of productivity, treats all firms as being part of the same industry. In effect, Figure 3.5(B) ranks firms in the pooled sample and accounts for worker reallocation both within and across industries.

Regardless of the methodology used, both series in Figure 3.5 present clear evidence of the cleansing and sullying effects that were apparent in worker reallocation across the firm productivity ladder using revenue per worker. For the within-industry rankings, we find that, on average, poaching flows to higher-paying firms lead to an increase in firm premium by 0.09 log points per quarter whereas nonemployment flows lead to a 0.09 log point decrease in the firm premium per quarter. These estimates are quite similar qualitatively and quantitatively to the patterns using the direct measure of within industry productivity differences to rank firms. We also find the cyclical patterns of these components are very similar to those using the direct measure of

Figure 3.5: Decomposition of Growth in the Indirect Measure of Productivity Over the Cycle

(A) Firms are Ranked within Industry

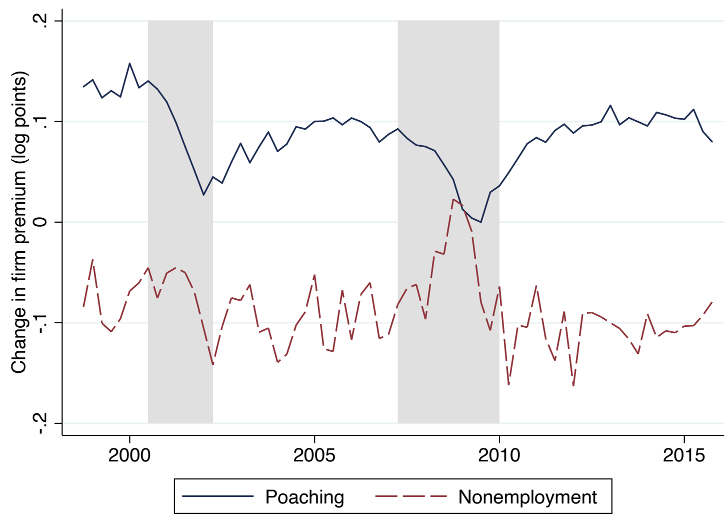

(B) Firms are Ranked within the Full Sample

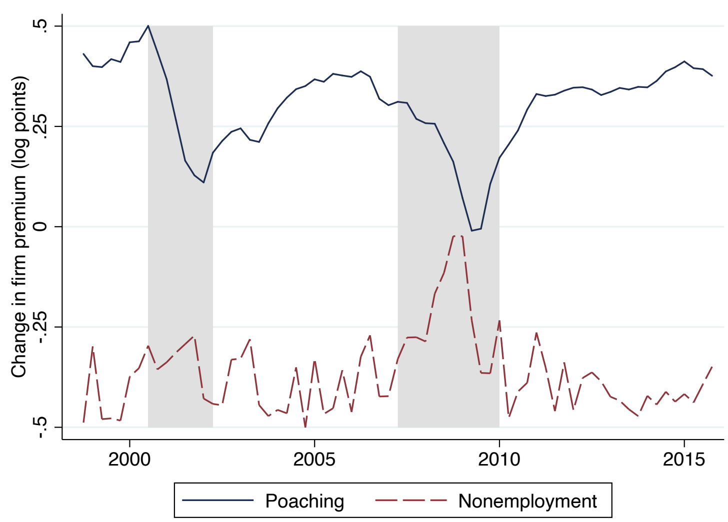

Notes: This figure presents the components of the indirect measure of productivity growth–i.e., the AKM firm fixed effect–that are attributable to worker reallocation between high- and low-productivity firms through poaching and nonemployment flows. Firms are ranked based on their AKM firm fixed effects. In panel (A) firms are ranked within their 4-digit NAICS industry codes. In panel (B) firms are ranked within the full sample. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

productivity. Figure 3.5(B) presents estimates based on ranking firms in the pooled sample (not within industry) and finds that, on average, poaching flows contribute 0.3 log points to the growth in the AKM firm premium each quarter whereas nonemployment flows lead a decline by 0.4 log points per quarter. The estimates are about three times as large as those for the within industry based tabulations. While much larger in magnitude, the qualitative patterns are very similar.

Columns 2 and 3 of Table 3.3 quantify the cyclicality of the contributions to this indirect measure of firm performance. As in the first column, we find that the cleansing contribution outweighs the sullying contribution in response to an increase in the unemployment rate while the opposite is true in response to a positive deviation of the unemployment rate from trend. These findings hold regardless of whether we rank firms within industries (column 2) or in the pooled sample (column 3). Since the AKM firm premium abstracts from worker heterogeneity, this suggests our main results by firm productivity are not being driven by variation in the patterns of sorting of heterogeneous workers across heterogeneous firms over the cycle. The larger coefficients in column (3) relative to column (2) suggest that cleansing and sullying effects are larger in magnitude when inter-industry productivity differentials are accounted for but that the relative contribution of these effects on productivity growth are qualitatively similar.

3.5 Implications for Earnings

Whereas Section 3.4 focused on implications for productivity growth, the current section asks if moving up the firm productivity ladder benefits workers in the form of higher earnings. We use the same methodology described in equation 3.4, but we replace the productivity differentials, $\tilde{R}^h_{t-1} - \tilde{R}^l_{t-1}$, with earnings differentials. We measure firm-level earnings in two ways: (i) average log earnings of all workers, and (ii) the AKM firm fixed effect. The earnings differential is the difference between the employment-weighted average of earnings at high- and low-productivity firms. Consistent with our core analysis, we deviate earnings from the industry average when calculating the differentials. Note that the exercise is distinct from the analysis discussed in Section 3.4.3, since we are investigating the earnings implications of worker flows across firms ranked by the revenue per worker measure of productivity.

Figure 3.6 shows that worker reallocation up the firm productivity ladder through poaching and nonemployment flows has meaningful implications for the earnings of workers. Figures 3.6(A) and 3.6(B) present results in which the earnings differentials are measured with average log earnings and the AKM firm premium, respectively. The results in the two figures are quantitatively and qualitatively similar. Job-to-job transitions add an average of 0.05 log points to average earnings in any given quarter. The earnings contribution of job-to-job flows is lower during recessions. During the 2007-2009 recession, the earnings contribution of job-to-job flows fell to a series low of 0.01. The contribution of nonemployment transitions to earnings growth is in the opposite direction and similar in magnitude to the contribution of job-to-job flows.[33] In the average quarter, worker movements into and from nonemployment subtract an average of 0.03 log points from earnings. This negative effect of nonemployment transitions is less present during recessions. During the 2001 recession, the contribution of nonemployment transitions is close to zero. During the 2007-2009 recession, the contribution of nonemployment transitions to earnings growth is briefly positive.

Comparing these results to Figure 3.4, we can draw conclusions about the extent to which the gains from productivity-enhancing reallocation are realized by workers. The earnings and productivity implications of employment transitions have similar signs, with job-to-job transitions providing gains, while nonemployment transitions subtract from each. The magnitudes, however, differ. The proportionate changes in earnings are approximately half the magnitude of the analogous changes in productivity. This suggests incomplete pass-through of the gains in revenue productivity from worker flows into earnings.

3.6 Conclusion

Consistent with the existing literature on firm heterogeneity, we find evidence of large differences in productivity across firms within the same industry. We also find that more productive firms in the same industry are more likely to grow and less productive firms more likely to contract and exit. The dispersion of productivity across firms

[33] Hahn, Hyatt, and Janicki (2021) consider the implications of employment transitions for earnings growth without considering productivity. They report that job-to-job and nonemployment transitions move in opposite directions and are roughly similar in magnitude, and follow opposite cyclical patterns.

Figure 3.6: Earnings Growth from Worker Reallocation up the Firm Productivity Ladder

(A) Average Earnings

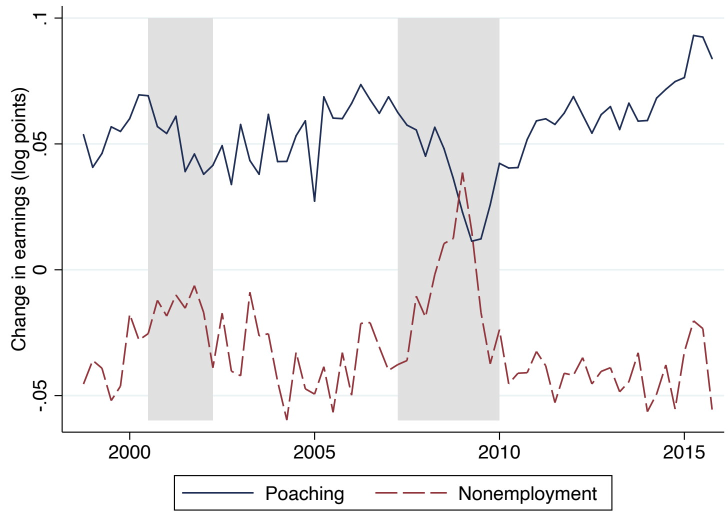

(B) Firm Fixed Effect

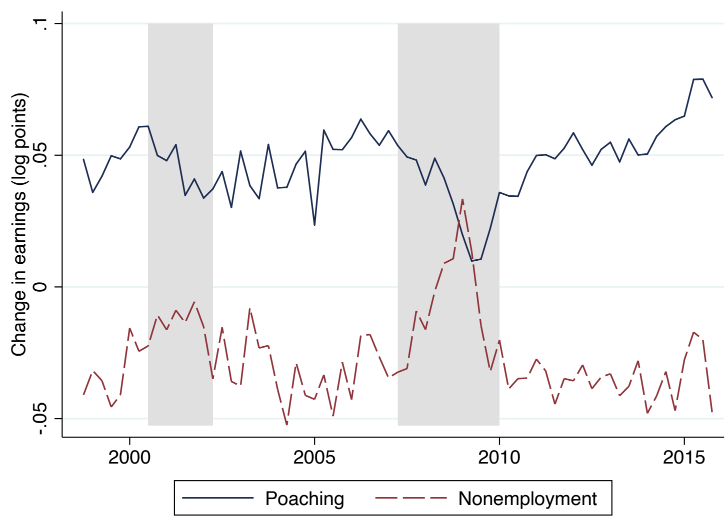

Notes: Panel (A) shows the components of earnings growth that attributable to worker reallocation between high- and low-productivity firms through poaching and nonemployment flows. We use the differences between the average earnings of workers at high and low productivity firms in order to quantify the implications of changes in the share of workers at high-productivity firms. Panel (B) presents the components of firm premium growth that are attributable to worker reallocation between high- and low-productivity firms through poaching and nonemployment flows. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

is large in magnitude contributing to a high pace of reallocation of workers across firms. Using a decomposition of net job flows into those accounted for by job-to-job flows and those accounted for net flows from nonemployment, we find that much of the overall reallocation of employment from less productive to more productive firms is accounted for by job-to-job flows. The pace at which workers move up the productivity job ladder is highly procyclical. The collapse of the productivity job ladder is consistent with a sullying effect of recessions. In recessions, we find that the reallocation of workers away from less productive firms via nonemployment flows increases. This occurs through a spike in separations to nonemployment along with a decline in hires from nonemployment at low productivity firms. Thus, we also find evidence that this component of reallocation is consistent with a cleansing effect of recessions.

The timing of the cleansing and sullying effects differs across stages of the cycle. The cleansing effect peaks relatively early in a downturn coincident with the relatively early spike in separations. The sullying effect peaks later in a downturn but lingers into the early stages of a recovery when unemployment is falling but remains well above trend.

Our findings are robust to using a direct measure of productivity based on relative differences in revenue per worker across firms within the same industry and an indirect measure of productivity/firm performance based on using the AKM firm premium to rank firms and quantify the contribution of worker flows to improved aggregate performance. When we focus on relative differences in the AKM firm premium across firms within the same industry, the results are very similar qualitatively and quantitatively to those using the direct measure of relative productivity. Since the AKM firm premium abstracts from worker heterogeneity, this suggests our results are not being driven by variation in the patterns of sorting of heterogeneous workers across heterogeneous firms over the cycle. This is not to suggest that the latter is unimportant but rather that there may be additional effects of cleansing and sullying from sorting above and beyond those we have quantified. We recognize that any conclusions about the role of sorting over the cycle for productivity fluctuations are tentative at best. Further theoretical and empirical work is needed in this area.

Chapter A: Appendix Material for Chapter 1

A.1 Additional Empirical Results

Figure A.1: Earnings Before and After Entry

(A) Average Earnings

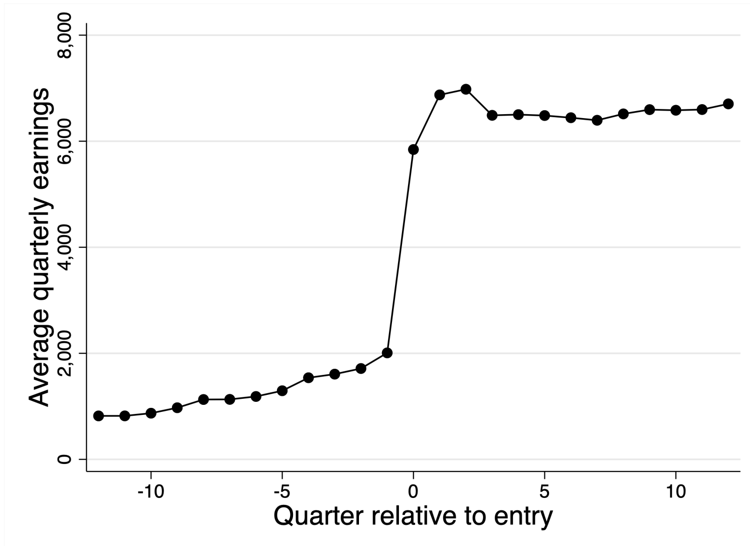

(B) Earnings Categories

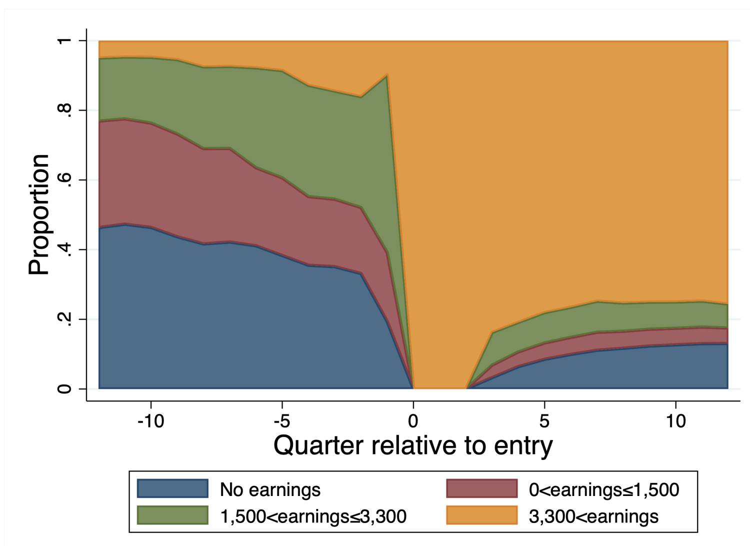

Notes: Both figures plot earnings in the 12 quarters before and after entry. Panel A plots the average quarterly earnings and Panel B plots the proportion of individuals with quarterly earnings in one of four mutually exclusive categories. All statistics are calculated using sample weights.
Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Figure A.2: Age-Earnings Profile by Age of Entry

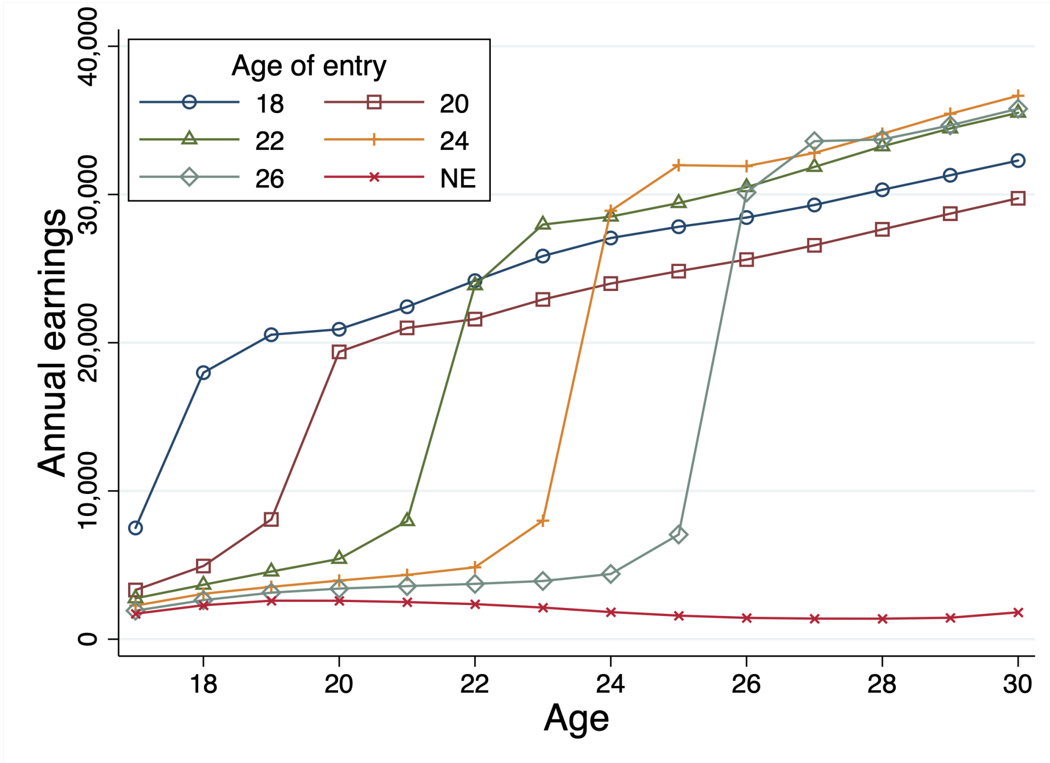

Notes: The figure plots the average annual earnings by age for different groups of workers defined by the age they were when they entered the labor market. The category, NE, is a group of workers that never entered the labor market. The sample includes all children who turned 30 by 2016 and all statistics are calculated using sample weights.
Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Figure A.3: Age of Entry

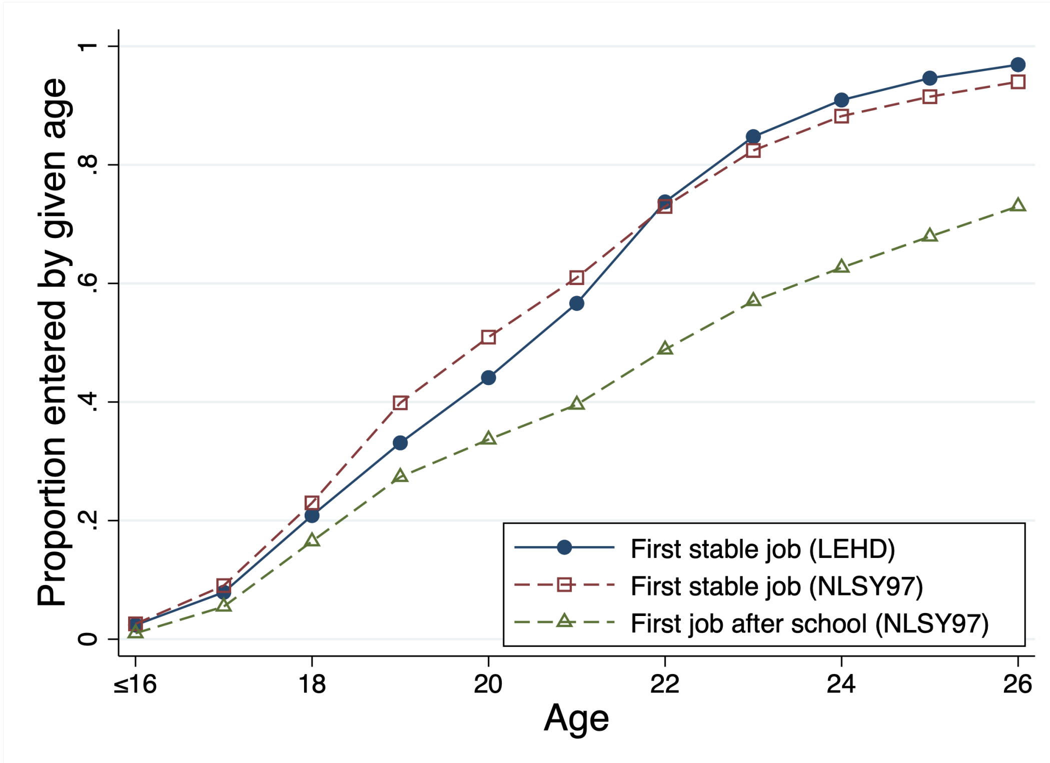

Notes: The figure plots the cumulative proportion of children that have entered the labor market by the age indicated on horizontal axis. For comparison, I also plot results using alternative measures of entry constructed from the NLSY97. These measures include the first stable job (working at least 35 hours for 36 consecutive weeks) and the first stable job after all schooling is completed. All statistics are calculated using sample weights.
Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics (LEHD) and 2000 Decennial Census files and data from the National Longitudinal Survey of Youth 1997 cohort (NLSY97).

Table A.1: Intergenerational Transmission of Employers and Education

	works for parent's employer				
	(1)	(2)	(3)	(4)	(5)
less than high school	0.055***	0.061***	0.048***	0.023*	0.055***
	(0.003)	(0.004)	(0.006)	(0.009)	(0.002)
high school	0.043***	0.054***	0.047***	0.042***	0.048***
	(0.002)	(0.002)	(0.002)	(0.003)	(0.001)
some college	0.020***	0.027***	0.024***	0.025***	0.024***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
parental earnings quartile	first	second	third	fourth	all
observations	180,000	183,000	177,000	165,000	705,000

Notes: Each column presents estimates from a separate regression. The outcome variable is an indicator equal to one if the first stable job is at the employer of either parent. The main independent variables include indicator variables for the highest level of education: less than high school, high school or equivalent, and some college or Associate degree. Bachelor's degree or advanced degree is the omitted educational category. Each regression controls for the interaction between the sex of the individual and the percentile of the parental earnigns distribution. All results are based on the sample of individuals who respond to the American Community Survey after they turn 25. Columns 1 through 4 present estimates based on the sample of individuals whose parents are in the first through fourth quartiles of the parental earnigns distribution, respectively. Column 5 includes all indviduals.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics, 2000 Decennial Census files and responses to the American Community Survey.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Table A.2: Intergenerational Transmission of Employers and Unemployment

	works for parent's employer			
	(1)	(2)	(3)	(4)
unemployment rate	-0.068**	0.128***	0.191***	0.064*
	(0.022)	(0.024)	(0.032)	(0.031)
covariates				
age of entry		X	X	X
quarter of entry			X	
county				X
observations	17,010,000	17,010,000	17,010,000	17,010,000

Notes: Each column presents estimates from a separate regression. The outcome variable is an indicator equal to one if the first stable job is at the employer of either parent. The main independent variable is the county-level unemployment rate, which ranges from zero to one, measured in the year in which the child enters the labor market. The different columns include additional covariates as indicated by the rows below the estimates. The covariates include fixed effects for: the age of entry, the quarter of entry, the county in which the individual entered the labor market. Standard errors are two-way clustered at the county and quarter of entry.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files and unemployment data from the U.S. Bureau of Labor Statistics.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Figure A.4: Intergenerational Transmission of Employers and Industry Wage Premiums

(A) Daughters

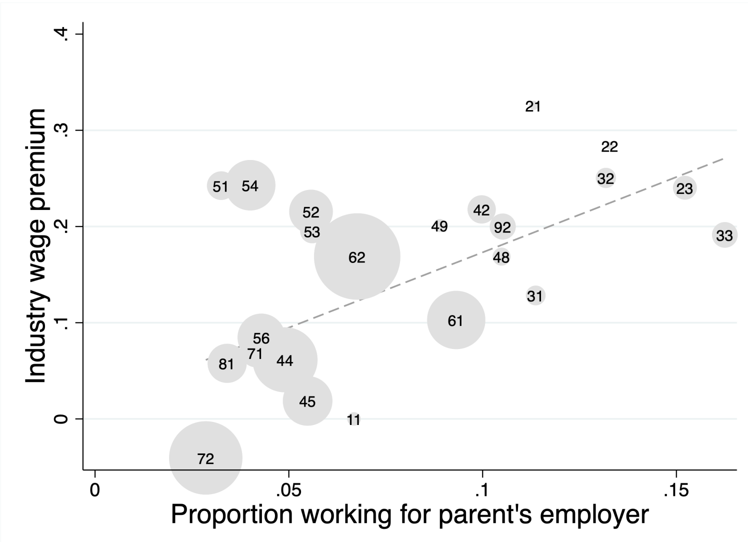

(B) Sons

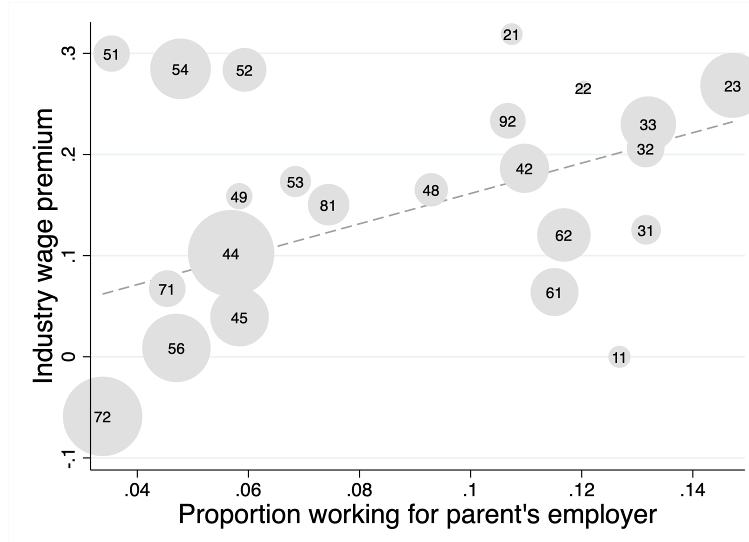

Notes: Panels A and B present results for daughters and sons, respectively. Each point on the plot presents information related to an industry, measured as the two-digit North American Industry Classification System code. The horizontal axis represents the proportion of individuals in a given industry who work for the employer of either parent at their first stable job, where this proportion is calculated separately for sons and daughters. The vertical axis is the industry-level wage premium. The wage premium is estimated using data from the Current Population Survey (CPS) by regressing log wages on an set of industry dummies (Agriculture, Forestry, Fishing, and Hunting is the omitted industry category) and controlling for year fixed effects, a third order polynomial in potential experience and fixed effects for the level of education. The sample from the CPS includes all employed individuals between the ages of 26 and 35 and regressions are estimated separately for males and females. The size of the marker is proportional to the number of individuals whose first stable job is in that industry.
Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files and publicly available data from the Current Population Survey.

Figure A.5: Intergenerational Transmission of Employers and Union Membership

(A) Daughters

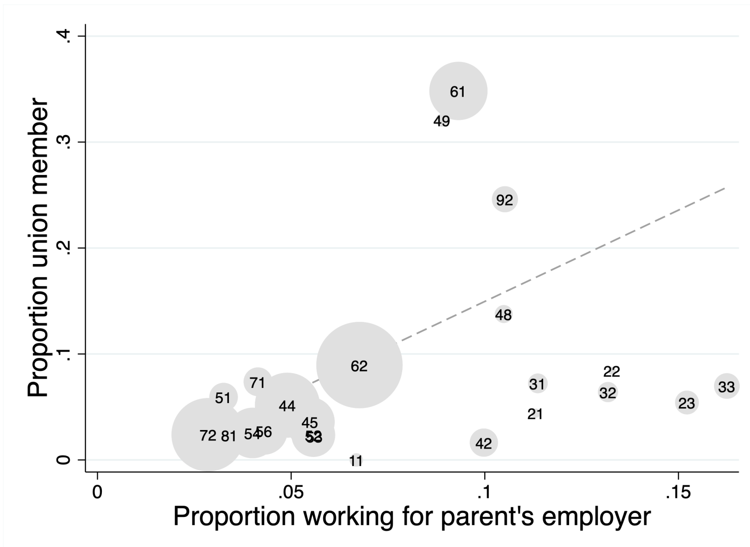

(B) Sons

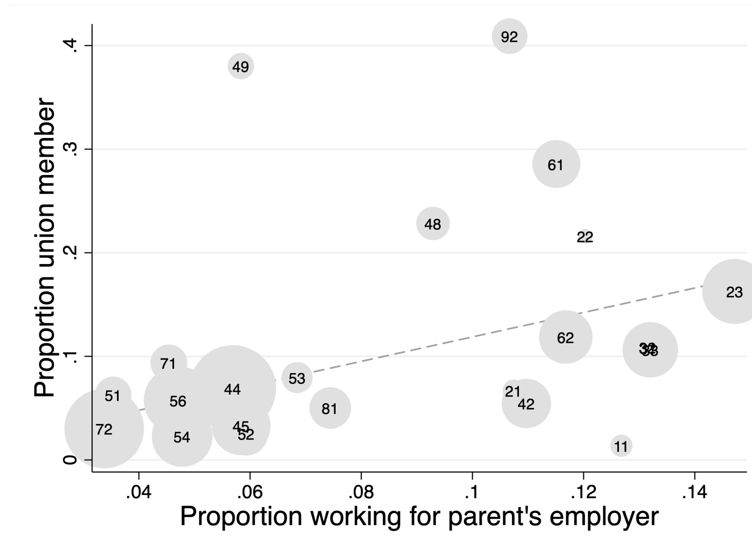

Notes: Panels A and B present results for daughters and sons, respectively. Each point on the plot presents information related to an industry, measured as the two-digit North American Industry Classification System code. The horizontal axis represents the proportion of individuals in a given two-digit industry who work for the employer of either parent at their first stable job, where this proportion is calculated separately for sons and daughters. The vertical axis presents proportion of individuals within that industry that are a member of a union. This statistic is calculated separately for males and females using data from the Current Population Survey, which include a sample of employed individuals between the ages of 26 and 35. The size of the marker is proportional to the number of individuals whose first stable job is in that industry.

Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files and publically available data from the Current Population Survey.

Table A.3: Intergenerational Transmission of Employers by Sex

	works for employer of			
	neither parent (1)	father (2)	mother (3)	both parents (4)
A. Daughters				
	0.940	0.013	0.040	0.006
A. Sons				
	0.922	0.042	0.026	0.010

Notes: Panels A and B present results for daughters and sons, respectively. Columns 1 through 4 present the proportion of individuals who find their first stable job at the same employer as neither parent, the father, the mother, and both parents, respectively. The proportions are calculated separately by the sex of the child. All statistics are calculated using sample weights. Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Figure A.6: Correlates with Parental Earnings

(A) Parent is Employed

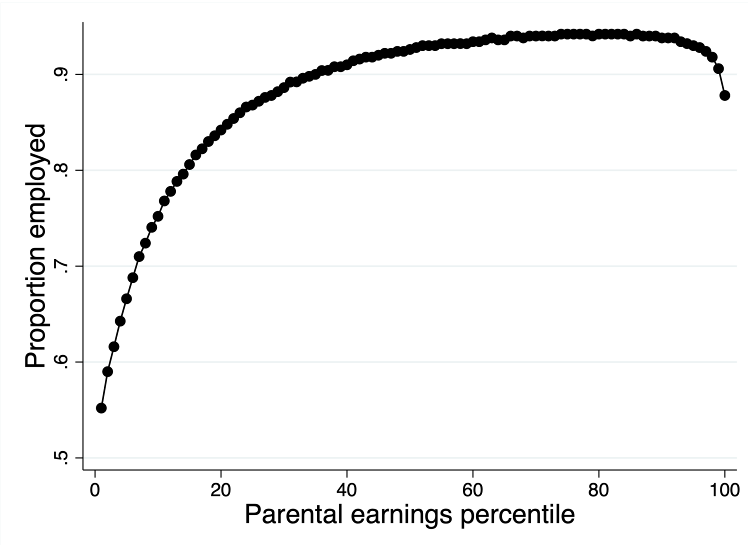

(B) Parent is Top Earner within Employer

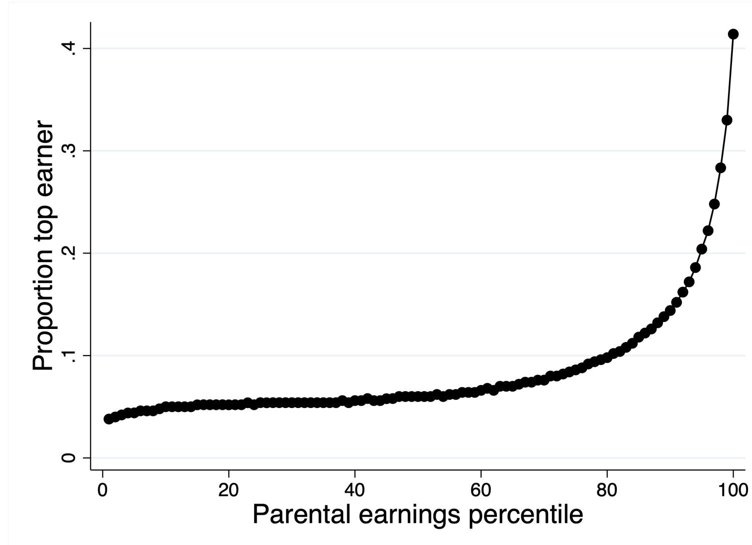

Notes: This figure describes the labor market outcomes of the parent who is the primary earner in the quarter in which their child enters the labor market. Panel A presents the proportion of parents (primary earner) that are employed. Panel B presents the proportion of parents (primary earner) whose earnings are in the top percentile of the within employer earnings distribution. All statistics are presented separately for each percentile of the parental earnings distribution and are calculated using sample weights.

Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Figure A.7: Long-Run Measures of the Intergenerational Transmission of Employers

(A) Daughters

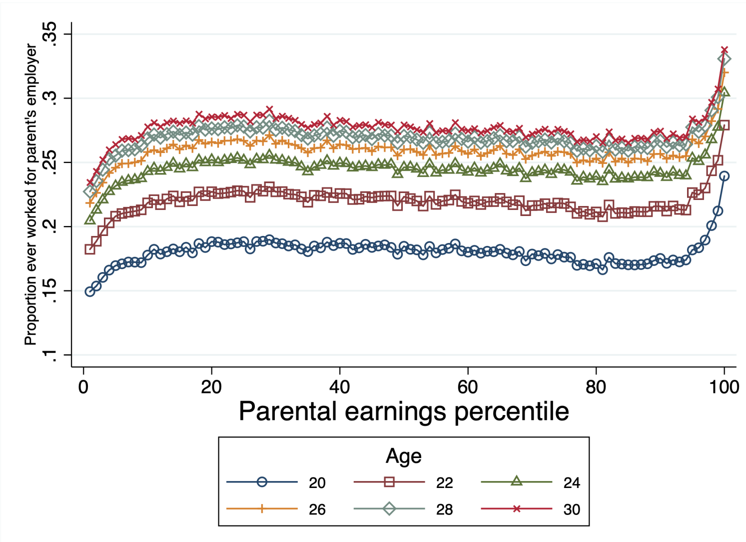

(B) Sons

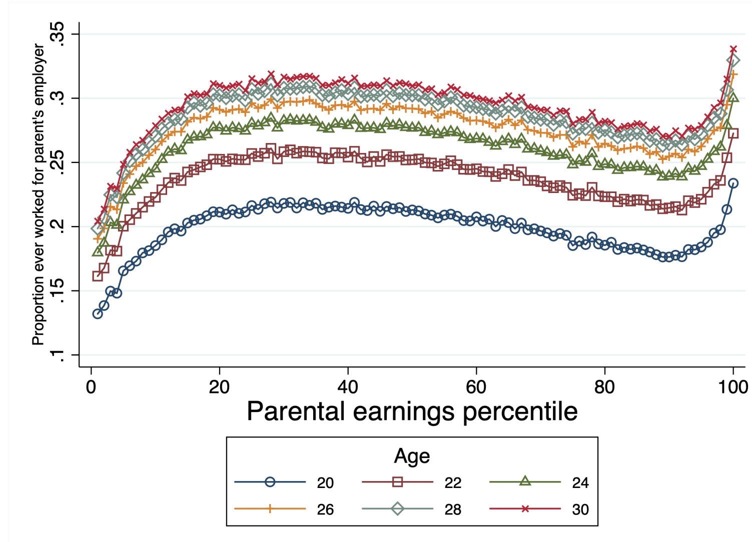

Notes: Each line in the plot the proportion of individuals who have ever worked for an employer of either parent between the ages of 18 and the age indicated in the legend. Each statistic is reported separately by the percentile of the parental earnings distribution. Panels A and B present results for the sample of daughters and sons, respectively. The results are based on a subsample that include children who turned 30 by 2016. All statistics are calculated using sample weights.

Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Table A.4: Control for Changes in Offer Wages

	log of quarterly earnings				
	(1)	(2)	(3)	(4)	(5)
works for parent's employer	0.307*** (0.029)	0.299*** (0.029)	0.312*** (0.028)	0.342*** (0.031)	0.307*** (0.030)
added covariates	none	new hire earnings	earnings growth of parent	earnings growth of all employees	employment growth
observations	11,460,000	11,460,000	11,460,000	11,460,000	11,460,000

Notes: Each column presents estimates from a separate regression estimated by two-stage least squares. The outcome variable is the log of the first full-quarter of earnings at the first stable job. The endogenous variable is an indicator equal to one if the child works for their parent's employer (primary earner) at the first stable job. The excluded instrument is the average quarterly hiring rate at the parent's employer in the four quarters prior to the quarter in which the individual enters the labor market. Column 1 reproduces the main results and columns 2-5 extend the baseline specification to include controls for the log of average earnings of stable new hires in the year before entry, the average quarterly earnings growth of the primary earner in the year before entry, the average annual earnings growth of all workers in the year before entry, and the average quarterly employment growth rate in the year prior to entry, respectively. All specifications include a fixed effect for the parent's employer; a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry, a fixed effect for the cohort of the child, and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the United States. Standard errors are clustered at the level of parent's employer and are presented in parentheses.

Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Table A.5: Household Fixed Effects

	log of quarterly earnings	
	(1)	(2)
works for parent's employer	0.199***	0.155***
	(0.040)	(0.045)
fixed effect	employer	household
control mean	8.757	8.757
observations	4,476,000	4,476,000

Notes: Each column presents estimates from a separate regression estimated by two-stage least squares. The outcome variable is the log of the first full-quarter of earnings at the first stable job. The endogenous variable is an indicator equal to one if the individual works for their parent's employer (primary earner) at the first stable job. The excluded instrument is the average quarterly hiring rate at the parent's employer in the four quarters prior to the quarter of entry. The specification in column 1 includes a fixed effect for the parent's employer whereas the specification in column 2 includes a fixed effect for the parent's employer by household. Both specifications are estimated on the same sample (which drop singleton observations) and include a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry, a fixed effect for the cohort of the child and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the United States. Standard errors are clustered at the level of parent's employer and are presented in parentheses. Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Table A.6: Effect on Timing of Entry

	average in three years prior to entry		
	quarterly earnings (1)	quarters worked (2)	quarter of entry (3)
works for parent's employer	-84.870	-0.066**	-3.973***
	(61.840)	(0.020)	(0.570)
control mean	1,269	0.612	13.170
observations	11,460,000	11,460,000	11,460,000

Notes: Each column presents estimates from a separate regression estimated by two-stage least squares. The outcome variables in columns 1 and 2 are average quarterly earnings and employment in the three years pior to entry, respectively. The outcome variable in column 3 is the quarter of entry relative to the expected quarter of high school graduation (based on birth cohort). The endogenous variable is an indicator equal to one if the child works for their parent's employer (primary earner) at the first stable job. The excluded instrument is the average quarterly hiring rate at the parent's employer in the four quaters prior to the quarter in which the individual enters the labor market. All specifications include a fixed effect for the parent's employer; a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the United States. The specificaitons in columns 1 and 2 also include a fixed effect for the cohort of the child. Standard errors are clustered at the level of parent's employer and are presented in parentheses.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Table A.7: Robustness to Measuring Hiring Rate at Alternative Times

	log of quarterly earnings								
	(1)	(2)	(3)	(4)	(5)	(6)	(6)	(8)	(9)
works for parent's employer	0.721***	0.590***	0.603***	0.414***	0.307***	0.250***	0.293***	0.347***	0.494***
	(0.129)	(0.102)	(0.074)	(0.053)	(0.029)	(0.013)	(0.013)	(0.015)	(0.023)
quarters used to calculate average hiring rate	[-8,-6]	[-7,-4]	[-6,-3]	[-5,-2]	[-4,-1]	[-3,0]	[-2,1]	[-1,2]	[0,3]
first stage F-statistic	177	275	461	798	1,434	4,308	7,725	5,445	2,219
observations	11,460,000	11,460,000	11,460,000	11,460,000	11,460,000	11,460,000	11,460,000	11,460,000	11,460,000

Notes: Each column presents estimates from a separate regression estimated by two-stage least squares. The outcome variable is the log of the first full-quarter earnings at the first stable job. The endogenous variable is an indicator equal to one if the child works for their parent's employer (primary earner) at the first stable job. In columns 1 through 9 the excluded instrument is the average quarterly hiring rate at the parent's employer in the four quarters prior to six quarters before through three quarters after the quarter in which the individual enters the labor market. All specifications include a fixed effect for the parent's employer; a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry, a fixed effect for the cohort of the child and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the United States. Standard errors are clustered at the level of parent's employer and are presented in parentheses.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Table A.8: Effect on Industry and Employer Pay Premium

	Industry Pay Premium				Employer Pay Premium
	two-digit	three-digit	four-digit	six-digit	
	(1)	(2)	(3)	(4)	(5)
works for parent's employer	0.167***	0.200***	0.208***	0.230***	0.304***
	(0.013)	(0.015)	(0.016)	(0.016)	(0.024)
control s.d.	0.178	0.208	0.222	0.232	0.366
observations	11,460,000	11,460,000	11,460,000	11,460,000	11,460,000

Notes: Each column presents estimates from a separate regression estimated by two-stage least squares. The outcome variables in columns 1-4 are the estimated pay premiums associated with the two-digit, three-digit, four-digit and six-digit industry codes, respectively. The outcome variable in column 5 is the estimated pay premiums associated with the employer. The endogenous variable is an indicator equal to one if the child works for their parent's employer (primary earner) at the first stable job. The excluded instrument is the average quarterly hiring rate at the parent's employer in the four quarters prior to the quarter in which the individual enters the labor market. All specifications include a fixed effect for the parent's employer; a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry, a fixed effect for the cohort of the child and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the United States. Standard errors are clustered at the level of parent's employer and are presented in parentheses.

Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Figure A.8: Long-Run Effects

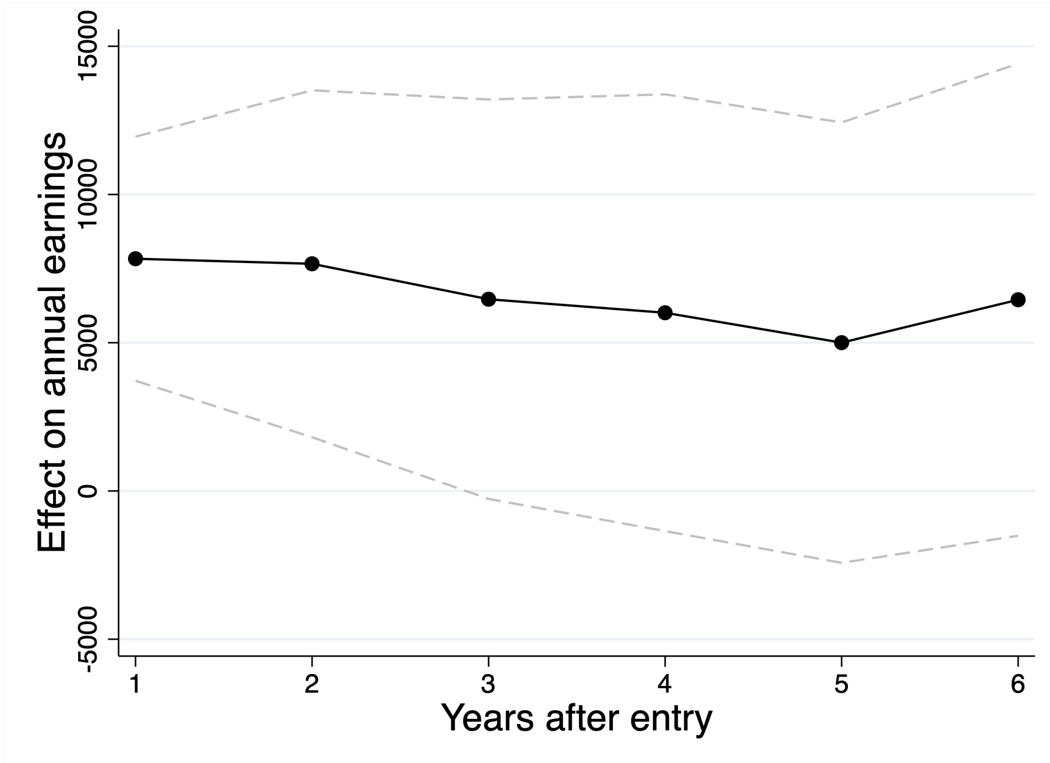

Notes: Each point on the figure represents an estimate from a separate regression. The outcome is the annual earnings x years after entry, where x refers to the coordinate on the horizontal axis. The endogenous variable is an indicator equal to one if the child work for their parent's employer (primary earner) at the first stable job. The excluded instrument is the average quarterly hiring rate at the parent's employer in the four quarters prior to the quarter in which the individual enters the labor market. All specifications include a fixed effect for the parent's employer; a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry, a fixed effect for the cohort of the child and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the U.S. Standard errors are clustered at the level of parent's employer and are used to construct the 95% confidence interval, which is denoted by the dashed lines. All regressions are estimated on a sample of 3,441,000 individuals who are expected to graduate high school in 2004 or earlier and who entered the labor market between the year in which they were expected to graduate high school and six and a half years later. The F-statistic from the first stage is 364.
Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Table A.9: Effect of Working for Employer of Secondary Earner

	works for parent's employer	log of quarterly earnings	
	(1)	(2)	(3)
hiring rate	0.071***	0.021***	
	(0.004)	(0.006)	
works for parent's employer			0.291***
			(0.081)
estimator	OLS	OLS	2SLS
F-statistic	365		
mean	0.042		
control mean		8.762	8.762
observations	4,447,000	4,447,000	4,447,000

Notes: Each column presents results from a separate regression. The outcome variable in column 1 is an indicator equal to one if the child works for their parent's employer (secondary earner) at their first stable job and the outcome variable in columns 2-4 is the log of the first full-quarter earnings at the first stable job. The main independent variable in column 1 is the average quarterly hiring rate at the parent's employer and the main independent variable in columns 2-4 is an indicator equal to one if the individual worked for the employer of their parent. The results in columns 1-3 are estaimted by Ordinary Least Squares (OLS) and the results in column 4 are estimated by two-stage least squares (2SLS), where the instrument is the average quarterly hiring rate at the parent's employer in the four quarters prior to the quarter in which the child enters the labor market. All specifications include a fixed effect for the parent's employer; a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry, a fixed effect for the cohort of the child and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the United States. Standard errors are clustered at the level of parent's employer and are presented in parentheses.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Table A.10: Effect of Working for Father's and Mother's Employer

	log of quarterly earnings			
	(1)	(2)	(3)	(4)
works for father's employer	0.522***	0.326***		
	(0.102)	(0.039)		
works for mother's employer			0.281***	0.336***
			(0.056)	(0.075)
sample	daughters	sons	daughters	sons
first stage F-statistic	387.900	760.900	760.900	391.800
observations	3,511,000	3,691,000	3,691,000	4,168,000

Notes: Each column presents results from a separate regression. The outcome variable is the log of the first full-quarter earnings at the first stable job. Columns 1-2 estimate the effect of working for the mother's emloyer and columns 3-4 estimate the effect of working at the father's employer. The main independent variable is an indicator equal to one if the individual worked for the employer of their parent. Each specification is estimated by two-stage least squares, where the excluded instrument is the average quarterly hiring rate at the parent's employer in the four quarters prior to the quarter in which the child enters the labor market. All specifications include a fixed effect for the parent's employer; a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry, a fixed effect for the cohort of the child and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the United States. Standard errors are clustered at the level of parent's employer and are presented in parentheses.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Figure A.9: Residualized Hiring Rate

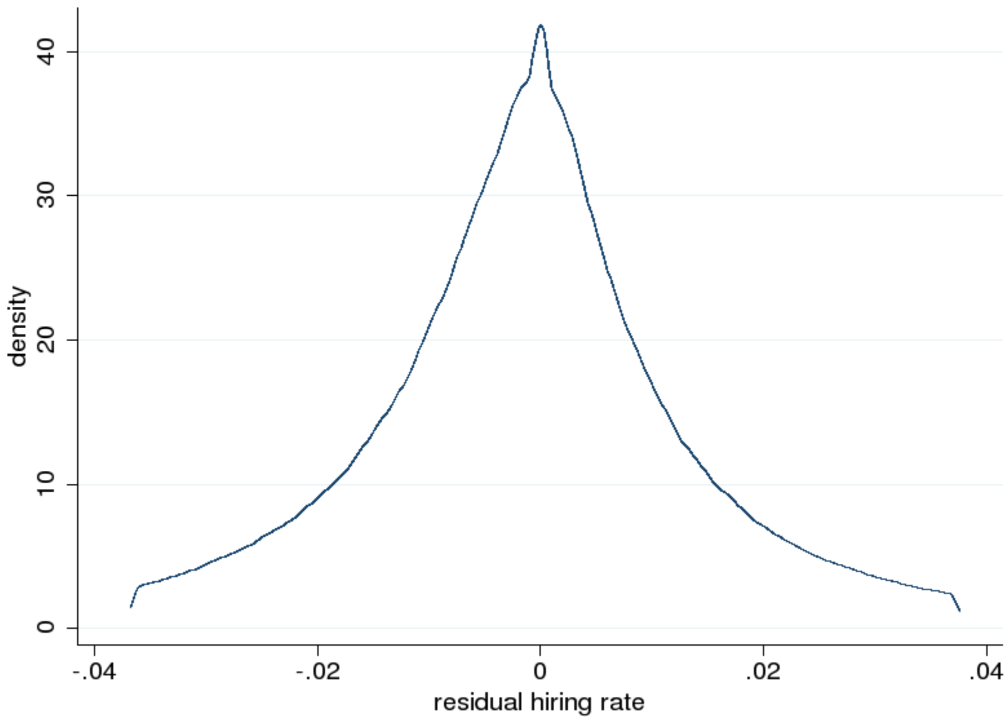

Notes: This figure presents the kernel density of the residuals from a regression of the average quarterly hiring rate at the parents' (primary earner) employer in the four quarters prior to entry on a fixed effect for the parents' employer; a fixed effect for the year of entry by two-digit industry code of parents' employer by state of parents' employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry, a fixed effect for the cohort of the child and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the United States. Standard errors are clustered at the level of parents' employer. The distribution is winsorized at the 5^{th} and 95^{th} percentiles according to the Census Bureau's rules.
Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Table A.11: Heterogeneity by Residualized Hiring Rate

	log of quarterly earnings		
	(1)	(2)	(3)
works for parent's employer	0.436***	0.310***	0.228*
	(0.048)	(0.029)	(0.114)
estimation sample			
first tercile	X	X	
second tercile	X		X
third tercile		X	X
first stage F-stat	999	1,429	212
observations	7,304,000	7,606,000	7,308,000

Notes: Each column presents estimates from a separate regression estimated by two-stage least squares. The outcome variable is the log of the first full-quarter earnings at the first stable job. The endogenous variable is an indicator equal to one if the child works for their parent's employer (primary earner) at the first stable job. The excluded instrument is the average quarterly hiring rate at the parent's employer in the four quarters prior to the quarter in which the individual enters the labor market. The sample is partitioned into terciles based on the residualized hiring rate. The row below the estimates indicates whether observations from a given tercile are included in the estimation sample. All specifications include a fixed effect for the parent's employer; a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry, a fixed effect for the cohort of the child and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the United States. Standard errors are clustered at the level of parent's employer and are presented in parentheses.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Table A.12: Characteristics of Compliers

	works for parent's employer		Characteristics of Compliers		
	no	yes	IV(p25)	IV(p50)	IV(p75)
	(1)	(2)	(3)	(4)	(5)
A. Individual					
male	0.50	0.60	0.54	0.54	0.51
White non-Hispanic	0.74	0.74	0.73	0.74	0.74
Black non-Hispanic	0.09	0.08	0.10	0.11	0.10
Asian non-Hispanic	0.03	0.02	0.03	0.03	0.03
Hispanic	0.11	0.13	0.12	0.10	0.10
other	0.03	0.03	0.03	0.03	0.03
born in United States	0.96	0.96	0.97	0.98	0.97
B. Parent and their Employer					
skilled services	0.49	0.38	0.49	0.48	0.66
unskilled services	0.15	0.26	0.20	0.16	0.10
manufacturing/production	0.35	0.36	0.31	0.36	0.25
tenure of parent	23.96	22.63	24.52	25.27	26.71
earnings rank within employer	68.49	77.93	63.97	51.65	65.40
parental earnings rank	55.47	54.40	58.48	66.39	60.33
Sample Size					
proportion of full sample	0.94	0.06	0.04	0.03	0.15

Notes: Each row presents estimates for the variable defined in the first column.Columns 1 and 2 present the average value of the variable for the sample of individuals who do not and do work for the employer of their parent at their first stable job, respectively. Colummns 3-5 present the average characteristics of the compliers for the case in which the instrumental variable is a binary variable equal to one if the residualized hiring rate exceeds the 25th, 50th, and 75th percentile, respectively. The complier characteristics are estimated using the methodology described by Abadie (2003). I winsorize the estimates of κ at the 1st and 99th percentiles to reduce the influence of outlier values.

Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Table A.13: Intergenerational Elasticity of Earnings Using Long-Run Earnings

	earnings of child in 2016		
	(1)	(2)	(3)
Panel A. Including Zero Earnings			
log of parental earnings	0.378	0.417	0.396
	(0.002)	(0.002)	(0.001)
sample	daughters	sons	all
observations	8,416,000	8,591,000	17,010,000
Panel B. Excluding Zero Earnings			
log of parental earnings	0.2499	0.2203	0.2348
	(0.0006)	(0.0005)	(0.0004)
sample	daughters	sons	all
observations	7,412,000	7,706,000	15,120,000

Notes: Columns 1 through 3 present results based on a sample of daughters, sons, and all children, respectively. The estimates in Panel A are the coefficients from a regression in which the independent variable is the log of parental earnings and the dependent variable is the inverse hyperbolic sine of the earnings of the child in 2016. The samples used to estimate the regressions in Panel A include children who have zero earnings in 2016. The estimates in Panel B are the coefficients from a regression in which the independent variable is the log of parental earnings and the dependent variable is the log of the earnings of the child in 2016. The samples used to estimate the regressions in Panel B do not include children who have zero earnings in 2016. All regressions are estimated via weighted least squares using sample weights.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Figure A.10: Log Earnings of Parents and Children

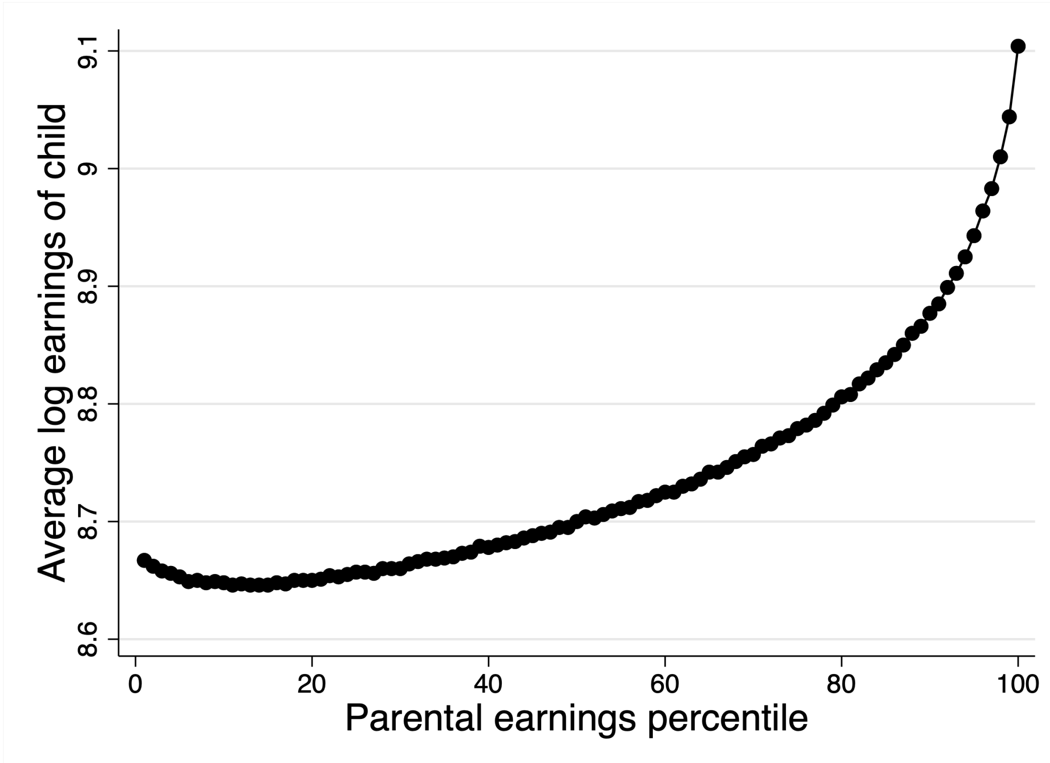

Notes: The figure plots the average log earnings of the children against the average log earnings of the parents. Each point represents the average outcome of individuals and their parents for a given percentile of the parental earnings distribution. The horizontal and vertical axes correspond to the average value of log parental earnings and the average value of the log of the first full-quarter of earnings at the first stable job of the child, respectively. All statistics are calculated using sample weights.
Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Table A.14: Intergenerational Elasticity of Earnings with Homogeneous Treatment Effects

	sample		
	daughters (1)	sons (2)	all (3)
A. Observed			
IGE	0.1565	0.1298	0.1430
	(0.0002)	(0.0003)	(0.0002)
B. No Transmission with Primary Earner			
percent change in IGE	-0.95%	-1.14%	-1.05%
	(0.09)	(0.11)	(0.10)
C. No Transmission with Either Parent			
percent change in IGE	-1.94%	-2.34%	-2.14%
	(0.18)	(0.22)	(0.20)
observations	8,416,000	8,591,000	17,010,000

Notes: The results in columns 1-3 corresopnd to daughters, sons and all children, respectively. Panel A presents the observed intergenerational elasticity of earnings (IGE), which is denoted $\rho(y_{ijt}, y_p)$ and is estimated with sample weights via weighted least squares. Panels B and C present the percent by which the IGE estimates in Panel A would change if no children were to work for the employer of the parent who is the primary earner or either parent, respectively. The percent change is defined as, $\frac{\rho(y_{ijt},y_p)-\rho(y_{i(j0)t},y_p)}{\rho(y_{ijt},y_p)} \times 100$. The treatment effects used to construct the counterfactual estimates are estimated via two-stage least squares and are estimated for the entire sample, pooling sons and daughters and children from all five quintiles of the parental earnings distrbution. Standard errors are presented in parentheses and are calculated using the delta method and take into account the uncertainty in the estimated earnings consequences.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Figure A.11: Intergenerational Transmission of Employers by Parental Earnings

(A) Daughters

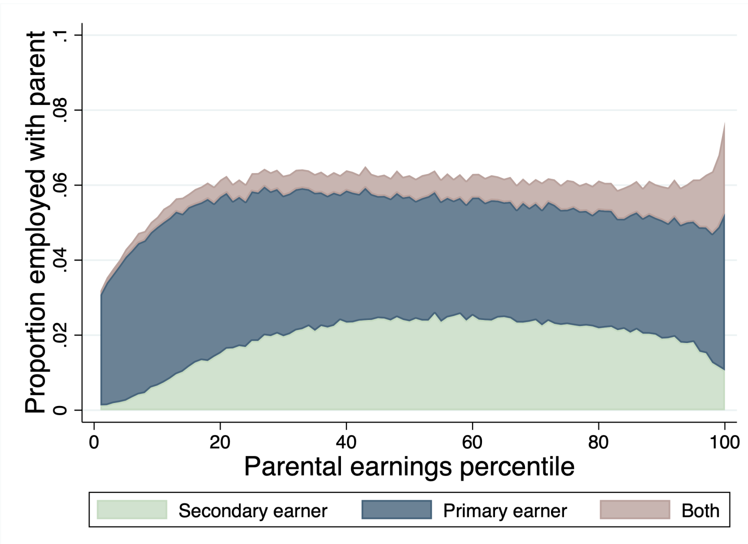

(B) Sons

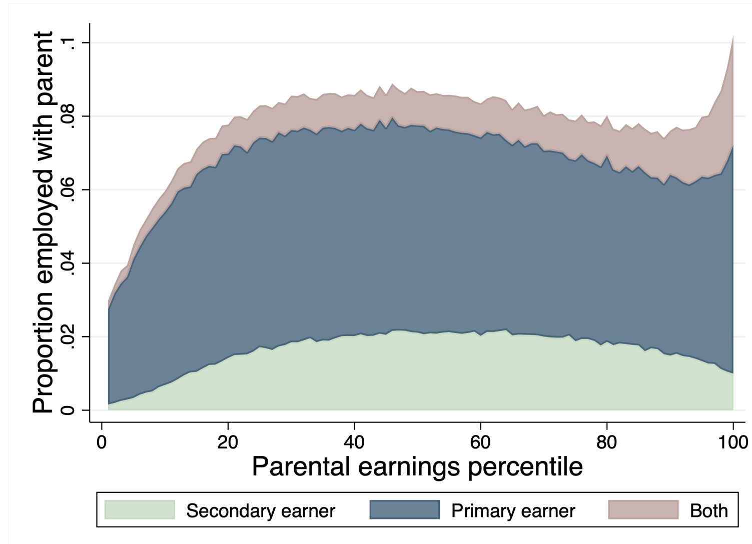

Notes: The figures plot the proportion of children whose first stable job is at the same employer as the secondary earner only, primary earner only, or both parents, respectively. Each statistic is reported separately by the percentile of the parental earnings distribution. Panels A and B present results for the sample of daughters and sons, respectively. All statistics are calculated using sample weights.

Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

Table A.15: Effect on Earnings Rank by Sex and Parental Earnings

	percentile earnings rank					
	(1)	(2)	(3)	(4)	(5)	(6)
A. All						
works for parent's employer	14.30*	20.64***	15.93***	28.96***	20.63***	21.64***
	(05.88)	(04.70)	(04.42)	(05.34)	(05.46)	(01.97)
F-statistic	238.9	358.0	446.8	330.4	316.2	1434.1
observations	1,350,000	1,987,000	2,297,000	2,462,000	2,487,000	11,460,000
B. Daughters						
works for parent's employer	25.25*	31.27***	27.32*	39.76***	26.58*	31.13***
	(12.49)	(09.68)	(12.06)	(12.55)	(12.30)	(04.02)
F-statistic	64.3	131.2	100.0	106.1	130.9	679.8
observations	586,000	876,000	1,029,000	1,128,000	1,152,000	5,387,000
C. Sons						
works for parent's employer	08.23	15.95**	19.61***	29.59***	27.04***	21.72***
	(08.23)	(06.06)	(05.37)	(06.70)	(07.31)	(02.46)
F-statistic	97.7	198.3	245.7	176.9	161.3	854.2
observations	600,000	909,000	1,067,000	1,149,000	1,148,000	5,501,000
Sample Description						
parental earnings quintile	first	second	third	fourth	fifth	all

Notes: This table presents estimates based on subsamples defined by the interaction between the quintile of parental earnings (defined by the column) and sex (defined by the panel). The outcome varable is the percentile rank of the individual's earnings at their first stable job. The endogenous variable is an indicator equal to one if the individual works for their parent's employer (primary earner) at the first stable job. The excluded instrument is the average quarterly hiring rate at the parent's employer in the four quarters prior to the quarter of entry. All specifications include a fixed effect for the parent's employer; a fixed effect for the year of entry by two-digit industry code of parent's employer by state of parent's employer; and a vector of covariates that includes log annual earnings of the parent in the year prior to entry, a fixed effect for the cohort of the child and interactions between the sex of the child and race, ethnicity, and an indicator equal to one if born in the United States. Standard errors are clustered at the level of parent's employer and are presented in parentheses.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.
*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

Figure A.12: Conditional Expected Rank

(A) Observed and Counterfactural

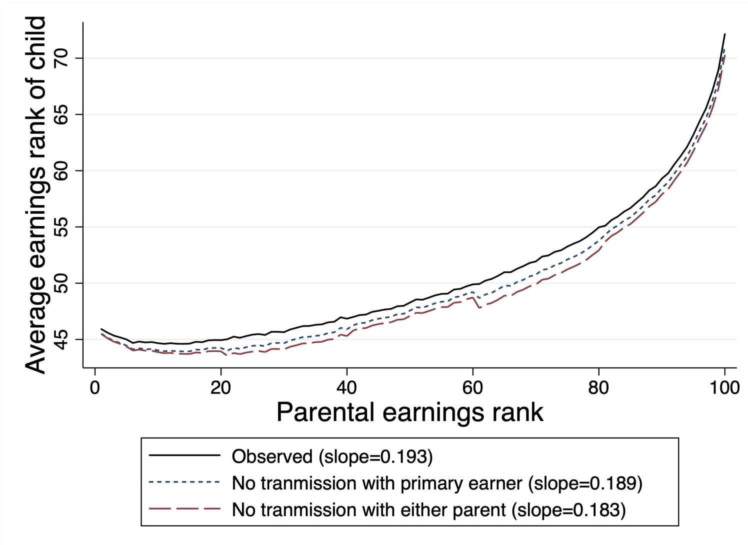

(B) Difference

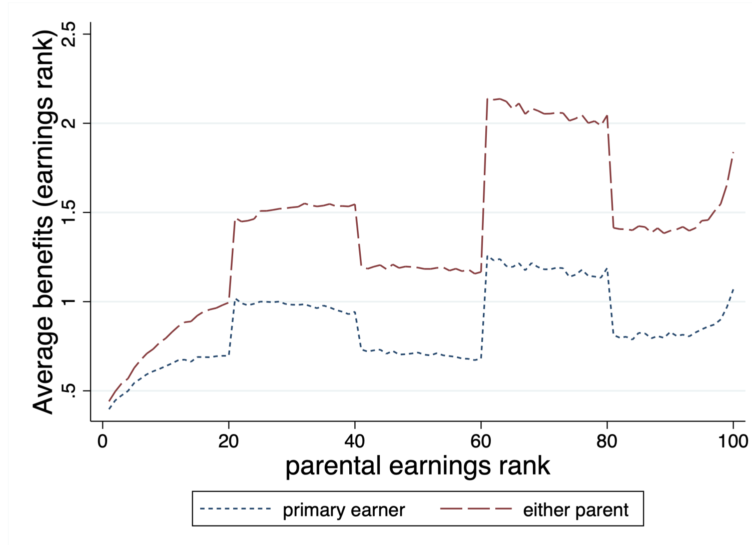

Notes: The solid line in Panel A presents the conditional expected rank measure, which is the average percentile rank the of earnings at the first stable job of the child for each percentile of the parental earnings distribution. The dashed lines represent the counterfactual measures that correspond to the two different scenarios in which no individual works for the employer of the primary earner or either parent. Panel B plots the difference between the observed and counterfactual measure for each percentile of the earnings distribution. The treatment effects used to construct the counterfactual estimates are estimated via two-stage least squares and are estimated separately by the quintile of the parental earnings distribution. All statistics, aside from the two-stage least squares estimates, are calculated using sample weights.

Source: Author's calculations based on matched data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

A.2 Details on Data

A.2.1 Sample Frame

The Hundred Percent Census Edited File (HCEF) is an edited version of the Hundred Percent Census Unedited File, which contains all household and person records included in the 2000 Decennial Census. Edits are applied to remove duplicate observations and to ensure consistency between the long and short-form files. While the Decennial Census surveys aim to interview everyone who resides in the United States, in practice, the sample frame considered in my paper does not include all children (within the appropriate age range) living in the United States in 2000. In addition to coverage issues in the 2000 Decennial Census discussed in the text and by Mulry (2007) and the technical report "Coverage Evaluation of Census 2000: Design and Methodology", some children do not live with their parents. Specifically, 91% of individuals younger than 18 lived with their parents in 2000. The remaining 9% individuals will be excluded from my sample since I require that the parent is the head of household.[1]

Panel A and B of Figure A.13 depict the share of individuals whose relationship to the household head is defined as a child by age in 2000 and race/ethnicity, respectively. While my sample frame excludes some individuals for these two reasons, it does include the vast majority of children who fall within the age range. Nevertheless, I point out that the results in this paper aim to be representative of the sample frame and I make no attempts to adjust for additional differences between the sample frame and other populations.

[1]This statistic is based on the authors own calculations using a 5% sample of the 2000 Decennial Census made available through IPUMS, see Ruggles et al. (2019).

Figure A.13: Relationship to Head of Household

(A) By Age in 2000

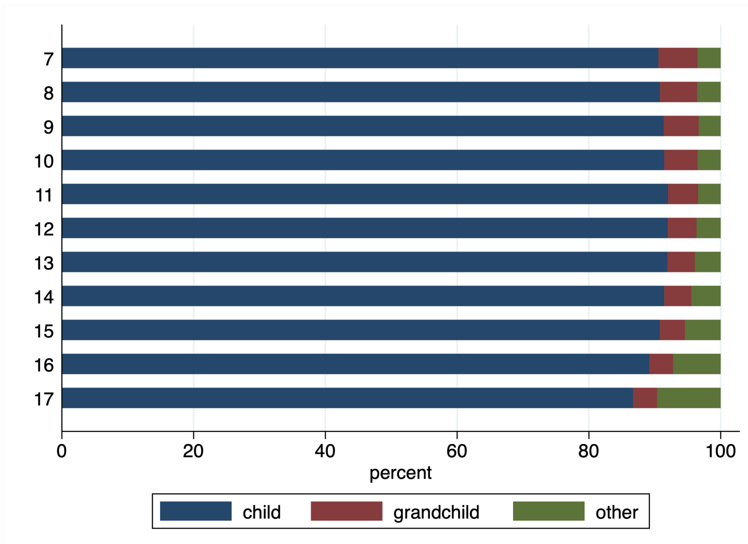

(B) By Race/Ethnicity

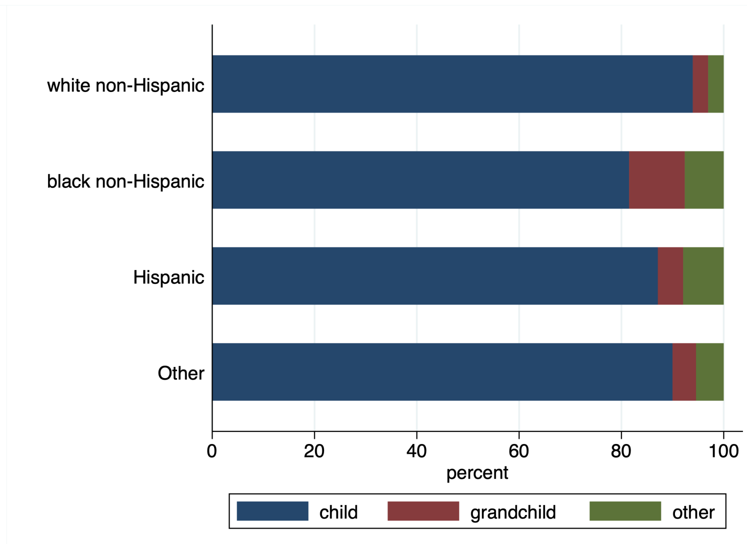

Notes: The figures present the proportion of children born between 1982 and 1992 whose relationship to the head of household in the 2000 Decennial Census was defined as: child, grandchild, or other. Panel A breaks out the results by the age of the child at the time of the Decennial Census and Panel B breaks out the results by the race/ethnicity of the child.

Source: Author's calculations based on a 5% sample from the 2000 Decennial Census obtained from IPUMS, see Ruggles et al. (2019).

A.2.2 Sample Restrictions

I make several key sample restrictions in the move from the sample frame to the analysis sample, all of which are summarized in Table A.16. First, I implement a number of restrictions to ensure that I can accurately link the records of the children from the HCEF to the data from the Longitudinal Employer-Household Dynamics (LEHD) program. Individuals are identified by a Protected Identification Key (PIK), which the Census Bureau generates using personally identifiable information.[2] I use the PIK to link person records between the HCEF and the LEHD and to attach employer characteristics to jobs. Various types of measurement error in the HCEF may prevent a PIK from being accurately assigned to an individual. In order to ensure that each child is accurately assigned a PIK, I require that a unique PIK be assigned to the individual and the year and month of birth recorded in the Individual Characteristic File (ICF) match those recorded in the HCEF.[3] The decision to retain only observations with unique non-missing PIKs and matching year and month of birth between the HCEF and the LEHD is conservative, in the sense that it may drop some individuals who could accurately be linked across the two datasets. The justification for doing this is to limit measurement error in intergenerational relationships, which would arise if PIKs were incorrectly assigned to the child or either parent. While these restrictions reduce sample size, they do not introduce bias to the extent that the sample weights account for the selected nature of the sample. 79% of the children in the sample frame satisfy these restrictions.

Second, I implement a number of restrictions to ensure that I accurately measure the relationship between children and parents and link parental records to the LEHD.

[2]See Wagner and Layne (2014) for a description of the methodology by which PIKs are assigned to individual observations.

[3]The ICF contains a record for every individual that ever appears in the LEHD and contains basic observable characteristics such as race, sex, and date of birth. The primary source for the date of birth variable is the Person Characteristic File (PCF), which is drawn from information recorded from transactions with the Social Security Administration.

To ensure that the relationship between children and parents is accurately measured in the HCEF, I require that the household contains no more than 15 individuals in the HCEF. To ensure that I am able to link the records of the parents to the LEHD files, I require that a unique PIK be assigned to both parents and the year and month of birth recorded in the ICF match those recorded in the HCEF for both parents.[4] 62% of the children in the sample frame satisfy the restrictions in this and the preceding paragraph.

I construct sample weights in order to address the possibility that the first two sample restrictions produce a selected sample. Specifically, using a dataset that includes every child in the sample frame, I estimate the propensity score as the probability of satisfying the first two sample restrictions as a function of observable characteristics that include: sex, relationship to head of household (biological child, adopted child or step child), race (White, Black, Native American, Asian, or other), Hispanic ethnicity, number of parents in the household in 2000, and a vector of observable characteristics of the census tract in which the household resided in at the time of the 2000 Decennial Survey (share of parents that are single parents, median household income, poverty rate, proportion of residents who were living in the same house five years ago, urban/rural, proposition of households receiving public assistance). The sample weights are the inverse of the estimated propensity score.

Third, I implement a set of restrictions to ensure that the measurement of key labor market outcomes are not impacted by coverage issues in the LEHD. Since much of the analysis focuses on the labor market outcomes associated with first stable jobs, I drop children if their first stable job is likely to not be covered in the LEHD. Specifically, I identify the state in which children reside in in the year they are expected to graduate from high school and retain observations only if the state was participating in the LEHD for more than a year prior to that year and the

[4]Parents are defined as the household head and either their spouse or unmarried partner. Note that edits applied to the HCEF imply that there are at most two parents in each household.

Table A.16: Sample Restriction Criteria

Exclusion Criteria	Observations Remaining	
	number	percent
none (sample frame with no restrictions)	37,120,000	100%
child not assigned a unique PIK or the year and month of birth recorded in the HCEF does not match the date of birth in the Social Security Administration transaction file	29,165,000	79%
head of household and spouse (or unmarried partner) is not assigned a unique PIK, the year and month of birth recorded in the HCEF does not match the date in the LEHD or there are more than 15 individuals in the household	23,169,000	62%
the state in which the child resided in began reporting to the LEHD less than a year prior when they are expected to graduate high school or the year child entered the labor market, or if parental earnings is below the 5^{th} percentile	21,321,000	57%
child did not enter the labor market by the end of 2016	17,010,000	46%

Notes: This table describes the sample restrictions applied to the sample frame. The first column describes the criteria and the second column presents the rounded number of observations that remain after dropping the observations that meet the criteria. These numbers represent a cummulative count after the all sample restrictions described in preceding rows are applied. The third column presents this infomration as a percent of the total sample frame.
Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

year child entered the labor market. Since an important dimension of the project is to study differences across the parental earnings distribution, I also drop parents for whom I cannot reliable measure earnings. Specifically, I construct a long-run measure of parental earnings (discussed in detail in Appendix Section A.2.4) and I drop parents whose earnings is below the 5^{th} percentile. The percentile is calculated on a dataset with all previously discussed sample restrictions and also conditional on the child entering the labor market. For parents below this threshold, it is difficult to distinguish between low earnings and earnings missed in the LEHD and I find that measures of earnings and other economic indicators (such as the poverty rate of median value in the census tract in which the household lived in 2000) start to diverge for these households. These two sets of restrictions drop an additional 1.9 million children, which leaves 57% of the sample frame.

Lastly, much of the analysis is restricted to a set of children who enter the labor market. I define entry as the first quarter in which the individual earns at least $3,300 per quarter for three consecutive quarters and receives positive earnings from the same employer for those three quarters. 46% of the children in the sample frame satisfy the restrictions in this and the preceding paragraphs.

A.2.3 Edits to Individual Earnings Records

Earnings data in the LEHD come from Unemployment Insurance (UI) records, which report total amount paid to each worker per employer per quarter. In measuring quarterly earnings, I sum earnings records across employers within a quarter for each individual to construct a measure of total individual earnings per quarter. While the administrative data are not subject to various types of measurement error that plague survey data, they are not error free. A key issue is that data errors can produce very large outlier observations. Researchers typically deal with these by winsorizing the data–editing or dropping earnings records above some percentile of the distribution.

The issue with this methodology is that it incorrectly impacts the earnings of workers who truly have earnings in the top percentiles.

In order to retain top earners in my sample, I use an alternative methodology to deal with outliers. The methodology, which I have also employed in Fallick et al. (2019), is based on the fact that outliers often appear in the form of a large spike for a single quarter for an individual. Let $z_i = max\{median(y_{it}), 10000\}$ be the greater of the median of earnings observed for individual i over the entire sample and 10,000.[5] Then define earnings growth as:

$$\Delta_{it} = \frac{y_{it} - z_i}{\frac{1}{2}(y_{it} + z_i)} \tag{A.1}$$

where t is the quarter and y is the earnings. The growth rate, Δ_{it}, captures the extent to which earnings in a given quarter exceeds the typical earnings of that individual. The choice to set a minimum value of z is motivated by the desire to avoid editing the earnings of low earners, since the outliers are driven by very large levels of earnings.

I define outliers as earnings records that produce growth rates that exceed the 95^{th} percentile of the distribution. Let $\Delta(p95)$ denote the 95^{th} percentile, then the earnings variable used in this paper is defined as:

$$\tilde{y}_{it} = \begin{cases} y_{it} & \text{if } \Delta_{it} < \Delta(p95) \\ z_i * \frac{1+\frac{1}{2}\Delta(p95)}{1-\frac{1}{2}\Delta(p95)} & \text{if } \Delta_{it} > \Delta(p95) \end{cases} \tag{A.2}$$

This methodology edits outlier observations so that if the growth rate were calculated on the edited value it would be equal to the 95^{th} percentile. The advantage of this methodology over the traditional winsorization method is that it retains the earnings records of individuals who consistently have high levels of earnings.

[5]The median is calculated from a sample that contains strictly positive earnings.

A.2.4 Measuring Parental Earnings

The ideal dataset would contain earnings data for each worker over their entire working life, and lifetime earnings would simply be calculated as the sum of all observed earnings. However, the LEHD fall short of the ideal data because some sources of earnings are not included in the data and because they do not cover the full working life of all parents in the sample. Thus, I require an alternative method to estimate lifetime earnings.

A common approach in the literature is to calculate parental earnings as the average earnings over a limited number of years. For example, recent work by Chetty et al. (2014) measure parental earnings as the average earnings measured across five years. Even using comprehensive income data derived from the 1040 tax forms, there are various issues with their approach (see Mazumder 2016 for a detailed discussion). The first is related to the number of years over which the earnings are averaged. A large literature inspired by Solon (1992) and Zimmerman (1992) finds that measuring parental earnings over a short time periods introduces measurement error and leads to artificially low estimates of the intergenerational relationship in economic outcomes. Mazumder (2005) suggest that even fifteen years of data may not be enough to accurately measure lifetime earnings. The second issue, is that parental earnings measured at different points in the life cycle may not be comparable (see Jenkins 1987; Solon 1992; Grawe 2006; Bohlmark and Lindquist 2006; Haider and Solon 2006). For example, two individuals aged 35 and 55 might have similar earnings in a given year but very different levels of lifetime earnings.

There are also a number of additional issues that are specific to the LEHD. The main challenge is that it is not clear how to interpret missing data because it is difficult to distinguish between zero earnings and missing earnings. There are two main reasons why earnings data from the LEHD might be missing for a given individual in a given quarter. First, data availability in the LEHD varies on a state-by-state basis. While all states are currently reporting, coverage is less complete for years further in the past. Figure A.14 illustrates when the different states entered the program. While the residential data in the LEHD can be used to identify whether workers are living in a state that participates in the LEHD, imperfect coverage of these data and workers who commute across state boundaries make it difficult to accurately flag workers whose earnings are missing due to a lack of state reporting.

Second, while most earnings (96% of salary employment) are covered under the UI system, the LEHD systematically misses some sources of earnings. Measurement issues at the bottom of the wage earnings distribution are of particular concern. Figure A.15 demonstrates this point by using data from the CPS to plot average total household income by source against percentiles of parental wage earnings distribution. For most of the distribution, wage earnings (which are accurately measured in the LEHD) are the primary source of both income and earnings. However, this is not true at the bottom of the distribution. Below the vertical line marks the set of households with no wage earnings (12% of household in this sample have no reported wage earnings). Below the 25^{th} percentile, alternative sources of income start becoming an increasingly more important source of total household income, so much so that households with zero reported wage earnings actually have higher average total income relative to households who have positive, but little, wage earnings. Most importantly, since my focus is on earnings, self-employment (not captured in the LEHD) is a main source of earnings for parents at the bottom of the wage earnings distribution. Wage earnings is the primary source of income for households with total income (as opposed to total wage earnings) that is above the 10th percentile. The same is not true for households with income below the 10th percentile, for whom transfer income is relatively more important. While Figure A.15 seems to indicate that wage earnings represent the primary source of earnings at the top of the distribution, Smith et al.

(2019) find that non-wage earnings become increasingly important in the top 1% of earners. Taken together, the measure of parental earnings constructed using earnings data from the LEHD should be seen as representative of working families, which excludes roughly the bottom 10% and top 1% of earners.

In order to address the measurement issues in the LEHD, I use an estimation procedure that leverages all of the available data. In particular, I estimate the following regression:

$$y_{it} = \alpha_i + \beta^g X_{it} + u_{it} \tag{A.3}$$

where is is the individual, t is the quarter, y is total quarterly earnings, α is an individual fixed effect and X is vector that consists of a third order polynomial in age. To allow for a flexible age earnings profile, I estimate this specification separately for groups, g, defined by the interaction between sex, race/ethnicity (White non-Hispanic, Black non-Hispanic, Asian non-Hispanic, Hispanic, and other), and state of residence in 2000. The data are a panel that include all strictly positive earnings records between 2000 and 2016 for the parents in the sample. I further restrict the panel to individuals between the ages of 30 and 60 and drop individuals that have fewer than 4 quarters of strictly positive earnings over the entire time period.

I use the estimates from this model to construct a measure of lifetime earnings for each parent. I predict the value of earnings for each quarter between the ages of 35 and 55 and define lifetime earnings as the average of these values. Individuals with either missing or negative values are assigned a lifetime earnings of zero. For single-headed households parental earnings is simply the lifetime earnings of the parent. For two-parent households, parental earnings is the average of the lifetime earnings of both parents.[6]

Much of the analysis relies on percentile ranks of parental earnings. Thus, it is

[6]The choice to take the average earnings across parents is in line with the assumptions made by Chetty et al. (2014).

Figure A.14: States Participating in the LEHD Program

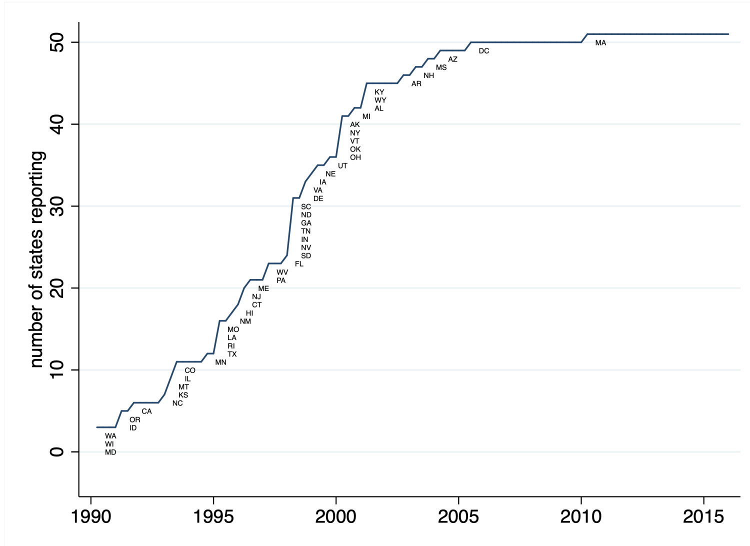

Notes: The figure plots the number of states that are reporting to the Longitudinal Household-Employer Dynamics (LEHD) program in a given year. The abbreviations below the solid line represent the states that begin reporting in that year.

Figure A.15: Source of Earnings Across the Wage Earnings Distribution

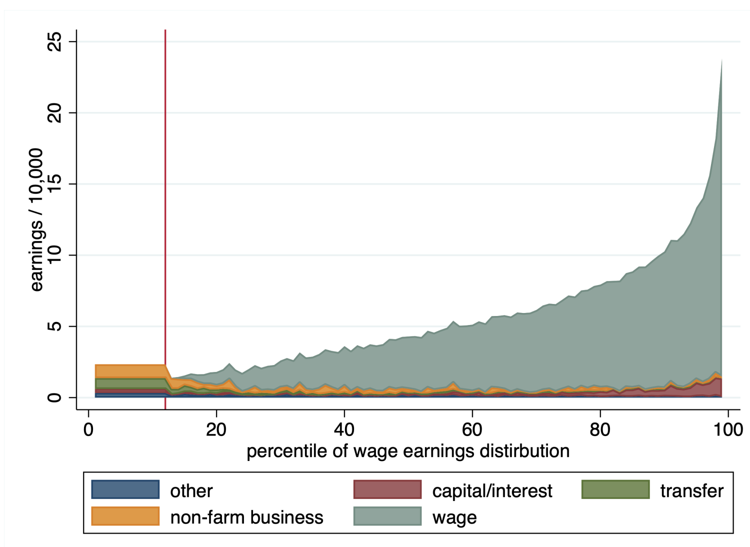

Notes: The figure presents the average household earnings by the percentile of total household wage earnings. Income is broken out into five sources that include: capital/interest, transfer, non-farm business, other and wages. Percentiles below the vertical line have zero wage earnings. The sample includes all households that have at least one child present and excludes the households in the top percentile of the wage earnings distribution due to outlier values.

Source: Author's calculations based on data from the the 2000 March supplement to the Current Population Survey (CPS) and were obtained from IPUMS, see Ruggles et al. (2019).

Figure A.16: Parental Earnings and Neighborhood Characteristics

(A) Poverty

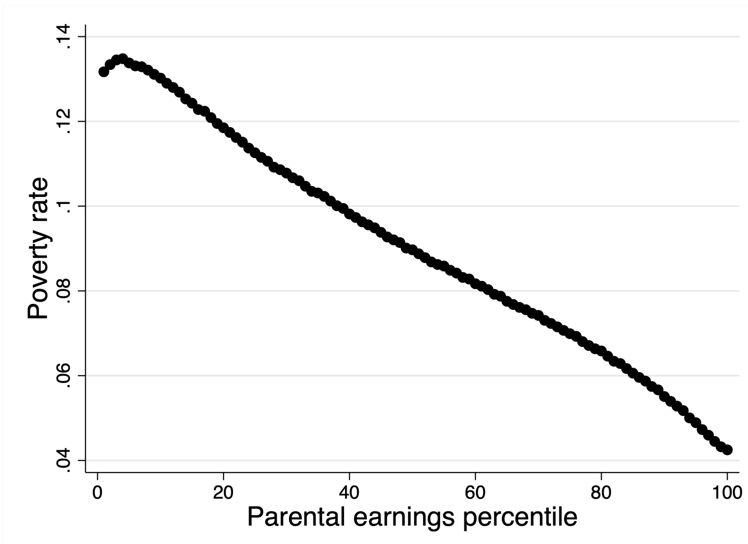

(B) Median Income

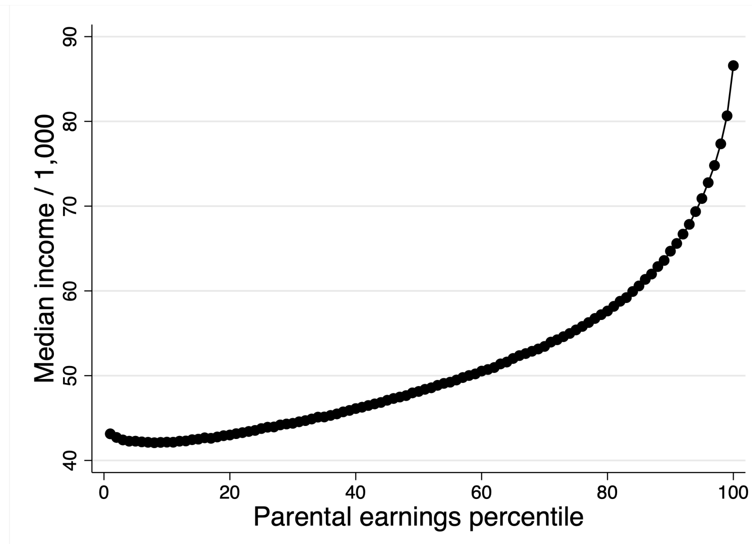

Notes: The figure plots the average characteristic of the census block group of residence in 2000 for each percentile of the parental earnings distribution. The characteristics in Panel A and B are poverty rate and median income, respectively.

Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

critical that the estimates of lifetime earnings preserve the rank of the true values of lifetime earnings. While I do not have an objective measure of lifetime earnings against which to validate my measure, I do have other proxies. In particular, I use the HCEF to identify the census block group in which all households reside in 2000 and measure characteristics of those neighborhoods. I focus on poverty rate and median income, since these are likely to be correlated with lifetime earnings. Figure A.16 plots the average value of these neighborhood level variables against the percentile of the lifetime earnings distribution (percentiles are calculated within cohorts of children). If all measures are proxies of lifetime earnings then there should be a monotonic relationship between the variables. The figure illustrates that this is true for most of the distribution. The one exception is that very bottom of the distribution, where parental earnings may be measured with more error. But overall, the figure indicates a strong relationship between the measure of parental earnings used in this paper and other measures of economic status and thus should alleviate concerns related to measurement error.

If the imputed measure of parental earnings is a multiple of the true lifetime earnings value, then the estimates of IGE will be unaffected. However, if the error is not multiplicative, or differs across individuals, then measurement error may affect the estimates of IGE. A main concern is that my measure is unable to account for differences in labor force participation. By failing to account for periods of nonemployment, my measure will produce artificially high levels of lifetime earnings for parents who have many periods of zero earnings. This may reduce the elasticity of the initial earnings of a child with respect to parental earnings in the lower parts of the distribution. For this reason, it is useful to compare the results with elasticities to those using percentiles. It is worth pointing out that the issue of measuring the earnings of low-income households is not unique to my setting. For example, Chetty et al. (2014) find that their estimates of the IGE are sensitive to the inclusion of the households below the 10^{th} percentile.

A.2.5 Grouping Industries into Sectors

I group two-digit North American Industry Classification System (NAICS) industry codes into three distinct sectors, which are defined below. The unskilled service sector includes: retail trade (44,45); administrative and support and waste management and remediation services (56); arts, entertainment and recreation (71); accommodation and food services (72); and other services (81). The skilled service sector includes: information (51); finance and insurance (52); real estate and rental and leasing (53); profession, scientific and technical services (54); management of companies and enterprises (55); educational services (61); health care and social assistance (62); and public administration (92). The manufacturing/production sector includes: agriculture, forestry, fishing and hunting (11); mining, quarrying, and oil and gas extraction (21); utilities (22); construction (23); manufacturing (31,32,33); wholesale trade (42); and transportation and warehousing (48,49).

A.2.6 Employer and Industry Pay Premiums

In order to estimate the earnings-premium associated with specific employers, I use the methodology developed by Abowd et al. (1999), or commonly referred to as the AKM model. Specifically, I estimate the following specification,

$$y_{it} = \alpha_i + \Psi_{j(i,t)} + X_{it}\beta + \epsilon_{it} \tag{A.4}$$

where i is the individual; t is the year; y is the log of average quarterly earnings; X_{it} is a vector of time varying controls that include a fixed effect for the year and a third order polynomial in age interacted with sex and education; α_i is an individual fixed effect; $\Psi_{j(i,t)}$ is a fixed effect for the employer of i in time t; and ϵ_{it} is a regression residual.[7] The estimate, $\hat{\Psi}_{j(i,t)}$, is a time-invariant measure of the employer pay premium (measured in quarterly earnings).

I estimate this specification using a national sample that includes all earnings records from the LEHD measured between the years 2000 and 2016 and workers between the ages of 25 and 40. I retain jobs that provide over half of the earnings for that year and calculate quarterly earnings as the average of full-quarter earnings for a given employer within the year.[8] Due to computational constraints, I estimate the specification separately within 15 mutually exclusive samples defined by the 9 census divisions and the six largest states (CA, TX, FL, NY, PA, IL). As is standard in the literature, I restrict the sample to the largest connected set within each of these samples. In order to account for the fact that the level of firm pay premiums are not comparable across estimates from distinct samples, I follow Gerard et al. (2018) and normalize all employer fixed effects by subtracting the mean value of the fixed effect for employers in the accommodation and food services industry. Intuitively, this normalization assumes that employers in this industry offer a pay premium of zero, on average.

I am unable to compute the employer pay premium for employers that lie outside of the largest connected set within each of the 15 mutually exclusive samples. In practice this happens in a very small fraction of cases. In order to avoid disclosure issues related to releasing results on multiple samples, I impute missing data with the mean value of individuals who do not work at the employer of a parent and include a control for imputed values in the empirical specification.

[7] Identification of the age and time effects in the presence of individual fixed effects is achieved by following Card et al. (2013) and omitting the linear age term in for each sex by education group and using a cubic polynomial in age minus 40. This normalization assumes that the age-earnings profile is flat at age 40. While the normalization affects the estimates of the individual fixed effects and the covariate index $X_{it}\beta$, the employer fixed effects are invariant to the normalization used. Data on education comes from the individual characteristics file and is sourced from various surveys and is imputed for many observations.

[8] Outliers in the earnings data are dealt with using the same methodology described in Appendix Section A.2.3.

I estimate the industry-level premium using the same data and methodology except I replace the employer fixed effect with a fixed effect for the industry code. Because all industries are connected through worker mobility, estimation is performed on the national sample but to ease computational burden, I take a random 10% subsample of workers. I am able to estimate an industry-level pay premium for all industries, and thus there are no missing data for this variable.

A.2.7 Employer- and Firm-Level Variables

A.2.7.1 Poaching Hires

For each employer I calculate the share of new stable hires that are acquired through poaching flows as opposed to nonemployment flows. In order to explain how poaching rates are constructed, it is useful to establish the following terminology. Each worker with positive earnings in quarter t can have one of four types of employment spells defined in Table A.17, where "+" denotes positive earnings and "0" denotes zero earnings at the employer at quarter t.

Table A.17: Classification of Employment Spells

	earnings at employer		
	t-1	t	t+1
beginning of quarter	+	+	0
end of quarter	0	+	+
middle of quarter	0	+	0
full quarter	+	+	+

A worker with a beginning of quarter employment spell is relatively attached to the employer at the start of quarter t but separates from the employer at some point during quarter t. Similarly, a work with an end of quarter employment spell joins the employer at some point during quarter t and experiences a stable spell of employment that continues into the following quarter. Middle of quarter employment spells represent spells that begin and end within the quarter and, following the conventions used to construct the Job-to-Job Flows statistics, I do not use them when constructing poaching rates.

Workers who experience an end of quarter employment spell in quarter t are defined as stable new hires. These workers begin their employment spell at some point during quarter t, and I define the hire as a poaching hire if the worker also left their previous employer in quarter t. In other words, a poaching hire is an individual who switches employers and begins their new job no later than one quarter after leaving their old job. In practice, I identify poaching hires as individuals who experience an end of quarter employment spell in quarter t and experience either a full quarter or end of quarter employment spell (at a different employer) in quarter t-1. All stable new hires that do not meet these criteria are defined as hires from nonemployment.

For each employer, I calculate the total number of stable hires made through poaching and nonemployment flows between 2000 and 2016. I then calculate an employer-level poaching rate as the proportion of stable new hires made through poaching flows over the entire period. Lastly, I rank employers from 0 to 100 based on their poaching hire rate, where the ranks are calculated using average employer size as weights.

A small fraction of employers have insufficient observations to calculate this measure. In order to avoid disclosure issues related to releasing results on multiple samples, I impute missing data with the mean value of individuals who do not work at the employer of a parent and include a control for imputed values in the empirical specification.

A.2.7.2 Average Earnings

I calculate average earnings at the employer using full quarter employment spells. Specifically, using data between 2000 and 2016, I retain all workers who experience a full quarter employment spell and take the log of their earnings (I top code earnings at \$1,000,000 to mitigate the impact of outliers). The employer-level average of log earnings is simply the average of the quarterly earnings records. I rank employers from 0 to 100 based on their average log earnings, where the ranks are calculated using average employer size as weights. There are no missing data for any of the employers in the sample.

A.2.7.3 Productivity

The firm-level measure of productivity is based on data from the Revenue Enhanced Longitudinal Business Database (RE-LBD). The RE-LBD supplements the LBD with revenue data from the Census Business Registrar (BR). The BR contains annual measures of revenue measured at the tax reporting or employer identification number (EIN) level. Haltwanger et al. (2016) describe how the revenue data and the employment data from the LBD are combined to construct firm level measures of log revenue per worker, which represent the measure of productivity.

There are two limitations of this particular measure of productivity. First, the coverage is not universal since the employment and revenue data for some firms cannot be linked and since the coverage excludes non-profit firms and firms in the Agriculture, Forestry, Fishing and Hunting (NAICS=11) and Public Administration (NAICS=92) industries. Haltwanger et al. (2016) show that the revenue data cover about 80% of firms in the LBD and patterns of missing productivity data are only weakly related to observable firm characteristics. Second, the revenue per worker measure fails to account for differences in intermediate inputs across industries, which imply that this measure cannot be used to compare productivity of firms that are located in different industries.

In order to overcome the latter limitation, I follow Haltiwanger et al. (2017) and construct a time invariant measure of productivity. Specifically, after attaching firm productivity to the employer-level dataset, I calculate average productivity for each employer as the employment-weighted average of log revenue per worker observed across all periods. From each employer I then subtract the employment-weighted average of productivity at the level of the four-digit NAICS industry code. Thus, this measure of productivity is a time invariant measure that captures the productivity of an employer relative to other employers in the same industry. Productivity ranks that range from 0 to 100 are calculated within four-digit industry codes and are employment weighted, where employment refers to the average number of employees at the employer observed over the sample period.

A.2.7.4 Firm Age and Size

Measures of firm age and firm size are derived from the Longitudinal Business Database (LBD).[9] The LBD is an annual dataset that covers the universe of establishments and firms in the US non-farm business sector with at least one paid employee. Establishment-level employment is measured as the number of workers on payroll in the pay-period that covers the 12th day of March in the previous year. Firm size is simply the sum of employment at all establishments within the firm. Firm age measures the number of years since the firms formation and accounts for changes in firm identifiers as well as mergers and acquisitions.[10]

[9] See Jarmin and Miranda (2002) for a detailed description of the LBD and Haltiwanger et al. (2014) for a description of how firm-level outcomes from the LBD are linked to the employers int he LEHD.

[10] See Davis et al. (2007) for a detailed description of how the firm age variable is constructed.

A.3 Approximation Methodology

By definition, $\text{cov}(D_i\beta_i, y_p) = \mathbb{E}[D_i\beta_i y_p] - \mathbb{E}[D_i\beta_i]\mathbb{E}[y_p]$. By iterated expectations,

$$\mathbb{E}[D_i\beta_i] = \mathbb{E}\big[\mathbb{E}[D_i\beta_i|D_i]\big] = \mathbb{E}[D_i]\mathbb{E}[\beta_i|D_i = 1] \tag{A.5}$$

and

$$\mathbb{E}[D_i\beta_i y_p] = \mathbb{E}\big[\mathbb{E}[D_i\beta_i y_p|r_p]\big] \tag{A.6}$$

where r_p is the percentile rank of parental earnings. Because the Pearson correlation coefficient is bounded between -1 and 1, it follows that,

$$cov(D_i\beta_i, y_p|r_p)^2 \leq var(D_i\beta_i|r_p) \times var(y_p|r_p) \tag{A.7}$$

In practice, I condition on r_p, but one could think to condition on more detailed ranks. As the number of ranks approaches the sample size, $var(y_p|r_p)$ approaches zero and the covariance term therefore approaches zero. Thus,

$$\begin{aligned} \mathbb{E}[y_p D_i\beta_i|r_p] &= \mathbb{E}[y_p|r_p] \times \mathbb{E}[D_i\beta_i|r_p] + cov(D_i\beta_i, y_p|r_p) \\ &\approx \mathbb{E}[y_p|r_p] \times \mathbb{E}[D_i\beta_i|r_p] \end{aligned} \tag{A.8}$$

where equation A.7 suggests that $cov(D_i\beta_i, y_p|r_p)$ will be close to zero when conditioned on parental earnings ranks that are defined at a sufficiently high level of detail. Combing these pieces yields the approximation in equation 1.3.

I assess the performance of the approximation methodology by using the same methodology to approximate the observed IGE. By definition, $\rho(y_{ij}, y_p) = \frac{\text{cov}(y_{ij}, y_p)}{\text{var}(y_p)}$. The variance term, $\text{var}(y_p)$, is directly observed and I use the following approximation for the covariance term,

$$\text{cov}(y_{ij}, y_p) \approx \mathbb{E}\Bigg[\mathbb{E}[y_p|r_p] \times \mathbb{E}[y_{ij}|r_p]\Bigg] - E[y_p] \times E[y_{ij}] \tag{A.9}$$

Where this approximation relies on the same assumption used to derive equation 1.3. Table A.18 compares the estimates of the IGE from the micro data, in Panel A, to the approximated values, in Panel B. The approximated values are virtually identical to the actual values, which suggests that the methodology performs well in this context.

Standard errors for the counterfactual estimates in Table 1.7 are estimated via the delta method. Specifically, let

$$\begin{aligned}\Gamma(\vec{B}) &= \frac{\rho(y_{ij}, y_p) - \rho(y_{ij(0)t}, y_p)}{\rho(y_{ij}, y_p)} \times 100 \\ &= (\frac{100}{\rho(y_{ij}, y_p) var(y_p)}) \sum_{q=1}^{5} \hat{\beta}^q [\frac{1}{100}(\sum_{k=(q-1)*20+1}^{q*20} \mathbb{E}[y_p|r_p = k]\mathbb{E}[D_i|r_p = k]) - \mathbb{E}[y_p]\mathbb{E}[D_i]/5]\end{aligned} \tag{A.10}$$

where $\vec{B} = [\hat{\beta}^1, \hat{\beta}^2, \hat{\beta}^3, \hat{\beta}^4, \hat{\beta}^5]$ is a 1×5 vector where the components are the effects conditional on parental earnings for the five parental earnings quintiles. Then we have,

$$\frac{\partial\Gamma(\vec{B})}{\partial\beta^q} = \frac{100}{\rho(y_{ij}, y_p) var(y_p)} \times \left[\frac{1}{100} \sum_{k=(q-1)*20+1}^{q\times 20} \big(E[y_p|r_p = k]E\big[D_i|r_p = k\big]\big) - E[y_p]E[D_i]/5\right] \tag{A.11}$$

Assuming independence between the β^k estimates, leads to the following expression by the delta method,

$$se(\Gamma(\vec{B})) = \sum_{k=1}^{5} var(\beta^k) \times [\frac{\partial\Gamma(\vec{B})}{\partial\beta^k}]^2 \tag{A.12}$$

where $var(\beta^k)$ is simply the square of the standard error from Table 1.6.

Table A.18: Approximation of the Intergenerational Elasticity of Earnings

	(1)	(2)	(3)
A. Individual-Level Data			
$\rho(Y_i, Y_{p(i)})$	0.157	0.130	0.143
B. Approximation			
$\rho(Y_i, Y_{p(i)})$	0.155	0.131	0.143
sample	daughters	sons	all

Notes: The results in columns 1-3 corresopnd to daughters, sons, and all children, respectively. Panel A presents the estimated coefficient from a regression of the log of the first full-quarter of earnings at the first job of the child on the log of parental earnings. The regression is estimated via weighted least squares with sample weights applied. Panel B presents the approximations of the values in Panel A.

Source: Author's calculations based on data from the Longitudinal Employer-Household Dynamics and 2000 Decennial Census files.

A.4 Stylized Model

This section synthesizes my findings by developing a stylized model that describes how the intergenerational transmission of employers affects intergenerational mobility. Relative to other models of intergenerational mobility, the novel features of my model are that I: (1) incorporate a employer-specific component into individual earnings and (2) explicitly model the choices that lead individuals to work for a parent's employer. The key insights from the model include:

1. parents affect the earnings of their children not only by shaping the development of their human capital, but also by providing access to higher-paying employers;

2. if working at the parent's employer is determined by choices made by the employer and the child, then there are conditions under which the instrumental variables estimator identifies the average treatment effect for the population that works for their parents employer, even in the presence of selection bias and selection on gains;

3. the effect of the intergenerational transmission of employers on intergenerational mobility is theoretically ambiguous;

4. if parents adjust their investments in the human capital of their children based on their expectations of whether the child will work for their employer, this could either amplify or dampen the implications for intergenerational mobility.

Let y_{ij} denote the log earnings of individual i at employer j. Assume that log earnings are additive in the log of the human capital (h_i), the employer pay premium (f_j), and an idiosyncratic error terms (u_i). Thus,

$$y_{ij} = h_i + f_j + u_i \tag{A.13}$$

The individual component is defined independent of where the individual is employed and employer transmission affects earnings entirely through its effect on the employer pay premium.

Using the notation of the potential outcomes framework, let $j(1)$ denote the parent's employer and let $j(0)$ denote the employer that represents the outside option. The employer pay premium can be written as,

$$f_j = f_{j(0)} + D_i \beta_i \tag{A.14}$$

where D_i is an indicator equal to one if the individual works for their parent's employer and zero otherwise and $\beta_i = f_{j(1)} - f_{j(0)}$ is the effect of working for a parent's employer.

An individual's outside option is related to their human capital. Specifically, the labor market exhibits sorting between workers and firms, characterized by the following equation:

$$f_{j(0)} = \lambda h_i + \nu_i \tag{A.15}$$

where ν_i is an idiosyncratic error term and $\lambda > 0$ indicates that individuals with higher levels of human capital tend to match to employers that offer higher pay premiums. The same matching process applies to parents, but I abstract from the possibility that parents might work for the employers of their parents.[11] Furthermore, the relationship between the human capital of the child and earnings of the parent is characterized by,

$$h_i = x + \theta y_{pj(1)} + \eta_i \tag{A.16}$$

where p denotes the parent of i, η_i is an idiosyncratic error term and $\theta > 0$ implies that human capital is increasing in parental earnings.

[11] Formally, I assume that $D_p = 0$, where p denotes the parent of i. This assumption simplifies the analysis and allows me to write the earnings benefits associated with working for the parent's employer as function of parental earnings and unobserved error terms $\beta_i = (\frac{\lambda}{1+\lambda} - \lambda\theta) y_{pj(1)} + [\lambda/(1+\lambda)](\lambda \nu_p - u_{pj(1)}) - [\lambda x + \lambda \eta_i + \nu_i]$.

Whether a child works for the employer of their parent depends on choices made by both the employer and the child. Let O_i be equal to one if the parents' employer makes a job offer to the child and zero otherwise. The offer decision depends on the instrument, $z_i \in \{z', z''\}$ with $z' > 0 > z''$, and the human capital of the parent and the child. Specifically, $O_i = \mathbb{1}\{\phi h_p + \gamma h_i > z_i\}$, where ϕ and γ could be positive or negative.[12] Let A_i be equal to one if the child would accept a job offer from the parent's firm. The child will choose to accept the offer if the earnings gains, β_i, exceed any costs, c, such that $A_i = \mathbb{1}\{\beta_i > c\}$. The child will work with their parent only if they receive a job offer and it is optimal for them to accept,

$$D_i = \mathbb{1}\{\phi h_p + \gamma h_i > z_i\} \times \mathbb{1}\{\beta_i > c\} \tag{A.17}$$

Unlike the standard selection models, equation A.17 illustrates that selection into treatment depends on the choices of multiple agents.

Combining equations A.13, A.14, A.15 and A.16 yields the following relationship between the earnings of the child, the earnings of the parent and the effect of the transmission of employers,

$$y_{ij} = \alpha_1 + \alpha_2 y_{pj(1)} + D_i\beta_i + \epsilon_i \tag{A.18}$$

where $\epsilon_i = \nu_i + (1+\lambda)\eta_i + u_i$ is an unobserved error term, and where $\alpha_1 = (1+\lambda)x$ and $\alpha_2 = (1+\lambda)\theta$. Equation A.17 illustrates that D_i is related to ϵ_i through the unobserved error terms, implying that estimating equation A.18 via OLS will produce biased estimates with a sign that is theoretically ambiguous.[13]

[12] ϕ might be positive if higher-ability parents have more control over the hiring process because they hold leadership positions, or negative if lower-ability parents work at firms that rely more heavily on networks in the hiring process. γ may be positive if firms are more likely to make a job offer to high ability workers, or negative if parents exert more effort to procure job opportunities for low ability children.

[13] To more clearly see the relationship between D_i and ϵ_i note that the offer and acceptance decisions can be re-written as: $O_i = \mathbb{1}\{(\frac{\phi}{1+\lambda} + \gamma\theta)y_{pj(1)} + \gamma x - \frac{\phi}{1+\lambda}(\nu_p + u_p) + \gamma(x + \eta_i) > z_i\}$ and

Under the assumption that the instrument is orthogonal to the unobserved components of the individual's earnings ($z_i \perp\!\!\!\perp \eta_i, \nu_i, u_i$) and parent's earnings ($z_i \perp\!\!\!\perp \nu_p, u_p$), an instrumental variables estimator that uses z_i as an instrument identifies a local average treatment effect (LATE), which is defined as $\mathbb{E}[\beta_i | D_i(z') < D_i(z'')]$. In the standard one-agent selection framework the LATE will depend on the value of the instruments since the decision-making process directly links the benefits and instruments.

In my context, in which selection into treatment is determined by two agents, this link is potentially broken. The implication is stated in the following proposition,

Proposition 1 *If $\phi = 0$ and $\gamma = 0$, then $O_i \perp\!\!\!\perp \beta_i$ and*

$$\underbrace{E[\beta_i | D_i = 1]}_{ATT} = \underbrace{E[\beta_i | D_i(z') < D_i(z'')]}_{LATE} \tag{A.19}$$

Proof 1 *If $\gamma = 0$ and $\phi = 0$ then $O_i = \mathbb{1}\{0 > z_i\}$ and it follows that $O_i \perp\!\!\!\perp \beta_i$. For any two values of the instrument, $z' > 0 > z''$, it follows that,*

$$\begin{aligned}
\underbrace{E[\beta_i | D_i = 1]}_{ATT} &= E[E[\beta_i | A_i = 1] | O_i = 1] \\
&= E[E[\beta_i | A_i = 1] | O_i(z') < O_i(z'')] \\
&= \underbrace{E[\beta_i | D_i(z') < D_i(z'')]}_{LATE}
\end{aligned} \tag{A.20}$$

where the first and third inequalities hold by the law of iterated expectations and the second inequality holds as a result of $O_i \perp\!\!\!\perp \beta_i$.[14]

If the offer decision is unrelated to the human capital of the parent ($\phi = 0$) and the human capital of the child ($\gamma = 0$), then the offer decision and the earnings gains

$A_i = \mathbb{1}\{(\frac{\lambda}{1+\lambda} - \lambda\theta)y_{pj(1)} + (\frac{\lambda}{(1+\lambda)})(\nu_p/\lambda - u_p) > c + \lambda x + \lambda\eta_i + \nu_i\}$. See Appendix Section A.4.2 for details.

[14]It also exploits the fact that $O_i \perp\!\!\!\perp A_i$, which follows directly from $O_i \perp\!\!\!\perp \beta_i$.

will be independent ($O_i \perp\!\!\!\perp \beta_i$). Under these conditions, the instrument affects the treatment status of a random sample of individuals who would accept job offers at their parent's employer and the LATE is equivalent to the ATT. This equivalence, which may hold even in the presence of selection bias and selection on gains, is possible because treatment status is determined by the choices of multiple agents.

While the empirical evidence suggests that the intergenerational transmission of employers reduces mobility, the relationship is theoretically ambiguous. This is formalized in the following proposition, which states that the counterfactual IGE corresponding to a wold in which no one worked for a parent's employer could be greater or small than the observed IGE.

Proposition 2 *Consider a deterministic case of the model by letting z_i, η_i, ν_i and u_i be equal to zero and let $c \geq 0$. Then the following statements are true:*

- *if $\frac{1}{1+\lambda} > \theta$ and $\phi > -\theta\gamma(1+\lambda)$ then $\rho(y_{ij}, y_{pj(1)}) > \rho(y_{ij(0)}, y_{pj(1)})$*
- *if $\frac{1}{1+\lambda} < \theta$ and $\phi < -\theta\gamma(1+\lambda)$ then $\rho(y_{ij}, y_{pj(1)}) < \rho(y_{ij(0)}, y_{pj(1)})$*

Proof 2 *To prove the results it is useful to start by noting the implications of the deterministic setting (η_i, ν_i, u_i and z_i are set to zero) for the following expressions,*

$$
\begin{aligned}
O_i &= \mathbb{1}\{(\frac{\phi}{1+\lambda} - \theta\gamma)y_{pj(1)} > 0\} \\
A_i &= \mathbb{1}\{(\frac{\lambda}{1+\lambda} - \lambda\theta)y_{pj(1)} - \lambda x > c\} \\
\beta_i &= (\frac{\lambda}{1+\lambda} - \lambda\theta)y_{pj(1)} - \lambda x
\end{aligned}
\tag{A.21}
$$

It is straightforward to show that $cov(\beta_i, y_{pj(1)}) = (\frac{\lambda}{1+\lambda} - \lambda\theta)var(y_{pj(1)})$. In the first case, when $\frac{1}{1+\lambda} > \theta$ and $\phi > -\theta\gamma(1+\lambda)$, it immediately follows that $\frac{\partial\beta_i}{\partial y_{pj(1)}} > 0$, $\frac{\partial O_i}{\partial y_{pj(1)}} > 0$, $\frac{\partial A_i}{\partial y_{pj(1)}} > 0$ and $\frac{\partial D_i}{\partial y_{pj(1)}} > 0$. Under the assumption that $c \geq 0$, D_i and β_i are both increasing in $y_{pj(1)}$, and it follows that $D_i\beta_i$ is a monotonic transformation of

β_i. Thus, $cov(\beta_i, y_{pj(1)})$ and $cov(D_i\beta_i, y_{pj(1)})$ have the same sign, which implies that, $cov(D_i\beta_i, y_{pj(1)}) > 0$. The proof for the second case uses the same logic.

Proposition 2 highlight highlights two competing forces. On the one hand, the transmission of employers will reduce mobility if high income parents are best able to procure high-paying job offers for their children. On the other hand, the transmission of employers will increase mobility if children from low income households have lower levels of human capital and are more reliant on their parents to find work. In contrast to previous theoretical work by Corak and Piraino (2012) and Magruder (2010), which does not model selection into the parent's employer, reasonable arguments can be made that the transmission of employers could either increase or reduce intergenerational mobility, making this relationship theoretically ambiguous. Thus, while my empirical evidence suggests that employer transmission reduces mobility, this conclusion might differ in other contexts depending the characteristics of the labor market and the human capital accumulation process.

A.4.1 Extension with Parental Investment in Human Capital

Within economics, virtually all of the theoretical work on intergenerational mobility builds on the framework of Becker and Tomes (1976, 1986), in which the persistence of economic outcomes across generations is driven by investments human capital that are determined by optimizing behavior on the part of the parents. Even the two papers that have studied the role of parental labor market networks from theoretical perspective, Corak and Piraino (2012) and Magruder (2010), have used this approach. In contrast, I have ignored the decisions related to human capital investment and have instead focused on the component of earnings attributable to employer pay premiums. I refer to these effects on the employer pay premium, which are conditional on the human capital of the children, as the "direct effects." While I argue that this is most important feature to focus on, these channels are not mutually exclusive and may interact in interesting ways. I explore this possibility in this section by extending the stylized model to allow for parents to shape the human capital of their children through investments. I refer to the effects mediated by parental investment decisions as the "indirect effect" of the intergenerational transmission of employers.

I consider a model in the vein Becker and Tomes (1976, 1986) in which parents make decisions regarding the optimal investments of the human capital of their children. For tractability I focus on the deterministic setting (z_i, η_i, ν_i and u_i are equal to zero) and assume that children only accept job offers from their parents when the earnings benefits are positive ($c \geq 0$). Furthermore, I maintain the assumptions underlying equations A.13, A.14 and A.15. However, I do not impose the assumption stated in equation A.16, because the goal of this section is to derive the relationship between parental earnings and the human capital of the child as the result of optimizing behavior on the part of the parents. For notation, I use lower case letters to denote the log of upper case variables (for examples, $h_i = log(H_i)$).

Parents care about their current period consumption, C_p, and the total financial resources of their children, which depends on the earnings of the children, Y_{ij}, and bequests, B_i, plus interest accrued at rate R. Parents solve the following problem:

$$\max_{C_p, C_i, B_i} \{v(C_p) + u(Y_{ij} + RB_i)\} \text{ subject to } C_p + S_i + B_i \leq Y_{pj(1)} \tag{A.22}$$

where S_i represents investment in the human capital of the children and $u(\cdot)$ and $v(\cdot)$ are continuous functions that both have the following properties: $u'(\cdot) > 0$, $u''(\cdot) < 0$ and $u'(0) = \infty$. This setup assumes that there are no credit constraints, as bequests may be negative.

While there are a number of ways to generate intergenerational persistence in earnings in the absence of credit constraints, I follow Becker et al. (2018) and assume that there are complimentarities between the human capital of the parent and the production of human capital of the child. Specifically, investment translates into human capital according to the following production function, $H_i = H_p^{\sigma} S^{\alpha}$. Intuitively, this captures the fact that investments in human capital might be more productive if made by parents with higher ability. I also assume that $\alpha(1+\lambda) < 1$ which implies that there are diminishing returns to parental investment. The optimal level of investment in human capital is defined by the level at which the marginal rate of return is equal to the interest rate, $\frac{\partial Y_{ij}}{\partial S_i} = R$. Combining terms, we can rewrite the expression determining optimal investment as follows,

$$\alpha(1+\lambda)H_p^{\sigma(1+\lambda)} S_i^{\alpha(1+\lambda)-1} exp\{D_i\beta_i\} + H_p^{\sigma(1+\lambda)} S_i^{\alpha(1+\lambda)} \frac{\partial exp\{D_i\beta_i\}}{\partial S_i} = R \qquad \text{(A.23)}$$

where the left-hand side represents the marginal returns to investments in human capital and the right-hand side represents the marginal returns to bequests.

To understand how the transmission of employers shapes the investment decision it is useful to consider three cases. As a starting point consider the case in which parents do not account for employer transmission when making investment decisions ($exp\{D_i\beta_i\} = 1$ and $\frac{\partial exp\{D_i\beta_i\}}{\partial S_i} = 0$). Under these conditions is it straight forward to show that the optimal level of investment is given as:

$$S_i' = [\frac{R}{\alpha(1+\lambda)}]^{1/[\alpha(1+\lambda)-1]} H_p^{\sigma(1+\lambda)/[1-\alpha(1+\lambda)]} \qquad \text{(A.24)}$$

Thus, the optimal level of parental investment is increasing in the human capital of the parent and decreasing in the interest rate and it produces the following relationship between the human capital of the child and the earnings of the parent, $h_i = x + \theta y_{pj(1)}$, where $x = \frac{-\sigma}{1-\alpha(1+\lambda)} log\left(\frac{R}{\alpha(1+\lambda)}\right)$ and $\theta = \frac{\sigma/(1+\lambda)-(1-\alpha)}{1-\alpha(1+\lambda)}$. Note that this linear relationship is exactly the one assumed in Section A.4.

How will this relationship change if parents consider the possibility of helping their child to secure a job within their employer when making investment decisions? In a step towards answering this question, consider a second case in which parents account for the fact that the transmission of employers might affect the level of earnings ($exp\{D_i\beta_i\} \neq 1$) but they do not account for the fact that investments might affect the gains associated with transmission ($\frac{\partial exp\{D_i\beta_i\}}{\partial S_i} = 0$). Under these assumptions, the optimal level of investment is defined as, $S_i'' = S_i' \times exp\{\frac{D_i\beta_i}{1-\alpha(1+\lambda)}\}$ and it follows that,

$$s_i' - s_i = \frac{D_i\beta_i}{1-\alpha(1+\lambda)} \geq 0 \tag{A.25}$$

Because $\exp\{D_i\beta_i\} \geq 0$ and $\alpha(1+\lambda) < 0$, this mechanism leads to an increase in parental investment. Intuitively, the transmission of employers provide access to firms that pay higher wages and thus parents who expect their children to work with them will expect a higher rate of return on investments in human capital.[15]

In the third case I allow for the investment decisions of parents to also depend on the anticipated effects of a rise in human capital on the gains of working for a parent's employer ($\frac{\partial exp\{D_i\beta_i\}}{\partial S_i} \neq 0$).[16] Because $\frac{\partial exp\{D_i\beta_i\}}{\partial S_i} < 0$, it is immediately apparent that if we were to plug in S_i'' into equation A.23 the sum of the terms of the left hand side would be less than the interest rate on the right hand side. Furthermore, under the assumption that $\gamma < 0$, both $\alpha(1+\lambda)H_p^{\sigma(1+\lambda)}S_i^{\alpha(1+\lambda)-1}exp\{D_i\beta_i\}$ and $H_p^{\sigma(1+\lambda)}S_i^{\alpha(1+\lambda)}\frac{\partial exp\{D_i\beta_i\}}{\partial S_i}$ are (weakly) decreasing in S_i, and it follows that the optimal level of investment in case 3 is less than the optimal level in case 2, $S_i''' < S_i''$. In the mechanism highlighted in this case, the intergenerational transmission of employers reduces the incentive to invest in human capital because the earnings gains associated with working the parents' employer are declining in the human capital of the child (both along intensive and extensive margins).

Taken together, the total indirect effect of the intergenerational transmission of

[15]Different assumptions could lead to alternative conclusions. For example, both Corak and Piraino (2012) and Magruder (2010) assume that the effect of networks on earnings is additive in levels, which leads them to conclude that parental investment decisions are unaffected by the presence of parental labor market networks.

[16]As in case 2, I continue to allow for the possibility that $exp\{D_i\beta_i\} \neq 0$.

employers on the level of parental investment is theoretically ambiguous.[17] On the one hand, the transmission of employers will increase the marginal returns to human capital investments by providing access to high-paying firms. On the other hand, the marginal returns are pushed down by the fact that higher-ability children are less likely to work with their parents and gain less conditional on doing so.

The implications for intergenerational mobility are similarly ambiguous. For simplicity, consider the case in which $\theta(1+\lambda) < 1$ and $\phi > -\theta\gamma(1+\lambda)$, which implies that the direct impact of employer transmission will increase IGE. Because these conditions imply that $D_i\beta_i$ is increasing in parental earnings, children from high income families will tend to be the greatest beneficiaries of working with their parents (being more likely to do so and experiencing greater benefits conditional on doing so). The mechanism highlighted in case 2 will amplify the disparities between children from high and low income households while the mechanism highlighted in case 3 will mitigate these differences. The total indirect effect on intergenerational mobility will depend on which force dominates.

A.4.2 Sign of Selection Bias

In order to highlight the empirical challenges created by the unobserved components of earnings, start by decomposing the following estimator into a causal effect and selection bias,

$$\begin{aligned}\underbrace{E[y_{ij}|D_i=1,y_{pj(1)}]-E[y_{ij}|D_i=0,y_{pj(1)}]}_{\text{estimator}} &= E[y_{ij(1)}-y_{ij(0)}|D_i,y_{pj(1)}]+E[y_{ij(0)}|D_i=1,y_{pj(1)}]-E[y_{ij(0)}|D_i=0,y_{pj(1)}]\\ &= \underbrace{E[\beta_i|D_i=1,y_{pj(1)}]}_{\text{ATT}}+\underbrace{E[\epsilon_i|D_i=1,y_{pj(1)}]-E[\epsilon_i|D_i=0,y_{pj(1)}]}_{\text{selection bias}}\end{aligned} \tag{A.26}$$

where $\epsilon_i = (1+\lambda)\eta_i + \nu_i + u_i$. From inspection, ϵ_i will generate selection bias if and only if $cov(\epsilon_i, D_i) < 0$.

[17] This follows from the fact that I have shown that $S_i' \leq S_i''$ and $S_i''' < S_i''$. Thus the total effect (difference between S_i' and S_i''') will depend on whether the mechanism highlighted in case 2 or 3 is stronger.

In order to sign the selection bias term we must rewrite D_i as a function of parental earnings and the idiosyncratic error terms. The assumption that parents do not share an employer with their own parents (the employer of p is $j(1)$) in conjunction with equations A.13, A.14 and A.15 implies that $f_{ij(1)} = \frac{\lambda y_{pj(1)} - \lambda u_i + \nu_p}{1+\lambda}$. Combining this expression with equation A.16 yields,

$$A_i = \mathbb{1}\{(\frac{\lambda}{1+\lambda} - \lambda\theta)y_{pj(1)} + \frac{\lambda}{1+\lambda}(\nu_p/\lambda - u_p) > c + \lambda x + \lambda\eta_i + \nu_i\} \tag{A.27}$$

Equation A.15 implies $h_p = \frac{y_{pj(1)} - \nu_p - u_p}{1+\lambda}$ and combing this with equation A.16 yields,

$$O_i = \mathbb{1}\{(\frac{\phi}{1+\lambda} + \gamma\theta)y_{pj(1)} - \frac{\phi}{1+\lambda}(\nu_p - u_p) + \gamma(x + \eta_i) > z_i\} \tag{A.28}$$

Thus, because $D_i = O_i \times A_i$, we have written D_i as a function of parental earnings and the idiosyncratic error terms.

Inspection of how η_i and ν_i enter equations A.27 and A.4.2 illustrates two potential sources of selection bias. First note that ν_i and O_i are independent while A_i and ν_i are negatively correlated. Thus, ν_i and D_i will be negatively correlated and ν_i will generate negative selection bias. Intuitively, children who receive job offers at low-paying firms will be more willing to accept offers at their parents employers and this will lead us to underestimate the benefits of employer transmission. Second, η_i and A_i are negatively correlated, which again will tend to produce negative selection bias. Intuitively, low-ability children will have more limited outside employment opportunities and will be more willing to work at their parents' employer. However, the relationship between η_i and O_i is ambiguous and will depend on the sign of γ. If $\gamma < 0$ then η_i and O_i will be negatively correlated, which will produce negative selection bias because the low ability children will be more likely to receive job offers. However, if $\gamma > 0$ then η_i and O_i will positively correlated. In this latter case, the effect of η_i on the selection bias term will be ambiguous and will depend on the relative importance of its effect on O_i and A_i.

Chapter B: Appendix Material for Chapter 2

B.1 Additional Figures

Figure B.1: PHA Size in Existing Research and HOPE VI Sample

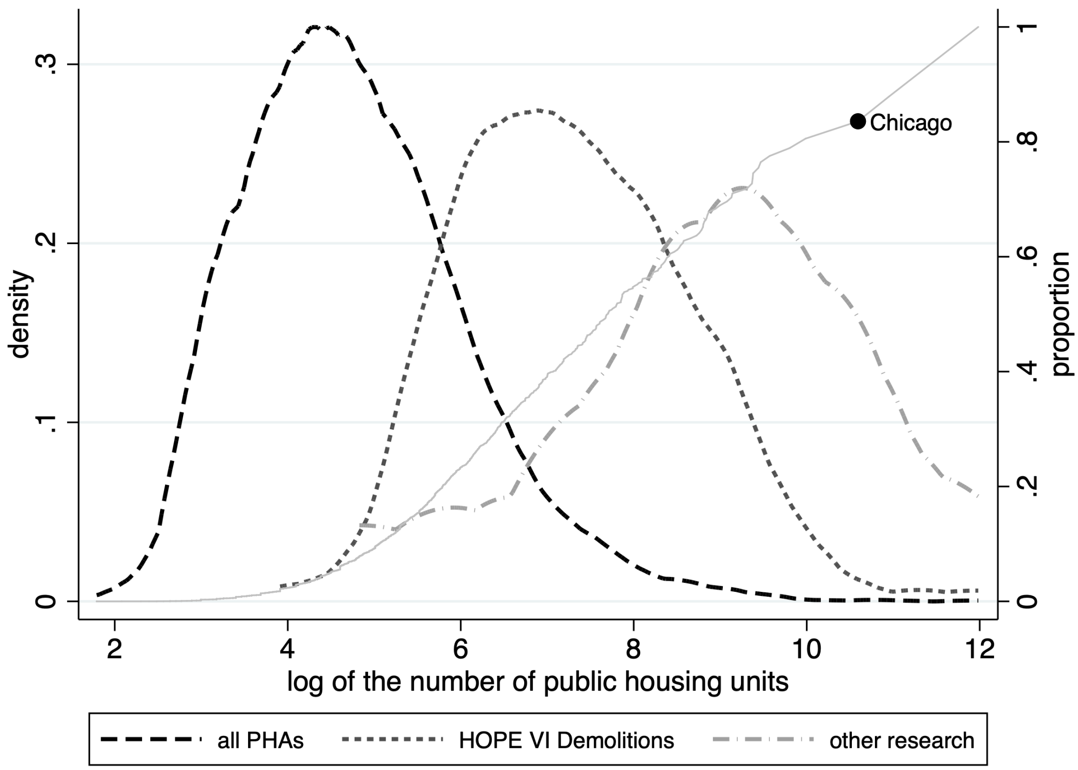

Notes: This figure plots the kernel density of the log of the total number of public housing units (measured in 1997) within each Public Housing Authority (PHA) in the U.S. The densities are plotted for three groups: all PHAs, PHAs that received a HOPE VI demolitions grant, and PHAs that participated in other randomized controlled studies. The PHAs from the other studies include: Baltimore MD, Boston MA, Chicago IL, Los Angeles CA, and New York NY from the Moving to Opportunity experiment (Ludwig et al. 2013); Chicago IL from the Gautreaux program (Rosenbaum 1995); and Atlanta GA, Augusta GA, Fresno CA, Houston TX, Los Angeles CA, and Spokane WA from the Effects of Housing Choice Voucher on Welfare Families project (Mills et al. 2006). The solid grey line presents the proportion of total public housing units in the U.S. that are located in a PHA with fewer than the number of housing units indicated on the horizontal axis. The black marker indicates the size of the Chicago PHA, which is the setting for Jacob (2004) and Chyn (2018).
Source: Author's calculations based on data from HUD USER.

Figure B.2: Baseline Differences between HOPE VI and non-HOPE VI Observations

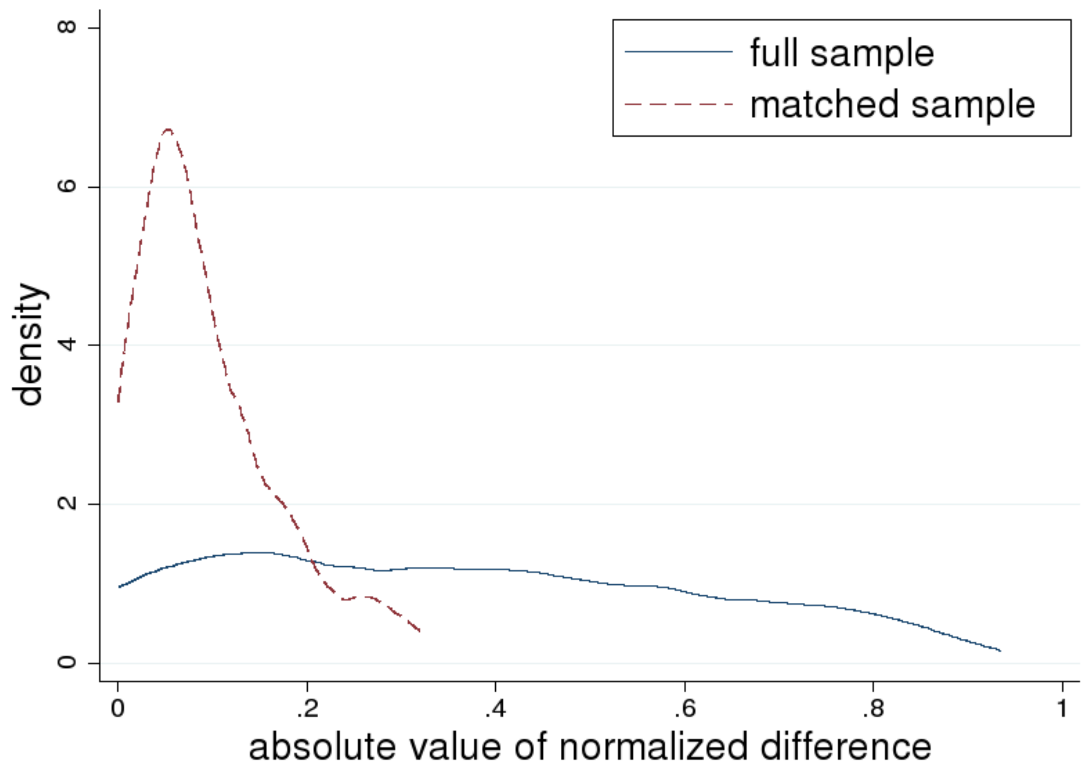

Notes: The figure presents kernel densities of normalized differences of all 92 baseline covariates calculated from the full and matched samples (see Appendix C for a description of the variables). To account for the clustered nature of the data within projects, we collapse data to the average value at the project-level to calculate the normalized difference for each variable. The normalized difference is defined as $(\bar{x}_1 - \bar{x}_0)/(\sqrt{(s_1^2 + s_0^2)/2})$, where $\bar{x}_d$ and s_d^2 is the sample average and variance for the HOPE VI ($d = 1$) and non-HOPE VI ($d = 0$) samples, respectively.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

Figure B.3: Baseline Differences within Strata

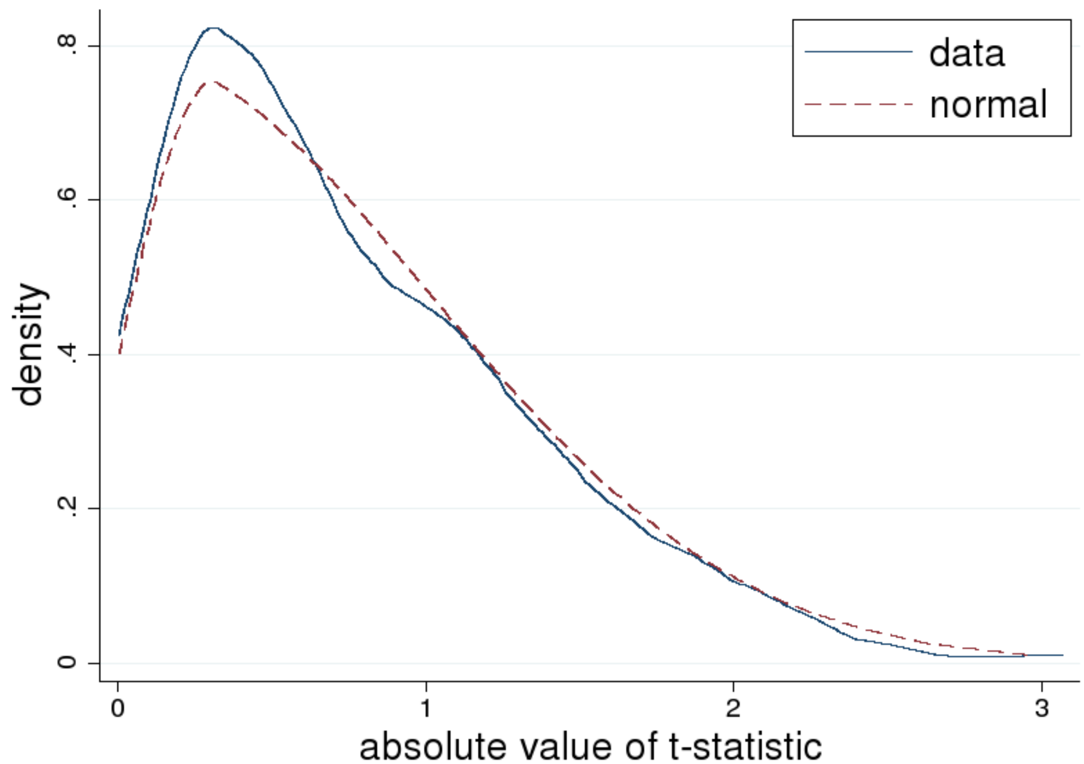

Notes: The figure presents the distribution of the absolute value of t-statistics obtained from regressing a baseline variable on an indicator for HOPE VI within each stratum. The t-statistics are calculated using the household-year dataset when the the baseline variable is measured at the household-year level and using the child-year dataset for all other variables. Standard errors are clustered at the project level. With 92 baseline variables (see Appendix C for a description of the variables) and 7 strata, the figure summarizes the distribution of 644 t-statistics. To aid interpretation, we also plot the distribution of the absolute value of t-statistics from a normal distribution. All statistics are calculated on the matched sample.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

Figure B.4: Evolution of Effect of HOPE VI on Earnings in Large PHAs

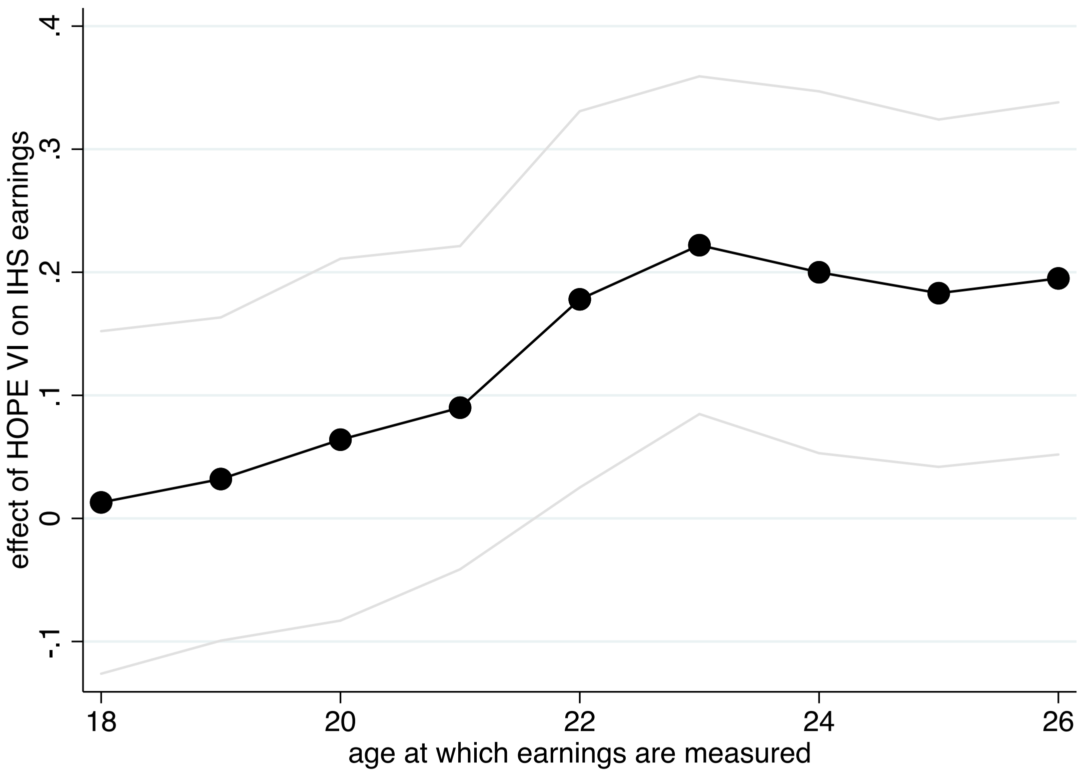

Notes: The figure presents estimates of the effect of HOPE VI on the inverse hyperbolic sine (IHS) of annual earnings measured in the year in which the child turns 18-26. Effects on earnings are estimated using the stratification with regression estimator where all stratum-level regressions control for the base year in which the child appears in public housing, the year in which earnings are measured and standard vector of individual- and project- level characteristics. Standard errors are clustered at the project level and the gray line indicates the 95% confidence interval. Estimates are for large PHAs only.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

Figure B.5: Baseline Differences between HOPE VI and Failed Applicants

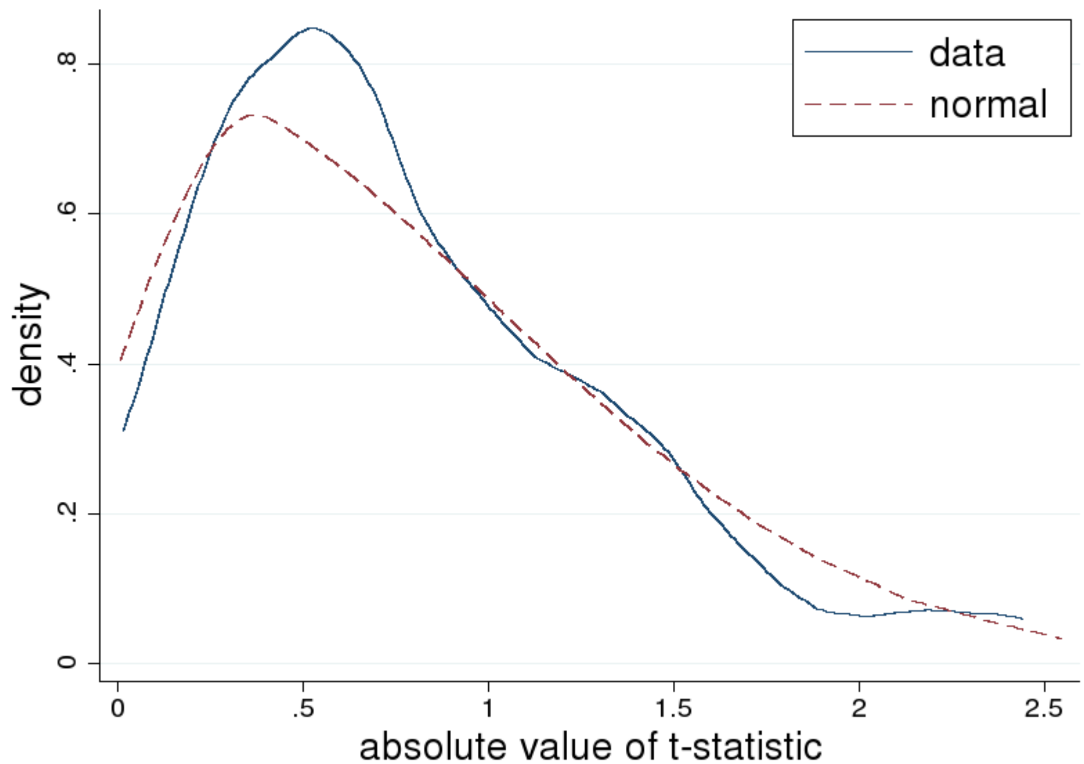

Notes: The figure presents the distribution of the absolute value of t-statistics obtained from regressing a baseline variable on an indicator for HOPE VI. Regressions are estimated separately for large and small PHAs and the set of non-HOPE VI projects includes only the failed applicants. The t-statistics are calculated using the household-year dataset when the the baseline variable is measured at the household-year level and using the child-year dataset for all other variables. Standard errors are clustered at the project level. With 92 baseline variables (see Appendix C for a description of the variables) and regressions with large and small PHAs, the figure summarizes the distribution of 184 t-statistics. To aid interpretation, we also plot the distribution of the absolute value of t-statistics from a normal distribution.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

Figure B.6: Distribution of Earnings for Household Heads and Children

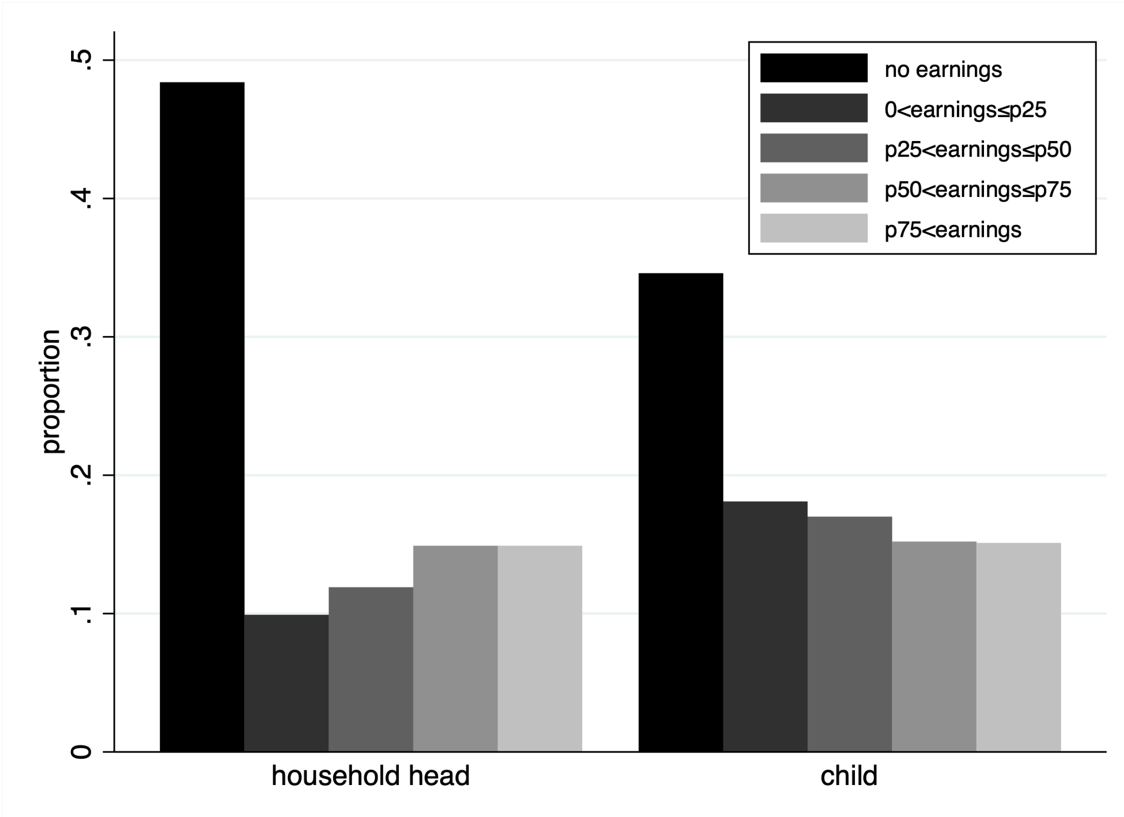

Notes: The figure presents the proportion of household heads and children whose earnings are zero or within a given quartile of the overall distribution of positive earnings. In the legend, the notation p25 denotes the 25th percentile. Parental earnings are measured 10 years after the reference year whereas the earnings of children are measured in the year they turn 26.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

B.2 Additional Tables

Table B.1: Construction of the Sample of HOPE VI Projects

			Number of Projects	
	Criteria to Drop from Sample	**Justification**	**Dropped**	**Remaining**
1	located in U.S. territory	Our interest is in understanding how HOPE VI demolitions affected the long-term outcomes for children in the U.S.	4	281
2	defined as a scattered-site in HOPE VI award list	The majority of scattered-sites demolished under HOPE VI had fewer than 5 units demolished, although there were some scatter-sites with substantially more units demolished. In total, less than 2% of the total units demolished under the HOPE VI program were scattered sites. While there is no formal definition of scatter-site housing, the two key characteristics are low density and low concentration (Hogan 1996).	16	265
3	previously received a HOPE VI demolitions grant	Some projects received more than one grant and we limit our focus to the first awards.	19	246
4	unable to assign a project ID	The project ID is required to link records to the administrative data.	3	243
5	data irregularities due to MTW	HUD's Moving to Work (MTW) demonstration program exempted participating local housing authorities from HUD reporting requirements see Abravanel et al. (2004).	≈50	≈190
6	fewer than 15 occupied units or more than 3000 occupied units	The lower threshold limits the sample to larger public housing projects while the upper threshold ensures that we do not mistakenly group together spatially disparate projects.	≈20	≈170
7	senior housing	Given our focus on children, we drop projects if they are senior housing (over 80% of residents are above 55 years of age) or if they have no children ages 10-18 residing in them.	≈10	≈160

Notes: This table describes the sample selection criteria that reduce the full set of 285 HOPE VI demolition awards to the (approximately) 160 awards studied in this paper. The columns present the restriction applied to the sample, the justification for imposing this restriction, the number of projects affected and the number of projects remaining. The symbol, ≈, denotes that the count is rounded according to disclosure avoidance rules of the U.S. Census Bureau.

Table B.2: Sample Size within Strata

	PHA		Project		Households		Children	
	HOPE VI	Control	HOPE VI	Control	HOPE VI	Control	HOPE VI	Control
Panel A. Large PHAs								
stratum I	20	20	30	40	4,000	8,500	7,000	15,000
stratum II	≤ 15	20	20	40	1,500	7,000	2,500	13,000
stratum III	≤ 15	50	20	100	1,500	17,000	2,000	32,000
total	40	90	70	180	7,000	26,200	11,500	60,000
Panel B. Small PHAs								
stratum IV	40	60	50	60	2,500	4,000	4,000	7,000
stratum V	≤ 15	40	≤ 15	40	500	3,500	800	6,000
stratum VI	≤ 15	80	≤ 15	90	400	7,500	700	13,000
stratum VII	20	200	30	200	800	13,000	1,500	22,000
total	80	380	90	390	4,200	28,000	7,000	48,000

Notes: The table summarizes the sample size within each stratum. Projects in large and small Public Housing Authorities (PHAs) are divided into three and four distinct strata, respectively, resulting in a total of 7 strata. Panel A and B present counts for the large and small PHA's, respectively. The even and odd columns present the counts of HOPE VI and non-HOPE VI projects, respectively. Columns 1 and 2 present the unique number of PHAs, columns 3 and 4 present the unique number of projects, columns 5 and 6 present the unique number of households, and columns 7 and 8 present the unique number of children. Numbers below: 100 are rounded to the nearest 10, 999 are rounded to the nearest 50, 9,999 are rounded to the neaest 100 and 99,999 to the neaest 500. The total row is calculated using the numbers presented in the table and we use 10 to impute the count of cells with fewer the 15 observations (except for the count of HOPE VI projects in small PHAs, since we released the total count of HOPE VI projects as 160).
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

Table B.3: Covariate Balance After Matching

	Without Strata		Across Strata		Within Strata	
	number of tests (1)	$\|t\| > 1.645$ (2)	number of tests (3)	$\|t\| > 1.645$ (4)	number of tests (5)	$\|t\| > 1.645$ (6)
Large PHAs						
Project	16	0.000	16	0.000	48	0.146
PHA	19	0.105	19	0.105	57	0.140
Neighborhood	35	0.029	35	0.000	105	0.114
Child	11	0.091	11	0.091	33	0.091
Household Head	11	0.091	11	0.091	33	0.061
All Variables	92	0.054	92	0.043	276	0.116
Small PHAs						
Project	16	0.000	16	0.000	64	0.078
PHA	19	0.053	19	0.158	76	0.092
Neighborhood	35	0.029	35	0.057	140	0.043
Child	11	0.182	11	0.182	44	0.068
Household Head	11	0.091	11	0.182	44	0.045
All Variables	92	0.055	92	0.098	368	0.062

Notes: Panel A and B present results for large and small Public Housing Authorities (PHAs), respectively. Each row summarizes a number of balance tests for the category of baseline variables defined by the row label (project, PHA, neighborhood, child and household head). Appendix C defines all 92 baseline variables that appear in the five categories. The even numbered columns summarize the proportion of t-statistics that are greater in absolute value than 1.645. The odd numbered columns present the number of t-statistics that contribute to calculating this proportion. All t-statistics are calculated by estimating a model in which the outcome is a baseline variable and the regoressor is an indicator for HOPE VI. The t-statistic is equal to the estimated coefficent divided by the standard error, where standard errors are clustered at the level of the project. Ordinary Least Squares (OLS) is used to estimate the t-statistics summarized in column 2. These regressions are estimated within the large and small PHA samples, but do not account for the stratification structure. The stratification with regression estimator is used to estimate the t-statistics summarized in column 4. Column 6 summarizes the t-statistics obtained from estimating OLS regressions within each of the strata. All estimates are calculated using a child-level dataset, except for the variables the correspond to the head of household household, which are estimated using the head of household-level dataset.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.

Table B.4: Placebo Test Using Failed Applicants

	Large PHAs		Small PHAs	
	qrtrs worked (1)	IHS earnings (2)	qrtrs worked (3)	IHS earnings (4)
Failed Applicant	-0.058	-0.142	-0.028	-0.070
	(0.037)	(0.091)	(0.027)	(0.071)
control	2.200	6.410	2.220	6.420
	[1.720]	[4.490]	[1.730]	[4.490]
observations	99,000	99,000	124,000	124,000

Note: Failed applicants are projects that applied for but never received HOPE VI funding (either the Revitalization or Demolition program). Columns 1-2 and 3-4 present estimates from the stratification with regression estimators from large and small Public Housing Authorities (PHAs), respectively. The outcome variables in the odd and even numbered columns are annual labor market outcomes measured in the year in which the child turns 26 including number of quarters worked and the inverse hyperbolic sine (IHS) of earnings, respectively. All stratum-level regressions control for the base year in which the child appears in public housing, the year in which the child turns 26, and standard vector of individual- and project-level characteristics. Standard errors are clustered at the project-level and are presented in parentheses. The mean and standard deviation, presented in brackets, of the outcome for the control group are a weighted aggregate of stratum-level statistics, where the weights are proportional to the number of treated individuals in a strata.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$.

Table B.5: Pseudo Outcomes Analysis

	effect of HOPE VI on row variable		
	(1)	(2)	(3)
Household-Level Variables			
age	0.242	0.493	0.333
	(0.270)	(0.303)	(0.204)
black	0.026	0.054	0.036
	(0.042)	(0.040)	(0.030)
Hispanic	0.026	0.004	0.018
	(0.045)	(0.025)	(0.030)
white	0.001	-0.049**	-0.017
	(0.014)	(0.024)	(0.013)
other non-Hispanic	-0.014	0.013	-0.004
	(0.021)	(0.012)	(0.014)
dependents	-0.001	0.018	0.006
	(0.023)	(0.020)	(0.016)
disability	-0.005	0.006	-0.001
	(0.007)	(0.009)	(0.005)
household size	0.037	0.013	0.028
	(0.025)	(0.018)	(0.017)
female	0.004	-0.003	0.002
	(0.007)	(0.005)	(0.005)
married	0.008	0.002	0.006
	(0.005)	(0.006)	(0.004)
income	-0.936***	-0.790***	-0.883***
	(0.293)	(0.240)	(0.206)
Child-Level Variables			
age	0.087**	0.018	0.062**
	(0.039)	(0.042)	(0.029)
black	0.035	0.057	0.043
	(0.043)	(0.040)	(0.031)
Hispanic	0.022	0.001	0.015
	(0.038)	(0.026)	(0.026)
white	0.001	-0.050**	-0.017
	(0.013)	(0.024)	(0.012)
other non-Hispanic	-0.024	0.015	-0.010
	(0.025)	(0.012)	(0.017)
disability	0.003	-0.005	0.000
	(0.005)	(0.003)	(0.003)
female	0.000	-0.005**	-0.002
	(0.002)	(0.002)	(0.001)
Sample of PHAs	large	small	all

Notes: Columns 1-3 present estimates from the stratification with regression estimator for large, small, and all Public Housing Authorities (PHAs), respectively. Each row presents the results from a specification in which the variable listed in the row is the pseudo outcome. For each pseudo outcome, the entire matching procedure is implemented but the pseudo outcome (or any variable constructed using this variable) is omitted from the process. The results presented in the table are coefficients from a stratification with regression estimator, which regresses the pseudo outcome on an indicator for HOPE VI and the set of standard covariates (we omit the pseudo outcome from the covariates). Note that there are four mutually exclusive race/ethnicity categories, including: white (non-Hispanic), black (non-Hispanic), Hispanic, and other non-Hispanic. Standard errors are clustered at the project level and are presented in parentheses.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** p≤0.01, ** p≤0.05, * p≤0.10.

Table B.6: Sensitivity to Stratification and Covariates

	IHS of Earnings at Age 26			
	(1)	(2)	(3)	(4)
Panel A. Large PHAs				
HOPE VI	0.059	0.191**	0.157	0.195***
	(0.102)	(0.076)	(0.107)	(0.073)
stratification			X	X
covariates		X		X
observations	149,000	149,000	149,000	149,000
Panel B. Small PHAs				
HOPE VI	-0.090	0.015	0.005	0.045
	(0.099)	(0.087)	(0.101)	(0.087)
stratification			X	X
covariates		X		X
observations	109,000	109,000	109,000	109,000

Notes: Panels A and B present estimates from the stratification with regression estimator for large and small Public Housing Authorities (PHAs), respectively. The outcome variable is the inverse hyperbolic sine (IHS) of annual earnings at age 26. The rows below the point estimates indicate whether the stratification structure was used (if not, Ordinary Least Squares is used) and whether the standard vector of individula- and project-level controls are included in the regression. All specifications include a fixed effect for the base year in which the child appears in public housing as well as a fixed effect for the year in which the child turns 26. Standard errors are clustered at the project-level and are presented in parentheses.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$.

Table B.7: Ordinary Least Squares Estimates

	IHS of Earnings at Age 26			
	(1)	(2)	(3)	(4)
Panel A. Large PHAs				
HOPE VI	-0.101	0.154**	0.088	0.288**
	(0.091)	(0.078)	(0.134)	(0.111)
estimated with controls	no	yes	no	yes
non-HOPE VI sample	same PHA	same PHA	applicants	applicants
observations	338,000	338,000	19,000	19,000
Panel B. Small PHAs				
HOPE VI	-0.206**	-0.021	-0.037	0.087
	(0.100)	(0.086)	(0.121)	(0.104)
estimated with controls	no	yes	no	yes
non-HOPE VI sample	same PHA	same PHA	applicants	applicants
observations	92,000	92,000	13,000	13,000

Notes: Panels A and B present estimates from large and small Public Housing Authorities (PHAs), respectively. Each estimate is from a separate regression estimated by Ordinary Least Squares in which the dependent variable is the inverse hyperbolic sine (IHS) of annual earnings measured in the year the child turns 26 and the main independent variables is an indicator equal to one if the project received a HOPE VI grant. All regressions contain controls for the year in which the individual appears in public housing as well as the year in which the individual turns 26. The row below the point estimates indicates whether the standard set of additional project- and individual-level covariates are included in each specification. The sample of HOPE VI projects is identical across all specifications and the row above the observation counts indicates whether the set of non-HOPE VI projects includes projects in PHAs that were awarded HOPE VI funding (same PHA) or projects that applied but did not receive HOPE VI funding (applicants to the Revitalization or Demolition program). Standard errors are clustered at the project-level and are presented in parentheses.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** p≤0.01, ** p≤0.05, * p≤0.10.

Table B.8: Housing Outcomes

	Housing Subsidy Type			Co-Residence with Parents		
	public (1)	voucher (2)	non-subsidized (3)	census tract (4)	household (5)	incarcerated (6)
Panel A. Large PHAs						
HOPE VI	-0.039**	0.018*	0.021	-0.008	-0.012	0.001
	(0.02)	(0.01)	(0.019)	(0.007)	(0.007)	(0.003)
control mean	0.683	0.089	0.228	0.523	0.519	0.047
observations	149,000	149,000	149,000	149,000	149,000	149,000
Panel B. Small PHAs						
HOPE VI	-0.013	0.025***	-0.012	-0.013	-0.006	0.005
	(0.018)	(0.008)	(0.02)	(0.009)	(0.008)	(0.003)
control mean	0.612	0.072	0.316	0.542	0.507	0.042
observations	109,000	109,000	109,000	109,000	109,000	109,000

Notes: Panels A and B present estimates from the stratification with regression estimator for large and small Public Housing Authorities (PHAs), respectively. All outcomes are measured in 2010. In columns 1-3 the outcome is an indicator equal to one if individual lived in public, voucher, or other housing, respectively (cateogries are mutually exclusive). In columns 4 and 5 the outcome is an indicator equal to one if the child lives in the same census tract or household as their parent, respectively. The outcome variable in column 6 is an indicator equal to one if the child was incaracerated in the 2010 census. All stratum-level regressions control for the base year in which the child appears in public housing, the year in which the child turns 26, and the standard vector of individual- and project-level characteristics Standard errors are clustered at the project-level and are presented in parentheses. The mean and standard deviation, presented in brackets, of the outcome for the control group are a weighted aggregate of stratum-level statistics, where the weights are proportional to the number of treated individuals in a strata.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$.

Table B.9: Dispersion of Project Co-Residents

	share of former residents living within			avg. log
	one mile	three miles	five miles	distance
	(1)	(2)	(3)	(4)
Panel A. Large PHAs				
HOPE VI	-0.01	-0.006	-0.006	0.043
	(0.006)	(0.012)	(0.016)	(0.086)
control mean	0.065	0.193	0.315	3.29
observations	149,000	149,000	149,000	149,000
Panel B. Small PHAs				
HOPE VI	-0.021***	-0.02	-0.016	0.047
	(0.007)	(0.018)	(0.021)	(0.088)
control mean	0.086	0.22	0.316	3.41
observations	109,000	109,000	109,000	109,000

Notes: Panels A and B present estimates from the stratification with regression estimator for large and small Public Housing Authorities (PHAs), respectively. In columns 1-3 the outcome is the share of former public housing residents that live within a one-, three- and five-mile radius, respectively. The outcome in column 4 is the average log distance (measured in miles) between the individual and each of the former public housing residents. The longitude and lattitude of residence correspond to the the centroid of the census tract in which the child resides in 2010. All stratum-level regressions control for the base year in which the child appears in public housing, the year in which the child turns 26, and the standard vector of individual- and project-level characteristics. Standard errors are clustered at the project-level and are presented in parentheses. The mean and standard deviation, presented in brackets, of the outcome for the control group are a weighted aggregate of stratum-level statistics, where the weights are proportional to the number of treated individuals in a strata.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** p$\leq$0.01, ** p$\leq$0.05, * p$\leq$0.10.

Table B.10: Relationship between Poverty, Job Proximity, and Rent

	1990 Characteristics		2010 Characteristics	
	avg. commute (1)	poverty (2)	job proximity index (3)	poverty (4)
Panel A. Large PHAs				
median rent	-0.0189	-0.5301***	0.0062	-0.5616***
	(0.0420)	(0.0534)	(0.0300)	(0.0310)
R-squared	0.0003	0.1843	0.0000	0.3140
observations	13,934	12,865	13,834	13,834
Panel B. Small PHAs				
median rent	0.1978***	-0.4124***	-0.1216***	-0.5227***
	(0.0716)	(0.0337)	(0.0337)	(0.0229)
R-squared	0.0355	0.1085	0.0147	0.2722
observations	4,639	4,324	4,595	4,595

Notes: Panels A and B present results for large and small PHAs, respectively. The sample includes all census tracts located in counties that contain a HOPE VI project. Columns 1-2 present results in which the characteristics of the census tracts are measured in 1990. Column 3-4 present results in which the characteristics of the census tracts are measured in 2010. Each variable corresponds to a percentile rank that is calculated within counties and weighted by the census tract population. Within each panel each column presents results from a separate regression in which the outcome variable is listed in the column header and the independent variable is listed in the row. Standard errors are clustered at the level of the county.
Source: Authors' calculations based on publicly available data from the 1990 and 2010 Decennial Files and HUD.
*** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.10$.

Table B.11: Neighborhood Job Density

	log job density (1)	log population density (2)	log job/population (3)
Panel A. Large PHAs			
HOPE VI	0.045	-0.617***	0.220**
	(0.028)	(0.167)	(0.087)
control mean	0.07	-1.4	0.256
observations	149,000	149,000	149,000
Panel B. Small PHAs			
HOPE VI	-0.005	0.081	-0.001
	(0.014)	(0.171)	(0.056)
control mean	0.053	-2.77	0.354
observations	109,000	109,000	109,000

Notes: Panels A and B present estimates from the stratification with regression estimator for large and small Public Housing Authorities (PHAs), respectively. In columns 1-3 the outcome variable is a characteristic (measured in 2010) of the census tract in which the project was located, including: the log of the job density, the log of the population density and the log of the ratio of jobs to population, respecitvely. All stratum-level regressions control for the base year in which the child appears in public housing, the year in which the child turns 26, and the standard vector of individual- and project-level characteristics Standard errors are clustered at the project-level and are presented in parentheses. The mean and standard deviation, presented in brackets, of the outcome for the control group are a weighted aggregate of stratum-level statistics, where the weights are proportional to the number of treated individuals in a strata.
Source: Authors' calculations from matched Longitudinal Employer-Household Dynamics, Department of Housing and Urban Development, and Decennial Census files.
*** p≤0.01, ** p≤0.05, * p≤0.10.

B.3 Additional Robustness Checks

Revitalization Program: While the focus of our paper is on the HOPE VI Demolitions program, there is some overlap with the HOPE VI Revitalizations program by which some projects received both Demolition and Revitalization grants. Typically, the award of the Revitalization grants and their implementation took place well after the Demolition grant, but it is possible that our estimates are affected by the Revitalizations program. To investigate this, we start by estimating two specifications. First, we estimate a specification in which we interact the indicator for the HOPE VI Demolitions award with an indicator for the Revitalizations program. Second, we estimate a specification in which we drop all projects that received a HOPE VI Revitalization grant.[1] We conduct both sets of analyses separately within large and small PHAs and the inverse hyperbolic sine of age 26 earnings is the outcome variable. For the first specification, the coefficient on the interaction term is statistically insignificant at the 10 percent level for both the large and small PHA samples. In other words, there is not a statistically significant difference between the treatment effects for projects awarded both a Revitalization and Demolition and projects that were only awarded a Demolitions grant. The results from both the first and second specifications indicate that the effect of the HOPE VI Demolition is positive and statistically significant at the 10 percent level in large PHAs and is positive but statistically insignificant in small PHAs. Taken together, these analyses suggest that our main estimates are not affected by the HOPE VI Revitalization program. Another important result in our paper is that the effect of HOPE VI is largest for projects that were located in neighborhoods characterized by high poverty rates and limited job accessibility (proxied by population density and commute time). To see if these results are sensitive to

[1]Note that in constructing the sample we already drop non-HOPE VI Demolition projects that receive a Revitalization grant. Thus, this restriction only drops projects that received both a Demolition and Revitalization grant.

the presence of the Revitalization program we estimate a model in which we interact the HOPE VI indicator with both an indicator for the Revitalization program and baseline characteristics of the neighborhood. Our main findings with respect to heterogeneous effects by neighborhood characteristics are robust to controlling for the Revitalization program. Specifically, the interaction term between HOPE VI and the baseline characteristics is positive and statistically significant at the 10 percent level in large PHAs for all three baseline characteristics, whereas this interaction term is never statistically significant in small PHAs.

Early Demolitions: We define the year of the demolition as two years prior to the award year. However, the microdata begins in 1997 and we therefore define the year of the demolition as 1997 for HOPE VI projects that were awarded funding before 1999. We conduct two sets of analyses to investigate whether the timing of these early demolitions impact our results. First, we estimate a specification in which we interact the HOPE VI indicator with an indicator equal to one if the award was made prior to 1999. Second, we estimate a specification in which we drop all HOPE VI projects that received funding before 1999. We conduct both sets of analyses separately within large and small PHAs and the inverse hyperbolic sine of age 26 earnings is the outcome variable. For the first specification, the coefficient on the interaction term is not statistically significant at the 10 percent level for both the large and small PHA samples. In other words, the treatment effect for HOPE VI projects awarded funding prior to 1999 is not statistically different from the treatment effect for HOPE VI projects awarded funding in 1999 or later. The results from both the first and second specifications indicate that the effect of HOPE VI is positive and statistically significant at the 5 percent level in large PHAs and is positive but statistically insignificant in small PHAs. Taken together, these results suggest that our main findings are not sensitive to how we treat the early demolitions that occur prior 1999.

Partial Demolitions: Some projects were only partially demolished. We argue that all residents were potentially affected by the program, since the demolitions could change the neighborhoods in which the projects were located. Thus, our main results estimate the effect of HOPE VI on all residents of the public housing project, as opposed to just the ones that were forced to move. To investigate the importance of partial demolitions, we measure the growth rate in project size between the year of the demolition and five years after. We then estimate a specification in which we interact this growth rate with the indicator for HOPE VI. We conduct the analysis separately within large and small PHAs and the inverse hyperbolic sine of age 26 earnings is the outcome variable. We find no evidence of heterogeneous effects by changes in project size. Specifically, the interaction terms between HOPE VI and the project growth rate is positive but not statistically significant at the 10 percent level for both the large and small PHA samples. The results provide no evidence that residents who experienced partial demolitions faired differently relative to those exposed to more complete demolitions.

Tuning Parameters: While our methodology is in some ways data-driven, we do select tuning parameters that govern this process. While we use standard values for most tuning parameters, we choose custom values for (1) the number of matched used in the trimming procedure and (2) the thresholds for selection of covariates to be included in the propensity score. The choice of tuning parameters was based on their ability to eliminate baseline differences between HOPE VI and non-HOPE VI projects within stratum. However, we also assess the sensitivity of our results to alternative choices of these parameters. First, we try implementing the entire stratification procedure using 3 through 8 matches (our baseline specification uses 5 matches) when creating the trimmed sample. Second, we implement the stratification procedure with alternative thresholds for the likelihood ratio test by using values for the threshold firs the first-order terms of 1.5 to 5 by intervals of 0.5 and setting the threshold value for the second-order terms to 1.71 plus the threshold for the first-order term (our main specification uses 2.5 and 4.21 for the tuning parameters related to the first-order and second-order terms, respectively). Using the resulting stratification structures, we then estimate the effect of HOPE VI on the inverse hyperbolic sine of age 26 earnings as the main outcome and do this separately for large and small PHAs. For small PHAs the effect of HOPE VI is never statistically significant at the 10 percent level across all fourteen specifications. For large PHAs, the point estimate is always positive and estimates in nine, twelve and thirteen of the fourteen estimates are statistically significant at the 1 percent, 5 percent and 10 percent level, respectively. Thus, for large PHAs only one of the fourteen estimates is not statistically significant at the 10 percent level. Thus, our main findings are quite robust to alternative choices of the tuning parameters.

Duplicate Observations: Control observations (individual-year observations) may appear in the data multiple times, as they are included for each year they appear in public housing. We cluster standard errors at the project level, which accounts for these duplicate observations within projects. However, if individuals in the control projects move to new projects, they will appear multiple times in the data and the clustering will not adequately account for the correlation in their outcomes. To assess the degree to which this is a problem we drop all individuals who appear in more than one project in our sample. We then re-estimate our main results focusing on the inverse hyperbolic sine of age 26 earnings. In large PHAs the effect of HOPE VI is positive and statistically significant at the 5 percent level and in small PHAs the effect of HOPE VI is positive but is not statistically significant. Thus, our main findings are not sensitive to how we account for individuals who appear in multiple projects within the sample period.

Heterogeneous Effects: Table 3 presents estimates of the effect of HOPE VI interacted with different individual characteristics for the large PHA sample. We replicate this analysis for the small PHA sample. With one exception, none of the interaction terms are statistically significant at the 10 percent level. The one exception is that, in the specification that interacts HOPE VI with race, the estimated coefficient on HOPE VI*black is negative and statistically significant at the 10 percent level whereas the estimated coefficient on HOPE VI is positive and statistically significant at the 10 percent level. Thus, there is some evidence that HOP VI may have been beneficial for white non-Hispanic children in small PHAs.

Alternative Controls: We compare our preferred estimates to estimates from two other estimators. We use OLS to estimate specifications in which we regress the IHS of annual earnings at age 26 on an indicator for HOPE VI and a set of covariates. We estimate these regressions separately for the large and small PHA samples and experiment with using alternative samples for the control group, including: (1) non-HOPE VI projects that are located in a PHA that received some HOPE VI funding, and (2) the sample of projects that applied for but never received HOPE VI funding. The results, presented in columns 2 and 4 of Table A.7, indicate that these alternative approaches produce qualitatively similar conclusions: HOPE VI led to substantial long-run labor market benefits in large PHAs (points estimates of 0.154 and 0.288) but had no discernible impact in small PHAs (point estimates of -0.021 and 0.087). The robustness of our main findings bolsters confidence in our conclusions. Columns 1 and 3 present results from OLS regressions that do not include controls. Here we find no evidence that HOPE VI had positive effect on earnings in either large or small PHAs. The difference between the estimates from specifications that do and do not include covariates in the mode underscores the importance of adjusting for baseline differences between the HOPE VI and non-HOPE VI projects.

B.4 Description of Variables

Description	Unit of Measurement	Used in Nearest-Neighbor Matching	Used in Propensity Score Estimation
average total household income	public housing project	X	X
proportion of household heads who are black non-Hispanic	public housing project	X	X
proportion of household heads who are white non-Hispanic	public housing project	X	X
log of the number of occupied units	public housing project	X	X
proportion of household heads who are disabled	public housing project	X	X
proportion of household heads who are married	public housing project	X	X
proportion of household heads who are female	public housing project	X	X
average count of individuals per housing unit	public housing project		X
proportion of household heads who are Hispanic	public housing project		X
average age of household heads	public housing project		X
proportion of household heads who are over 55	public housing project		
proportion of household heads who are other race/ethnicity (not white, not black not Hispanic)	public housing project		
average number of dependents per household	public housing project		
proportion of children with disability	public housing project		
proportion of residents who are children (under age 18)	public housing project		

average gross household rent per month	public housing authority	X	X
proportion of households with majority of income from wages and or business income	public housing authority	X	X
log of the number of available units	public housing authority		X
proportion of units that are occupied	public housing authority		
proportion of household reporting	public housing authority		
average household size	public housing authority		
average federal spending per unit per month	public housing authority		
average total household income	public housing authority		
proportion of households with majority of income from welfare	public housing authority		
average of household's income as a percent of local median income	public housing authority		
proportion of household heads (or spouse) who are under 25 years old	public housing authority		
proportion of household heads who are older than 62 years old	public housing authority		
proportion minority	public housing authority		
proportion black, non-Hispanic	public housing authority		
proportion Hispanic	public housing authority		
proportion household heads (with children) married	public housing authority		
proportion of household heads (with children) single parents	public housing authority		
proportion over-housed with more bedrooms than people	public housing authority		
average assets	public housing authority		

average number of months since manager reported on household	public housing authority		
Median Family Income or (area median income), on which HUD bases income limits on	county 1990	X	X
unemployment rate in 1996	county	X	X
average pay in 1996	county		X
unemployment rate in 1990	county		
average pay in 1990	county		
unemployment rate in 1991	county		
average pay in 1991	county		
unemployment rate in 1992	county		
average pay in 1992	county		
unemployment rate in 1993	county		
average pay in 1993	county		
unemployment rate in 1994	county		
average pay in 1994	county		
unemployment rate in 1995	county		
average pay in 1995	county		
poverty rate	1990 census tract	X	X
proportion of households with wage or salary income	1990 census tract	X	X
proportion of households with public assistance income	1990 census tract	X	X
median rent	1990 census tract	X	X
median year housing structure built	1990 census tract	X	X
proportion of adults with high school education only	1990 census tract		X
proportion of housing units vacant	1990 census tract		X
proportion of population living in rural area	1990 census tract		X

proportion of households with social security income	1990 census tract		X
proportion of mothers who are single	1990 census tract		X
proportion black, non-Hispanic	1990 census tract		X
proportion Hispanic	1990 census tract		X
proportion born in the US	1990 census tract		
proportion of adults with graduate degree	1990 census tract		
median income	1990 census tract		
proportion of households with interest, dividend or net rental income	1990 census tract		
median gross rent as percent of household income	1990 census tract		
proportion white, non-Hispanic	1990 census tract		
proportion senior citizen	1990 census tract		
median home value	1990 census tract		
age	child		
black, non-Hispanic	child		
white, non-Hispanic	child		
Hispanic	child		
other race/ethnicity (not white, not black not Hispanic)	child		
total household income	child		
head of household is married	child		
female	child		
disabled	child		
number of dependents in the household	child		
total number of people living in the household	child		
age	head of household		
black, non-Hispanic	head of household		
white, non-Hispanic	head of household		
Hispanic	head of household		

other race/ethnicity (not white, not black not Hispanic)	head of household		
total household income	head of household		
married	head of household		
female	head of household		
disabled	head of household		
number of dependents in the household	head of household		
total number of people living in the household	head of household		

Notes: This table defines all of the 92 baseline variables mentioned throughout the paper. The third and fourth column indicate whether the variable was included in the matching specification used to trim the sample and considered in the propensity score estimation, respectively. Variables at the level of the public housing project are created using the microdata from the HUD-PIC files. For recipients of the HOPE VI demolition grants, the values correspond to the year of the demolition, for other projects the values correspond to the average values between 1997 and 2001. The variables at the level of the public housing authority are based on public use data made available through the HUDUSER web tool. Family median income is also based on public use data obtained through the HUDUSER web tool. The unemployment rate and average earnings are from public use data made available by the Bureau of Labor Statistics. The characteristics of the census tracts in 1990 are derived from the 1990 Decennial Census and are from public use data provided by IPUMS (see Ruggles et al. 2019). The variables at the child and head of household level are from the HUD-PIC files.

Chapter C: Appendix Material for Chapter 3

C.1 Assessing Measurement Issues

C.1.1 Worker Flows

There is a residual term, $\tilde{\epsilon}_t^h$, in equation 3.2. Figure C.1 plots this residual and illustrates that it is smaller in magnitude than the poaching and nonemployment flows and does not exhibit clear a cyclical pattern. One possible explanation for this residual is that the flows in this section only include cases in which both the origin and destination firms are located in our 28-state sample. In unreported results, we directly measure the out-of-sample worker flows and find that they are not the source of the residual term. To provide evidence that out-of-sample worker flows are not affecting our results, Figure C.1 also plots the poaching and nonemployment flows that include all hires into firms in our sample (including hires from firms located in states outside of our sample) and all separations from firms in our sample (including separations to firms located in states outside of our sample). The patterns are quite similar regardless of whether out-of-sample flows are included, suggesting the migration in and out of our 28-state sample is unlikely to affect our main results.

C.1.2 Productivity

To put our numbers in perspective we compare them to an aggregate measure of productivity growth calculated from the RE-LBD. We calculate aggregate productivity growth by limiting the sample of firms to those that appear in our sample and then calculating the log difference between total revenue per worker within each 4-digit NAICS industry code. Aggregate productivity growth for this purpose is measured as the employment-weighted average of these industry-level growth rates.[1] Figure C.2 compares our measure of productivity growth to other widely used measures

[1]The employment used for weighting is the average of employment in the current and previous year.

Figure C.1: Differential Flows between High and Low Productivity Firms with Residual

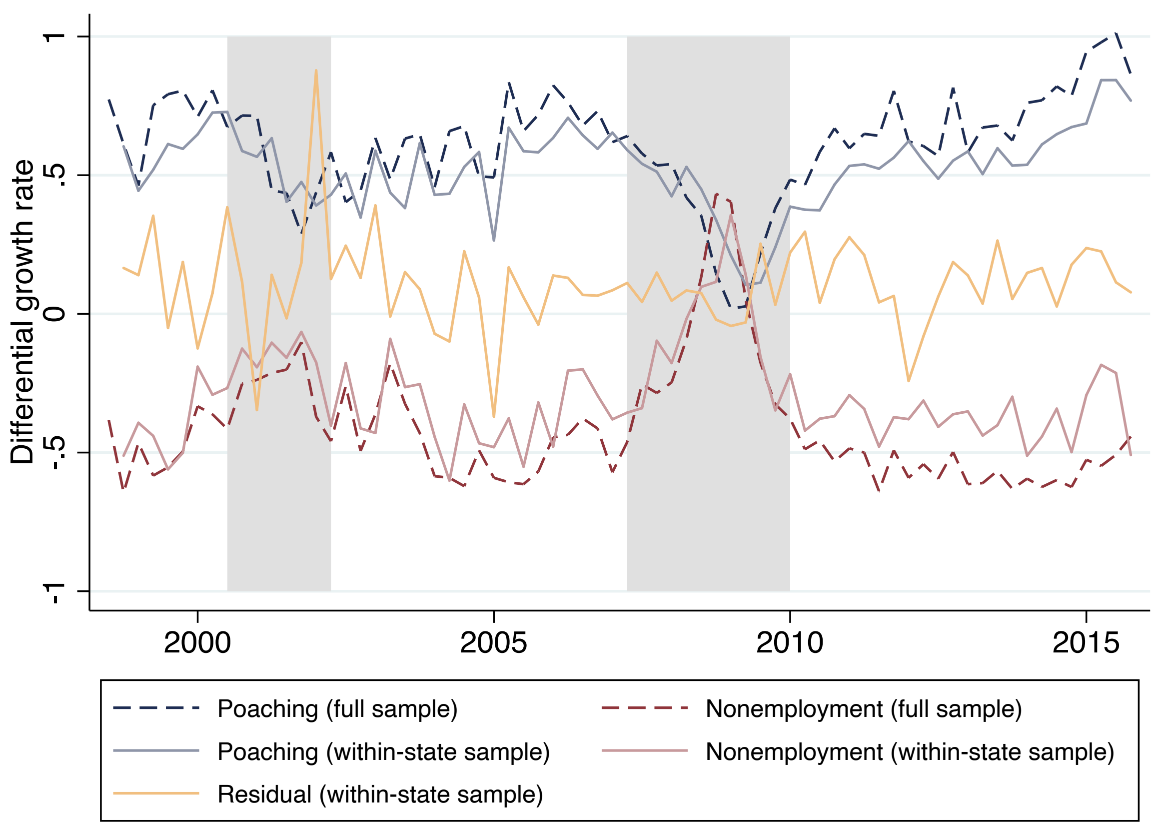

Notes: Differential growth rates are the difference in quarterly employment growth rates between high- and low-productivity firms. Results are presented separately for poaching, nonemployment, and residual flows. The poaching and nonemployment flows depicted by the dashed lines are based on the full sample. The poaching and nonemployment flows depicted by the solid lines are based on a sample that is limited to worker flows in which both the origin and destination employers are in one of the 28 states in our sample. The residual growth rate is the difference between the observed changes in the share of employment at high- and low-productivity firms and what is predicted by the poaching and nonemployment flows. Data are seasonally adjusted using X-12.

from the BLS and the BEA. The measure labeled BLS(industry) is the closest conceptually to our measure. It is an employment-weighted average of industry-level (4-digit NAICS) labor productivity growth rates. The industry-level growth rate is based on the growth rate in what BLS calls sectoral output per worker. Sectoral output is gross output less intrasectoral transactions. At the 4-digit level the adjustment for intrasectoral transctions is modest. The mean of the RE-LBD based measure is 1.3 log points while the mean of the BLS(industry) measure is 1.7 log points. The correlation of these two series is 0.85.

The BLS(sector) and BEA measures are growth rates of value added per worker for the private, non-farm business sector. These measures are distinct conceptually from the RE-LBD and BLS(industry) measures. The latter only capture within-industry contributions to aggregate productivity growth while the value-added measures not only use a conceptually different output measure but also reflect shifts in employment from low- to high-productivity industries. Still these alternative measures exhibit similar patterns to the RE-LBD and BLS(industry) measures. Interestingly, the BLS(sector) average growth rate is 2 log points which is larger than the BLS(industry) measure at 1.7 log points. While appropriate caution is needed to compare these measures given conceptual differences, this pattern is consistent with between industry effects contributing positively to aggregate productivity growth.

Figure C.3 compares aggregate productivity growth to the growth attributable to worker reallocation. Data in the RE-LBD are reported at an annual frequency so we sum the quarterly components of productivity growth from our decomposition within each calendar year. Across all years in the sample, average aggregate productivity growth is 1.3 log points per year and the components attributable to poaching and nonemployment flows contribute, on average, 0.4 and -0.3 log points per year, respectively. Thus, at the annual frequency our decomposition captures a quantitatively important aspect of productivity growth. The aggregate series exhibit a relatively

Figure C.2: Aggregate Productivity Growth from Alternative Sources

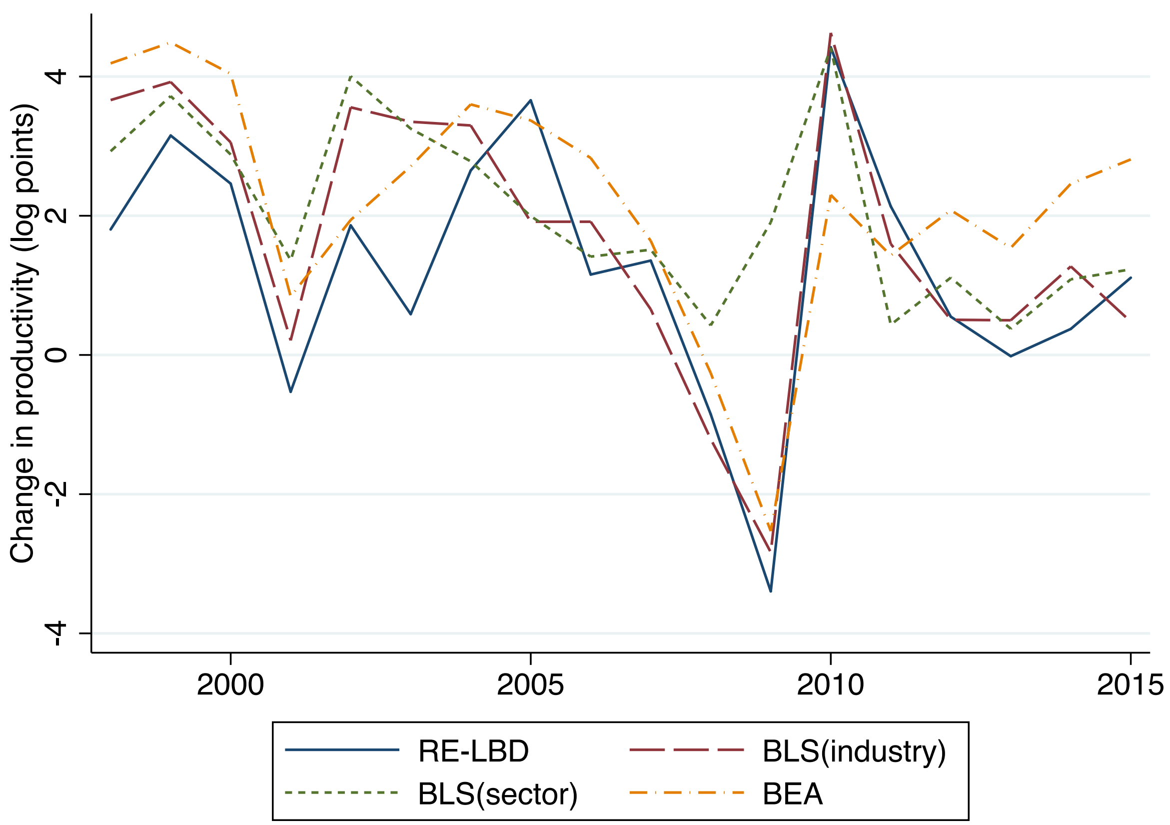

Notes: This figure presents annual aggregate productivity growth based on four different sources that include: (1) confidential data from the RE-LBD that characterizes the growth of all firms in our data, (2) publicly available data from the BLS based on industry-level productivity growth, (3) publicly available data from the BLS based on sector-level productivity growth, and (4) publicly available data on industry-level measures of value added from the BEA. To estimate aggregate productivity growth from the RE-LBD, we follow a methodology similar to the BLS and calculate the log difference in total revenue per total number of workers between the current and subsequent year for each industry, then take the employment weighted average across industries.

larger decline in productivity growth during recessions. For example, between 2007 and 2009 annual productivity growth declined by 4.8 log points. In contrast, the decline in productivity growth attributable to poaching flows can only account for 0.3 log points. However, part of the movements in aggregate productivity reflect measurement error due to cyclical changes in factor utilization.

One concern with our measure of productivity is that it does not account for intertemporal variation in factor utilization. During a recession a firm may decide to cut back on production (possibly by reducing workers' hours or worker intensity), which could lead to a decline in log revenue per worker without any real changes in productivity. This issue could affect both our decomposition results and the our aggregate measure of productivity growth.

Intertemporal variation in factor utilization does not appear to meaningfully affect the decomposition results. This is because, to the extent that log revenue per worker is subject to this concern, it affects high- and low-productivity firms to an equal extent. Figure C.4(a) plots, $\tilde{R}^h_{t-1} - \tilde{R}^l_{t-1}$, which is the productivity differential between high- and low-productivity firms. The productivity differential is orders of magnitudes larger than the short-term variation, which may be driven by intertemporal variation in factor utilization. To illustrate that the short-term variation in these differentials does not affect the decomposition exercise, we construct a smoothed series by fitting a linear time trend to the productivity differentials. We then use this smoothed series to implement the productivity growth decomposition. Figure C.4(b) presents the results and shows that the decomposition using the actual and smoothed productivity series yield essentially the same results. To quantify this we regress the component of productivity growth attributable to poaching flows using the observed productivity differentials on the series using the smoothed differentials. The R-squared is 0.99. The analogous R-squared for the nonemployment flows is 0.999.

Variation in factor utilization over the business cycle is a greater issue for our mea-

Figure C.3: Annual Productivity Growth from Worker Reallocation

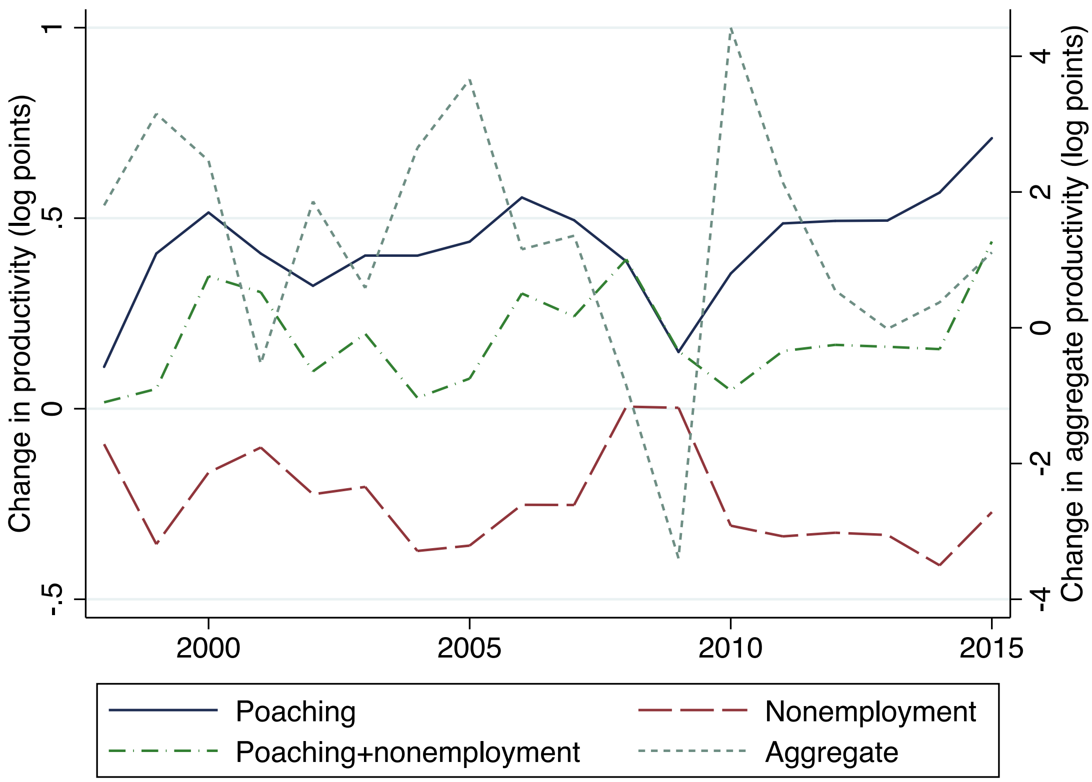

Notes: The figure presents the components of annual productivity growth that are attributable to worker reallocation between high- and low-productivity firms through poaching and nonemployment flows as well as the sum of these two components. Annual productivity growth is the sum of the quarterly growth within a calendar year for these components. The figure also presents a measure of aggregate productivity growth that is calculated from the RE-LBD micro data.

Figure C.4: Decomposition and Intertemporal Variation in Factor Utilization

(A) Productivity Differential

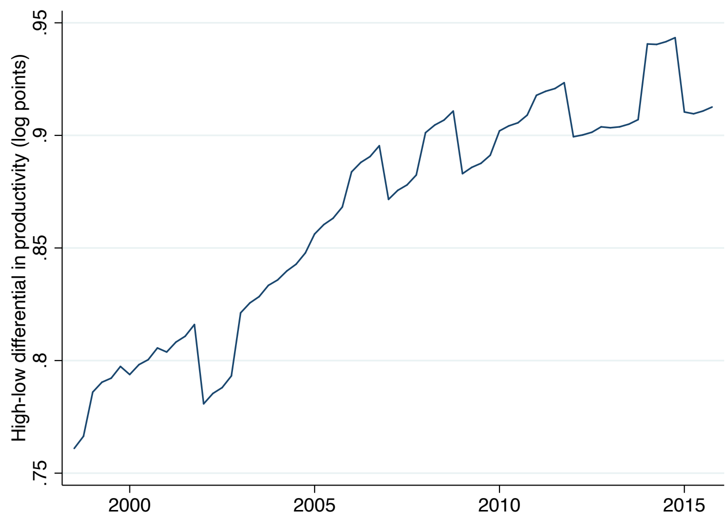

(B) Productivity Decomposition

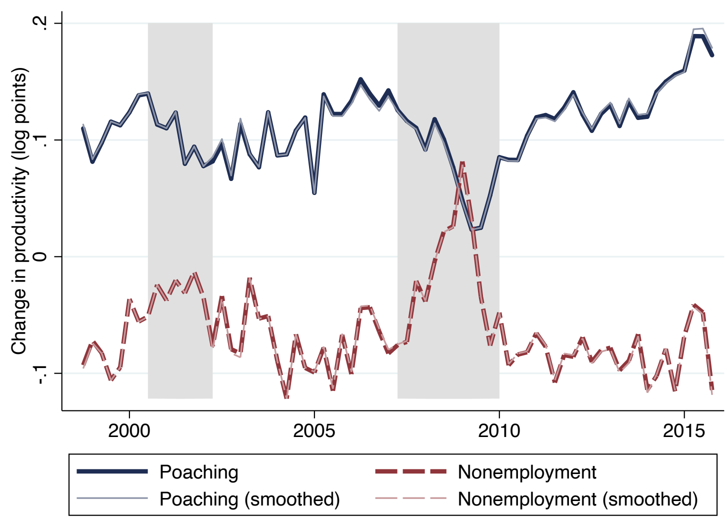

Notes: Panel (a) plots the difference between the average productivity at high- and low-productivity firms. Panel (b) presents the components of productivity growth that attributable to worker reallocation between high- and low-productivity firms through poaching and nonemployment flows. The results in Panel (b) implement the decomposition using both the observed productivity differentials (depicted in Panel (a)) as well as the productivity differentials from a smoothed series generated by fitting the observed productivity differentials with a linear time trend. Data are seasonally adjusted using X-12.

sure of aggregate productivity growth. Fernald (2014) produces a series of growth in business sector TFP that adjusts for factor utilization. Figure C.5(a) presents both the adjusted and unadjusted growth rates from these data. The unadjusted series exhibit a larger decline in TFP during recessions relative to the adjusted series. To get a sense of how sensitive our measure of productivity growth is to changes in factor utilization, we use the difference between the unadjusted and adjusted series from Fernald (2004) to adjust for variation in factor utilization.[2] Figure C.5(b) compares the productivity growth attributable to worker reallocation to the adjusted and unadjusted measures of aggregate productivity growth.

[2]Specifically, we calculate the difference between the adjusted and unadjusted growth rates in 2001 and 2008. We adjust our growth rate by adding these differences to the growth rates in 2001 and 2008. Because our data also exhibit a large decline in 2009, we also add the 2008 difference from the Fernald (2014) series to the 2009 growth rate.

Figure C.5: Aggregate Productivity Growth with Factor Utilization Adjustment

(A) Aggregate Productivity Growth from Fernald (2014)

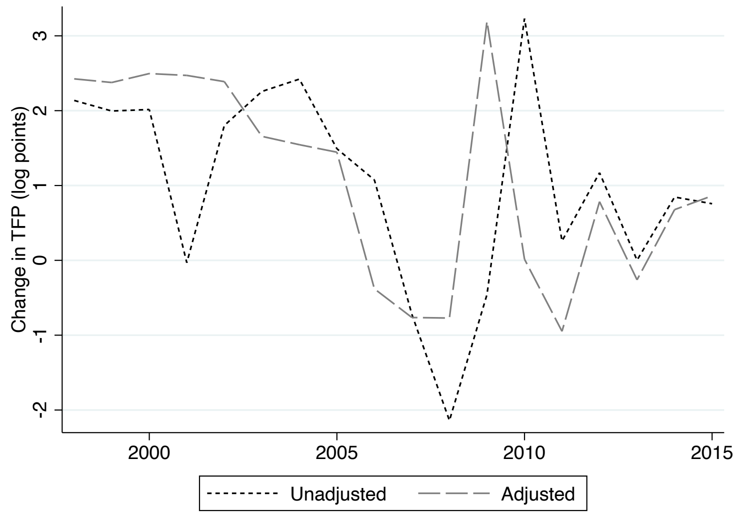

(B) Productivity Decomposition with Utilization Adjustment

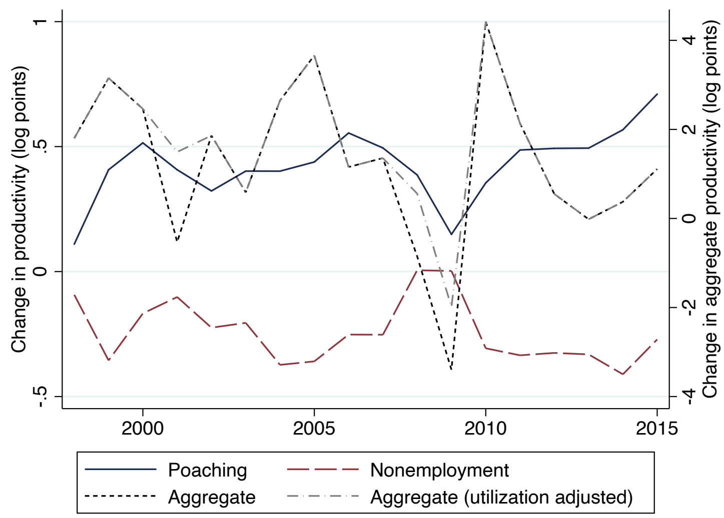

Notes: Panel (a) presents data from Fernald (2014) on the growth in business sector total factor productivity (TFP) as well as a measure that implements an adjustment for variation in factor utilization. Panel (b) presents the annual productivity growth from worker reallocation through poaching and nonemployment flows as well as aggregate productivity growth. In addition, Panel (b) includes a series that adjusts aggregate productivity growth using the difference between the unadjusted and adjusted measures of growth from Fernald (2014).

C.2 Decomposition Methodology

C.2.1 Employment

Equation 3.2 is an exact accounting identity. To see why this equation holds, begin by assuming that there is a closed system in which all changes in employment are accounted for by hires and separations. This assumption implies that,

$$\frac{\Delta E_t^i}{E_{t-1}^i} = \lambda_t^i + \delta_t^i \text{ for } i \in \{h, l\} \tag{C.1}$$

Then we can write,

$$\begin{aligned}
\Delta\theta_t^h &= \frac{E_t^h}{E_t} - \frac{E_{t-1}^h}{E_{t-1}} \\
&= \frac{E_t^h E_{t-1} - E_{t-1}^h E_t}{E_t E_{t-1}} \\
&= \frac{E_{t-1}^l \Delta E_t^h - E_{t-1}^h \Delta E_t^l}{E_t E_{t-1}} \\
&= (\frac{\Delta E_t^h}{E_{t-1}^h} - \frac{\Delta E_t^l}{E_{t-1}^l})\theta_{t-1}^h \theta_{t-1}^l (\frac{E_{t-1}}{E_t}) \\
&= \left(\lambda_t^h - \lambda_t^l + \delta_t^h - \delta_t^l\right)\theta_{t-1}^h \theta_{t-1}^l (\frac{E_{t-1}}{E_t}) \\
&= \tilde{\lambda}_t^h + \tilde{\delta}_t^h
\end{aligned} \tag{C.2}$$

where

$$\tilde{x}^h = (x^h - x^l)\theta_{t-1}^h \theta_{t-1}^l (\frac{E_{t-1}}{E_t}) \text{ for } x \in \lambda, \delta \tag{C.3}$$

In practice, the system is not closed and hires and separations do not perfectly predict changes in employment. Thus, define the residual term as $\tilde{\epsilon}_t^h = \Delta\theta_t^h - (\tilde{\lambda}_t^h + \tilde{\delta}_t^h)$.

C.2.2 Productivity

Assumptions A1 and A2 allow us to isolate the component of productivity growth that is attributable to changes in the share of workers at high-productivity firms. First, note that using an accounting identity, we can rewrite the unobserved measure of productivity as,

$$\begin{aligned} P_t &= \sum_k \theta_t(k) P_t(k) \\ &= \sum_k \theta_t(k)[\theta_t^l(k) P_t^l(k) + \theta_t^h(k) P_t^h(k)] \end{aligned} \tag{C.4}$$

Then by assumption A1,

$$\begin{aligned} P_t &= \sum_k \theta_t(k)[\theta_t^l(k) R_t^l(k) + \theta_t^h(k) R_t^h(k) + U_t(k)] \\ &= \sum_k \theta_t(k)[\theta_t^l(k) \tilde{R}_t^l(k) + \theta_t^h(k) \tilde{R}_t^h(k) + \bar{P}_t(k)] \\ &= \sum_k [\theta_t(k)\bar{P}_t(k)] + \sum_{i\in\{l,h\}} \left[\sum_k [\theta_t(k)\theta_t^i(k)] \sum_k [\theta_t(k)\tilde{R}_t^i(k)] + \sum_k [cov(\theta_t^i(k), \tilde{R}_t^i(k))] \right] \end{aligned} \tag{C.5}$$

where $cov(\theta_t^i(k), \tilde{R}_t^i(k)) \equiv \sum_k [\theta_t(k)\theta_t^i(k)\tilde{R}_t^i(k)] - \sum_k [\theta_t(k)\theta_t^i(k)] \sum_k [\theta_t(k)\tilde{R}_t^i(k)]$ and $\bar{P}_t(k) \equiv [P_t^l(k) + P_t^h(k)]/2$. Equation C.5 in combination with assumption A2 allows to rewrite productivity growth as,

$$\begin{aligned} \Delta P_t &= \Delta\left(\sum_k [\theta_t(k)\bar{P}_t(k)] + \sum_{i\in\{l,h\}} \left[\sum_k [\theta_t(k)\theta_t^i(k)] \sum_k [\theta_t(k)\tilde{R}_t^i(k)] \right] \right) \\ &= \Delta\left(\sum_k [\theta_t(k)\bar{P}_t(k)] + \sum_{i\in\{l,h\}} \theta_t^i \tilde{R}_t^i \right) \\ &= \sum_{i\in\{l,h\}} \left[\tilde{R}_{t-1}^i \Delta\theta_t^i + \theta_t^i \Delta\tilde{R}_t^i \right] + \Delta\left(\sum_k [\theta_t(k)\bar{P}_t(k)] \right) \\ &= (\tilde{R}_{t-1}^h - \tilde{R}_{t-1}^l)\Delta\theta^h + \theta_t^l \Delta\tilde{R}_t^l + \theta_t^h \Delta\tilde{R}_t^h + \Delta\left(\sum_k [\theta_t(k)\bar{P}_t(k)] \right) \end{aligned} \tag{C.6}$$

The first term is the component of productivity growth that is attributable to worker reallocation between high- and low-productivity firms.

C.2.3 Empirical Assessment of Assumptions

Assumption A1 states that log revenue per worker is a valid measure of productivity that is comparable across industries up to an additive constant. Section 3.2.1 provides some evidence to support this assumption by showing that log revenue per worker deviated from the industry average is predictive of employment growth and survival. To further assess the plausibility of this assumption we plot the average industry-deviated log revenue per worker, $\tilde{R}^i(k)$, for each two-digit NAICS industry code. The results, presented in Figure C.6, illustrate that there are no outliers in terms of the dispersion of log revenue per worker within industries. The industry with the least dispersion in productivity is Health Care and Social Assistance (NAICS=62) and the difference between the average log revenue per work at high- and low-productivity firms is 76 log points. The industry with the most dispersion in log revenue per worker is Finance and Insurance (NAICS=52) and the difference between the average log revenue per work at high- and low-productivity firms is 145 log points. While there are clearly differences in the dispersion of log revenue per worker within different industries, the lack of outliers lend some support to the plausibility of assumption A1.

Assumption A2 states that the covariance between the share of employment at high-productivity firms and the dispersion of log revenue per worker does not change over time. This assumption could be violated if industries that experience an increase in productivity growth also experience an increase in the dispersion of productivity across firms. While this is possible, we argue that any violation of this assumption is not quantiatively important. The term $\Delta\big(cov\big(\theta_t^i(k), \tilde{R}_t^i(k)\big)\big)$ would show up on the right-hand-side of equation 3.3 if it were nonzero and we can show that the

Figure C.6: Log Revenue per Worker Deviated from Industry Average

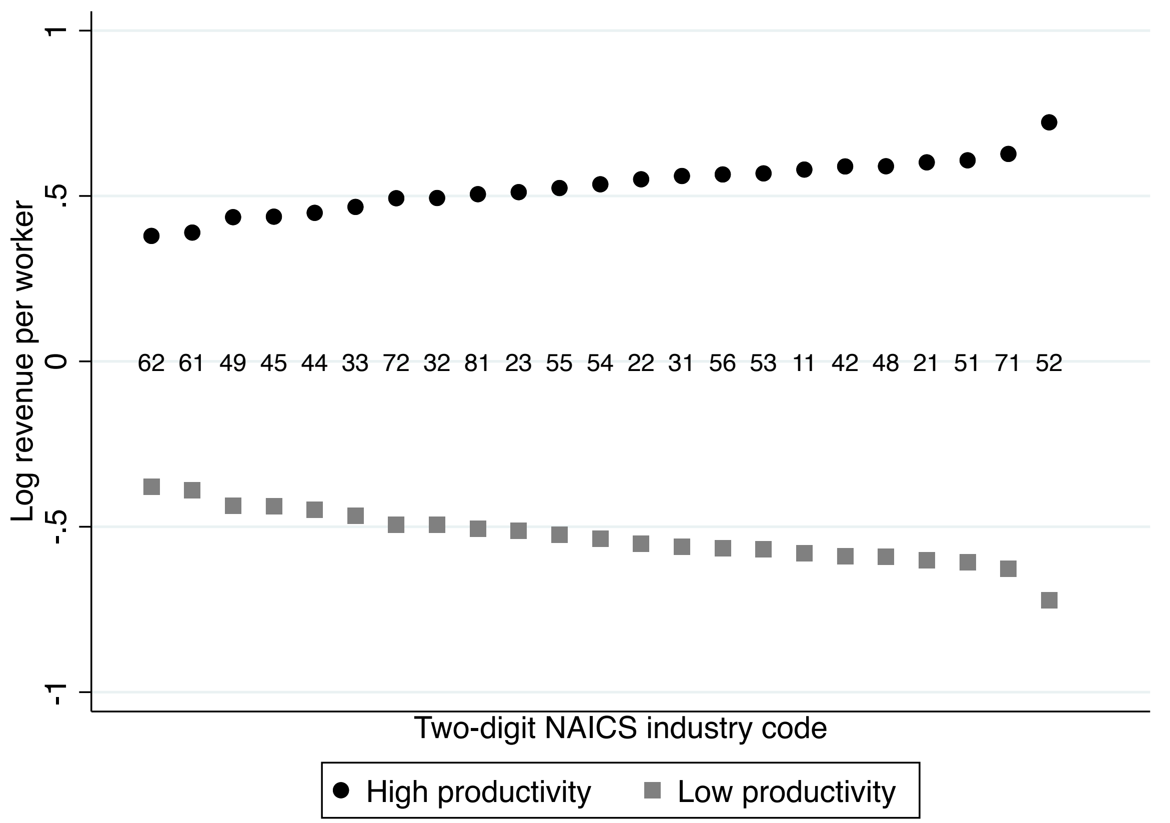

Notes: This figure presents the average value of log revenue per worker deviated from the average at the four-digit NAICS industry code for each two-digit NAICS industry code.

Figure C.7: Magnitude of Change in Covariance Term

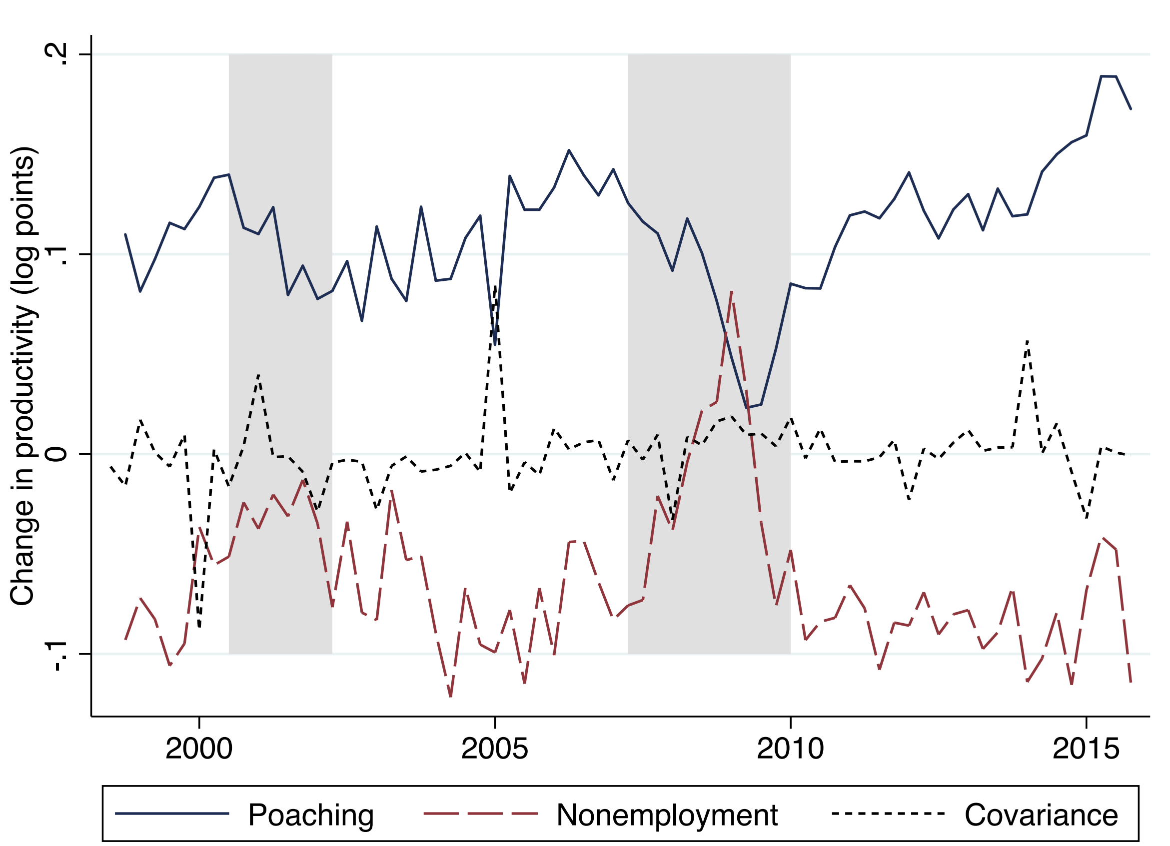

Notes: This figure presents the components of productivity growth that attributable to worker reallocation between high- and low-productivity firms through poaching and nonemployment flows as well as the covariance term, $\Delta\big(cov\big(\theta_t^i(k), \tilde{R}_t^i(k)\big)\big)$. Assumption A2 states that the covariance term is zero. This term is numerically equivalent for high- and low-productivity firms. Data are seasonally adjusted using X-12.

term is substantially smaller than the components of productivity growth that are attributable to worker reallocation. Specifically, the average value of the absolute value of the productivity growth attributable to worker reallocation through poaching and nonemployment flows is nine and six times larger than the average value of the absolute value of the covariance term, respectively. Figure C.7 makes this same point in more detail by plotting the covariance term as well as the components of worker reallocation through poaching and nonemployment flows over time. Taken together, assumption A2 appears to be a reasonable assumption.

References

Abadie, A. (2003) "Semiparametric instrumental variable estimation of treatment response models," *Journal of Econometrics.*

Abowd, John, Francis Kramarz, and David Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67(2): 251-333.

Abowd, John, John Haltiwanger, and Julia Lane. 2004. "Integrated Longitudinal Employee-Employer Data for the United States." *American Economic Review* 94(2): 224-229.

Abowd, John, Bryce Stephens, Lars Vilhuber, Fredrik Andersson, Kevin McKinney, Marc Roemer, and Simon Woodcock. 2009. "The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators." In *Producer Dynamics: New Evidence from Micro Data*, 68, Studies in Income and Wealth, ed. Timothy Dunne, J.Bradford Jensen and Mark J. Roberts, 149-230. Chicago: University of Chicago Press.

Acemoglu, D. and Autor, D. (2011) "Skills, tasks and technologies: Implications for employment and earnings," *Handbook of Labor Economics.* doi: 10.1016/S0169-7218(11)02410-5.

Altonji, J. G., Kahn, L. B. and Speer, J. D. (2016) "Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success," *Journal of Labor Economics*, 34(S1), pp. S361-S401. doi: 10.1086/682938.

Andersson, Fredrik, John C. Haltiwanger, Mark J. Kutzbach, Giordano Palloni, Henry O. Pollakowski, and Daniel H. Weinberg. 2018a. "Childhood Housing and Adult Earnings: A Between-Siblings Analysis of Housing Vouchers and Public Housing." *NBER Working Paper 22721.*

Andersson, Fredrik, John Haltiwanger, Mark Kutzbach, Henry Pollakowski, and Daniel Weinberg. 2018b. "Job Displacement and the Duration of Joblessness:

The Role of Spatial Mismatch." *The Review of Economics and Statistics* 100(2): 203-218.

Arellano-Bover, J. (2020) "Career Consequences of Firm Heterogeneity for Young Workers: First Job and Firm Size."

Bagger, J. and Lentz, R. (2019) "An Empirical Model of Wage Dispersion with Sorting," *Review of Economic Studies*. Oxford University Press, 86(1), pp. 153-190. doi: 10.1093/restud/rdy022.

Bayer, P., Ross, S. L. and Topa, G. (2008) "Place of work and place of residence: Informal hiring networks and labor market outcomes," *Journal of Political Economy*, 116(6), pp. 1150-1196. doi: 10.1086/595975.

Beaman, L. A. (2012) "Social networks and the dynamics of labour market outcomes: Evidence from refugees resettled in the U.S.," *Review of Economic Studies*, 79(1), pp. 128-161. doi: 10.1093/restud/rdr017.

Becker, G. S., S. Kominers, K Murphy and J. Spenkuch (2018) "A theory of intergenerational mobility", *Journal of Political Economy*, 126, pp. S7–S25. doi: 10.1086/698759.

Becker, G. S. and Tomes, N. (1986) "Human Capital and the Rise and Fall of Families", *Journal of Labor Economics*, 4(3), pp. 1–39. doi: 10.1086/298118.

Becker, G. S. and Tomes, N. (1979) "An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility", *Journal of Political Economy*, 87(6), pp. 1153–1189. doi: 10.1086/260831.

Barlevy, Gadi. 2002. "The Sullying Effect of Recessions." *Review of Economic Studies* 69(1): 65-96.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta. 2013. "Cross Country Differences in Productivity: The Role of Allocative Efficiency." *American Economic Review* 103(1): 305-334.

Bertheau, Antoine, Henning Bunzel, Rune Vejlin. 2020. "Employment Reallocation over the Business Cycle: Evidence from Danish Data." IZA Discussion Paper # 13681.

Bingley, P., Corak, M. and Westergard-Nielsen, N. (2011) "The Intergenerational Transmission of Employers in Canada and Denmark." doi: 10.1086/656371.

Black, S. E., Devereux, P. J. and Salvanes, K. G. (2005) "Why the apple doesn't fall far: Understanding intergenerational transmission of human capital," *American Economic Review*, pp. 437-449. doi: 10.1257/0002828053828635.

Black, S. E. and Devereux, P. J. (2011) "Recent developments in intergenerational mobility", *Handbook of Labor Economics*, pp. 1487-1541. doi: 10.1016/S0169-7218(11)02414-2.

Bohlmark, A. and M. J. Lindquist (2006) "Life-cycle variations in the association between current and lifetime income: Replication and extension for Sweden", *Journal of Labor Economics*, 24(4), pp. 879–896. doi: 10.1086/506489.

Burbidge, John B., Lonnie Magee, and A. Leslie Robb. 1988. "Alternative Transformations to Handle Extreme Values of the Dependent Variable" *Journal of the American Statistical Association* 83(401): 123-127.

Blackwood, Glenn, Lucia Foster, Cheryl Grim, John Haltiwanger, and Zoltan Wolf. 2020 (forthcoming). "Macro and Micro Dynamics of Productivity: From Devilish Details to Insights." *American Economic Journal: Macroeconomics*.

Burdett, Kenneth, and Dale Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39(2): 257-273.

Caballero, Ricardo J., and Mohamad L. Hammour. 1994. "The cleansing effect of recessions." *The American Economic Review*, pp.1350-1368.

Crane, Leland, Henry Hyatt, and Seth Murray. 2020. "Cyclical Labor Market Sorting." Unpublished draft, U.S. Census Bureau.

Caldwell, S. and Harmon, N. (2019) "Outside Options, Bargaining, and Wages: Evidence from Coworker Networks," *Working Paper*.

Carlson, Deven, Robert Haveman, Tom Kaplan, and Barbara Wolfe. "Long-term earnings and employment effects of housing voucher receipt." *Journal of Urban Economics* 71, no. 1 (2012): 128-150.

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women." *The Quarterly Journal of Economics*, 131(2), pp.633-686.

Card, D. et al. (2018) "Firms and labor market inequality: Evidence and some theory", *Journal of Labor Economics*, 36(1), pp. 13-70. doi: 10.1086/694153.

Card, D., J, Heining, and P. Kline (2013) "Workplace Heterogeneity and the Rise of West German Wage Inequality", *Quarterly Journal of Economics*. Oxford University Press, 128(3). doi: 10.1093/qje/qjt006.

Chetty, R., N. Hendren, Kline and E. Saez (2014) "Where is the land of opportunity? The geography of intergenerational mobility in the United States", *Quarterly Journal of Economics*, 129(4), pp. 1553–1623. doi: 10.1093/qje/qju022.

Chetty, R. et al. (2017) "The fading American dream: Trends in absolute income mobility since 1940," *Science*, 356(6336), pp. 398-406. doi: 10.1126/science.aal4617.

Chetty, Raj, Nathaniel Hendren, and Lawrence Katz. 2016. "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment." *American Economic Review* 106(4): 855-902.

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

Chetty, R. et al. (2020) 'Race and economic opportunity in the United States: An intergenerational perspective," *Quarterly Journal of Economics*, 135(2), pp. 711-783. doi: 10.1093/qje/qjz042.

Chyn, Eric, 2018. "Moved to opportunity: The long-run effects of public housing demolition on children." *American Economic Review* 108(10): 3028-3056.

Cingano, F. and Rosolia, A. (2012) "People I Know: Job Search and Social Networks," *Journal of Labor Economics*. University of Chicago Press, 30(2), pp. 291-332. doi: 10.1086/663357.

Collinson, Robert and Peter Ganong. 2018. "How Do Changes in Housing Voucher Design Affect Rent and Neighborhood Quality." *American Economic Journal: Economic Policy* 10(2): 62-89.

Corak, M. and Piraino, P. (2012) "Intergenerational Earnings Mobility and the Inheritance of Employers."

Corak, M. and Piraino, P. (2011) "The Intergenerational Transmission of Employers," *Journal of Labor Economics*, 29(1), pp. 37-68. doi: 10.1086/656371.

Davis, Steven and John Haltiwanger. 1990. "Gross job creation and destruction: Microeconomic evidence and macroeconomic implications." *NBER Macroeconomics Annual*, 5, pp.123-168.

Davis, S. J., J. Haltiwanger, R. Jarmin and J. Miranda (2007) "Volatility and dispersion in business growth rates: Publicly traded versus privately held firms." *NBER Macroeconomics Annual*, 21, pp. 107–156.

Davis, Steven, John Haltiwanger, and Scott Schuh. 1996. "Small business and job creation: Dissecting the myth and reassessing the facts." *Small Business Economics* 8(4): 297-315.

Davis, Steven, John Haltiwanger, and Scott Schuh (1996) *Job Creation and Destruction*, Cambridge: MIT Press.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2017. "Declining dynamism, allocative efficiency, and the productivity slowdown." *American Economic Review*, 107(5), pp.322-26.

Decker, Ryan, John Haltiwanger, Javier Miranda, and Ron Jarmin. 2020. "Changes in Business Dynamism and Productivity: Shocks vs. Responsiveness" *American Economic Review* 110(12): 3952-90.

Dunn, T. and Holtz-Eakin, D. (2000) "Financial Capital, Human Capital, and the Transition to Self-Employment: Evidence from Intergenerational Links," *Journal of Labor Economics*, 18(2), pp. 282-305. doi: 10.1086/209959.

Dustmann, C. et al. (2016) "Referral-based job search networks," *Review of Economic Studies*, 83(2), pp. 514-546. doi: 10.1093/restud/rdv045.

Eliason, M. et al. (2019) "Social Connections and the Sorting of Workers to Firms."

Eriksen, Michael and Amanda Ross. 2013. "The impact of housing vouchers on mobility and neighborhood attributes." *Real Estate Economics* 41(2): 255-277.

Fallick, Bruce, Erika McEntarfer, John Haltiwanger and Matthew Staiger (2019) "Job-to-Job Flows and the Consequences of Job Separations", *SSRN Electronic Journal*. doi: 10.2139/ssrn.3503543.

Farber, H. et al. (2018) "Unions and Inequality Over the Twentieth Century: New Evidence from Survey Data." *NBER Working Paper 24587*. Available at: http://www.nber.org/pape

Fernald, John. 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Federal Reserve Bank of San Francisco Working Paper 2012-19.

Foster, Lucia, John Haltiwanger, and C. J. Krizan. 2001. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." In *New Developments in Productivity Analysis*, 63, Studies in Income and Wealth, ed. Charles Hulten, Edwin Dean, and Michael Harper, 303-372. Chicago: University of Chicago Press.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review* 98(1): 394-425.

Fradkin, A., Panier, F. and Tojerow, I. (2019) "Blame the parents? How parental unemployment affects labor supply and job quality for young adults," *Journal of Labor Economics*, 37(1), pp. 35-100. doi: 10.1086/698896.

Gallagher, Megan, and Beata Bajaj. 2007. "Moving on: Benefits and challenges of HOPE VI for children." HOPE VI: Where do we go from here.

Gerard, F., L. Lagos, E. Severnini and D. Card (2018) "Assortative Matching or Exclusionary Hiring? The Impact of Firm Policies on Racial Wage Differences in Brazil", *National Bureau of Economic Research*. doi: 10.3386/w25176.

Gibbons, R. and Waldman, M. (2006) "Enriching a Theory of Wage and Promotion Dynamics inside Firmsm" *Journal of Labor Economics*, 24(1), pp. 59-107. doi: 10.1086/497819.

Glitz, A. (2017) "Coworker networks in the labour market," *Labour Economics*. Elsevier B.V., 44, pp. 218-230. doi: 10.1016/j.labeco.2016.12.006.

Goetz, E. G. 2010. "Desegregation in 3D: Displacement, dispersal and development in American public housing. Housing Studies" 25(2): 137-158.

Graham, Matthew, Mark Kutzbach, and Danielle Sandler. 2017. "Developing a Residence Candidate File for Use with Employer-Employee Matched Data." *U.S. Census Bureau, Center for Economic Studies Discussion Paper* 17-40.

Grawe, N. D. (2006) "Lifecycle bias in estimates of intergenerational earnings persistence", *Labour Economics*, 13(5), pp. 551–570. doi: 10.1016/j.labeco.2005.04.002.

Guimaraes, Paulo, and Pedro Portugal. 2010. "A simple feasible procedure to fit models with high-dimensional fixed effects." *The Stata Journal*, 10(4), pp.628-649.

Hahn, Joyce, Henry Hyatt, and Hubert Janicki. 2021. "Job Ladders and Growth in Earnings, Hours, and Wages." *European Economic Review* 133: 103654.

Haider, S. and G. Solon (2006) "Life-cycle variation in the association between current and lifetime earnings", *American Economic Review*, 96(4), pp. 1308–1320. doi: 10.1257/aer.96.4.1308.

Haltiwanger, John, Henry Hyatt, Lisa Kahn, and Erika McEntarfer. 2018. "Cyclical Job Ladders by Firm Size and Firm Wage." *American Economic Journal: Macroeconomics* 10(2): 52-85.

Haltiwanger, John, Henry Hyatt, Erika McEntarfer. 2018. "Who Moves Up the Job Ladder?" *Journal of Labor Economics* 36(S1): S301-S336.

Haltiwanger, J., H. Hyatt, E McEntarfer, L Sousa and S. Tibbets. (2014) "Firm Age and Size in the Longitudinal Employer-Household Dynamics Data." *SSRN Electronic Journal.* doi: 10.2139/ssrn.2423452.

Haltiwanger, J., R. Jarmin and R. Kulick (2016) "High Growth Young Firms: Contribution to Job, Output and Productivity Growth." doi: 10.2139/ssrn.2866566.

Haltiwanger, John, Ron Jarmin, Robert Kulick, and Javier Miranda. 2017. "High Growth Young Firms: Contribution to Job, Output, and Productivity Growth." In: *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*, 75, Studies in Income and Wealth, ed. John Haltiwanger, Erik Hurst, Javier Miranda, and Antoinette Schoar, Chicago: University of Chicago Press: 11-62.

Haltiwanger, John, Javier Miranda, and Ron Jarmin. 2013. "Who Creates Jobs? Small versus Large versus Young." *Review of Economics and Statistics* 95(2): 347-361.

Haltiwanger, J., Hyatt, H. and Mcentarfer, E. (2017) "Do Workers Move Up the Firm Productivity Job Ladder?"

Heath, R. (2018) "Why do firms hire using referrals? Evidence from Bangladeshi garment factories," *Journal of Political Economy*, 126(4), pp. 1691-1746. doi: 10.1086/697903.

Hellerstein, Judith K., Melissa McInerney, and David Neumark. 2011. "Neighbors and Co-Workers: The Importance of Residential Labor Market Networks." *Journal of Labor Economics* 29(4): 659-695.

Hellerstein, Judith, Mark Kutzbach, and David Neumark. 2019. "Labor market networks and recovery from mass layoffs: Evidence from the Great Recession period," *Journal of Urban Economics* 113.

Hilger, N. G. (2016) 'Parental job loss and children's long-term outcomes: Evidence from 7 million fathers' layoffs," *American Economic Journal: Applied Economics*, 8(3), pp. 247-283. doi: 10.1257/app.20150295.

Hsieh, Chang-Tai and Peter Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics* 124(4): 1403-1448.

Imbens G. 2015. "Matching Methods in Practice." *Journal of Human Resources* 50(2): 373-419.

Imbens, G. W. and Angrist, J. D. (1994) "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62(2), p. 467. doi: 10.2307/2951620.

Imbens, Guido, and Donald Rubin. 2015. "Causal Inference in Statistics, Social, and Biomedical Sciences." New York: Cambridge University Press.

Imbens, Guido W., and Jeffrey M. Wooldridge. 2009. "Recent Developments in the Econometrics of Program Evaluation." *Journal of Economic Literature* 47(1): 5-86.

Ioannides, Y. M. and Loury, L. D. (2004) "Job Information Networks, Neighborhood Effects, and Inequality," *Journal of Economic Literature*, 42(4), pp. 1056-1093. doi: 10.1257/0022051043004595.

Jacob, Brian A. 2004. "Public Housing, Housing Vouchers, and Student Achievement: Evidence from Public Housing Demolitions in Chicago." *American Economic Review* 94(1): 233-258.

Jacob, Brian A., Jens Ludwig, and Douglas L. Miller. 2013. "The effects of housing and neighborhood conditions on child mortality." *Journal of Health Economics* 32(1): 195-206.

Jones, Roderick W., and Derek J. Paulsen. 2011. "HOPE VI Resident Displacement: Using HOPE VI Program Goals to Evaluate Neighborhood Outcomes." *Cityscape: A Journal of Policy Development and Research* 13(3): 85-102.

Jarmin, R. S. and Miranda, J. (2012) "The Longitudinal Business Database." *SSRN Electronic Journal*. doi: 10.2139/ssrn.2128793.

Jenkins, S. "Snapshots versus movies:'lifecycle biases' and the estimation of intergenerational earnings inheritance."*European Economic Review* 31, no. 5 (1987): 1149-1158.

Jovanovic, B. and Nyarko, Y. (1997) “Stepping-stone Mobility,” *Carnegie-Rochester Conference Series on Public Policy*, 46, pp. 289-325. doi: 10.1016/S0167-2231(97)00012-2.

Kahn, Lisa (2010) “The long-term labor market consequences of graduating from college in a bad economy,” *Labour Economics*, 17(2), pp. 303-316. doi: 10.1016/j.labeco.2009.09.002.

Kain, John F. 1968. “Housing Segregation, Negro Employment, and Metropolitan Decentralization.” *Quarterly Journal of Economics* 82: 32-59.

Katz, L. F. and Summers, L. H. (1989) “Industry Rents: Evidence and Implications,” Brookings Papers on Economic Activity. Microeconomics, 1989, p. 209. doi: 10.2307/2534722.

Kingsley, Thomas G., Jennifer Johnson, and Kathryn Pettit. 2003. “Patterns of Section 8 Relocation in the HOPE VI Program.” *Journal of Urban Affairs* 25(4): 427–47.

Kramarz, F. and Skans, O. N. (2014) “When strong ties are strong: Networks and youth labour market entry,” *Review of Economic Studies*, 81(3), pp. 1164-1200. doi: 10.1093/restud/rdt049.

Lachowska, M. et al. (2019) “Do firm effects drift? Evidence from Washington Administrative Data,” *NBER Working Paper*.

Lens, Michael. 2014. “Employment Accessibility Among Housing Subsidy Recipients.” *Housing Policy Debate* 24(4): 671-691.

Lens, Michael, Kirk McClure, and Brent Mast. 2019. ”Does Jobs Proximity Matter in the Housing Choice Voucher Program?” *Cityscape* 21(1): 145-162.

Lise, Jeremy, and Jean-Marc Robin. 2017. “The Macrodynamics of Sorting between Workers and Firms.” *American Economic Review* 107(4): 1104-1135.

Loury, L. D. (2006) “Some contacts are more equal than others: Informal networks, job tenure, and wages,” *Journal of Labor Economics*, 24(2), pp. 299-318. doi: 10.1086/499974.

Ludwig, Jens, Greg Duncan, Lisa Gennetian, Lawrence Katz, Ronald Kessler, Jeffrey Kling, and Lisa Sanbonmatsu. 2013. “Long-term Neighborhood Effects on Low-Income Families: Evidence from Moving to Opportunity.” *American Economic Review* 103(3): 226-231.

Magruder, J. R. (2010) “Intergenerational networks, unemployment, and persistent inequality in South Africa,” *American Economic Journal: Applied Economics*, 2(1), pp. 62-85. doi: 10.1257/app.2.1.62.

Manning, A. (2011) “Imperfect competition in the labor market,” *Handbook of Labor Economics*. doi: 10.1016/S0169-7218(11)02409-9.

Mazumder, B. (2016) “Estimating the intergenerational elasticity and rank associationin the United States: Overcoming the current limitations of Taxdata”, *Research in Labor Economics*, 43, pp. 83–129. doi: 10.1108/S0147-912120160000043012.

Mazumder, B. (2005) “Fortunate sons: New estimates of intergenerational mobility in the united states using social security earnings data”, *Review of Economics and Statistics*, 87(2), pp. 235–255. doi: 10.1162/0034653053970249.

McCarty, Maggie. 2005. ”HOPE VI: background, funding, and issues. Report No. RL32236, Congressional Research Service.” The Library of Congress, Washington, DC.

McClure, Kirk. 2018. ”Length of Stay in Assisted Housing.” *Cityscape* 20(1): 11-38.

Melitz, Marc, and Sašo Polanec. 2015. “Dynamic Olley-Pakes productivity decomposition with entry and exit.” *The RAND Journal of Economics* 46(2): 362-375.

Mills, Gregory, Daniel Gubits, Larry Orr, David Long, Judie Feins, Bulbul Kaul, Michelle Wood, Amy Jones & Associates, Cloudburst Consulting, and the QED Group. 2006. The Effects of Housing Vouchers on Welfare Families. Washington, DC: U.S. Department of Housing and Urban Development, Office of Policy Development and Research.

Montgomery, J. D. (1991) “Social Networks and Labor-Market Outcomes: Toward an Economic Analysis," *The American Economic Review*, 81(5), pp. 1408-1418. doi: 10.1016/j.jcorpfin.2007.04.005.

Mortensen, D. T. and Vishwanath, T. (1995) "Personal contacts and earnings: It is who you know!" *Labour Economics*, 2(1), pp. 103-104. doi: 10.1016/0927-5371(95)80020-x.

Mortensen, D. T. and Burdett, K. (1998) "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39(2), pp. 257-273. Available at: http://www.jstor.org/stable/2527292.

Moscarini, Giuseppe, and Fabien Postel-Vinay. 2013. "Stochastic Search Equilibrium." *Review of Economic Studies* 80(4): 1545-1581.

Moscarini, Giuseppe, and Fabien Postel-Vinay. 2016. "Did the job ladder fail after the Great Recession?" *Journal of Labor Economics*, 34(S1), pp.S55-S93.

Mortensen, Dale, and Christopher Pissarides. 1994. "Job Creation and Destruction in the Theory of Unemployment." *Review of Economic Studies* 61: 397-415.

Mulry, M. (2007) "Summary of accuracy and coverage evaluation for the U.S. Census 2000", *Journal of official statistics*, 23(3), pp. 345–370.

Murphy, LaShonia. 2012. "The Allocation of Funds with HOPE VI: Applicants and Recipients." National Commission on Severely Distressed Public Housing (NCS-DPH). 1992. The Final Report: A Report to the Secretary of Housing and Urban Development. Washington D.C. August.

Olley, Steven and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64(6): 1263-1297.

Oreopoulos, Philip. 2003. "The Long-Run Consequences of Living in a Poor Neighborhood." *Quarterly Journal of Economics* 118(4): 1533-1575.

Oreopoulos, P., von Wachter, T. and Heisz, A. (2012) "The short- and long-term career effects of graduating in a recession," *American Economic Journal: Applied Economics*, 4(1), pp. 1-29. doi: 10.1257/app.4.1.1.

Patterson, Rhiannon, Michelle Wood, Ken Lam, Satyendra Patrabansh, and Gregory Mills. 2004. "Evaluation of the Welfare to Work Voucher Program: Report to Congress." Abt Associates Inc. Cambridge, MA.

Popkin, Susan J., Diane K. Levy, Laura E. Harris, Jennifer Comey, Mary K. Cunningham, and Larry F. Buron. 2004. "The HOPE VI Program: What about the Residents?" *Housing Policy Debate* 15: 385-414.

Popkin, Susan J., Diane K. Levy, and Larry F. Buron. 2009. "Has HOPE VI Transformed Residents' Lives? New Evidence from the HOPE VI Panel Study". *Housing Studies* 24(4): 477-502.

Postel-Vinay, F. and Robin, J. M. (2002) "Equilibrium wage dispersion with worker and employer heterogeneity," *Econometrica*, 70(6), pp. 2295-2350. doi: 10.1111/1468-0262.00377.

Rosenbaum, James E. 1995. "Changing the Geography of Opportunity by Expanding Residential Choice: Lessons from the Gautreaux Program." *Housing Policy Debate* 6(1): 231-269.

Rosenbaum, Paul, and D. Rubin. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrika* 70: 41-55.

Rosenbaum, Paul, and D. Rubin. 1984. "Reducing the Bias in Observational Studies Using Subclassification on the Propensity Score." *Journal of the American Statistical Association* 79: 516-524.

Roy, A. D. (1951) "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3(2), pp. 135-146. doi: 10.1111/j.2044-8317.1966.tb00372.x.

Ruggles, S., S. Flood, R. Goeken, J. Grover, E. Meyer, J. Pacas and M. Sobek. IPUMS USA: Version 10.0 [dataset]. Minneapolis, MN: IPUMS, 2020. https://doi.org/10.18128/D010.V1

Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas and Matthew Sobek. IPUMS USA: Version 9.0 [dataset]. Minneapolis, MN: IPUMS, 2019. https://doi.org/10.18128/D010.V9.0.

San, S. (2020) "Who Works Where and Why? Parental Networks and the Labor Market."

Sandler, Danielle H. 2017. "Externalities of Public Housing: The Effect of Public Housing Demolitions on Local Crime." *Regional Science and Urban Economics* 62: 24-35.

Schmieder, J. F., Von Wachter, T. and Heining, J. (2020) "The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany." Available at: http://www.econ.ucla.edu/tvwachter/papers/Jobloss_wp_2018_3_20.pdf.

Schmutte, I. M. (2015) "Job Referral Networks and the Determination of Earnings in Local Labor Markets," *Journal of Labor Economics*, 33(1), pp. 1-32. doi: 10.2139/ssrn.1727374.

Shea, J. (2000) "Does parents' money matter?," *Journal of Public Economics*, 77, pp. 155-184. Available at: www.elsevier.nl.

Shen, Qing. 1998. "Location Characteristics of Inner-City Neighborhoods and Employment Accessibility of Low-Income Workers," *Environment and Planning B: Planning and Design* 25(3): 345–365.

Shumpeter, Peter. 1939. *Business Cycles* New York: McGraw-Hill.

Smith, M., D. Yagan, W. Zidar, E. Zwick (2019) "Capitalists in the twenty-first century", *Quarterly Journal of Economics.* Oxford University Press, 134(4), pp. 1675–1745. doi: 10.1093/qje/qjz020.

Smith, Tony E. and Yves Zenou. 2003. "Spatial mismatch, search effort, and urban spatial structure." *Journal of Urban Economics* 54(1): 129-156.

Solon, G. (1992) "Intergenerational income mobility in the United States", *American Economic Review*, 82(3), pp. 393–408. doi: 10.2307/2117312.

Solon, G. (1989) "Biases in the Estimation of Intergenerational Earnings Correlations," *The Review of Economics and Statistics*, 71(1), p. 172. doi: 10.2307/1928066.

Solon, G. (2002) "Cross-country differences in intergenerational earnings mobility," *Journal of Economic Perspectives*, 16(3), pp. 59-66. doi: 10.1257/089533002760278712.

Solon, G., Haider, S. J. and Wooldridge, J. M. (2015) "What are we weighting for?", *Journal of Human Resources*, 50(2), pp. 301-316. doi: 10.3368/jhr.50.2.301.

Stinson, M. and Wignall, C. (2018) "Fathers, Children, and the Intergenerational Transmission of Employers." Available at: www.census.gov/ces.

Susin, Scott. 2002. "Rent vouchers and the price of low-income housing." *Journal of Public Economics* 83(1): 109-152.

Syverson, Chad. 2004. "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy* 112(6): 1181-1222.

Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature* 49(2): 326-365.

Topa, G. (2011) "Labor markets and referrals," *Handbook of Social Economics*, 1(1 B), pp. 1193-1221. doi: 10.1016/B978-0-444-53707-2.00005-0.

Turner, Margery Austin, Mark Woolley, G. Thomas Kingsley, Susan J. Popkin, Diane Levy, and Elizabeth Cove. 2007. "Severely Distressed Public Housing: The Costs of Inaction." Washington, DC: The Urban Institute.

U.S. Bureau of Labor Statistics. 1997. BLS Handbook of Methods. Office of Publications and Special Studies, Washington, DC.

U.S. Bureau of Labor Statistics. 2017. BLS Handbook of Methods. Division of Information Services, Washington, DC. Last modified April 14, 2017.

U.S. Census Bureau (2004) "Coverage Evaluation of Census 2000: Design and Methodology." Available at: www.census.gov/dmd/www/refroom.html.

U.S. Department of Housing and Urban Development (HUD). 2007. HOPE VI Program Authority and Funding History. https://www.hud.gov/sites/documents/DOC_9838. (accessed 19 September 2019).

U.S. General Accounting Office (GAO). 2003. Public Housing: HUD's Oversight of HOPE VI Sites Needs to be More Consistent. US General Accounting Office.

Vilhuber, Lars. 2018. “LEHD Infrastructure S2014 files in the FSRDC.” *U.S. Census Bureau, Center for Economic Studies Discussion Papers*, CES 18-27r.

Von Wachter, T. and Bender, S. (2006) “In the right place at the wrong time: The role of firms and luck in young workers’ careers,” *American Economic Review*, 96(5), pp. 1679-1705. doi: 10.1257/aer.96.5.1679.

Wagner, D. and Layne, M. (2014) “The Person Identification Validation System (PVS): Applying the Center for Administrative Records Research and Applications (CARRA) Record Linkage Software.”

Wang, Fahui. 2007. “Job Access in Disadvantaged Neighborhoods in Cleveland, 1980-2000: Implications for Spatial Mismatch and Association With Crime Patterns.” *Cityscape* 9(3): 95-121.

Zimmerman, D. J. (1992) “Regression Toward Mediocrity in Economic Stature,” *American Economic Review*, 82(3), pp. 409-429.