Essays in Labor Economics

by

Jacob French

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved April 2022 by the
Graduate Supervisory Committee:

Basit Zafar, Co-Chair
Esteban Aucejo, Co-Chair
Daniel Silverman
Berthold Herrendorf

ARIZONA STATE UNIVERSITY

May 2022

ABSTRACT

Are heterogeneous labor market outcomes a product of markets efficiently allocating resources or the result of structural market failures which should be corrected through well-crafted policy? In order to address this fundamental question in modern economics, we must first understand the forces which shape individuals' earnings, employment, and occupational choices. This collection of essays provides new evidence to support several novel channels which influence labor markets. First, I evaluate the connection between technological change and labor market outcomes by bringing new data and methods to study the mechanization of American agriculture in the early 20th century. Using an instrumental variables estimation strategy, I find that exogenous increases in exposure to technological change generated occupational displacement for incumbent laborers, increased income inequality, and had important impacts on intergenerational mobility for the children of affected workers. Additionally, I investigate the connection between low-opportunity neighborhoods and public housing residents' labor market outcomes. Leveraging quasi-random variation in neighborhood quality due to a public housing demolition, I find that residents' wages increased after moving to higher-opportunity neighborhoods and that more intense supportive services improved post-move employment. Taken together, these essays provide new evidence that both large-scale factors like new technologies and local factors like neighborhood quality contribute to heterogeneity in labor market outcomes both historically and up to the present day.

ACKNOWLEDGMENTS

I am especially grateful to Basit Zafar, Esteban Aucejo, Dan Silverman, and Berthold Herrendorf for their guidance and support. I thank Gustavo Ventura, Domenico Ferraro, Alvin Murphy, Jakob Dowling, Paola Ugalde, Spencer Perry, Marco Escobar, Zachary Tobin, Marco Mangini, Jorge Mesias, and participants in the ASU General Economics Workshop for helpful conversations and suggestions. I also thank Urban Strategies, Inc. and Esther Shin for making available the data used in this study. I have no relevant or material financial interests that relate to the research described in this paper and this research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

TABLE OF CONTENTS

Page

LIST OF TABLES vi

LIST OF FIGURES viii

CHAPTER

1 INTRODUCTION 1

2 TECHNOLOGICAL CHANGE, INEQUALITY, AND INTERGENERATIONAL MOBILITY: THE CASE OF EARLY 20TH CENTURY AGRICULTURE 4

2.1 Historical Setting 12

2.2 Data and Measurement 16

2.2.1 Data Sources 16

2.2.2 Measurement 17

2.3 Empirical Strategy 22

2.4 Empirical Results 28

2.4.1 Sample Characteristics 29

2.4.2 Occupational Mobility 30

2.4.3 Average Product of Agricultural Labor 32

2.4.4 Technological Change, Employment, and Occupational Wages 35

2.4.5 Family Structure and Migration 39

2.4.6 Second Generation 40

2.5 Discussion of Empirical Results 44

2.6 Model 48

2.6.1 Firm Problem 49

2.6.2 Worker Problem 50

2.6.3 Young Generation 52

2.6.4 Equilibrium . . . 53

2.7 Model Calibration . . . 54

2.8 Model Results . . . 57

2.9 Conclusion . . . 59

2.10 Tables . . . 61

2.11 Figures . . . 71

3 MOVES OF OPPORTUNITY, SUPPORTING SERVICES, AND LABOR MARKET OUTCOMES . . . 77

3.1 Background . . . 84

3.1.1 Setting . . . 84

3.1.2 Data . . . 86

3.2 Methodology . . . 89

3.2.1 Identification of Neighborhood Effects . . . 89

3.2.2 Empirical Strategy . . . 92

3.3 Results . . . 96

3.3.1 Relocation Shifted Neighborhood Quality . . . 96

3.3.2 Neighborhood Transitions Impacted Labor Market Outcomes 97

3.3.3 Relocation Effects Persist Over Time . . . 101

3.3.4 Heterogeneous Effects . . . 103

3.3.5 Wage and Employment Transitions . . . 105

3.4 Mechanisms . . . 108

3.4.1 Proximity to Job Opportunities Did Not Change After Moving 108

3.4.2 Case Management Prevent Job Loss over the Move . . . 110

3.4.3 Robustness . . . 114

CHAPTER Page

3.5 Conclusion 115

3.6 Tables 117

3.7 Figures 125

REFERENCES 134

APPENDIX

A APPENDIX TO CHAPTER 2 138

B APPENDIX TO CHAPTER 3 145

LIST OF TABLES

Table Page

2.1 Estimates of Equipment Investment and Non-Agricultural Occupational Mobility . . . 62
2.2 Observable Characteristics of Sample in 1910 . . . 63
2.3 Correlation of Instrument and Farm Wages Over Time . . . 64
2.4 Estimates of Equipment Investment and Occupational Mobility . . . 65
2.5 Estimates of Equipment Investment and Average Product of Labor . . . 66
2.6 Estimates of Equipment Investment and Occupational Wage . . . 67
2.7 Effects on Children of Incumbent Agricultural Workers: 1940 . . . 68
2.8 Calibrated Model Parameters . . . 68
2.9 Model Joint Calibration Fit . . . 69
2.10 Calibrated Model Parameters . . . 70
3.1 Sample Balance . . . 118
3.2 Shifts in Neighborhood Characteristics Over Move . . . 119
3.3 Relocation Effect Estimates . . . 120
3.4 Relocation Wage Effects and Neighborhood Shifts . . . 121
3.5 Heterogeneous Relocation Effects . . . 122
3.6 Distance and Time to Work . . . 123
3.7 Baseline Characteristics By Risk Score . . . 123
3.8 Risk Score Interaction . . . 124
A.1 Other Outcomes . . . 139
A.2 Displaced Worker Mechanisms . . . 141
A.3 Effects on Second Generation Skills . . . 142
A.4 Human Capital Transition Cost Parameterization . . . 142
B.1 Neighborhood Shift Over Relocation: Census Tract Observables . . . 146

Table Page

B.2 Shifts in Neighborhood Characteristics Over Move by Demograph Group 147

B.3 Average Within Indiviudal Change in Transfer Receipt 147

B.4 Employment Transitions 148

B.5 Relocation Effects: 1-2 Adult Households 148

B.6 Effect of Relocation on Adjusted Labor Market Outcomes 149

LIST OF FIGURES

Figure Page

2.1 Farm Equipment Productivity ... 72

2.2 Instrument Construction ... 73

2.3 Farm Equipment Per Worker ... 74

2.4 Equipment Investment and New Equipment Characteristics ... 75

2.5 Second Generation Wage Distribution ... 76

3.1 Distribution of Moving Dates ... 126

3.2 Neighborhood Poverty and Relocation ... 127

3.3 Wage Distribution Pre- and Post-Move ... 128

3.4 Wage and Hours Effects over Time ... 129

3.5 Employment Effects Over Time ... 130

3.6 Discontinuity in Case Manager Attention ... 131

3.7 Risk Score Distribution Pre- and Post-Relocation ... 132

3.8 Regression Discontinuity Bandwidths: Employment ... 133

A.1 First Generation Wage Distribution ... 140

A.2 Selection of the Investment Cutoff to Define Agricultural Regions in Model ... 143

A.3 Cumulative Distribution Functions of Changes in Equipment Value by Ruggedness ... 144

B.1 Neighborhood Characteristics over Time ... 150

B.2 Visual Regression Discontinuity Strategy ... 151

Chapter 1

INTRODUCTION

Income and wealth inequality are key aspects of the modern economy and are often cited as problems that policymakers should address. However, whether and how to address these issues largely depends on the underlying determinants of inequality. Understanding heterogeneity in labor market outcomes is a key step towards characterizing these determinants. This dissertation focuses on characterizing two factors that impact both the ways in which individuals interact with the labor market and the returns they receive for this interaction.

Chapter 2 focuses on the connection between labor-augmenting technologies and incumbent workers' future labor market outcomes by bringing new data and methods to evaluate the mechanization of early 20th-century agriculture in the United States, an episode of rapid technological change impacting a large proportion of the economy. Using an instrumental variables estimation strategy, I find that increased exposure to technological change caused incumbent workers to leave agriculture. These moves were disproportionately into lower-paying occupations as compared to the typical post-agricultural occupation. On the other hand, incumbent farmers faced no significant occupational displacement while also experiencing a significant increase in the average product of agricultural labor in their county. These effects did not attenuate over time and were transmitted into a second generation. The children of farmers from counties that experienced more technological change had higher non-agricultural wages in adulthood, while the children of wage-workers who left agriculture from the same regions had lower wages in adulthood as compared to their peers. These empirical results are used to discipline a dynamic occupational sorting model which indicates that 16% of workers had lower lifetime welfare due to technological change, and the total consumption equivalent cost to these individuals was 11% of the total surplus generated by the technological shock. This chapter highlights the way in which new technologies can both increase surplus and have long-lasting negative effects for some individuals.

Chapter 3 provides new evidence that neighborhood quality impacts labor market outcomes using novel panel data from a public housing demolition in Memphis, Tennessee. Residents at a large public housing site were required to relocate into the surrounding metropolitan area due to the demolition of their housing units, a move that significantly improved their neighborhood environment. This relocation is estimated to have increased hourly wages by $0.69, more than 7% of pre-move wages. Crucially, the impacts of relocation on employment were heterogeneous by age and education, with both more educated and younger adults avoiding the relocation associated job loss their peers experienced. This result suggests that some demographic groups unambiguously benefited from the relocation. Evidence suggests that these positive outcomes may have been modulated by the personalized case management services residents received over the course of relocation and for several subsequent years. Exploiting variation from a discontinuity in the intensity with which these services were offered, more attention from a case manager is estimated to have prevented post-move job loss. Finally, data on employers shows that post-move jobs were not physically inaccessible from the initial public housing site, suggesting proximity to job opportunities alone is not able to explain the wage increases.

The following essays provide new empirical evidence that both large macroeconomic factors, such as the emergence of new production technologies, and hyper-local factors, such as the quality of one's neighborhood environment, both play a measurable role in determining the returns individuals receive for providing labor services.

The details of these analyses provide valuable insights about specific mechanisms which generate heterogeneity in labor market outcomes. Taken as a whole, this work demonstrates that exogenous forces in strikingly different settings and timescales can have lasting impacts on individuals' outcomes, suggesting that inequality is, at least in part, a product of incomplete markets.

Chapter 2

TECHNOLOGICAL CHANGE, INEQUALITY, AND INTERGENERATIONAL MOBILITY: THE CASE OF EARLY 20TH CENTURY AGRICULTURE

This paper examines the mechanization of early 20th-century agriculture as a historical example of rapid labor-augmenting technological change. Commercially viable standard purpose tractors were introduced around 1910 and ushered in an era of significant improvements to the size and productivity of farm equipment. The setting of early 20th-century agriculture provides several unique advantages in evaluating the impact of new technologies. More than 30% of working-age men in 1910 worked in an agricultural occupation, and so the mechanization of American agriculture was a phenomenon impacting a wide swath of the economy. Additionally, population-level census data and recent advances in linking census records provide the opportunity to study channels such as intergenerational mobility that are often infeasible to evaluate in other settings.

I use two complementary methodologies to evaluate the impact that the rapid mechanization of agriculture had on the US labor market. First, causal estimates of the effect of more intense technological adoption on incumbent workers' subsequent labor market outcomes are derived using an instrumental variables (IV) estimation strategy. The proposed instrument is farmland ruggedness, which I show was unrelated to agricultural productivity before the widespread mechanization of agriculture but provided a significant impediment to the adoption of tractors and larger, more productive equipment. Conceptually, these estimates provide a comparison of the outcomes of ex-ante similar agricultural workers, one of which lived in a county better suited to adopt new agricultural machinery. In this sense, they provide inference about the effects of increased technological change in a given county within a partial equilibrium framework. To evaluate the welfare implications of technological change in early 20th century American agriculture in a general equilibrium setting, the second part of this paper develops an occupational sorting model using estimates derived from the IV specification to identify the magnitude of technological shock needed to explain the estimated effect of additional investment on incumbent workers' subsequent occupational mobility along the economies' transition path.

In many settings, technological change occurs along dimensions that are not directly observable. One further advantage of the 20th-century agricultural setting is that agricultural census data provides an account of the stock of agricultural equipment capital in each US county over time, providing a direct measure of the degree of technological change occurring in each county in the US. Evidence from a separate panel of detailed equipment counts and productivities for specific types of farm equipment indicates that changes in equipment value per acre of farmland in the agricultural census are strongly correlated with purchases of more novel (new to the region) and productive (measured in hours of labor per acre) pieces of equipment.[1] Based on this evidence, changes in equipment value per acre of farmland, *i.e.* investment net of depreciation, provides a measure of the amount of technological change a given region experienced.

The core IV estimation strategy used throughout the paper argues that the ruggedness of farmland, defined as the two-dimensional sum of square differences in elevation at a resolution of 10 meters, acted as a barrier to the adoption of tractors and the larger scale farm equipment which embodied agricultural technological change between 1910 and 1920.[2] This strategy is conceptually similar to the one used in Boone

[1] The specific data comes from McKibben *et al.*, 1939 and was digitized as part of this project

[2] Ruggedness measured at 10m intervals is aggregated to a county-level measure using the average

and Wilse-Samson, 2021 to study how the shock of the great depression differed based on the level of agricultural mechanization within a region. In contrast, this paper examines pre-depression mechanization and focuses on the labor market outcomes of incumbent workers and their children. Further, the underlying ruggedness measure implemented here is quite different from the measure in Boone and Wilse-Samson, 2021.

Primary source documentation qualitatively supports the relevance of this ruggedness instrument and standard first stage tests of relevancy reject ruggedness as a weak or irrelevant instrument. Further, two key pieces of evidence support the exclusion restriction required for the IV estimates to provide a causal estimate. First, ruggedness was unrelated to agricultural productivity measured as the average product of agricultural labor (APL) in 1900 but became strongly negatively correlated with APL by 1920, suggesting that i) there was a structural break between the relationship of ruggedness as measured in this paper and agricultural productivity concurrent with the initial adoption of tractors and increased productivity of farm equipment and ii) before this adoption began, regions with less rugged farmland were no less productive than regions with less rugged farmland. This second implication indicates that ruggedness is unlikely to violate the exclusion restriction due to a deep structural connection to agricultural productivity. The second piece of evidence is that the IV model estimates a null effect of additional investment in agricultural technology on non-agricultural workers' occupational mobility, indicating that the channel through which the instrument operates is specific to agricultural workers and not due to general economic shocks to different regions (a violation of the exclusion restriction).

The empirical analysis yields several insights. First, I find that increased technological adoption caused some workers to leave agriculture for other occupations, *i.e.*

value in the most fertile land within a county.

technological change caused occupational displacement. A one standard deviation increase in the rate of technological adoption (investment) within a county is estimated to have caused a 4 percentage point increase in the likelihood an incumbent worker left agriculture for an alternative occupation between 1910 and 1920. This occupational displacement was largely concentrated within wage-earning workers (as opposed to farm owner-operators and their family members) for whom a one standard deviation increase in technological adoption increased occupational mobility by nearly 20% (10pp).

Second, I find that increased investment in new technologies is estimated to have impacted workers' labor earnings. While wages are only observable in the 1940 census, regions that quasi-randomly received additional investment saw large increases in the average product of labor within agriculture, whereas evidence suggests that wage-workers who left agriculture from high-investment regions were more likely to transition into lower-paid occupations (as measured using 1940 wages) compared to the average move out of agriculture. While this latter estimate is conditional on an ex-post outcome (leaving agriculture), and thus potentially conflates a direct impact of technological change with selection on unobservable ability, the limited measures of human capital in the 1910 census suggest that the marginal wage-worker who left agriculture from a high-investment region had an identical level of literacy as a worker who left from a low-investment region, indicating that these results are unlikely to be explained by selection along measures closely associated with literacy. Interestingly, additional investment is estimated to have had no impact on geographic mobility out of workers' baseline county or state. In other words, displaced workers did not leave for regions with more opportunities, but remained in their incumbent counties and took lower-paid local jobs. This result suggests that while there was a significant amount of geographic mobility at the time, the marginally displaced worker still likely faced significant mobility frictions. Further evidence indicates that some workers were better situated to benefit from technological change. These "winners" were largely farmers and their family members who saw an increase in productivity and little impact to their occupational mobility. Conversely, some wage workers were pushed out of agricultural occupations due to technological change, resulting in the average worker moving into an occupation with lower than expected average earnings.

Third, this paper provides some of the first evidence that technological change can impact intergenerational mobility. Specifically, children of farmers from high-investment regions earned more as adults in 1940 than their peers from low-investment regions. The first generation estimates indicate that the fathers of these children were the likely beneficiaries of technological change, in that they did not experience significant occupational displacement and reaped the benefits of increased labor productivity. On the other hand, the children of agricultural workers who left agriculture by 1920 from high-investment regions had lower wages than their peers from low-investment regions. As with the first generation results, the interpretation of this second result as causal is slightly more complicated, as conditioning on fathers who left agriculture does not rule out selection on unobservable parental ability. However, these results indicate that either i) displaced fathers were very negatively selected, for which there is no evidence in the data, or ii) parental displacement due to technological change had a negative impact on children's wages. Together these results imply that children's outcomes mirrored their fathers', in that the sons of winners themselves were better off than their peers, while weaker evidence implies the sons of technological losers earned less than expected compared to the children of workers who left agriculture from low-investment regions.

Earnings inequality, measured as the variance of labor earnings, was higher within children from less-rugged (and thus higher investment) counties. This difference is particularly large among children of fathers who left agriculture. Interestingly, this increase in inequality is driven by both ends of the income distribution, with a higher share of children from less-rugged counties earning very little and much more than the median child as compared to the distribution of children from more-rugged counties. While it's difficult to isolate the mechanisms which transmitted the impact of technological change into a second generation, evidence suggests that education may play a role. Children of farmers from high-investment regions worked in occupations with a higher share of educated workers. Estimated impacts on years of education further support a human capital channel, with investment increasing educational attainment for sons of farmers and decreasing attainment for sons of workers who left agriculture, though these estimates are quite imprecise. Taken together, these empirical results provide valuable evidence that technological change in early 20th-century agriculture had important effects on occupational mobility, the distribution of labor income, and intergenerational mobility.

In order to compare how the magnitude of the costs of technological adoption highlighted by the empirical results compare to the aggregate benefits of additional productivity within a consistent framework, empirical estimates are used to discipline the transition path of an equilibrium dynamic occupational sorting model based on those in Dvorkin and Monge-Naranjo, 2019 and Garcıa-Couto, 2021. The model allows the realized path of wages and occupational mobility to be compared to a counterfactual world in which technological change was never realized. Additionally, the model allows for an evaluation of wages over time, rather than indirect measures such as occupational wages in years besides 1940. Within this model, agents dynamically sort across occupations based on relative productivity, and technological change is realized as an increase in labor productivity for agricultural laborers in one of the two agricultural regions. By matching the difference in occupational mobility between agricultural regions based on the productivity shock to the empirical estimates, the model is able to replicate the observed impact of technological change.

Overall, 84% of agents had higher lifetime utility in the economy which experiences the technological shock as compared to remaining in the initial steady state. However, the agents made worse off by technological change were significantly affected; the gross consumption equivalent costs of technological change are 11% of the gross surplus generated among all agents alive at baseline. The model further finds that the intergenerational channel identified in the empirical analysis is economically relevant when accounting for the welfare implications of technological change, with 20% of the welfare losses accruing to children. However, failing to account for children's welfare would also overestimate the relative size of the welfare losses relative to the gains, as children enter the labor market after some of the costly occupational transitions were already made and so are relatively less impacted than their fathers. These estimates highlight how new technologies can both "expand the pie" and have significant adverse effects on a relatively small proportion of the population.

This paper makes several contributions to the existing literature. First, it provides a consistent evaluation of the costs and the benefits of a major episode of technological change and demonstrates that while aggregate surplus is increased, a relatively small share of the population potentially experienced significant welfare losses to achieve this aggregate surplus. These findings contribute to several areas of research, perhaps most directly to the broad literature of skill-biased technical change, which has more recently focused on the adoption of computers and wage inequality (for example Autor, 2019), but has also been studied over a longer horizon (Katz and Margo, 2013, Acemoglu, 2002) and for specific historical technologies (Feigenbaum and Gross, 2020, Fiszbein *et al.*, 2020).

The second way this paper contributes to the existing literature is by documenting that welfare losses from technological change can be persistent and may transmit to a second-generation who did not directly experience the technological change, a channel rarely discussed in the context of technological change. One notable exception to this is Feigenbaum and Gross, 2020 who evaluate intergenerational effects across cohorts, rather than within family. There is also a body of work that connects parental job loss to children's outcomes such as educational attainment, as in Hilger, 2016. However, to my knowledge, this is the first paper to evaluate the direct impact of additional technological adoption experienced by a parent on children's subsequent outcomes in adulthood.

Third, this paper documents an explicit connection between the adoption of a new labor-augmenting technology and subsequent occupational displacement for a segment of the incumbent workforce. There is relatively scarce empirical evidence establishing a causal connection between technological change and subsequent outcomes for incumbent workers with Humlum, 2019, Bessen *et al.*, 2019, and Feigenbaum and Gross, 2020 as notable examples. Within the context of agricultural mechanization, Eisner, 2021 uses variation in crop shares to estimate how the introduction of tractors impacted the occupational choices of Midwestern farmers, with an emphasis on heterogeneous effects by age. This paper uses a different, more general, source of identifying variation, evaluates effects of technological investment in both machinery and tractors, focuses on heterogeneity by occupation (farmer vs farm laborer), and evaluates the intergenerational effects of technological change. Given the degree to which new technologies continue to shape our world, additional data on the connections between technological adoption and worker outcomes are key to improving our ability to predict what future labor markets will look like.

Finally, this project provides valuable historical insights by bringing modern methods to a major event in the history of the United States' economic development which is often cited in discussions about the adoption of new technologies (Acemoglu and Restrepo, 2019, for example, uses agricultural mechanization as a historical example of automation). Olmstead and Rhode, 2001 and Binswanger, 1986 provide additional insights into the mechanization of agriculture. Mostly closely related with this paper, Boone and Wilse-Samson, 2021 examine how the great depression modulated the structural change induced by agricultural mechanization.

Together, these contributions have important implications for our understanding of how wealth and income are distributed and open new questions about the extent to which modern technologies may be shaping the outcomes of younger generations even today.

2.1 Historical Setting

This section will present several facts about the US agricultural sector and agriculture technology in the late 19th and early 20th centuries which will guide the remainder of the analysis.

The early 20th century was a time of rapid changes in American agriculture. In some 30-40 years, American agriculture changed from being nearly entirely animal-powered to utilizing a relatively robust rural electricity grid, tractors, and gasoline-powered vehicles. Over the same period, significant accomplishments in breeding high-yield crop varieties, implementing modern land management practices, and the development of chemical fertilizers also had important impacts on agricultural productivity. However, this project will focus on the specific role that the increased adoption of more productive farm equipment, broadly defined as machinery used in agricultural production, played in altering the demand for agricultural labor.

The widespread adoption of more productive farm equipment had a substantial impact on labor productivity. Figure 2.1 Panel A plots the change in productivity, measured as the operating hours needed to complete the 5 categories of tasks central to crop farming, within a consistent panel of farms surveyed as part of the 1936 Works Progress Administration report, "Changes in Farm Power and Equipment" (see McKibben *et al.*, 1939 for details).[3] In just 10 years, from 1909 to 1919, the average labor hours required to complete all of the tasks fell 13% from 8.7 to 7.5, indicating significant changes in the quality of capital stocks. We further see that the rate of advancement slowed significantly over the following decades, suggesting that the largest improvement to the quality of agricultural capital occurred earlier in the century.

To what extent can the trends in Panel A of Figure 2.1 be attributed to technological innovation rather than the increased accumulation of established technologies? In one very tangible sense, tractors were effectively non-existent in 1909 but were being rapidly adopted in 1919, when as many as 1 in 5 farms had a tractor in some regions. A more direct way to evaluate the productivity gains from the adoption of the tractor is to evaluate the increased productivity of farm equipment made possible by the adoption of tractors. Panel B of Figure 2.1 plots the share of newly purchased equipment that represented a novel design to an agricultural region over time. Here, the number of newly purchased machines is measured as any positive change in the aggregate number of a specific size and type of machinery reported within the seven agricultural regions designated in McKibben *et al.*, 1939. We see that for some categories of machinery, namely Tillage and Plowing, more than 15% of newly purchased machines were of a size/type completely new to the region in which they

[3]The hours reported in Panel A of Figure 2.1 summarize the productivity of the capital stock of farmers over time, but do not accurately represent the actual labor required to successfully harvest a crop, as not all tasks required machinery and some tasks may be repeated several times over a growing season.

were purchased, and even within the other categories, a non-trivial proportion of new machinery represented demonstrably new technologies, at least within the sample of farms surveyed. It's important to note that the number of entirely new machines introduced successfully over the period in question was limited. Rather, improvements to productivity came from the refinement of existing classes of machinery, particularly the development of larger, more technologically advanced versions of existing designs, likely in large part driven by the increasing availability of tractor power. In this sense, technological change in early 20th-century agriculture is best characterized as a series of incremental improvements to existing designs that nevertheless cumulatively amounted to a significant improvement in productivity. To summarize, the Works Progress report on farm equipment provides compelling evidence that farm equipment became significantly more productive in the early 20th century, with many of the new pieces of machinery representing larger versions of existing designs not previously observed in the region in which they were purchased.

While the Works Progress report provides quantifiable evidence of technological progress, it has a limited ability to tie this progress to the labor market trends of incumbent agricultural workers. Specifically, it only covers the years 1909-1936 and provides evidence for a relatively small number of counties (35 counties, aggregated into 7 regions). Figure 2.3 provides evidence about the rate of capital accumulation within the agricultural sector over a much longer time frame using data from the full decennial and agricultural censuses. Specifically, the y-axis represents the average value of farm equipment per agricultural worker in constant dollars across all counties in the United States.[4] We see that in the Post-Civil War era, the rate of capital

[4]Here, and for the remainder of the paper, agricultural worker refers to anyone aged 16-65 whose primary occupation is recorded as farmer, farm manager, farm foremen, and farm laborer (wage and unpaid family workers)

accumulation steadily increased with a significant uptick beginning around 1900.[5] The single largest increase in capital per worker occurred between 1910 and 1920, with the rate of capital accumulation per worker slowing after 1920. These trends are consistent with the evidence on equipment productivity in Figure 2.1, in that the period with the fastest accumulation of capital per worker corresponds to the period with the greatest improvements in machinery productivity within the Works Progress panel. The fact that the 1910-1920 period saw the fastest rise in machinery value per worker and the largest gains in machinery productivity within the Works Progress panel informed the decision to focus the bulk of the following analysis on this period, hence the shading in Figure 2.1, although other time-points are also be evaluated to provide supporting evidence.

While Figure 2.1 demonstrates the increasing relevance of machinery in agricultural production over time, it's not clear the extent to which these trends are related to the technological progress on display in Figure 2.3. However, evidence suggests that the two are in fact intricately linked. Figure 2.4 demonstrates that within the 35 counties surveyed as part of the Works Progress Panel, larger increases in equipment value (in constant dollars) per acre between 1910 and 1920 as measured in the agricultural census is associated with: (i) increases in the quantity of new equipment, (ii) increases in the share of new equipment which is novel to a region, and (iii) and increased average productivity of new equipment.[6] Based on this evidence, increases in equipment value per acre will be used throughout this paper as an intensive margin measure of technological change.

[5]Records for the 1890 census have been lost.

[6]The unit of observation for each panel in Figure 2.4 is a county implement type, with the same implement types as in Figure 2.3. Standard errors derived from block-bootstrap at the county level.

2.2 Data and Measurement

2.2.1 Data Sources

The data for this paper was aggregated from several sources. Broadly, the data sources fall into two categories, historical census data at the individual or county level and geographic information system (GIS) data which was then aggregated to a useful geographic unit, such as a county.

First, I will discuss the primary sources for historical data. Micro-data from decennial censuses was obtained from IPUMS full count census data (Ruggles *et al.*, 2021. Linkages that allow individuals to be followed across censuses were taken from the Census Linking Project (Ran *et al.*, 2020). Together, these data sources provide information on individuals' occupational mobility, as well as a number of other outcomes of interest, both in the cross-section and within-individual over time. The full decennial census was also used to measure aggregate agricultural labor supply at the county and national level with a high degree of accuracy, as well wage data for the 1940 census.

Key data for this project was also obtained from the agricultural census, specifically from Haines *et al.*, 2018. Measures of the value of farm equipment, acres of farmland, and acres of improved farmland by county were taken from this dataset. This dataset also provided the key inputs for estimates of the returns to agricultural labor, namely the value of aggregate output, intermediate inputs, and agricultural rents by county.

Additional historical data was obtained for the 1936 Works Progress Administration report on Changes to Farm Equipment McKibben *et al.*, 1939. The original data for this report was collected from a sample of 3,363 farms across 35 counties in 1936 and retrospectively evaluated the number and productivity of farm equipment at these farms between 1909 and 1936. The aggregated data from the report was digitized as part of this research project and is publicly available from the authors' website. Specific sample sizes by county and information about the adoption of tractors and trucks were obtained from the sister report McKibben and Griffin, 1938.

Next, I will discuss the primary sources for GIS data. First, a nationwide elevation profile was obtained from the USGS 3D Elevation Program. Data on soil and environmental characteristics within a county were constructed from the USGS SSURGO dataset. Specifically, this dataset provided information on soil fertility (the National Commodity Crop Productivity Index), soil composition (the USGS soil taxonomy classification), root zone depth, available water within the root zone, the susceptibility of the soil to drought, the number of warm days per year, and soil erosion tolerance. County-level indexes were constructed as an area-weighted average of the spatial data.

2.2.2 Measurement

This section provides an overview of how several of the key variables used in the later analysis are constructed.

Ruggedness

As will be discussed in Section 2.3, terrain ruggedness will be used as an instrumental variable for the adoption of technology within a region. The specific measure of ruggedness was first described in Riley *et al.*, 1999, and captures the local variation in elevation at a specific resolution.[7] The index is constructed from rasterized elevation data and is defined as the square root of the sum of squared differences in

[7] See Nunn and Puga, 2012 for an example application of the Riley, DeGloria, Elliot ruggedness index.

elevation between a point and the 8 adjacent points. If e_{ij} is the elevation in grid point (i, j), then the ruggedness index at that location is calculated as:

$$Index_{ij} = \sqrt{\sum_{l=i-1}^{i+1} \sum_{k=j-1}^{j+1} (e_{lk} - e_{ij})^2}$$

Figure 2.2 provides a visualization of the ruggedness measure in two representative counties. Conceptually, this index captures the two-dimensional variation in elevation at a fixed distance. The resolution for most analysis will be fixed at $\frac{1}{3}$ of an arc-second (approximately 10 meters), the finest resolution currently available through the 3DEP program. At this resolution, rugged terrain is likely to be hilly or rough, but both smaller resolution features such as rockiness and larger resolution features such as large slopes or mountains will have minimal influence on the index. This fine resolution measure of ruggedness captures the sort of steep, sloped terrain which was not easily adapted to the use of larger, largely tractor-driven farm equipment.

The disaggregated index values were converted into a county-wide ruggedness measure by taking the average of all grid points within a county at which soil fertility was above a cutoff value, as visualized in Figure 2.2 Panel B. The cutoff was determined such that the total area of included points matched that area of all farmland in 1910. By conditioning on fertility, the final ruggedness index attempts to capture the ruggedness of representative farmland within a county, and avoid areas within a county that were unlikely to be used for agriculture at the turn of the 20th century.

Occupational Wage

Within the available full-census datasets, only the 1940 census contains information on individuals' wages. For analysis with outcomes measured in 1940, this is not a problem, but for years other than 1940, a measure of occupational wage will be used as a second-best approximation of the return to an individual's labor. Specifically, occupational wages are measured as the average labor income for men aged 16 to 65 observed in the 1940 census for a given occupation county pair. If a county occupation pair is empty, the geographically nearest non-empty value is used. For the small number of occupations for which the 1940 census did not provide a direct analog, the average occupational wage for agricultural workers from the same baseline (1910) county who transitioned out of agriculture is imputed. This measure captures the earnings a given worker in 1920 would have received if they were to remain in the same occupation for 20 years. Under the assumption that relative earnings across occupations were relatively constant between 1920 and 1940, this measure may be a reasonable approximation of the wage income a worker earned in 1920. Alternatively, occupational wage can be thought of as an index of the relative labor earnings of different occupations.

Average Product of Labor

Throughout this project, agricultural worker will refer to anyone of working-age whose primary occupation is recorded as farmer, farm manager, farm foremen, farm wage laborer, and unpaid family laborer. For farm managers, farm foremen, and farm wage laborers, data from the 1940 census may provide a reasonably accurate measure of their occupational wage. However, farmers and their unpaid family members often do not report meaningful wage data in the census. Therefore, an alternative measure of these workers' labor productivity, the average product of agricultural labor, will be used. This measure is calculated using the residual profit per agricultural worker after accounting for inputs into production and the user cost of capital. The average product of labor (APL) for county c is calculated as follows:

$$APL_c = \frac{LaborProduct_c}{CountAgLaborers_c} = \frac{ValueOutputs_c - ValueInputs_c - UserCostK_c}{CountAgLaborers_c}$$

The agricultural census provides direct estimates of $LaborProduct_c$ and counting the total number of agricultural workers observed in the full census data provides an accurate estimate of $CountAgLaborers_c$, which can account for unpaid family workers, farmers, and hired labor. Estimates of $ValueInputs_c$ are also taken from the agricultural census and include expenditures on feed, commercial fertilizer, and liming materials.

Removing the user cost of capital from the product of labor allows us to separate the income a farmer receives for owning capital from the returns he receives for his labor. The average product of labor is what is relevant for determining how occupational mobility impacts income as a farmer who owns land and equipment can move into a manufacturing occupation but still rent out his capital. The specific calculation of user costs of capital is as follows:

$$UserCostK_c = TotalAcFarm_c RentPerAcre_c + EquipValue_c UserCostEquip$$

$TotalAcFarm_c$, $RentPerAcre_c$, and $EquipValue_c$ are all taken directly from the agricultural census, while $UserCostEquip$ is approximated at 15% based on a 3% interest rate as a 12% cost of ownership (including maintenance and depreciation), estimated in Butz and Lloyd, 1939. The measure of rent per acre includes rent for both the land itself and the structures on that land and is taken from farmers who reported renting their farms. Therefore, rent per acre allows for a reasonable approximation of the user costs of both the land and building capital within a county. $RentPerAcre_c$ is only available in the 1940 agricultural census, and so when calculating the average product of labor in other years, the 1940 value is adjusted by the national trend in agricultural rental prices taken from Lindert, 1988.

While treating the average product of labor as a measure of labor income is an imperfect approximation, the nature of the agricultural labor market, particularly historically, makes estimating labor income in agriculture difficult even with high-quality data. One test to determine if the estimates of the average product of labor yield reasonable estimates is to evaluate where they place agricultural workers in the income distribution. The average agricultural worker had a 1940 average product of labor of $559, less than the average wage of a general laborer ($673) and much less than the average engineer ($2,498).

Investment in Agricultural Capital

Using agricultural census data, a measure of the change in agricultural capital per farm acre between 1910 and 1920 will be used as a measure of technological change within a county. Section 2.1 discusses the interpretation of this measure in more detail, but the goal is to capture changes in equipment capital over time, accounting for the fact that farmland might be expanded or contracted within a county. In a slight abuse of standard definitions, changes in the value of equipment capital will be referred to as investment, although it is not taken net of depreciation. For a county c, the specific measure is calculated as follows:

$$g_c = 1 - \left(\frac{\frac{v_c^{1920}}{acres_c^{1920}} - \frac{v_c^{1910}}{acres_c^{1910}}}{\frac{v_c^{1910}}{acres_c^{1910}}} \right)^{\frac{1}{9}}$$

where v_c^t is the value of all farm equipment in county c in year t (adjusted to 1940 dollars using the consumer price index) and $acres_c^t$ is the total number of improved farmland acres in a county (farmland that could be used to grow crops). The investment is annualized over the 10 year period to produce a measure of annual percent growth of equipment value per acre. To aid in the interpretation of the results, this measure of investment will often be expressed in units of county standard deviations.

2.3 Empirical Strategy

This section will outline the main empirical strategy used throughout the paper. The goal of the empirical section of this paper is to estimate the causal relationship between the adoption of new technology, as measured by increases in the value of farm equipment at the county level, and subsequent outcomes for incumbent agricultural workers. The primary obstacle to achieving this goal is a lack of exogenous variation in agricultural investment. The endogeneity of investments in agricultural equipment to workers' occupational mobility (to name one of several outcomes of interest) could arise for a number of reasons. Perhaps most obviously, one should be concerned about reverse causality. If agricultural workers primarily left farming for higher-paid manufacturing jobs, then the agricultural sector may be forced to invest more heavily in equipment in order to increase the productivity of remaining workers and meet the demand for agricultural goods. In this case, one would expect to see a positive correlation between the likelihood a given worker left agriculture and the amount of investment in agricultural equipment. However, it would be incorrect to conclude that the adoption of new technologies caused the exodus from agriculture.

In order to avoid incorrectly conflating a correlation with a causal relationship, this paper will utilize an instrumental variables (IV) strategy to estimate the causal impact that the adoption of new agricultural technology had on key outcomes of interest for incumbent workers. The proposed instrument will be the terrain ruggedness of the most fertile land in a county, with a measure of ruggedness inspired by Nunn and Puga, 2012. Effectively, ruggedness can be thought of as the two-dimensional variation in terrain elevation at a fixed resolution. The specific regression model will be as follows:

$$g_{iacs} = \alpha Rugged_c^{10m} + \mathbf{X_c^{env}}\gamma + \mathbf{X_i^{indiv}}\theta + \lambda_s + \nu_a + \eta_{iacs} \tag{2.1}$$

$$Y_{iacs} = \beta\widehat{g_{iacs}} + \mathbf{X_c^{env}}\delta + \mathbf{X_i^{indiv}}\kappa + \mu_s + \xi_a + \epsilon_{iacs} \tag{2.2}$$

where i indexes the individual, a indexes 1910 age, c indexes 1910 county, and s indexes 1910 state. g represents the annualized percentage increase in equipment value per acre of farmland between 1910 and 1920, measured in county standard deviations (to provide estimates in a more meaningful unit), $\widehat{g}$ represents the predicted value of g estimated using Equation 2.1, $Rugged^{10m}$ represents the ruggedness of agricultural land at a 10m ($\frac{1}{3}$ arc-second) resolution, X^{env} represents county specific environmental controls, and X^{indiv} represents controls specific to the individual. Further, both the specifications in Equation 2.1 and Equation 2.2 include state and age fixed effects. Unless otherwise noted, the complete IV model specified in Equations 2.1 and 2.2 will be estimated using limited information maximum likelihood (liml) with standard errors clustered at the 1910 county level (c).

The inspiration to use topography as an instrumental variable comes from primary sources material on the adoption of tractors and the productivity gains of farm equipment. For example, McKibben and Griffin, 1938, p. 31 states; "it is an accepted fact that tractors are more adaptable to use on level or slightly undulating land than on rolling or rugged farmland." In their subsequent report, McKibben *et al.*, 1939 describe how rugged terrain reduces the productivity advantages of larger pieces of equipment: "if the machine is made up of units ... any difficulty - clogging, breakage or misadjustment - with one unit stops the entire machine... This situation has a bearing particularly in small fields and where the soil is stony or the topography rugged" (page 16).

Under a set of assumptions discussed below, the estimated β coefficient in Equation 2.2 has a casual interpretation. The sources quoted above make a compelling argument that the first assumption, relevancy, is satisfied in this setting. However, relevancy is also testable, and the first stage heteroskedasticity-robust F statistic (derived from standard errors clustered at the county level using the method described in Olea and Pflueger, 2013) is around 18 or 19 depending on the specification, above the critical value derived by Stock and Yogo, 2005 bounding the maximal bias of the IV estimate below 10% of the OLS bias.

The second assumption required for the proposed IV strategy to provide a casual interpretation is the exclusion restriction. Specifically, I require that conditional on controls, the instrumental variable (terrain ruggedness) only impacts an outcome of interest, for example, the likelihood an agricultural worker remains in an agricultural occupation, through the endogenous variable (the increase in equipment value in the worker's initial county). While this is a fundamentally untestable assumption, several pieces of evidence suggest it is likely reasonable. Many of the potential threats to the exclusion restriction fall into one of two broad categories, either the instrument is correlated with local economic shocks/trends outside of agriculture or the instrument is correlated with agricultural productivity (outside of the ability to adopt new technologies).

First, consider threats to identification due to local economic shocks or trends outside of agriculture. Terrain ruggedness could be correlated with changes to local economic conditions for a number of reasons. For example, more rugged regions may be less likely to experience additional investment in their railroad network or be more likely to have a mining boom that impacts the labor market decisions of all inhabitants. The common denominator behind this class of threat is the implication that while ruggedness impacts outcomes, it does so (at least partially) through a channel besides the adoption of new agricultural technology. If this is the case, one would expect to estimate similar effects within agricultural and non-agricultural workers. Therefore, one way to test for this category of threat would be to look at the outcomes of non-agricultural workers. Table 2.1 does just this and displays estimated impact of increased adoption of agricultural technology on the likelihood several different types of workers changed occupation using the IV model described above. Specifically, column 1 evaluates impacts to the occupational mobility of white-collar workers, whose labor market is likely largely separated from that of agricultural workers. Column 2 evaluates impacts to the occupational mobility of general laborers, whose labor market likely intersects significantly with that of agricultural workers. Finally, column 3 evaluates impacts to the likelihood a non-manufacturing laborer transitioned into a manufacturing industry, in order to test for endogeneity due to the local availability of manufacturing jobs. For all three samples of non-agricultural workers, Table 2.1 demonstrates that instrumented increases in technological investment are statistically uncorrelated with the occupational mobility of non-agricultural workers. Further, the magnitudes of the point estimates are economically small. These estimates stand in stark contrast to those for agricultural workers as discussed in Section 2.4. Overall, Table 2.1 indicates that any threat to the exclusion restriction which would invalidate the IV estimates must come from a factor specific to the agricultural sector.

Next, consider correlations with agricultural productivity. Given that the proposed instrument is by construction associated with the physical environment, it is natural to assume that it may be correlated with other aspects of the environment such as soil quality, weather, and hydrology which have a direct effect on agricultural productivity. On the other hand, the specific measure of ruggedness is somewhat narrow, in that it measures how much the terrain changes elevation at a very specific distance interval (10 meters) that is likely to impact the use of larger agricultural equipment. In principle, ruggedness at 10 meters may be only weakly correlated with ruggedness at 100 meters or 1 meter. To ensure that any potential link between my 10 meter ruggedness measure and relevant features of the physical environment is accounted for, the first and second stages of the IV model in Equations 2.1 and 2.2 include controls for key environmental factors which impact agricultural productivity, such as an index of the agricultural productivity of the soil within the county, ruggedness calculated at a 900m resolution, the average root zone depth for commodity crops, the amount of water storage within the root zone, the average number of warm days, and the erosion tolerance of the soil.[8] The IV model also includes state fixed effects, and so estimates are only identified using within-state variation, suggesting factors like crop type and growing zone are unlikely to explain results.

Table 2.3 tests the extent to which the control variables in Equations 2.1 and 2.2 can explain any correlation between the instrument and the returns to agricultural labor before and after the period of interest. If the exclusion restriction fails due to the second category of concern, intrinsic correlations with agricultural productivity, one would expect returns to agricultural labor to be correlated with the instrument, even in periods before qualitative evidence suggests the widespread adoption of tractors and larger farm equipment began. Columns 1 and 2 of Table 2.3 display a simple regression of log average product of labor in 1900 and 1920 respectively onto the instrument. We see that without any controls (outside of state fixed effects) the relationship is statistically indistinguishable from 0 in both periods. Further, I cannot reject the hypothesis that the coefficients in both models are equal. Columns 3 and 4 repeat this exercise but include all of the baseline controls described above. Here I see that the estimated relationship is over 10x as large in 1920 as in 1900, with only the 1920 estimate being statistically different from zero. The lack of an economically or

[8]Section 2.2.2 provides more detail on each of these measures.

statistically meaningful estimate in 1900 (columns 1 and 3) strongly suggests that the exclusion restriction is unlikely to be violated due to the instrument's correlation with intrinsic factor(s) of agricultural productivity. Further, the difference in estimates between 1900 and 1920 provides additional support for the instrument's relevancy, as it suggests that some structural shift happened between 1900 and 1920 in the relationship between terrain ruggedness and agricultural productivity. While many potential mechanisms could explain such a shift, the negative coefficient in column 4 is consistent with ruggedness restricting access to new technologies which did not exist in 1900. Finally, columns 5 and 6 repeat the estimates in columns 3 and 4 respectively, with additional controls for the type of soil within a county. Specifically, each model includes controls for the share of a county's area which is comprised of each of the 12 soil orders in the USDA soil taxonomy. We see that neither the 1900 or 1920 estimates change significantly relative to columns 3 and 4, suggesting that there is unlikely to be major omitted variables pertaining to the physical environment within a county in equations 3 and 4 which may explain the connection (or lack of connection) between ruggedness and agricultural productivity. The similar estimates on the coefficient for 10 meter ruggedness in columns 3 and 5 (as well as columns 4 and 6) also provides evidence that the baseline set of controls in columns 3 and 4 are sufficient to render 10 meter ruggedness as good as randomly assigned, the third assumption required for the estimates to have a causal interpretation. Soil orders are not included in the baseline model out of concern that they may also be indirectly correlated with the ability of a region's farmers to adopt new technologies, and there may attenuate the predictive power of the ruggedness index.

Under the additional assumption of monotonicity, the estimated β coefficient in Equation 2.2 can further be interpreted as the weighted average effect of increased technological adoption on the outcome of interest (*i.e.* the local average treatment effect, LATE), where the experiences of individuals in counties whose adoption of agricultural technology was more affected by terrain (complier counties) receive more weight. In this setting, monotonicity requires that ruggedness must either have zero impact or restrict the ability of a region to adopt new technology. If for some subset of counties, ruggedness increased the accumulation of technology, monotonicity would be violated. An example of how this assumption could be violated would be if machines specific to hilly, rugged terrain were developed over this period, and thus very flat and very rugged regions both adopted more technology. While a review of the primary source documents has not provided any evidence that this is the case, monotonicity is not directly verifiable. However, Angrist and Imbens, 1995 suggest that a necessary condition for monotonicity to hold is that the cumulative distribution functions (CDFs) of equipment investment must not cross for rugged and non-rugged regions. Figure A.3 indicates that the CDF for regions with above median ruggedness does not cross the CDF for regions with below median ruggedness, suggesting that monotonicity may be a reasonable assumption in this setting.

2.4 Empirical Results

This section presents empirical evidence on the causal impact that increased adoption of new agricultural technologies had on incumbent workers. A connection between the mechanization of American agriculture, increased agricultural productivity, and a decline in labor demand has been acknowledged for many decades. For example, McKibben *et al.*, 1939 states that "there has been a pronounced trend towards the use of larger implements [in agriculture]... The principal advantages of higher-capacity machines are decreased labor required per unit of output and increased timeliness in the performance of critical farm operations" (page 14). This section will attempt to quantify the causal impacts of this reduction in labor demand on incumbent workers.

Overall, evidence suggests that increased technological adoption caused some incumbent wage-workers (*i.e.* not farmer owner-operators) to leave agriculture. Further, the new technologies appear to have been labor-augmenting, with incumbent farmers experiencing significant increases in the average product of labor. Several pieces of evidence suggest that workers displaced by technological change experienced lower occupational wages as a consequence. Together, these results imply that technological change increased the dispersion in labor income among incumbent workers. Finally, by linking children's outcomes in adulthood to the technological change experienced by their fathers, I find that the children of workers who benefited from technological change had higher incomes than their peers, while the children of displaced workers had lower wages, providing new insights into the sources of inter-generational mobility.

2.4.1 Sample Characteristics

Table 2.2 provides descriptive statistics about the primary analysis sample. Column 1 displays the observable characteristics of all men aged 16 to 55 in 1910. Of particular note is the share of individuals who lived in a rural area (over 50%) and the high labor force participation rate (93.4%). The latter is likely influenced by the lack of a modern social safety net which began as part of the 1930s era New Deal legislation. Column 2 restricts the sample in column 1 to only men working in agricultural occupations. Unsurprisingly, the share of individuals living in a rural area is much higher within this sub-sample. The average agricultural worker also has nearly 0.5 more children, is slightly more likely to be married, and is less likely to be literate. Column 3 restricts column 2 by limiting the sample to only agricultural workers who are potentially eligible to be included in the primary analysis sample. Specifically, this column excludes men living outside the Contiguous United States, those living in urban counties (determined by the lowest 5% of counties by agricultural labor share), and men below 22 or above 55 years of age. Column 4 then presents observables for the share of individuals in column 3 with an exact match to an observation in the 1920 census using the standard exact matches provided by Ran *et al.*, 2020. The final column presents the test of equality between columns 3 and 4. As is well documented in other projects using linked census data, the likelihood a given individual is linked is not random. However, the actual degree of selection appears unlikely to be economically relevant along many dimensions such as age, family size, number of children, rural/urban household, etc. There are several dimensions on which selection into the matched sample appears potentially meaningful, such as literacy or race, where literate and white respondents appear more likely to be matched. Therefore, I expect my results to overweight the experiences of white or literate workers over their African American or illiterate countrymen. As all analysis will be conducted within the matched sample, any selection into this sample is unlikely to cause spurious results but does highlight the fact that some demographic groups may not be well represented by the conclusions ultimately drawn from the analysis sample.

2.4.2 Occupational Mobility

First, I turn to the connection between technological change and occupational mobility. Table 2.4 displays estimates of the impact of increased exposure to technological change, measured in (county) standard deviations of the change in equipment value per acre of farmland between 1910 and 1920, on the likelihood a given 1910 agricultural worker moved to a non-agricultural occupation in 1920. Column 1 presents estimates from a naive OLS regression. From this specification, I estimate a statistically insignificant 0.01pp increase in the likelihood an incumbent agricultural worker remained in an agricultural occupation if their exposure to technical change was increased by one standard deviation. Compared to the 31.7% likelihood a random worker left agriculture, this estimate appears to be a precisely estimated zero effect. However, column 2 demonstrates that this OLS estimate is indeed biased, as the equivalent IV estimate using the same controls finds a large positive effect of technological change on the likelihood a worker left agriculture. Specifically, a one standard deviation increase in technological change is associated with a 3.7pp increase in the likelihood a worker left agriculture. The heteroskedasticity-consistent first stage F-stat testing for the relevance of the ruggedness instrument indicates that at most, the IV estimate has less than 10% of the bias of the OLS estimate in column 1 (based on Stock and Yogo, 2005 critical values, see Andrews *et al.*, 2019 for a discussion of heteroskedasticity and weak instrument tests). Therefore, my instrumental variable estimation strategy discussed in Section 2.3 appears to satisfy conventional tests of instrument relevancy.

Agricultural workers are a broad category that includes farmers, unpaid family workers, farm managers, and wage laborers. We may expect that additional technology may have a differential impact depending on a worker's initial occupation type. Columns 3 and 4 of Table 2.4 separately estimate the impact of technology on occupational mobility for farmers (both farmers and unpaid family workers) and wage-workers respectively. While both groups are estimated to have had increased occupational mobility due to technological change, the estimated effect on farmers is not statically significant and is approximately $\frac{1}{4}$ the magnitude of the estimates for wage-workers. These results indicate that the new technology largely replaced the demand for wage labor, as opposed to reducing the demand for labor from a farmer's own family. This is an intuitive result, in that a farmer making employment decisions over the family farm will likely treat the labor supplied by the household as relatively more fixed than hired labor.

It's important to emphasize that the IV estimates in Table 2.4 do not imply that all transitions out of agriculture were driven by technological change. On the contrary, a very rough back-of-the-envelope calculation would suggest that, at most, 9% of all observed moves out of agriculture between 1910 and 1920 can be explained by reduced labor demand from technological change.[9] Evidence presented in Section 2.4.4 will also suggest that the average move out of agriculture was associated with a significant increase in occupational wage, implying that the representative move out of agriculture was likely beneficial, at least in terms of wages. However, the estimates in Table 2.4 directly imply that some workers who left agriculture by 1920 would have chosen to stay had the degree of adoption of new agricultural technologies in their area been reduced. This group of workers can be thought of as being directly displaced from their incumbent occupation by technological change. From a revealed preference standpoint, it is less clear that these displaced workers benefited from their induced occupational mobility.

2.4.3 Average Product of Agricultural Labor

The IV estimates in Table 2.4 are identified using quasi-random variation in the ability of a region to adopt new technologies, and thus under the exclusion restriction isolate a causal connection between technological change and occupational mobility to a demand, rather than supply, channel.[10] Given this evidence of a decrease in labor

[9] This calculation sums the number of workers estimated to have been displaced in each county by multiplying the β coefficient by the investment level in the county and the number of workers in that county. Together, the total number of displaced workers is only 9% of all of the moves out of agriculture in the 1910-1920 panel

[10] To see why, consider the case that some agricultural workers were "poached" from agriculture by higher-paid manufacturing jobs, and thus the remaining workers purchased equipment to make up for the shortfall in labor. Under this violation of the exclusion restriction, I would expect low-skilled

demand, equilibrium analysis would intuitively suggest that the average product of labor also decreased in areas with more technological adoption relative to areas with less adoption. However, Table 2.5 shows that the opposite occurred. Specifically, Table 2.5 display coefficients from IV regressions of average product of labor in agriculture using the baseline IV model described in Section 2.3. The sample and coefficient of interest are identical to those in column 2 of Table 2.4, so that the LATE estimates in Table 2.5 rely on the same identifying variation as in other Tables, even though the outcome varies at the county, not the individual level. In each column, the outcome variable is calculated using a different time period. We see that changes in equipment value per acre are uncorrelated with the average product of labor in the baseline period (1910). However, there is a positive relationship starting in 1920, which becomes more accurately estimated in 1940, the period for which non-agricultural wages can be observed.

---

non-agricultural workers to also be pulled into these better manufacturing jobs. However, as discussed in Section 2.3, IV estimates of the likelihood low-skilled laborers moved into a manufacturing industry were statically insignificant and of the opposite sign as the estimates on agricultural workers. Later analysis will argue that the technological change in agriculture identified using changes in equipment value per acre was likely labor-augmenting, and so given the evidence in support of the exclusion restriction, the estimates in Table 2.4 provide direct evidence that increased adoption of new labor-augmenting technologies can have a direct impact on the subsequent demand for labor. Further, these estimates imply that such a reduction in demand due to technological change can have significant ramifications for incumbent workers and that these effects may be heterogeneous by initial occupation. While the historical context cannot be ignored, neither of these findings are ex-ante obvious. It is completely plausible that the causation was reversed, that technological adoption in US agriculture was wholly driven by shocks to labor supply, rather than causing shifts to labor demand. Further, it is conceivable any such demand shocks were met by changes in occupational entry, rather than the exit of incumbent workers. In this sense, the evidence in Table 2.4 provides an important data point in my larger understanding of the effects of technological change.

As described in Section 2.2.2, the average product of labor is measured after removing the estimated returns to capital from land, buildings, and equipment. Therefore, an increase in aggregate productivity due to the accumulation of additional equipment capital does not mechanically increase the average product of labor. Another possible explanation for the positive coefficient in columns 2 and 3 of Table 2.5 is selection. If the marginal workers who were induced to leave agriculture because of technological growth (*i.e.* the compliers in the LATE framework) were less productive than the average worker, then you may expect to get similar estimates to those in panel A without any changes to workers' productivity. However, if selection was the primary mechanism explaining the increase in the average product of labor, then measuring returns holding the number of workers constant at the baseline 1910 level should yield null or negative estimates. Panel B demonstrates that holding the count of workers constant at the 1910 level results in nearly identical estimates as those in Panel A, indicating that selection is not a potential explanation for the increased average product of labor. Therefore, any explanation of the increased average product of labor in counties with high investment rates requires that the productivity of at least some incumbent workers increased. This strongly argues for the interpretation that the new agricultural machinery which embodied technological change was labor-augmenting.

It is important to note that while the average product of labor provides insight into labor incomes, it does not necessarily correspond to a full accounting of the effects on incumbents' welfare, as new technology may have changed the value of the capital incumbents' already owned. For example, technological change may have reduced the value of incumbents' existing equipment stock or increased the value of their land, neither of which is addressed in Table 2.5. While these are important margins to consider when evaluating the full costs and benefits of early 20th-century agricultural technology, it is the impact on workers' labor income that is relevant to evaluate the impact of occupational mobility on workers' welfare. Further, impacts to labor income are the most relevant factor when using this historical event as an analogy for more recent trends, as it is rare in more modern settings for workers to hold significant amounts of the specific capital used by their employer.

2.4.4 Technological Change, Employment, and Occupational Wages

Taken together the evidence on occupational mobility and the average product of labor imply that increased adoption of new agricultural technology decreased demand for wage labor while simultaneously increasing the productivity of the remaining workforce. While these results indicate that some of the incumbent workers who remained likely benefited from technological change, at least in terms of their labor income, it does not say anything about the outcomes of workers who were pushed out of their initial occupation (*i.e.* displaced workers) due to the increased adoption of new technology within their region.

Table 2.6 evaluates the impact technological change had on incumbent workers' labor market outcomes, specifically the average wages in their subsequent occupations. Each column of Table 2.6 presents estimates from the same specification, as described in Section 2.3 using different sample criteria, while each panel represents a different outcome variable. Panel A evaluates the impact of additional technology on the likelihood a worker was employed in 1920. Column 1, which includes all incumbent agricultural workers, demonstrates that additional technological adoption in a worker's baseline county had no statistical or economically meaningful impact on the likelihood that worker remained employed in 1920. Column 2 restricts the sample in column 1 to only 1910 farmers (and their family members), and I see very similar estimates for this sub-sample. Column 3 presents estimates from the opposite sample to Column 2 (1910 wage-workers), and again finds no statistical relationship between technological adoption and future employment.

Given that technological adoption caused agricultural wage-workers to be much more likely to leave agriculture, column 4 of Table 2.6 restricts the wage-worker sample in column 3 to only those workers who left agriculture by 1920. This restriction changes the interpretation of the estimate slightly, in that while the IV assumptions prevent factors like workers' outside options from influencing the results it does not address selection in who the marginal worker who left due to technological change was compared to the average worker who left agriculture. For example, if all workers who left agriculture due to technological change (the compliers to the IV instrument) were excluded from the sample, the IV estimation assumptions necessitate that the estimate in column 4 should be 0. However I cannot observe deterministic displacement at an individual level, and so if the marginal worker who left agriculture due to technological change was lower ability than those who left for other reasons, they may have had worse post-agricultural outcomes. Acknowledging these limitations, the results from column 4 imply there was no impact on the likelihood a worker remained employed in 1920. However, unlike in columns 1 to 3, the estimate in column 4 is of a potentially economically relevant magnitude and would imply that a 1 standard deviation increase in technological adoption decreased the likelihood a worker left the labor market by $\frac{1}{3}$ of the sample average likelihood, though the estimate is very imprecisely estimated.

Panel B of Table 2.6 evaluates the impact of technological change on workers' subsequent labor returns. As discussed in Section 2.2.2, individual wage data is not available before the 1940 census, and so the dependent variable in Panel B is occupational wage, corresponding to the average wage in the 1940 census for working-aged men in the same occupation county pair for both non-agricultural and wage-earning agricultural workers. For farmers (and their family members), the average product of labor in agriculture is used as an analogous measure of the returns to labor.[11] Each of the regression models in Panel B includes the additional control "1940 rental price index" for a worker's 1910 county, measured as the average rent in a county (in the 1940 census) relative to the national average rental rate, in order to control for price levels across baseline counties.

Column 1 of Panel B demonstrates that overall, the returns to labor for incumbent workers significantly increased in regions with more technological adoption. This result should not be surprising given the evidence of increased average product of labor in agriculture due to technological change and the general momentum behind occupational decisions (31.7% of the baseline sample moved occupations over the 1910-1920 period). Columns 2 and 3 indicate that virtually all of this increase in returns to labor was realized by workers who were farmers in 1910. A test of the equality of the estimates in columns 2 and 3 using an interaction model rejects the hypothesis that technological adoption had the same impact on wage-workers and farmers ($p = 0.025$). This is consistent with the divergence in outcomes between incumbent wage-workers and farmers observed for occupational mobility, and further supports the conclusion that the benefits to technology largely flowed to a specific subgroup of workers, thereby increasing income inequality. The results in column 3 do not indicate that wage-workers in high-investment regions were worse off than wage-workers in low-investment regions.

Comparing the average 1920 occupational wages for incumbent wage-workers in Panel B columns 3 and 4 of Table 2.6, I find that wages were much higher among

[11] To stay consistent with the non-agricultural occupational wages, the average product of labor in agriculture is also measured using 1940 agricultural census, though results are robust to using 1920 agricultural census.

the group that transitioned out of agriculture compared to the unconditional group ($1,045 vs $776), indicating that the average move out of agriculture was associated with a significant increase in occupational wage. However, the estimate in column 4 indicates that, among wage-workers who transitioned out of agriculture, workers who left high-investment counties and therefore are more likely to have been displaced by technological change moved into occupations with lower average wages.[12] As discussed above, the interpretation of the estimates in column 4 as casual is slightly more nuanced, as the sample in column 4 is conditioned on an ex-post outcome, leaving agriculture. However, they do suggest that either i) displaced workers were very negatively selected or ii) displacement due to technological change had a negative impact on occupational wages. Section 2.5 presents evidence that observable measures of ability do not support the conclusion that displaced workers were negatively selected on observable measures of ability. However, this does not rule out selection along unobservable dimensions.

Though the estimate in column 4 Panel B of Table 2.6 is noisy and only marginally statistically significant, analysis of this same sample who are matched in the 1940 census and still employed finds a statistically significant effect of approximately 40 log points on their actual wages (p=0.093) as measured in 1940. While the 20 year difference in time-frame means the 1940 sample is much smaller, over a longer horizon, and selected by age, the fact that effects persist throughout a worker's life-span in a consistent manner and appear meaningfully in actual earnings data rather than occupational wage indexes strongly supports the conclusion that displaced workers had lower post-agricultural wages than would be expected given the average move

[12]The estimate in Table 2.6 Panel B column 4 has the added benefit of being derived from only the average occupational wage measure as opposed to a combination of average product of labor and average wages within occupation.

out of agricultural wage work.

Overall, technological change appears to have had no impact on workers' subsequent likelihood of remaining employed. While the results for wages are both noisy and come with some non-trivial measurement issues, they tell a consistent story that one subset of workers likely benefited, while another likely experienced worse than expected labor returns after being displaced from their baseline occupation. The model described below provides an alternative way to measure the impact of technological change on workers' labor earnings which bypasses some of the limitations of the empirical analyses by using moments pertaining to actual wages along the economies' transition path, rather than the occupational wages utilized for this section.

2.4.5 Family Structure and Migration

While historical census data is relatively sparse by modern standards, there are a number of potentially interesting outcomes to evaluate in the context of technological change and incumbent workers, particularly in regards to family structure and migration. Appendix Table A.1 presents results for several of these outcomes separately for incumbent wage-workers and farmers. No statistically significant or economically meaningful impact on county or state migration is estimated (columns 1 and 2). This is a particularly interesting result as it suggests that the occupational mobility caused by technological change occurred within county, rather than through migration to larger metropolitan areas.

Columns 3, 4, and 5 evaluate the impact of technological adoption on family structure, in particular the likelihood a worker had a new child between 1910 and 1920, the likelihood an unmarried worker was married in 1920, and the likelihood a married worker was divorced in 1920. Technological adoption is estimated to have decreased the likelihood an incumbent farmer had a new child by 7.2% (on a base of

64%) per standard deviation, while the point estimate for farmers is approximately half as large and not statistically significant. Larger families would provide additional labor to work a family farm, and therefore a decline in fertility within farmers further supports the finding that technology impacted the demand for labor. Technological change is estimated to have had effectively zero impact on the marriage and divorce rates of farmers, while a one standard deviation increase in technological adoption is estimated to have increased the likelihood of divorce by 2.1% (from a very low base of 1%). Although this estimate is only marginally statistically significant at p=0.114, it is consistent with the occupational wages evidence that the occupational transitions among wage-workers caused by technological adoption may have been significantly stressful or costly.

2.4.6 Second Generation

Connections between parental income, occupation, and mobility and children's subsequent outcomes have long been established. These channels maybe even more relevant in a setting where a parent is displaced from their occupation due to a persistent shock in the demand for their occupation-specific skills, as appears to be the case for technological change in 20th century agriculture. This section directly evaluates the extent to which a father's experience of technological change had a causal impact on his son's future labor market outcomes. The specific methodology used to evaluate the impact of technological change on children closely mirrors the first generation analysis above. In particular, an equivalent IV model is used on a sample of agricultural workers' sons aged 12 or younger in 1910 who lived with their father and who are linked to a 1940 census record. Under equivalent identifying assumptions, the estimated coefficient on instrumented agricultural investment between 1910 and 1920 provides a consistent estimate of the impact of a father's increased exposure to technological change on his son's prime-age labor market outcomes.

A connection between technological change and children's outcomes could exist for a number of reasons; children of "winners" from technological change could experience more material investment from their parents and thus have higher unobservable skills. Conversely, children of displaced workers may be better positioned to learn skills in growing industries that provide higher returns in adulthood or may move to higher quality urban schools. Alternatively, as evidence suggests that displaced workers moved into low paid occupations, and there is a strong correlation between the occupations of fathers and sons, sons of displaced workers may be more likely to inherit skills that have a low overall return.

Table 2.7 presents IV estimation results for outcomes measured in 1940 within a sample of men who are matched to a 1910 census observation and whose father worked in an agricultural occupation in 1910. Panel A presents results separately for children whose father was an agricultural wage worker (columns 1-3) and whose father was a farmer (columns 4-6). Columns 1 and 4 demonstrate that the degree of technological change experienced by a father between 1910 and 1920 had no significant impact on a child's likelihood of employment for both children of wage-workers and farmers. Columns 2 and 4 demonstrate that, conditional on being employed, a father's experience of agricultural technological change had no impact on the likelihood a child worked in agriculture. Columns 4 and 6 evaluate impacts to children's labor income, which is known at the individual level for employed, non-agricultural workers in the 1940 census. Given that technological change had no statistical impact on the extensive margin of employment or the likelihood of working in agriculture, estimates on wages are unlikely to be biased by selection. Here I find that technological change had no statistical impact on wages for children of wage-workers. However, a one standard deviation increase in the degree of technological change experienced by a father working as a farmer is estimated to have increased his child's wages by 15 log points. This evidence is consistent with the estimates on the first-generation, in that this same group of fathers also experienced large increases in their average product of labor in agriculture, and thus the children of the "winners" from technological change also seem to have benefited.

Panel B of Table 2.7 mirrors the analysis in Panel A but splits the sample by the occupational mobility of fathers, rather than fathers' baseline occupation. As discussed in Section 2.4.4, by conditioning on an ex-post outcome, the interpretation of estimates in Table 2.7 are more nuanced, and do not rule out selection on parental ability. However, there is no evidence of this selection along observable measures in the data. As in Panel A, I find no statistical impact on children's employment or likelihood of working in agriculture for both children of fathers who left agriculture (columns 1 and 2) and children of fathers who remained in agriculture by 1920 (columns 4 and 5). Column 3 indicates that the children of fathers who experienced a higher degree of technological change and who left agriculture, *i.e.* were more likely to have been displaced by technological change, had much lower wages in adulthood compared to children whose father also left agriculture, but experienced a lesser degree of technological change.[13] Finally, Panel B column 6 indicates that, like children of fathers who were farmers (Panel A), children of fathers who remained in agriculture from high-investment regions had higher wages in adulthood, though the estimate is much less precise for the sub-sample in Panel B. Taken together Table 2.7 indicates

[13]One caveat to this result is that columns 1 and 2 of Table 2.7 Panel B, although very imprecisely estimated, indicate that the children of fathers who left agriculture from higher-investment regions may have had higher employment and been less likely to work in agriculture compared those from lower-investment regions. This implies that the high-investment sample may have had a degree of negative selection. However, even if the point estimates in columns 1 and 2 are taken as precisely estimated, their magnitude is unlikely to be able to explain the large point estimate in column 3.

that the degree to which a father experienced technological change had a significant impact on his son's labor earnings later some 20 years later and that in particular, the sons of "winners" from technological change had higher earnings than their peers, while the sons of workers who were more likely to have been displaced by technological change had lower labor earnings than their peers.

So far the empirical analysis has focused on average effects. However, these averages may be masking interesting distributional shifts. For example, it is possible that having a father who was displaced from an agricultural occupation due to technological change positioned some subset of children to benefit through any of the channels previously suggested, even if this was not the representative experience. Further, given the results indicating that technological adoption created both winners and losers, evaluating the distributional shifts in wages provides an analytical way to document increases in wage inequality. Figure 2.5 plots kernel density estimates for the residualized wages of children of agricultural workers separately for the upper and lower terciles of terrain ruggedness in the father's baseline (1910) county. For these plots, wages are residualized against the same controls as in Table 2.7 in order to mimic the IV estimation strategy in a reduced form analysis. Panel A includes a sample of all matched sons who are employed in a non-agricultural occupation in 1940, while Panels B and C present distributions for the subsample of sons of farmers and sons of fathers who left agriculture by 1920 respectively. All three panels indicate that the children of fathers who worked in regions with smoother terrain (*i.e.* who were more likely to have experienced more intense technological change) had a higher variation in wages. This difference is particularly stark within the sample whose fathers left agriculture by 1920 (Panel C). Here I find a clear flattening of the wage distribution for children of fathers from less-rugged regions with additional mass at both the left and right tails, which indicates that although the average effect of technological change on the children of workers who left agriculture was negative, it was beneficial for some subset.

2.5 Discussion of Empirical Results

The previous section outlined the main empirical results of the paper and found i) increased technological adoption in early 20th century agriculture caused some incumbent wage-workers to leave agriculture ii) owner/operators and their families likely saw an increase in their average product of labor due to technological change, iii) displaced wage-workers transitioned into lower paying occupations than would be expected for the typical move out of agriculture, iv) technological change impacted the labor earnings of children of incumbent workers in a direction which mirrored the impact on their fathers. This section will evaluate the extent to which these results and supporting evidence can be synthesized into a consistent framework that highlights specific causal mechanisms and which will provide a foundation for an equilibrium sorting model.

As discussed above, the IV estimation strategy isolates any impact on occupational mobility to a demand channel (rather than a supply effect). Then, the reduction in labor demand, concentration of mobility effects (displacement) within wage-workers, increased average product of labor, and decline in fertility among farmers are all broadly consistent with new agricultural technologies improving agricultural workers' productivity, thus reducing labor demand (result i) and increasing the labor returns of the residual workforce (result ii). Assuming that both of these results are driven by a single technological shock, they jointly imply that agricultural wage-workers and farmers provide different types of labor inputs into agricultural production, otherwise a demand shock for a single input would imply the opposite effect on the equilibrium price of that input.[14]

While the preponderance of evidence provides a relatively straightforward explanation of how a single shock to labor productivity can explain results i and ii, it is not ex-ante obvious what mechanisms underpin results iii and iv. On one hand, it is possible that job separations due to technological change are different than the typical occupational moves due to factors such as search frictions causing decreased match quality or a relatively inelastic demand function for non-agricultural low-skilled labor. On the other hand, it is possible that displaced workers are negatively selected in terms of ability or skill relative to non-displaced workers who also left agriculture. While far from conclusive, additional evidence to evaluate these mechanisms is consistent with a story of inelastic local demand and inconsistent with selection on observable measures of ability.

Table A.2 provides additional analysis to evaluate why wage-workers who left agriculture from higher-investment regions ended up in lower paid occupations than those who left from lower-investment regions. First, given the lack of evidence that displaced workers were more likely to migrate, it is possible that local labor markets were relatively inelastic, and so an increase in the low-skilled non-agricultural workforce depressed wages. As discussed in Section 2.2.2, the measure of occupational wages is taken at the occupation/county level, and so would capture differences in wages due to a persistent increase in the supply of low-skilled labor.[15] If inelastic labor demand is a first-order mechanism, one would have expected estimates to be significantly reduced in magnitude when occupational wages were measured nationally

[14]It's also worth noting that the fact labor-augmenting technological change reduced labor demand implies that within a constant elasticity of substitution model of aggregate agricultural production, the elasticity of substation between labor and other some other input (*e.g.* land) must be greater than 1, *i.e.* the two inputs are gross complements.

[15]Assuming these effects persisted until wages are measured in 1940

rather than locally. Table A.2 column 1 presents these estimates in a sample comparable to that in Table 2.6 Panel B column 4. We see that using national occupational wages reduces the point estimate by nearly 75%, which is consistent with differences in occupational wages across regions due to the increased supply of low-skilled labor driving the estimate in Table 2.6.

Columns 2-4 in Table A.2 evaluate the extent to which wage-workers who left agriculture from higher-investment regions were selected on observables relative to those who left from lower-investment regions. Column 2 provides the strongest evidence and finds no differences in 1910 literacy rates across workers from different investment levels. This would suggest that workers were not selected on literacy, the only directly observable measure of ability in the 1910 census. Columns 2 and 3 supplement this analysis by evaluating the extent to which wage-workers who left agriculture from high-investment regions were more likely to have moved into occupations with higher levels of education in the 1940 census, the first year with observable education data. Again I fail to reject a null effect, implying that workers who left agricultural from higher-investment regions moved into occupations with similar educational requirements as workers who left from lower-investment regions. Taken together, columns 2-4 in Table A.2 imply that displaced workers were unlikely to have been selected based on observable measures of ability or move into occupations that had higher education levels, although there are many dimensions of unobservable ability which may play a relevant role in explaining the impact of technological change on workers. In fact, selection based on relative ability is an explicit feature of the model described below.

Finally, technological change may have intergenerational effects through a number of potential mechanisms. Broadly, these mechanisms can be split into two groups; mechanisms that shift children's skills such as differential parental investment and non-skill-based mechanisms such as changes to parental social networks or a similar inelastic labor demand channel discussed for the first generation. In terms of this latter possibility, estimates of the impact of investment on the wages of non-agricultural laborers' children do not show any effect of investment on wages, suggesting that differences in the local supply of low-skilled labor is unlikely to explain why children of workers who left agriculture from high-investment regions had lower wages.[16]

Table A.3 presents evidence that this first set of mechanisms may be relevant in explaining why the children of technological winners and losers themselves won or lost. Specifically, Columns 1-3 display the effect of a father's increased exposure to technological change on three measures of a child's ability measured in 1940 for the sample of children of fathers who had left agriculture by 1920. Columns 4-6 display estimates for the same outcomes but within the sample of children of fathers who were farmers in 1910. While none of the estimates in columns 1-3 are statistically significant, the point estimates imply economically relevant differences in the years of schooling and the likelihood a child graduated 8th grade between children of fathers who left agriculture from high-investment regions and children of fathers who left agriculture from low-investment regions. Columns 4-6 provide slightly more precise estimates, given the much larger sample, and imply that a 1 standard deviation increase in investment in a farmer's agricultural region increased the educational attainment of his son by more than 0.5 years and increased the likelihood his son worked in an occupation with higher education requirements. Overall, while estimates of the impact of technological change on children's skills are imprecise, they suggest that skill accumulation may be an important channel to explain why the impacts of technological change carried through into a second generation.

[16]The children of laborers had slightly lower average wages as compared to the average child of a worker who left agriculture.

2.6 Model

The previous empirical sections established that quasi-random increases in the rate of investment in new technology between 1910 and 1920 increased incumbent agricultural workers' occupational mobility and had lasting impacts on their occupational earnings. Evidence also indicates that the impacts of additional exposure to new technologies had implications for inter-generational mobility. However, the widespread adoption of improved agricultural technologies likely also conferred aggregate benefits to society in the form of increased productivity. How large are the estimated costs to individual workers compared to the aggregate benefits to society? Over a sufficiently long horizon, the answer is likely small, as the majority of costs are likely to accrue over a finite transition period, while all future generations benefit from a productivity boost. However, from the perspective of the baseline population, the relative magnitude of labor reallocation costs relative to the aggregate benefits is less clear. Further, the evidence presented above on intergenerational mobility suggests that technological change may have had long-lasting impacts on the distribution of labor earnings.

This section presents a dynamic occupational sorting model inspired by (Dvorkin and Monge-Naranjo, 2019) and (Garcıa-Couto, 2021) which provides a framework through which the costs and benefits of technological change documented above can be consistently compared. The key components of the model are endogenous occupational selection based on comparative advantage, inter-generational skill transfer, and a one time shock to production technology calibrated to match the empirical evidence on occupational mobility. The resulting model allows for the simulation of welfare of identical agents under two versions of the world, one which experiences the initial shock to the agricultural technology and one which remains on the pre-shock balanced growth path.

2.6.1 Firm Problem

Time is discrete and runs forever. Competitive firms hire labor services from each of four occupation categories to produce a homogeneous final consumption good using a nested constant elasticity of substation production technology. The four occupation categories are high-skill non-agricultural (h), low-skill non-agricultural (l), agricultural laborer (w), and farmer (f). Where agricultural laborers correspond to agricultural wage-workers as in the empirical section, and farmers include both farmers and their unpaid family members who work in agriculture. Workers in each of the two agricultural occupations (w and f) are employed in one of two geographic regions which are ex-ante identical, but which may differ in production technology after the technological shock occurs. The agricultural regions will be indexed as high-tech (T) and low-tech (M). Therefore, firms make hiring decisions over six types of labor (h, l, Tf, Mf, Tw, Mw). This occupational structure allows for comparisons to be made between the occupational mobility of wage-workers in high-investment and low-investment regions, which will be used to discipline the magnitude of technological shock, as well as provide a more realistic outside option for agricultural workers to transition into. For ease of exposition, these types will be referred to simply as occupations for the remainder of the paper, acknowledging that they may jointly describe both the location and occupation of a given worker.

For an allocation of aggregate human capital input:

$$\mathbf{H} = [H^h, H^l, H^{Tf}, H^{Mf}, H^{Tw}, H^{Mw}]$$

firms' are able to produce $Y(\mathbf{H})$ units of the final good using the production technology:

$$Y(\mathbf{H}) = \left[\left(A^h H^h\right)^{\frac{\sigma-1}{\sigma}} + \left(A^l H^l\right)^{\frac{\sigma-1}{\sigma}} + F(H^{Tf}, H^{Mf}, H^{Tw}, H^{Mw})^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{2.3}$$

- where σ is the elasticity of substation between high-skilled, low-skilled, and agricultural occupations and A^j parameterizes the productivity of labor from occupation j. The function $F(\cdot)$ describes the technology with which firms combine agricultural labor across both geography and occupation. Specifically, they use a two-step CES aggregator with an elasticity of substitution across regions ρ and elasticity of substitution across occupations ϵ:

$$F(H^{Tf}, H^{Mf}, H^{Tw}, H^{Mw}) = \left(f^T(H^{Tf}, H^{Tw})^{\frac{\rho-1}{\rho}} + f^M(H^{Mf}, H^{Mw})^{\frac{\rho-1}{\rho}}\right)^{\frac{\rho}{\rho-1}} \tag{2.4}$$

$$f^i(H^{if}, H^{iw}) = \left(\left[A^{if} H^{if}\right]^{\frac{\varepsilon-1}{\varepsilon}} + \left[A^{iw} H^{iw}\right]^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}} \qquad for\ \ i \in T, M \tag{2.5}$$

In practice, the outer elasticity (σ) will be calibrated such that high-skilled, low-skilled, and agricultural labor are complementary inputs ($\sigma < 1$), and the elasticities of substitution between agricultural regions (ρ) and agricultural occupations (ε) will be jointed calibrated with the other estimated parameters.

Each period, firms take prices of inputs (wages) as given and decide on the optimal levels of each type of labor to hire in order to maximize profits. Therefore, the constant returns to scale production technology, along with free entry and exit results in a zero profit condition in equilibrium.

2.6.2 Worker Problem

The workers' side of the model is largely inspired by Dvorkin and Monge-Naranjo, 2019 and closely follows Garcıa-Couto, 2021. There are measure one workers in each period, with a random share δ dying and being replaced by new young workers. Workers have a per-period constant relative risk aversion (CRRA) utility specification with coefficient risk aversion γ and discount future periods with discount factor β.

Workers enter each period attached to a specific occupation/location j and with h units of human capital. At the start of the period, workers draw a vector of idiosyncratic productivity shocks $\epsilon = [\epsilon_h, \epsilon_l, \epsilon_{Tf}, \epsilon_{Mf}, \epsilon_{Tw}, \epsilon_{Mw}]$ which determine their comparative advantage for production. Next, they supply $h\epsilon_j$ units of human capital to their employer and then immediately consume their labor earnings $w_j h\epsilon_j$, where w_j is the wage in occupation j. Workers then make an occupation decision for the next period based on their productivity vector ϵ. Finally, based on their occupational decision (ℓ), worker's human capital evolves according to:

$$h' = h \times \epsilon_\ell \times \tau_{j,\ell} \tag{2.6}$$

- where $\tau_{j,\ell} \in (0, 1]$ captures the cost to human capital of transitioning from occupation j to occupation ℓ.

The Bellman equation for the worker is then:

$$V(j, h, \epsilon) = \frac{(w_j h\epsilon_j)^{1-\gamma}}{1-\gamma} + \beta\,(1-\delta) \max_{\ell \in [h, l, Tf, Mf, Tw, Mw]} (\chi_{j,\ell}\, \mathbb{E}_{\epsilon'}[V(\ell, h', \epsilon')]) \tag{2.7}$$

- where $\chi_{j,\ell}$ captures the non-pecuniary costs of occupational transition such as barriers to geographic mobility or tastes for different occupations. As pointed out by Dvorkin and Monge-Naranjo, 2019, equation 2.7 can be factorized if ϵ is drawn from a distribution satisfying a boundedness condition, resulting in the following characterization of the Bellman:

$$v(j, \epsilon) = \frac{(w_j \epsilon_j)^{1-\gamma}}{1-\gamma} + \beta\,(1-\delta) \max_{\ell \in [h, l, Tf, Mf, Tw, Mw]} (\chi_{j,\ell}\, \mathbb{E}_{\epsilon'}[v(\ell, \epsilon')](\tau_{j,\ell}\epsilon_\ell)^{1-\gamma}) \tag{2.8}$$

This formulation of the Bellman equation greatly simplifies the model, as it eliminates the need to know the distribution of human capital within an occupation when solving for the equilibrium prices and quantities. Under the specification above, Dvorkin and Monge-Naranjo, 2019 Theorem 1 proves that for $w \in \mathbb{R}^j_+$ a unique finite set of v^j exists which solve:

$$v^j = \mathbb{E}_\epsilon[v(j, \epsilon)] \tag{2.9}$$

Further, the proportion of workers switching from occupation j to occupation ℓ at the end of the period is:

$$\mu(j, \ell) = \frac{[\lambda_\ell \tau_{j,\ell}(\chi_{j,\ell} v^\ell)^{\frac{1}{1-\gamma}}]^\alpha}{\sum_{k\in[h,l,Tf,Mf,Tw,Mw]}[\lambda_k \tau_{j,k}(\chi_{j,k} v^k)^{\frac{1}{1-\gamma}}]^\alpha} \tag{2.10}$$

- where α and λ_j are the shape and scale parameters for the fréchet distribution of ϵ_j.

2.6.3 Young Generation

The young generation in each period are born into the model at age 17 to a random father with likelihood of childbirth equal to $\frac{\delta}{1-\delta}$ (deceased agents cannot produce children, leaving 1-δ agents to supply δ children to remain a constant population). Children are born at the end of a model period and will begin productive work in the next period, inheriting their father's current occupation.[17] A child's human capital is based on their father's human capital and income (consumption) through the following human capital production function:

$$h_{child} = h_0 \left((1-\kappa)\bar{h} + \kappa \left[\nu h_{father}{}^{\frac{\phi-1}{\phi}} + (1-\nu) I_{father}{}^{\frac{\phi-1}{\phi}} \right]^{\frac{\phi}{\phi-1}} \right) \tag{2.11}$$

[17]This is a simplifying assumption that is not central to the model results.

- where $h_0 \in (0,1)$ parameterizes the difference in human capital between children and adults, $\bar{h}$ parameterizes the minimum level of human capital for children, κ parameterizes the overall importance of parental inputs, ν parameterizes the relative importance of parental human capital vs parental income, and ϕ represents the elasticity of substitution between parental income (I_{father}) and parental human capital (h_{father}).

For the remainder the paper, the human capital production technology will be assumed to be Cobb-Douglas ($\phi = 1$). Then, substituting the value of a father's income into equation 2.11 ($I_{father} = \epsilon_{father\,occ}\ w_{father\,occ}\ h_{father}$) and rearranging terms, I find:

$$h_{child} = h_0 \left((1-\kappa)\bar{h} + \kappa h_{father}\left[\epsilon_{father\,occ} w_{father\,occ}\right]^{1-\nu}\right) \tag{2.12}$$

Therefore, each period the new generation provides aggregate human capital to occupation j equal to:

$$\begin{aligned} H_{children,j} &= \delta\theta_j h_0 \left((1-\kappa)\bar{h} + \kappa\,\mathbb{E}\left[\epsilon_j^{1-\nu}\right]\mathbb{E}\left[h_{father}\right] w_j^{1-\nu}\right) && (2.13)\\ &= \delta\theta_j h_0 \left((1-\kappa)\bar{h} + \kappa\lambda_j^{1-\nu}\Gamma\left(1 - \frac{1-\nu}{\alpha}\right)\frac{H_j}{\theta_j} w_j^{1-\nu}\right) && (2.14)\end{aligned}$$

- where $\Gamma(\cdot)$ represents the gamma function.

2.6.4 Equilibrium

Let θ_j be the share of workers in occupation j, and let θ^0 and H^0 be the vectors of occupational distribution and human capital for the workers born in this period. Then, Dvorkin and Monge-Naranjo, 2019 Proposition 1 proves that there is a unique invariant distribution of all workers:

$$\theta = \delta\theta^0[I - (1-\delta)\mu]^{-1} \tag{2.15}$$

And, if it exists, a unique stationary stock of human capital:

$$H = \delta H^0[I - (1-\delta)M]^{-1} \tag{2.16}$$

Where M represents a matrix of aggregate human capital transitions of which each element is equal to:

$$M_{j,\ell} = \Gamma(1 - \frac{1}{\alpha})\tau_{j,\ell}\lambda_\ell[\mu(j\ell)]^{1-\frac{1}{\alpha}} \tag{2.17}$$

Note that $\Gamma(\cdot)$ represents the gamma function. If a stationary stock of human capital does not exist, then a unique balanced growth path does, which converges to stable ratios $\frac{H_j}{H_i}$

Then, given an initial population of workers, a competitive equilibrium in this model consists of a set of wages w_j, workers' occupational decisions, and aggregate demand for human capital H_j such that workers' optimal decisions are consistent with v^j and μ and labor markets clear.

2.7 Model Calibration

In order to discipline the model, the model's occupations and regions must be mapped into actual data. The definitions of farmers and laborers follow directly from the IPUMS harmonized 1950 occupation codes used in the empirical results above. Further, high-skill and low-skill occupations are defined using the share of workers in an occupation with at least a high-school degree in the 1940 census, with high-skilled occupations being those with a share greater than the occupation of the median non-agricultural worker.

The assignment of counties into the high-investment or low-investment region follows a cutoff rule along equipment investment rates between 1910 and 1920. Because the calibration assumes the low-tech region does not undergo technological change, the determination of the cutoff fixes the model results to be relative to the high-investment counties getting the same amount of investment as the low-investment counties. In other words, some stance must be taken about what constitutes ”zero technological change.” In this sense, the model may underestimate the impacts of technological change, to the extent that the low-investment regions are still undergoing a meaningful degree of technological adoption. The cutoff value was determined using the McKibben *et al.*, 1939 farm survey data. Specifically, the cutoff was selected as the value which maximized the difference in average increases in productivity (declines in hours required to work one acre) between high-investment and low-investment counties in the survey sample.[18]

To reduce the number of parameters to be jointly calibrated, several normalizations will be made. First, the scale of the entire economy is fixed by setting $A^h{=}1$. Next, the non-pecuniary costs of occupational transition χ are normalized to be relative to leaving an occupation, such that $\chi_{j,\ell}{=}1$ if $j{\neq}\ell$ and $\chi_{j,\ell}{\in}(0,1)$ if $j{=}\ell$.[19] Next, the human capital transition parameters τ are normalized such that remaining in an occupation does not incur a human capital penalty, *i.e.* $\tau_{j,\ell}{\in}(0,1)$ if $j{\neq}\ell$ and $\tau_{j,\ell}{=}1$ if $j{=}\ell$. Third, the high-tech and low-tech sectors will be assumed to be ex-ante identical, such that all parameters indexed by Ti for $i{\in}[w,f]$ will be equal to the parameter for Mi, except of course the post-shock productivity parameter $\widehat{A^{Tw}}$, where the hat

[18] This analysis was carried out at the county implement level as is Figure 2.4, and only cutoffs with more than 30% of the sample on either side were considered. Because the median investment level was included in the range of possible values, it was selected to simplify bring the model to data. See Figure A.2 for a visualization of the selection of the cutoff value.

[19] χ less than one increases lifetime utility if $\gamma{>}1$, as will be the case in this paper.

notation differentiates the pre-shock parameter from the post-shock parameter. Several other assumptions about the transfer of human capital are made to reduce the number of parameters needed to construct τ, see Appendix Table A.4 for details.

Several parameters were set outside of the joint calibration procedure and are presented in Table 2.8. The remaining 24 model parameters were jointly calibrated to match data moments, including the estimated difference in the likelihood that an agricultural laborer had left agriculture by 1920 between high-tech and low-tech regions, using the same IV estimation strategy as in the empirical section.[20] To avoid attributing occupational wages to actual labor earnings, only wage moments in 1940 are targeted. Importantly, this required the transition path of the economy to be calculated within the generalized method of moments procedure. The estimation procedure required first solving the initial 1910 steady state and the second, post-shock steady state (using $\widehat{A^{Tw}}$ rather than A^{Tw}), and finally solving the path of wages and value functions between these two steady states in order to match moments along this transition path.

While the joint calibration does not directly attribute any one moment to a single parameter, Table 2.9 presents the calibration targets, the parameter most closely associated with each moment, and the corresponding model values. Baseline occupational employment shares were used to discipline the sectoral productivity parameters (A^j). The Fréchet scale parameters (λ_j) were set to match the relative wages of young workers (18 to 25 year olds in the data) to older workers (26 to 65 year olds) for high-skilled, low-skilled, and agricultural laborers. Because farmer wages are not observable, the ratio of farmers' average product of labor and average high-skilled wages

[20] In order to match the discretized technology adoption in the model, a binary indicator for technology adoption is used in these regressions, rather than the continuous measure used in the empirical section.

was targeted as an alternative. The likelihood that workers remained in their respective occupations between 1900 and 1910 was used to pin down the non-pecuniary costs of occupational transition (χ). The human capital costs of occupational transition (τ) were disciplined using the average 1940 wages of workers who moved from h to l occupations between 1930 and 1940 as compared to the average wages of workers who were working in an h occupation in both 1930 and 1940. Additionally targeted are the equivalent moments for low-skilled workers who moved into a high-skilled occupation and agricultural laborers who became farmers(or vice versa). The likelihoods of transitioning into and out of agricultural occupations between 1900 and 1910 were targeted to pin down the τ values associated with agricultural occupations. The elasticities of substitution within the agricultural production technology between regions and between occupations were disciplined to match the relative employment in high-tech and low-tech regions in 1920 and the IV estimated difference in attrition for farmers respectively. Finally, the key children's human capital production technology parameters (κ and ν) are disciplined to match the relative earnings of sons from high-skill fathers and low-skill fathers in 1940 and the difference in earnings for children from agricultural fathers in high-tech vs low-tech agricultural regions estimated using the instrumented investment level. The calibrated parameter values can be found in table 2.10.

2.8 Model Results

To evaluate the impact of technological change within the model, 30,000 individual workers were simulated twice, using the same sequence of productivity shocks, once along the transition path due to the technological shock and once while remaining in the initial steady state equilibrium. Simulated workers are required to be less than 55 years at baseline and are simulated forward for 45 years. The utility of each agent in each period is then constructed for both economies and the consumption equivalent value of technological change (CE) is calculated for each agent according to the following equality:

$$\frac{\left(c_0^1 + CE\right)^{1-\gamma}}{1-\gamma} + \sum_{t=1}^{54} \left(\beta\left(1-\delta\right)\right)^t U_t^1 \prod_{s=1}^{t} \chi_{o_{s-1}^1, o_s^1} \tag{2.18}$$

$$= U_0^0 + \sum_{t=1}^{54} \left(\beta\left(1-\delta\right)\right)^t U_t^0 \prod_{s=1}^{t} \chi_{o_{s-1}^0, o_s^0} \tag{2.19}$$

- where period 0 is the first post-shock year, k is an indicator for the economy with technological change, c_t^k is consumption in period t economy k, o_t^k is the occupation in period t economy k, and U_t^k is the utility in period t economy k.

Based on these consumption-equivalent costs of technological change, 83.5% of agents alive at baseline had higher lifetime utility in the economy which transitions to a new steady state due to the technological shock compared to remaining in the initial steady state (*i.e.* had a negative CE value). However, the adult agents made worse off by technological change were disproportionately affected, and so the gross consumption equivalent costs of technological change are 18% of the gross surplus generated among all adult agents alive at baseline.

Reinforcing the empirical findings, the calibrated model indicates that inter-general mobility is a relevant margin to consider when evaluating the welfare implications of technological change in early 20th century American agriculture. We find that, through the model's lens, the relative welfare loss generated by the technological shock is significantly lower at 10.9% of total benefits if the welfare of children alive at baseline is also considered. The reduced welfare cost is driven by the fact that children may enter the labor market after a significant amount of the occupational displacement has already occurred, and thus can reap the benefits of increased productivity without having to undergo the costly occupational transitions of their fathers. The model also indicates that the channel of intergenerational mobility is also economically relevant when accounting for the total loss of surplus generated by the technological shock, in that 19.5% of all of the welfare losses due to technological change were born by children. Together, these results indicate that, in the model's perspective, the causal link between technological change and intergenerational mobility identified in the empirical section is economically relevant when evaluating the welfare implications of this specific technological change event.

Technological change can impact welfare through two broad ways in the model. First, holding occupational decisions constant, it has a direct effect on wages. Second, because workers endogenously respond to these changes in wages through occupational mobility, it can both i) reduce human capital by incentivizing workers in some occupations to accept lower value moves in terms of the net effect to their human capital from ϵ_ℓ and $\tau_{j,\ell}$ and ii) reduce welfare directly through increased accumulation of χ transition costs. The χ parameters capture non-pecuniary occupational attachment such as the cost of geographic mobility, the value of local amenities, the cost of breaking existing social networks and establishing new ones.

Overall, the model results highlight how a technology can both "expand the pie" for everyone and have significant adverse effects on a relatively small proportion of the population. Future iterations of the model will explicitly discipline intergenerational mobility to empirical estimates and evaluate its importance in weighing the costs and benefits of technological change.

2.9 Conclusion

This paper studies the mechanization of early 20th century American agriculture and finds that quasi-random variation in the degree of technological adoption had persistent effects on incumbent workers' occupational mobility and occupational wages.

Further, the children of these incumbents had divergent outcomes based on how well their father was positioned to take advantage of the technological change. These results represent some of the only evidence about the long-term impacts of large-scale technological change on incumbent workers and document that impacts are not isolated to only the first-generation who directly experiences the change but spread into a second-generation who were children at the time of technological adoption. A dynamic occupational sorting model disciplined to match the empirical results for the first generation indicates that technological change reduced lifetime welfare for 16% of workers, and the total consumption equivalent cost to these individuals was 11% of the surplus created for agents alive at baseline. Further, the model highlights the relevance of the inter-generational channel identified in the empirical analysis when evaluating the welfare implications of this technological shock.

2.10 Tables

Table 2.1: Estimates of Equipment Investment and Non-Agricultural Occupational Mobility

	White Collar Workers: In Same Occ. '10-'20	Non-Ag. Laborers: In Same Occ. '10-'20	Non-Mfg. Laborers: In Mfg. '20
	IV	IV	IV
	(1)	(2)	(3)
Pct. Δ Equip. Value per Acre '10-'20 (SD)	0.0130	0.0142	-0.0208
	(0.0317)	(0.0382)	(0.0339)
Controls			
State and Age FEs	x	x	x
Standard Soil and Topo	x	x	x
Obs.	91,923	57,537	39,994
Y mean	0.310	0.168	0.164
Montiel Olea-Plueger F-stat	20.1	12.3	16.5

***:p<0.01, **: p<0.05, *:p<0.10, +:p<0.15

Notes: Table displays regression coefficients derived from separate IV regressions of the outcome/sample labeled at the top of each column onto the change in equipment value per acre (coefficient displayed), origin (1910) state fixed effects, age fixed effects, and the standard soil/topological controls (900m ruggedness, fertility index, usable water depth, root depth, average number warm days, and erosion tolerance). Standard errors in parenthesis clustered at the 1910 county level.

Table 2.2: Observable Characteristics of Sample in 1910

	Full (1)	Ag. Workers (2)	Potential Analysis (3)	Matched Analysis (4)	p-val (3)=(4)
Obs.	27,472,437	7,656,650	5,230,240	319,817	
Age	32.3	32.7	36.6	36.3	0.000
Male	1.000	1.000	1.000	1.000	.
Head of household	0.523	0.619	0.780	0.791	0.000
Race: white	0.893	0.835	0.835	0.909	0.000
Race: African American	0.097	0.158	0.159	0.088	0.000
Race: other	0.010	0.007	0.006	0.003	0.000
Married	0.547	0.597	0.750	0.767	0.000
Family size	3.2	3.8	3.7	3.8	0.000
Number children	1.173	1.648	2.138	2.139	0.870
Rural	0.509	0.966	0.981	0.986	0.000
Own home	0.414	0.558	0.552	0.607	0.000
In labor force	0.934	0.998	0.999	0.999	0.000
Literate	0.920	0.881	0.875	0.926	0.000

Notes: Table displays average levels of observable characteristics for several different samples derived from the 1910 decennial census. Agricultural workers are defined as anyone in the following occupations: farmer, farm manager, farm foremen, farm laborers (wage and unpaid family workers), and farm service laborers. The potential analysis column restricts the sample to those at least 22 years old, in the Continental United States, and not in a predominately urban county (defined as counties below the 5th percentile for agricultural labor share). Column 4 further restricts the sample in column 3 by conditioning on being matched in the Ran *et al.*, 2020 exact linkages between 1910 and 1920.

Table 2.3: Correlation of Instrument and Farm Wages Over Time

	Log Average Product of Labor in...					
	1900	1920	1900	1920	1900	1920
	(1)	(2)	(3)	(4)	(5)	(6)
10m Ruggedness (SD)	-0.0084	-0.0116	-0.0024	-0.0295**	0.0015	-0.0249*
	(0.0092)	(0.0087)	(0.0120)	(0.0120)	(0.0130)	(0.0130)
900m Ruggedness (SD)			-0.0104	0.0222+	-0.0028	0.0165
			(0.0124)	(0.0142)	(0.0130)	(0.0149)
Fertility Index (SD)			0.0158	0.0316***	0.0179+	0.0296***
			(0.0113)	(0.0100)	(0.0115)	(0.0101)
Water Zone (SD)			-0.0149	-0.0156	-0.0138	-0.0228*
			(0.0106)	(0.0109)	(0.0137)	(0.0134)
Root Zone (SD)			-0.0057	0.0186*	-0.0024	0.0245**
			(0.0098)	(0.0095)	(0.0100)	(0.0099)
Warm Days (SD)			0.0043	-0.0163*	0.0069	-0.0044
			(0.0098)	(0.0092)	(0.0118)	(0.0111)
Erosion Tolerance (SD)			-0.0028	-0.0028	-0.0084	-0.0126
			(0.0087)	(0.0083)	(0.0115)	(0.0105)
P-val: Rough10m Equal	(1)=(2) 0.779		(3)=(4) 0.041		(5)=(6) 0.059	
Controls						
State FEs	x	x	x	x	x	x
12 Soil Order Shares					x	x
Obs.	2,569	2,569	2,550	2,550	2,550	2,550
r2	0.425	0.475	0.427	0.479	0.431	0.484
Y mean	539.0	999.5	539.0	1000.5	539.0	1000.5

***:p<0.01, **: p<0.05, *:p<0.10, +:p<0.15

Notes: Table displays regression coefficients derived from separate IV regressions of the log average product of labor measured in either 1900 or 1920 (as labeled at the top of each column) onto the 10 meter ruggedness index used as an instrument within the main IV estimations, along with the indicated control variables and state fixed effects. The unit of observation is a county. Below the regression coefficients are p-values for the test of equality between the 10 meter ruggedness index in different columns. Columns 5 and 6 also include more detailed controls on the share of each of the 12 soil orders found within a county. Robust standard errors in parenthesis.
Farm wage winsorized below 10%.

Table 2.4: Estimates of Equipment Investment and Occupational Mobility

	Share Leave Ag. 1920			
	1910 Farmers and Wage-Workers	1910 Farmers and Wage-Workers	1910 Farmers	1910 Wage-Workers
	OLS	IV	IV	IV
	(1)	(2)	(3)	(4)
Pct. Δ Equip. Value per Acre '10-'20 (SD)	-0.0008	0.0377*	0.0245	0.0984**
	(0.0021)	(0.0219)	(0.0226)	(0.0489)
Controls				
State and Age FEs	x	x	x	x
Standard Soil and Topo	x	x	x	x
Obs.	277,196	277,196	224,908	30,431
Adjusted r2	0.009	-0.003	-0.002	-0.016
Y mean	0.317	0.317	0.288	0.498
Montiel Olea-Plueger F-stat		18.4	17.4	16.4

***:p<0.01, **: p<0.05, *:p<0.10, +:p<0.15

Notes: Table displays regression coefficients derived from separate IV or OLS regressions, as labeled above each column, of a dummy variable that a given 1910 agricultural worker remains in an agricultural occupation in 1920 onto the change in equipment value per acre (coefficient displayed), origin (1910) state fixed effects, age fixed effects, and the standard soil/topological controls (900m ruggedness, fertility index, usable water depth, root depth, average number warm days, and erosion tolerance). Standard errors in parenthesis clustered at the 1910 county level. Column 3 also includes more detailed controls on the share of each of the 12 soil orders found within a county. The samples include workers linked between the 1910 and 1920 censuses, aged 25-55 in 1910, who were in an agricultural population in 1910.

Table 2.5: Estimates of Equipment Investment and Average Product of Labor

	1910 (1)	1920 (2)	1940 (3)
Panel A: Farm Income, Using Contemporaneous Ag. Worker Counts			
Pct. Δ Equip. Value per Acre '10-'20 (SD)	19.4	220.6*	222.2***
	(85.6)	(114.9)	(83.2)
Obs.	357,970	357,970	357,970
Y mean	761.0	1082.0	598.1
Montiel Olea-Plueger F-stat	20.0	20.0	20.0
Panel B: Farm Income, Using 1910 Ag. Worker Counts			
Pct. Δ Equip. Value per Acre '10-'20 (SD)		242.8**	226.3***
		(118.3)	(83.2)
Obs.		357,970	357,970
Y mean		993.0	503.4
Montiel Olea-Plueger F-stat		20.0	20.0
Controls			
State FEs	x	x	x
Standard Soil and Topo	x	x	x

***:p<0.01, **: p<0.05, *:p<0.10, +:p<0.15

Notes: Table displays regression coefficients derived from separate IV regressions of the average product of labor as measured in the year at the top of each column in an agricultural workers' 1910 county onto the change in equipment value per acre (coefficient displayed), origin (1910) state fixed effects, and the standard soil/topological controls (900m ruggedness, fertility index, usable water depth, root depth, average number warm days, and erosion tolerance). Standard errors in parenthesis clustered at the 1910 county level. Panel B presents alternative measures of the average product of labor, holding the count of agricultural workers fixed at the 1910 level. The samples include workers linked between the 1910 and 1920 censuses, aged 25-55 in 1910, who were in an agricultural population in 1910.

Table 2.6: Estimates of Equipment Investment and Occupational Wage

	1910 Farmers and Wage-Workers	1910 Farmers	1910 Wage-Workers	1910 Wage-Workers 1920 Non-Ag. Workers
	(1)	(2)	(3)	(4)
Panel A: Dependent Variable: Employed '20				
Pct. Δ equip. value	-0.0009	-0.0007	0.0041	0.0314
per acre '10-'20 (SD)	(0.0104)	(0.0112)	(0.0190)	(0.0363)
Controls				
State FEs	x	x	x	x
Standard soil and topo	x	x	x	x
Y mean	0.948	0.947	0.956	0.910
Obs.	276,483	243,612	32,871	16,077
Clusters	2407	2407	2354	2195
Montiel Olea-Plueger F-stat	16.9	15.9	16.9	17.9
Panel B: Log Occ. Wage '20				
Pct. Δ equip. value	0.2066**	0.2380***	0.0360	-0.1035^{+}
per acre '10-'20 (SD)	(0.0810)	(0.0916)	(0.0584)	(0.0698)
Controls				
State FEs	x	x	x	x
Standard soil and topo	x	x	x	x
1910 cnty price index	x	x	x	x
Y mean	720.6[1]	712.9[1]	776.8[1]	1045.7[1]
Obs.	260,015	228,835	31,180	14,593
Clusters	2405	2405	2339	2143
Montiel Olea-Plueger F-stat	16.1	15.0	16.8	18.4

***:p<0.01, **: p<0.05, *:p<0.10, +:p<0.15

1: mean in 1940 $

Notes: Table displays regression coefficients derived from separate IV regressions of a dummy variable equal to one if a worker was employed in 1920 (Panel A) and the log occupational wage for a workers' 1920 occupation (Panel B) onto the change in equipment value per acre (coefficient displayed), origin (1910) state fixed effects, age fixed effects, and the standard soil/topological controls (900m ruggedness, fertility index, usable water depth, root depth, average number warm days, and erosion tolerance). Standard errors in parenthesis clustered at the 1910 county level. The samples include workers linked between the 1910 and 1920 censuses, aged 25-55 in 1910, who were in an agricultural population in 1910. Columns 2-4 further restrict this sample as indicated at the bottom of the table.

Table 2.7: Effects on Children of Incumbent Agricultural Workers: 1940

	Employed (0/1)	Ag. Worker (0/1)	Log Wages	Employed (0/1)	Ag. Worker (0/1)	Log Wages
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A:	Father: Ag. Wage Worker 1910			Father: Farmer 1910		
		Child: Employed 1940	Child: Employed Non-Ag 1940		Child: Employed 1940	Child: Employed Non-Ag 1940
Pct. Δ equip. value	-0.0036	-0.0234	0.0274	0.0121	0.0109	0.1539**
per acre '10-'20 (SD)	(0.0185)	(0.0340)	(0.0689)	(0.0105)	(0.0367)	(0.0651)
Y mean	0.886	0.244	1097.0[1]	0.919	0.372	1202.6[1]
Obs.	73,386	64,992	43,916	927,052	851,574	456,249
Clusters	2388	2385	2348	2418	2418	2418
Panel B:	Father: Left Ag. 1920			Father: Stay Ag. 1920		
		Child: Employed 1940	Child: Employed Non-Ag 1940		Child: Employed 1940	Child: Employed Non-Ag 1940
Pct. Δ equip. value	0.0527	-0.0459	-0.2728*	-0.0178	0.0592	0.1565+
per acre '10-'20 (SD)	(0.0433)	(0.0707)	(0.1512)	(0.0230)	(0.0569)	(0.1084)
Y mean	0.909	0.236	1292.9[1]	0.925	0.416	1208.5[1]
Obs.	21,560	19,598	12,788	64,661	59,808	29,689
Clusters	2329	2314	2203	2391	2391	2360
Controls						
Father State FEs	x	x	x	x	x	x
Standard soil and topo	x	x	x	x	x	x

***:p<0.01, **: p<0.05, *:p<0.10, +:p<0.15
1: mean in 1940 $
Notes: Table displays regression coefficients from a sample of children whose father worked in an agricultural occupation in 1910. Each panel displays estimates for a different partition of this sample, Panel A partitioning by the father's 1910 occupation, and Panel B by whether or not the father was still working in agriculture in 1920. Each panel/column displays coefficients derived from separate IV regressions of the outcome above each column onto the change in equipment value per acre in the child's father's 1910 county (coefficient displayed), 1910 state fixed effects, age fixed effects, and the standard soil/topological controls (900m ruggedness, fertility index, usable water depth, root depth, average number warm days, and erosion tolerance). Standard errors in parenthesis clustered at the 1910 county level.

Table 2.8: Calibrated Model Parameters

Discount factor	β	0.95	Duernecker and Herrendorf (2017)
Mortality rate	δ	0.04	
Coef. rel. risk aversion	γ	2.0	
Outer elasticity of substitution	σ	0.56	Duernecker and Herrendorf (2017)
Productivity shock shape	α	13.0	Dvorkin and Monge-Naranjo (2019)
Children's human capital level	h_0	0.266	Ratio of father to son income 1940 census

Table 2.9: Model Joint Calibration Fit

Moment	Data	Model
Sectoral productivity (A_j)		
1910 H emp share	0.300	0.819
1910 L emp share	0.333	0.409
1910 F emp share	0.287	0.569
Productivity growth scale (λ_j)		
1940 wage growth H	0.321	0.322
1940 wage growth L	0.441	0.442
1940 rel wage F vs H	0.456	0.850
1940 wage growth W	0.565	0.006
Non-pec. cost of occ. trans. ($\chi_{j,\ell}$)		
1900-1910 likelihood H stay	0.701	0.649
1900-1910 likelihood L stay	0.539	0.277
1900-1910 likelihood F stay	0.759	0.689
1900-1910 likelihood W stay	0.120	0.120
HC cost of occ. trans. ($\tau_{j,\ell}$)		
Avg H to L wage vs H stay	0.630	0.820
Avg L to H wage vs L stay	1.157	1.285
Avg W(F) to F(W) wage vs F(W) stay	0.859	1.177
1900-1910 share transition H to Ag	0.097	0.086
1900-1910 share transition L to Ag	0.179	0.179
1900-1910 share transition Ag to H	0.100	0.108
1900-1910 share transition Ag to L	0.132	0.132
1900-1910 share transition HTF(LTF) to LTF(HTF)	0.075	0.075
Productivity Shock (A_{Tw})		
IV Diff. W attrition HTF-LTF	**0.304**	**0.304**
Production Elasticities		
IV Diff. F attrition HTF-LTF	0.054	0.054
Rel Emp M vs T 1920	0.481	0.838
Children's HC Production		
Rel earn H son vs L son 1940	1.190	1.350
IV Diff. earn child T vs child M	-0.343	-0.343

Notes: Table displays the 24 jointly targeted moments from the model calibration, ordered according to the parameter most closely associated with the moment. The row in bold presents the moment which identifies the magnitude of technological shock.

Table 2.10: Calibrated Model Parameters

Basline Sectoral Productivity (A_j)

h	l	Tf	Mf	Tw	Tf
1.0	7.11	6.568	6.568	0.483	0.483

Post-Shock Sectoral Productivity (A_j)

h	l	Tf	Mf	Tw	Tf
1.0	7.11	6.568	6.568	34.09	0.483

Productivity Growth Scale (λ_j)

h	l	Tf	Mf	Tw	Tf
1.291	1.211	1.268	1.268	1.234	1.234

Non-Pec. Cost of Occ. Trans. ($\chi_{j,\ell}$)

	h_{t+1}	l_{t+1}	Tf_{t+1}	Mf_{t+1}	Tw_{t+1}	Tf_{t+1}
h_t	0.941	1.0	1.0	1.0	1.0	1.0
l_t	1.0	0.886	1.0	1.0	1.0	1.0
Tf_t	1.0	1.0	0.983	1.0	1.0	1.0
Mf_t	1.0	1.0	1.0	0.983	1.0	1.0
Tw_t	1.0	1.0	1.0	1.0	0.934	1.0
Tf_t	1.0	1.0	1.0	1.0	1.0	0.934

HC Cost of Occ. Trans. ($\tau_{j,\ell}$)

	h_{t+1}	l_{t+1}	Tf_{t+1}	Mf_{t+1}	Tw_{t+1}	Tf_{t+1}
h_t	1.0	0.97	0.785	0.785	0.785	0.785
l_t	0.933	1.0	0.951	0.951	0.951	0.951
Tf_t	0.05	0.701	1.0	0.678	0.764	0.764
Mf_t	0.05	0.701	0.678	1.0	0.764	0.764
Tw_t	0.05	0.701	0.764	0.764	1.0	0.678
Tf_t	0.05	0.701	0.764	0.764	0.678	1.0

Production Elasticities

θ	ρ
2.238	0.081

Children's HC Production

κ	ν
0.94	0.7

Notes: Tables display estimated model parameter values.

2.11 Figures

Figure 2.1: Farm Equipment Productivity

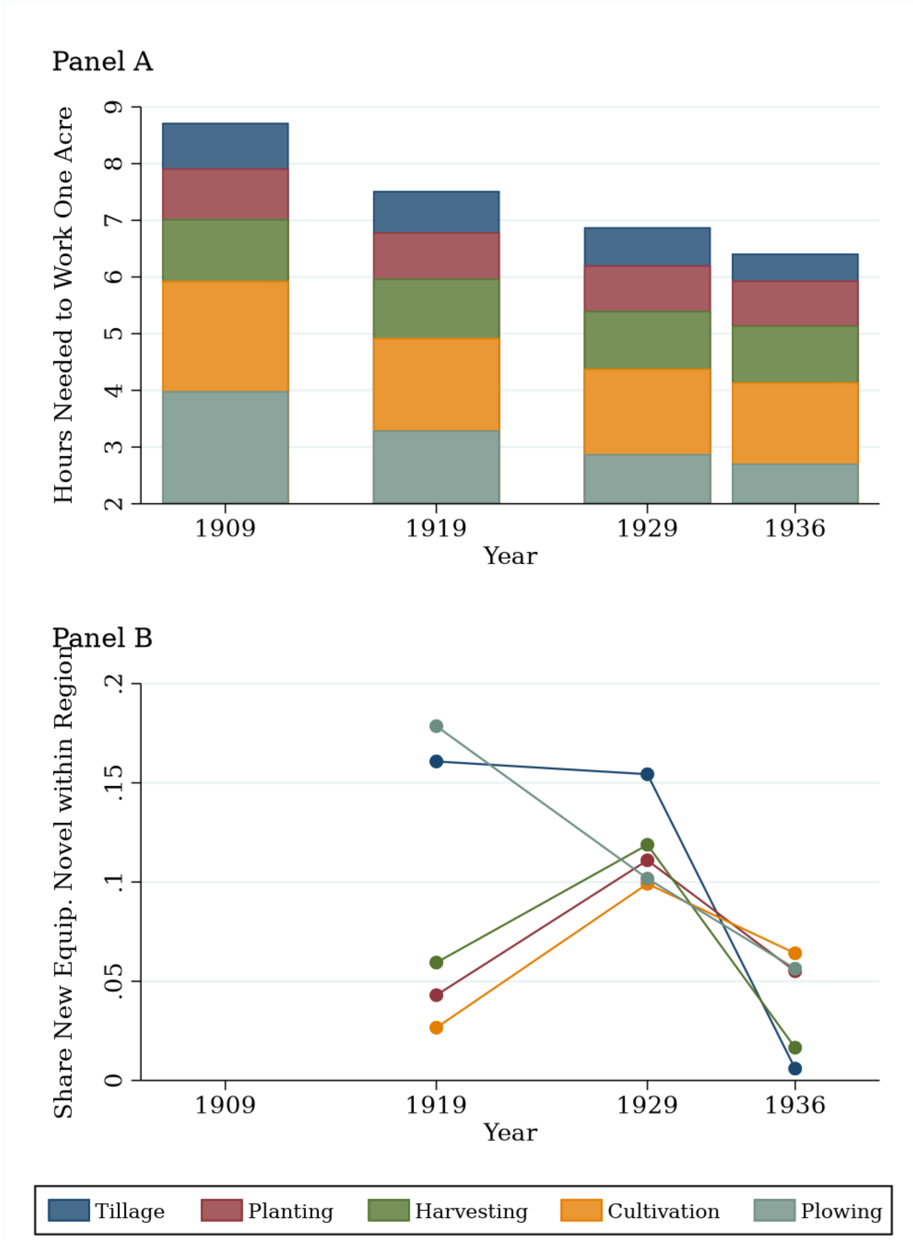

Notes: Data derived from McKibben *et al.*, 1939.
Panel A: Hours needed to complete each color coded task for one acre of land using the average technology (equipment) available over time within a geographically distributed panel of farms. *Panel B:* Share of new equipment which was not recorded in panel's previous period. New equipment of a given type measured as the difference between the number observed in period t minus the number observed in period $t-1$, or 0 if negative. Equipment type designated as novel if none were presented within region at the previous observation.

Figure 2.2: Instrument Construction

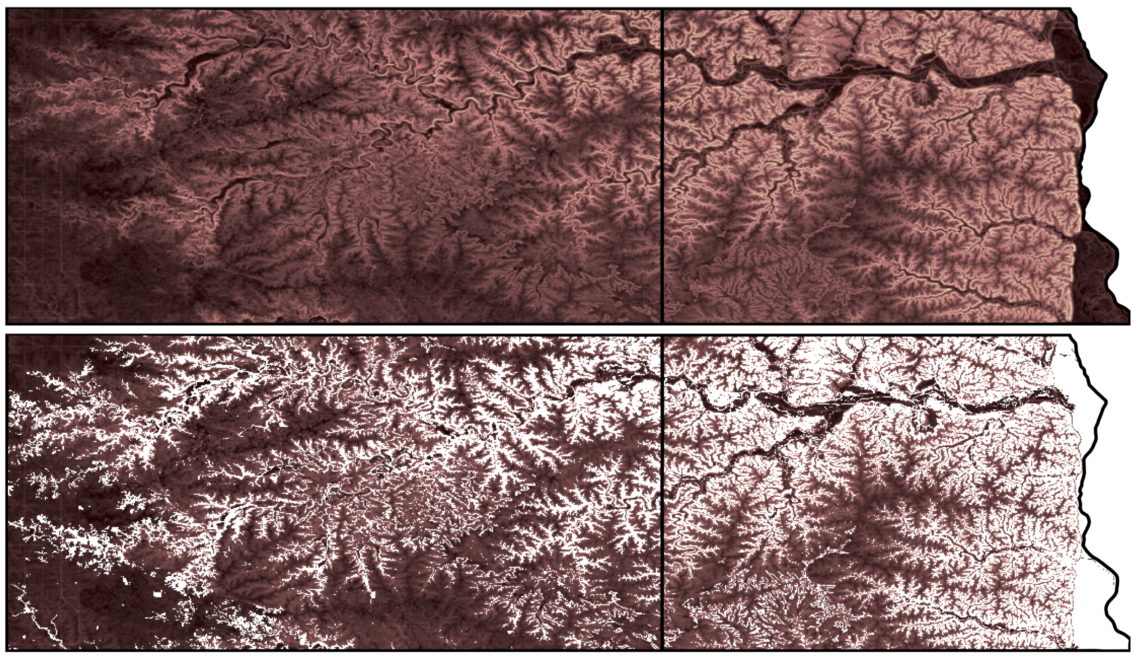

Notes: The upper image displays ruggedness at a $\frac{1}{3}$ arc second resolution for Fillmore County (left) and Houston County (right) Minnesota. Lighter pixels correspond to more rugged locations. The lower image displays the same ruggedness data, but only includes areas used to construct the average county ruggedness measure used as an instrumental variable. Included areas required a fertility index above a cutoff threshold set to match the total national land area under cultivation at baseline 1910. The constructed ruggedness index for Fillmore and Houston Counties are 1529.5 and 1137.9 respectively.

Figure 2.3: Farm Equipment Per Worker

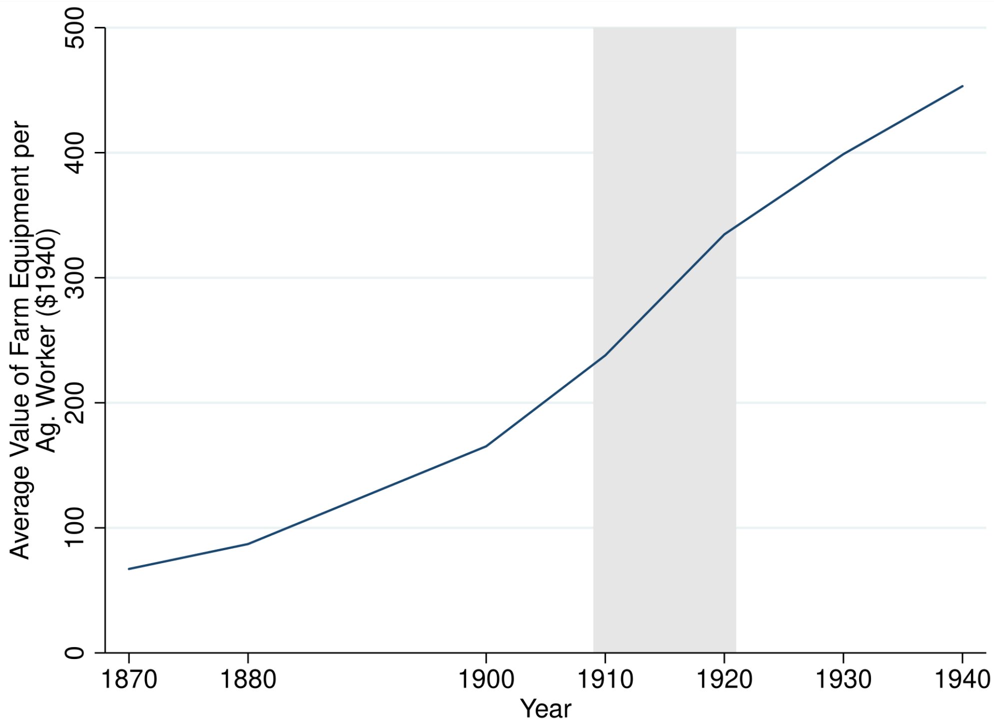

Notes: Average value of farm equipment per agricultural worker across US counties with significant agricultural employment share. Count of agricultural workers derived from full census, value of equipment derived from agricultural census estimates.

Figure 2.4: Equipment Investment and New Equipment Characteristics

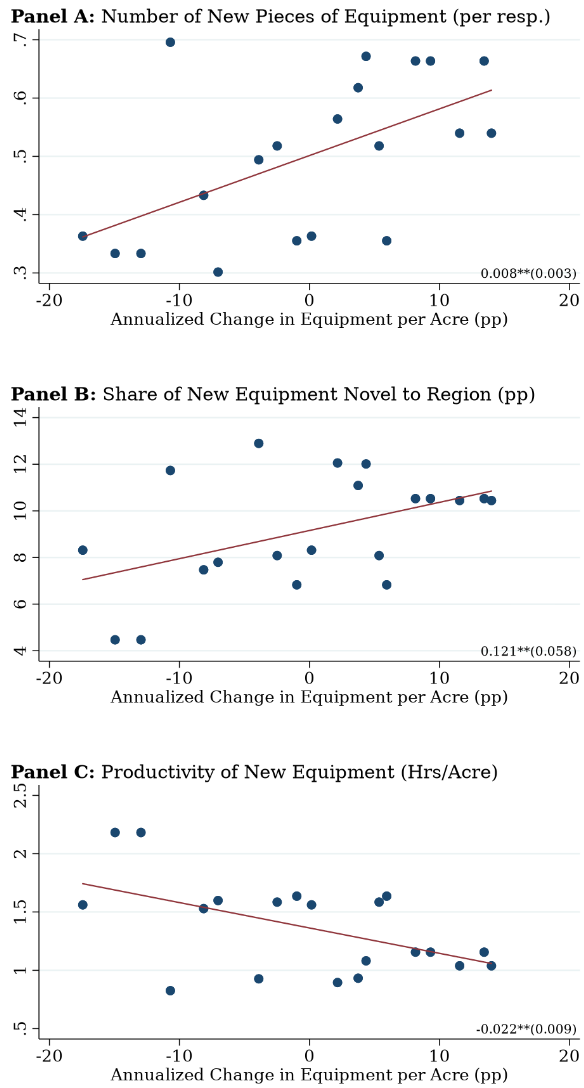

Notes: Data derived from McKibben *et al.*, 1939 and mapped to county level estimates using county sample shares within the equipment panel. Unit of observation is a county implement type. Standard errors derived from block bootstrap.
Panel A: Binscatter showing the relationship between growth in county level equipment (value) per acre and the number of new pieces of equipment the average survey respondent purchased between 1909/1910 and 1919/1920.
Panel B: Binscatter showing the relationship between growth in county level equipment (value) per acre and the share of new pieces of equipment novel to region. New equipment measured as the difference between the number observed in period t minus the number observed in period $t-1$, or 0 if negative. Equipment type designated as novel if 0 were present in the previous observation.
Panel C: Binscatter showing the relationship between growth in county level equipment (value) per acre and productivity of new equipment (defined as in Panel B), where lower hours per acre correspondents to more productive equipment.

Figure 2.5: Second Generation Wage Distribution

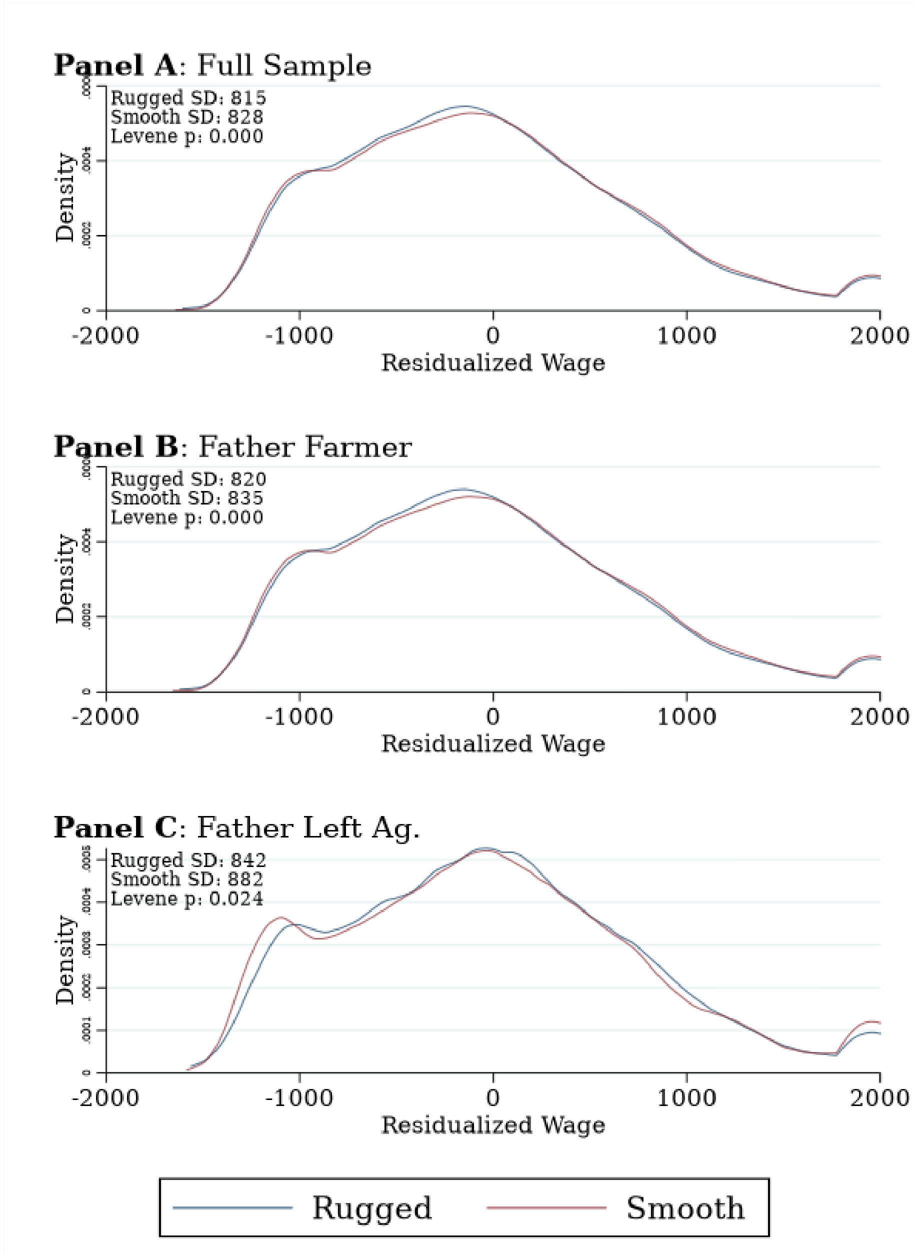

Notes: Figure displays kernel density plots for the residualized wages of children of agricultural workers separately for the upper and lower terciles of terrain ruggedness in the father's baseline (1910) county. Plots winsorized above \$2,000 for clarity, but standard deviations derived from underlying data. The sample in Panel A includes all matched sons who are employed in a non-agricultural occupation in 1940, and Panels B and C present distributions for the indicated sub-samples. Wages are residualized against the same controls as in Table 2.7. As a reference, the difference in the standard deviation of (unresidualized) wages between urban and rural men employed in non-agricultural occupations of the same age range in the 1940 census is \$127.

Chapter 3

MOVES OF OPPORTUNITY, SUPPORTING SERVICES, AND LABOR MARKET OUTCOMES

Together Federal, State, and Local Governments fund housing for over 1 million American households living in the public housing system.[1] Through this system, policymakers directly influence where 1 in 200 American households live. Often public housing sites are located in neighborhoods that may offer few job opportunities to residents and which may alter residents' productivity or preferences for supplying labor. This paper evaluates the extent to which moves out of such neighborhoods impact residents' labor market outcomes. A better understanding of the role low-opportunity neighborhoods play in restricting public housing residents' labor market outcomes is crucial to evaluating the costs and benefits of housing policy. Additionally, understanding the connection between neighborhood, employment, and productivity has been an explicit interest of researchers dating back to at least the mid 20th century when Kain (1968) tied racial housing segregation to minority unemployment, yet there is still little consensus about the magnitude or even existence of these links.

The goals of this paper are twofold. First, it aims to estimate if low-opportunity neighborhoods may have a causal impact on public housing residents' labor market outcomes. Second, it seeks to evaluate the extent to which such effects are modulated by the conditions surrounding a neighborhood transition, particularly the role of supportive services provided during the transition. Such a relationship may exist for a number of reasons. It may be that employers (correctly or incorrectly) infer

[1]Source: Department of Housing and Urban Development Resident Characteristics Report, September 2018. (https://pic.hud.gov/pic/RCRPublic/rcrmain.asp)

something about a potential employees' productivity from their neighborhood.[2] It is also possible that employment decisions of neighbors influence individuals' own decisions based on informal hiring, social preferences, or learning about the costs and benefits of employment.[3] It could also be that information frictions or commuting costs cause households to search for employment close to their residence, or high crime rates cause stress which reduces productivity.[4] Whatever the underlying mechanisms, any causal relationship between neighborhoods and adult labor market outcomes has remained difficult to estimate due to the endogenous sorting of households into neighborhoods based on unobservable factors potentially correlated with employment.

This paper addresses the endogeneity of neighborhood choice by evaluating an exogenous shock to neighborhood quality for a group of public housing residents. The specific identifying variation comes from a Choice Neighborhoods Implementation grant administered by the Department of Housing and Urban Development (HUD) to redevelop a large public housing site in Memphis, Tennessee. Due to this redevelopment, existing residents were offered a Section 8 voucher and required to find housing away from the site; a move that dramatically altered the type of neighborhood residents experienced. A novel aspect of the Memphis moves is that residents received ongoing case management services as part of the Choice Neighborhoods program, and thus the treatment they received was a combination of neighborhood relocation and ongoing case management. To disentangle the two treatments, relocated residents are compared to other Memphis public housing residents who were not required to relocate or given the option to take a housing voucher but received a similar case management program. Therefore, I am able to estimate the effect of relocation into a new neighborhood, conditional on receiving case management services over the relocation

[2]See for example Besbris *et al.* (2015)

[3]See for example Bayer *et al.* (2008) and Schmutte (2015)

[4]See for example Weinberg *et al.* (2004)

process. This distinction appears to matter, as households who were quasi-randomly assigned to receive more intense case management experienced less job loss post-move.

Using rich administrative data, the effects of relocation can be identified using only within-individual variation in relocation status through the use of an individual-level fixed effect model. Households were free to choose their destination housing, whether or not to take the voucher, and where to live within the metro area. However, a simple stylized model demonstrates that inference about the existence of neighborhood effects comes from the change in individuals' housing choice set created by the demolition and not from endogenous sorting into post-relocation neighborhoods. Specifically, finding that outcomes improved post-move provides strong evidence that neighborhoods impact labor market outcomes, even when the choice of new neighborhoods is endogenous, as the demolition required households to make new housing decisions under an expanded choice set within a relatively narrow window of time, thus changes in neighborhood environment are unlikely to correlated with time-varying unobservables within individual.

In a panel constructed from self-reported employment information, where the average relocation duration is approximately 12 months, relocation is estimated to have increased hourly wages by an average of \$0.69. Given the low baseline wage for this population (\$9.36 per hour), a \$0.69 wage increase represents a large increase in wages. No statistically significant effect on employment at the extensive or intensive margin was found on average over the period of relocation, although estimates for the extensive margin are noisy and do not rule out economically relevant impacts. Further, after separating the effect of relocation by duration in the new neighborhoods, full-time employment is estimated to have decreased after one year in the new neighborhoods, relative to the unmoved residents receiving case management services. More specifically, the combined effect of both receiving case management services and relocating appears to increase employment, although less drastically than receiving the case management alone. The share of households with at least 1 member working full-time is estimated to have been unaffected by relocation, suggesting that the moves may have shifted the labor allocation within households. However, all extensive margin results are relatively imprecise due to sample size limitations.

The effects of moving to better neighborhoods on employment become clearer after separating impacts by age or education. Relocation is estimated to have caused an economically and statistically significant reduction in employment for individuals older than 45 or with less than a high school degree, while their younger or more educated peers are estimated to have experienced no impact on employment and positive wage effects of approximately $0.90 per hour. Younger or more educated individuals earn higher wages post-move (and pre-move), which suggests that unambiguously positive neighborhood effects accrued to those with the highest value to employment measured by earnings potential and discounted value of work experience.

Key to the interpretation of these results is the assumption that relocation effects are well identified, *i.e.* that I can attribute any changes in outcomes coincident with relocation on the relocation itself and not some outside factor. Several pieces of analytical evidence support this assumption. First, the effect of relocation on wages is very similar for households who chose not to use the voucher (and thus move into a more modern public housing site), suggesting that voucher receipt itself is not behind the results.[5] Second, relocation caused large changes in observable measures of neighborhood quality. For example, on average residents chose to move into neighborhoods with 50% lower poverty rates, 40% lower rental rates, and 20% lower shares

[5]Note that even households who did not use the voucher experienced significant improvements in neighborhood quality because the baseline neighborhood had some of the worst observables (poverty rates) in the metro area.

of single-parent households compared to their baseline neighborhood (Census Block Group). Such large changes to an individual's neighborhood environment could very plausibly impact labor market outcomes through one or more of the many potential channels previously outlined. Finally, measures of post-move neighborhood quality are correlated with improved labor market outcomes, suggesting that the disruption of a non-voluntary move alone is unlikely to explain improved labor market outcomes.

It is important to emphasize that estimates are conditional on households receiving case management services over the relocation process and for several years afterward. However, I am able to net out the direct effect of these services by observing a separate population who were not given the opportunity to move, but received the same case management services. Evidence further suggests that this type of case management acts as a complement to relocation. Using a regression discontinuity design based on the way case managers assigned priority to households, I see that households who quasi-randomly received more intense case management over the move were 17 percentage points more likely to be employed after moving. This provides strong evidence that the amount of support households receive as they move out of low-opportunity public housing neighborhoods plays a pivotal role in the degree to which they benefit from such a move.

It is natural to expect that moving into neighborhoods of opportunity increased wages due to proximity to higher quality jobs. However, I fail to reject the hypothesis that post-move jobs were further away or less accessible from the baseline neighborhood as compared to pre-move jobs. Therefore, it appears physical barriers to better job opportunities cannot explain the observed wage increases.

Existing evidence on the impact of neighborhoods on adult outcomes is limited by the availability of potentially exogenous variation in neighborhood quality. Some of the first evidence that neighborhoods may have a relevant impact on adult labor market participation came from the Gautreaux Project, a natural experiment that moved low-income families to more prosperous, racially diverse neighborhoods. Rosenbaum (1995) finds adults relocated to suburban neighborhoods have higher employment five years after relocation relative to those moved to urban neighborhoods. Mendenhall *et al.* (2006) find this effect persisted for women assigned to more racially diverse neighborhoods up to 15 years after relocation.

Encouraged by the results of Gautreaux, researchers and policymakers designed the Moving to Opportunity (MTO) experiment, where volunteer families were randomly assigned housing vouchers with different restrictions based on neighborhood poverty. This experiment failed to find an impact of neighborhood poverty on adults' labor market outcomes but did find long-term benefits for children who were relocated to lower-poverty neighborhoods. (See Orr *et al.* (2003), Sanbonmatsu *et al.* (2011), and Chetty *et al.* (2016) for details.) More recently, van Dijk (2019) finds evidence quasi-exogenous moves into better neighborhoods due to public housing lotteries in the Netherlands increased household income.

Several studies have examined the effects of non-voluntary moves due to public housing demolitions and subsequent resident relocation due to HOPE VI, a predecessor to the Choice Neighborhoods program studied in this paper. In investigations focused on policy evaluation, Goetz (2002) and Clampet-Lundquist (2004) find no impact on adult labor market participation due to these relocations, while Anil *et al.* (2010) finds a small, but positive, effect on employment five years after relocation. In line with the MTO results, Chyn (2018) finds that children in households forced to relocate from disadvantaged neighborhoods due to public housing demolitions experienced positive effects on employment and earnings later in life.[6] There is also an

[6] Using the same Chicago demolitions as Chyn (2018), Jacob (2004) finds no effect of relocation on children's academic outcomes.

existing literature that addresses the consequences of exogenous mobility shocks from the perspective of evaluating the barriers to mobility. See for example Nakamura *et al.* (2019), and Bryan *et al.* (2014). This literature generally considers mobility across regions rather than neighborhoods and there is no consensus about the degree to which these results generalize to smaller geographies.

This paper contributes to the existing literature in at least three ways. First, it provides some of the first evidence that low-income households living in public housing can experience non-trivial impacts to their wages after moving to lower poverty neighborhoods, implying that some low-opportunity neighborhoods decrease the wages of residents. However, the evidence of a negative effect on employment for some demographic groups documents that non-voluntary moves may cause significant disruption to employment for some households. Therefore, any benefits from a policy that relocates households to improve outcomes should be carefully weighed against the potential costs to participating households.

Second, this paper builds on the emerging evidence that receiving additional support may complement the impact of moves to high opportunity areas for low-income households. This hypothesis is consistent with the results of Galiani *et al.* (2015), who use a structural model to estimate that voluntary take-up of the restricted vouchers in the MTO experiment would have been roughly half the rate observed had mobility counseling not been provided, implying that this much less intensive intervention may have been quite valuable to MTO participants. Bergman *et al.* (2020) also find evidence that interaction with an outside agent who provides information and helps reduce individual barriers during or before a move improves outcomes for low-income households. While the duration and type of case manager interaction are different than in Bergman *et al.* (2020), the results of this paper support the broad conclusion that providing guidance and support to low-income households as they transition into new neighborhoods improves outcomes.

Finally, this paper is one of the first to provide concrete evidence about the extent to which any causal impact of neighborhoods may be driven by physical barriers to better employment opportunities. This analysis provides useful evidence with which to disentangle the mechanisms linking neighborhood choice to labor market outcomes.

3.1 Background

3.1.1 Setting

Data for this paper comes from case management reports at several public housing sites in Memphis TN, recorded late-2015 to mid-2019. During this period, one of the sites underwent physical redevelopment through a Choice Neighborhoods implementation grant from the Department of Housing and Urban Development (HUD). This site will be referred to as the CN site. The CN site was selected as the target of this grant based on its physical infrastructure needs, not on the characteristics of its residents. However, all residents at the time of implementation were required to either take a section 8 voucher or find accommodation at another public housing site over the five-year redevelopment.

This paper will argue that residents at the CN site faced an exogenous shift in neighborhood quality due to the implementation grant, and then leverage this shift to evaluate the causal impact of neighborhoods on employment and wages. It is important to note that residents were free to choose where they relocated. However, as they all started from a significantly disadvantaged baseline, the relocation is associated with a significant improvement in average neighborhood quality along a number of observable dimensions.[7] The vast majority of relocations away from the CN public

[7] Only 2 households are observed moving into block groups with a higher poverty rate, with the increase being less than one percentage point.

housing site are observed to have occurred over a five-month window in late 2016. Figure 3.1 displays the distribution of relocation dates at the CN site.[8] Housing choice (Section 8) vouchers were made available to all residents at the start of this window, and all residents were required to relocate by the start of 2017 so that construction could begin at the CN site.

HUD required all CN grant recipients, to provide comprehensive case management services to current residents throughout the redevelopment process, in order to ease the burden of relocation and create lasting momentum at the site through residents who chose to return after redevelopment. The role of case managers at the CN site was to track the status of residents, connect residents to service providers, and help residents set and achieve goals for themselves and their households. Examples of potential service referrals include free GED prep courses, free or reduced-price legal services, or free after-school programs. In Memphis, these case management services were provided by a national not-for-profit organization, Urban Strategies, referred to as the NPO throughout this paper, who provided the data for this project. Towards the end of the Memphis CN relocation window, the NPO was contracted to provide a similar case management service to three other public housing sites in Memphis through a HUD ROSS grant. The goal of the ROSS grant was to improve resident self-sufficiency at the three target sites, which will be referred to as the non-CN sites, and the NPO provided essentially the same services as at the CN site. However, the non-CN sites were never required to relocate and never given the option to take a section 8 voucher.

[8]Relocation location is known for 67% of the CN individuals due to incomplete or missing data. Case managers recorded location information without using the full address, leading to ambiguities in geolocating addresses.

3.1.2 Data

To better target case management efforts and track efficacy, the NPO collected information on participating residents. Most of this information was recorded in the form of various status reports pertaining to individuals or households. The NPO provided access to selected pieces of these reports for both the Memphis CN and non-CN sites, which were then compiled into a monthly panel that tracks individuals' employment and wages from the start of case management until mid-2019.

The key report for this study recorded jobs held by residents, with the unit of observation being a job tied to a specific resident. This report provided starting and ending dates, starting and ending wages, starting and ending hours worked, and information about the employer. From this report, a monthly panel was constructed under the assumption that in the absence of a recorded job, an individual was unemployed. Throughout this paper, an individual will be considered unemployed if they do not have a recorded job. In a separate yearly survey of participants, this panel was consistent with participants' reported employment status 87.9% of the time, supporting this assumption. It is important to note this data only considers variation in wages and hours worked due to switching jobs, and so cannot speak to variation over time within the same job. This limitation may be less important given the high turnover rate in the specific population under consideration. Further, Appendix Section 3.4.3 displays results where hours worked and wages are adjusted for within job changes over time by assuming linear wage growth and leveraging the fact that ending wages are known for jobs that ended by the end of the panel.

Information from several other case management reports was combined with this employment data.[9] Survey responses were carried forward, with the understanding

[9]These include housing decisions, subjective neighborhood safety, adult and family subjective well being, the number and type of services a participant was referred to, family structure, and

the NPO considered them to reflect the current status of a household and updated a given report when it was no longer accurate. Data on quarterly risk scores constructed by the NPO for every participating adult was also made available and included in the panel. This measure was purportedly used to assess an individual adult's risk of experiencing adverse outcomes such as loss of housing, income, etc.

The combined panel of CN and non-CN residents consists of all adults participating in case management who are aged 18-65 over the period of observation. The NPO was only contracted to provide services at the non-CN site for two years, and so only observations in the first 24 months of case management are included in the main analysis panel. In the CN sample, the average relocation duration of relocated residents is 13.1 months and the average relocation distance is 4.9 miles away from the original CN site, for those with a known post-move location. To ensure that the pre-move baseline at the CN site is not biased by individuals anticipating the coming relocation and changing their employment status, the panel will also exclude observations from the second half of 2016 at the CN site, the period when the bulk of households began moving away from the site.[10]

Two distinct factors could lead to significant sample selection bias in this setting. One is that participation in case management was optional, and thus if non-participation was correlated with relevant unobservables, this non-response would bias the estimates. Second, due to the disruptive nature of relocation at the CN site, households could have preemptively left public housing before relocation, to avoid this process altogether. However, evidence suggests that both of these effects had a small impact on the overall sample. At the CN site, HUD mandated the housing authority provide one-to-one replacement of current housing and allow current residents the

---

income

[10]This includes June through December 2016.

option to return to the site upon completion of the revitalization process. In order to remain eligible to return to the CN site after renovation, residents were required to participate with the provided case management. Thus, the majority of families opted to participate. Of the 420 potential units at the CN site, there is at least one case management record for 342 unique families. Further, the CN site was not full at the time of relocation, as the housing authority had stopped placing residents there in anticipation of the redevelopment. Participation was similarly high at the non-CN sites, where 146 records exist for a potential 181 units. The fact that case management records exist for a large proportion of potential units at each site suggests that attrition due to turning down case management and preemptive relocation was minimal.

Several potential indicators of neighborhood quality were selected from the American Community Survey five year estimates 2012-2016 to capture a well-rounded picture of how the physical environment around residents shifted after relocation. Measures include the proportion of households under the poverty line, the proportion of minority residents, the proportion of adults with at least a high school degree, the proportion of households with a female head, and the proportion of households who rent their housing.

In order to rule out local economic shifts coinciding with the relocation period, monthly unemployment, average wage, and average hours worked for the Memphis metro area from the Bureau of Labor Statistics' Current Employment Statistics program were incorporated into the panel. These monthly estimates were then used to construct 12-month rolling averages.

3.2 Methodology

3.2.1 Identification of Neighborhood Effects

This section will present a stylized model of neighborhood choice to illustrate the causal relationship of interest and discuss identification in the specific Choice Neighborhood setting, specifically the post-move endogeneity of neighborhood choice.

First, let us focus only on the employment decision. Let $u_i(w, h; \Phi_n)$ represent the utility of individual i if she earns a wage w and works h hours. Vector Φ_n represents a set of factors relevant to her utility which may differ across her neighborhood, indexed by n. These factors may represent a number of different dimensions of neighborhood choice such as differential local amenities, the value of relative consumption, or differential access to consumption goods. Let $\Gamma_{i,n}$ be the menu of employment opportunities individual i receives if she lives in neighborhood n, where each option is defined as a combination of a wage and hours worked. Menus are allowed to vary across neighborhoods due to factors like commuting costs or discrimination. Then, within a neighborhood n, individual i picks the employment opportunity that maximizes her utility:

$$v_i(n) = \max_{\{w,h\} \in \Gamma_{i,n} \cup \{0,0\}} u_i(w, h; \Phi_n) \tag{3.1}$$

Where $v_i(n)$ is individual i's indirect utility function defined over her set of possible neighborhood choices Ω. Equation 3.1 makes explicit that each agent may choose not to work and earn zero wage. We will restrict ourselves to a static model and assume for ease of exposition that some aspects of neighborhood specific utility are additively separable from employment decisions:

$$u_i(w, h; \Phi_n) = u_i(w, h; \{\Phi_n^1, \Phi_n^2\}) = u_i^1(w, h; \Phi_n^1) + u_i^2(\Phi_n^2) \tag{3.2}$$

where u_i^1 and Φ_n^1 capture individual and neighborhood specific utility derived from the employment decision and $u_i^2(\Phi_n^2)$ captures individual-neighborhood specific match quality on factors independent of employment.

First notice that, because utility (u_i) and choice sets ($\Gamma_{i,n}$) are individual specific, the optimal neighborhood/employment decisions may differ across individuals. Further, because the factors in Φ_n and the choice sets ($\Gamma_{i,n}$) differ across neighborhoods, individuals may make different employment decisions in different neighborhoods. However, it is also entirely possible that individuals make the same employment decisions across all neighborhoods. With this in mind, one can say that the neighborhoods in the choice set Ω have a casual impact on wages for individual i if she would optimally pick employment opportunities with different wages in any two neighborhoods. Let Let $(w_i^*(n), h_i^*(n))$ denote the optimal decisions for the utility maximization problem (equation 3.1). For simplicity of notation, assume the utility problem has a unique solution in each neighborhood. Then neighborhoods causally impact wages if:

$$\exists\, n, n' \in \Omega \mid w_{i,n}^* > w_{i,n'}^* \tag{3.3}$$

How can I evaluate if the statement in 3.3 is true? The Memphis Choice Neighborhoods setting provides an ideal natural experiment in which to search for two such neighborhoods. Within this framework, the Choice Neighborhood relocation can be seen as removing a restriction on residents' choice set, Ω. Pre-move residents received a very large transfer (heavily subsidized housing) if they lived in their baseline neighborhood. After the move, residents received a voucher, and so this transfer was no longer contingent on living in a specific neighborhood (or in public housing). Therefore, I effectively observe two different sets of employment outcomes for each individual, one under a restricted neighborhood choice set and one with this restriction eased. We can then evaluate the inequality in 3.3 directly, at least for these two neighborhoods.

Given this framework, it is clear that the power to identify neighborhood effects in the Memphis Choice Neighborhoods setting comes from the removal of restrictions to housing choice, not from the cross-sectional sorting of households post-move. In fact, all of the specifications used to estimate the impact of relocation will include individual fixed effects, thus restricting identifying variation to within-individual. While a positive effect of relocation on outcomes would provide evidence that the statement in 3.3 holds in this context, there is no reason to believe that residents sorted in such a way as to maximize their labor market outcomes. The model makes this clear, as it could be that some neighborhoods with low earnings potential also offer large compensating amenities through the $u_i^2(\Phi_n^2)$ term. Therefore, I cannot say that my estimates of the impact of relocation (and thus switching neighborhoods) is the maximum possible within this context, but I can say that the maximum effect is *at least as big* as my estimates. Said differently, an estimate of economically relevant magnitude suggests that the maximum neighborhood effect is at least as large.

The Choice Neighborhoods setting is ideal because the moves were non-voluntary, and so the timing can be thought of as quasi-random within individual. Because of this, one may expect that any change in employment outcomes is not driven by unobserved factors within-individual. For example, in purely observational data I may be worried that some exogenous shock to $\Gamma_{i,n}$ or Φ_n, such as gaining more education, changing family structure, or a large wealth transfer may also cause individuals to move neighborhoods. In the Choice neighborhoods setting, I can be relatively confident that individuals' $\Gamma_{i,n}$ or Φ_n remained fixed over the move because the move was caused by a need to renovate the stock of public housing.

3.2.2 *Empirical Strategy*

Section 3.2.1 discussed how the Choice Neighborhood moves can be thought of as an experiment that repealed a restriction to individuals' housing choice set. This section will explore how to consistently estimate the average impact of relocation on residents' employment outcomes. As discussed in Section 3.2.1, the quasi-random nature of the Choice Neighborhoods moves means that observed changes in employment outcomes are unlikely to be caused by unobservable changes within-individual. Therefore, the focus of the empirical strategy will be addressing potential endogeneity due to the timing of moves relative to aggregate trends and the previously mentioned case management component of the Choice Neighborhoods program.

The relocation of residents at the CN site occurred over a planned period of case management intervention targeted at improving resident outcomes. Therefore, any empirical strategy to estimate the effect of relocation must address this intervention as well. Leveraging the availability of data on a group of public housing residents who received the same case management treatment, but were not forced to relocate and were not given the option of taking section 8 vouchers, the estimation strategy will identify the impact of relocation separately from the impact of ongoing case management. This approach is conceptually similar to a difference in difference strategy where the unit of treatment is an individual. Unlike the traditional diff-in-diff setting, where time is the running variable, here I will use the duration of case management.

The second potential issue, endogeneity of relocation due to the timing of moves relative to aggregate trends, will be addressed by directly controlling for a flexible function of short run local economic conditions. Specifically, 12-month rolling average metro unemployment for employment outcomes, 12-month rolling average Memphis metro wage for wage outcomes, and 12 month rolling average Memphis metro hours worked for hours worked outcomes. Further, given the often temporary labor market attachment of the population in question, outcomes may be more sensitive to seasonal variation than in other settings. Because relocation occurred over a 6 month period perfectly co-linear with seasonal variation, I will also control for seasonal fixed effects.

Therefore, the baseline regression model will be:

$$Y_{i,t} = \gamma_i + \beta R_{i,t} + \psi_{CM_{i,t}} + \kappa_{s_{i,t}} + g(E_t) + \epsilon_{i,t} \tag{3.4}$$

where $Y_{i,t}$ is the outcome for individual i at time (month) t, γ_i is an individual fixed effect, $R_{i,t}$ is an indicator that individual i has relocated by time t, $\psi_{CM_{i,t}}$ is a fixed effect for cumulative months of case management (denoted $CM_{i,t}$), $\kappa_{s_{i,t}}$ is a seasonal fixed effect where $s_{i,t}$ indexes the season, E_t is a 12-month rolling average labor market outcome for the metro area, $g(\cdot)$ is a quadratic function, and $\epsilon_{i,t}$ is an idiosyncratic error term.

The object of interest in this model is β, or the average effect of relocation on residents' labor market outcomes. The implicit identification assumption is that $R_{i,t}$ is exogenous conditional on the controls. It is worth emphasizing that, given the individual fixed effects (γ_i), all identification of β comes from within-individual variation in outcomes, and so this assumption need only hold within-individual over time.

Are the proposed controls sufficient to satisfy this assumption? We can think about threats to this assumption in two broad categories. First, it could be that the timing of moves correlated with unobserved individual-specific factors associated with employment, such as receiving a raise, changes to education, etc. This is unlikely given the quasi-random timing of the relocations. One caveat to this is the case management services individuals received over the move, which could cause unobservable changes. We are non-parametrically accounting for these potential changes through the $\psi_{CM_{i,t}}$ parameters. However, because the $\psi_{CM_{i,t}}$ parameters are only identified by observed changes in the non-CN group, it may be the case that the non-CN households do not provide a good counterfactual for the CN households.

One way to evaluate if the non-CN households form a suitable control group is to look at baseline observables in the two groups. The CN site was selected due to the physical infrastructure needs of the site and not on resident characteristics, and so I should not expect to see a difference in observables across sites if residents were not sorted a priori. Table 3.1 demonstrates that the two groups look quite similar along a number of dimensions. The one exception to this is family size and gender. Housing units at two of the three non-CN sites had an average of 2.93 bedrooms while units at the CN site had an average of 2.25.[11] This difference coincides with the disparity in family size across sites and also potentially explains the gender gap, as many larger families include older female relatives. Appendix Section 3.4.3 demonstrates that results are robust to restricting the sample to families with only 1 or 2 adults, for which there is no difference in family composition across sites.

Leaning into the similarities with a diff-in-diff approach, I can further test if the non-CN households form a reasonable counterfactual by examining the pre-move trends in outcomes, to determine if CN and non-CN individuals had similar responses to case management before the CN group moved. Given the similarity in observables, it's likely that the trends of pre-relocation CN residents and non-CN residents are similar. Figures 3.4 and 3.5 plot the estimated relocation effects in an event study framework, where any pre-trends can be directly observed. If the non-CN households responded to the case management services differently than the CN households, there should be a pronounced pre-trend. The absence of such a trend provides additional evidence that the CN and non-CN households likely responded to the case management treatment in similar ways, and thus the case management fixed effects are

[11]Data for the third non-CN site was not available

sufficient to address the endogeneity of $R_{i,t}$ to the duration of case management.

The second category of threats to identification concern shifts in local labor markets that may have been coincident with the relocation window. To address this concern I rely on the local economic controls $g(E_t)$ and seasonal fixed effects $\kappa_{s_{i,t}}$. The economic controls are adjusted using a polynomial (quadratic) function to allow a degree of curvature in the response of residents' outcomes to local shocks, given that the residents are pulled from a very selected part of the metro population. Results are robust to using a simpler linear term instead. One potential issue with these controls would be if an aggregate shock impacted the public housing population or low-income population, but not the average resident of Memphis. While this cannot be ruled out completely, the time period in question was relatively stable economically and there were no known policy changes that may have selectively improved public housing residents' outcomes. Further, as many residents began the case management program at approximately the same time, to the extent that any unobservable aggregate shock was public-housing specific, it should be reflected in the case management fixed effects. It is for this reason that, for the remainder of the paper, I will not ascribe a strong causal interpretation to the case management coefficients.

Finally, standard errors will be clustered at the individual level for all analysis, unless otherwise specified, to account for serial correlation in outcomes within-individual. While the first-best way to cluster may be at the unit of treatment assignment, the housing site, this is not possible for the Memphis Choice Neighborhood setting, as only one site was treated.

3.3 Results

3.3.1 Relocation Shifted Neighborhood Quality

Residents at the Choice Neighborhoods site experienced a substantial shift in neighborhood environment after relocation. Figure 3.2 plots the average census block group poverty level for the subset of CN households with location data.[12] We observe a substantial decline in neighborhood poverty level over the period of observation, from a baseline of 75% to an average of 35%. Not surprisingly, this shift in neighborhood observables is strongly correlated with the proportion of households relocated, shown as a dashed red line in Figure 3.2. The CN site was located in a neighborhood with an extremely high poverty rate, and thus, while there was no requirement for CN households to move into low poverty neighborhoods, relocation was correlated with a 40 percentage point reduction in neighborhood poverty rate.

The shift in neighborhood characteristics experienced by CN residents is reflected in many potential measures of neighborhood opportunity beyond poverty rates. Table 3.2 displays the average pre- and post-relocation levels for several observable characteristics of existing neighborhood residents: poverty, race, education, homeownership, and family structure.[13] In order to standardize across measures, the third column presents level differences in units of county standard deviations. For example, the average proportion of neighborhood families which rent their homes fell by 1.5 times the county standard deviation of block group rental share in Shelby County, Tennessee. All five of the measures moved in a direction associated with higher block group income, and four of the five measures improved by at least half a standard

[12]Relocation location is known for 67% of the CN individuals due to incomplete or missing data. Case managers recorded location information without using the full address, leading to ambiguities in geolocating addresses.

[13]Appendix Figure B.1 plots how these characteristics changed over time.

deviation. Subjective measures of perceived safety within the CN households also improved after moving, indicating that the households themselves identified a meaningful shift in neighborhood environment post-move. The proportion of households who reported feeling "very safe" or "extremely safe" in their home or neighborhood increased by 6.4 percentage points and 6.6 percentage points respectively.[14]

The only neighborhood characteristic which did not appreciably change was racial diversity, potentially due to lack of more diverse options within the majority-minority Memphis metro area. The fact that households did not move to more diverse neighborhoods separates the setting in this paper from several previously documented examples of exogenous neighborhood shifts, and thus the results here provide suggestive evidence about the relative importance of racial diversity in other contexts.

For robustness, Table B.1 displays the shift in neighborhood observables after relocation, but at the more aggregated census tract level. The change in neighborhood observables is still striking, even at this more aggregate level of measurement. Three of the five characteristics declined by more than one-third of a standard deviation, and two declined by more than one standard deviation. Overall, evidence suggests that the Memphis Choice Neighborhoods implementation grant caused a substantial change in residents' neighborhood characteristics along several relevant dimensions.

3.3.2 Neighborhood Transitions Impacted Labor Market Outcomes

The available data allows for labor market outcomes to be evaluated in three separate dimensions: employment status, hourly wages conditional on employment, and hours worked per week conditional on employment. Column 1 of Table 3.3 displays estimates of the relocation effect for each of these outcomes using the empirical strat-

[14]Before moving, 20.8% of households report feeling very safe or extremely safe in their home, and 16.6% of households report feeling very safe or extremely safe in their neighborhood.

egy discussed in Section 3.2.2. Relocation is estimated to have caused a positive and economically relevant increase in wages of approximately $0.70 per hour. This represents a 7% increase relative to the pre-move mean at the CN site.[15] It's important to emphasize that the inclusion of individual fixed effects limits identifying variation for the relocation coefficient to be within-individual, and so any extensive margin effects will not have a mechanical impact on the wage or hours worked estimations. We do not find statistically significant impacts on either the extensive or intensive margin of employment. However, the standard errors for these outcomes are sizable, and so economically relevant impacts cannot be ruled out, particularly for the extensive margin. The point estimate on employment implies relocation decreased employment by 3.7 percentage points or 12% of the pre-move baseline. As discussed further in Section 3.3.5, any decline in employment was due to both fewer individuals transiting into employment and more individuals transitioning out of employment, compared to the non-CN benchmark. Despite this, the total employment level within the CN sample increased over the panel, given that the case management regime appears to have successfully improved employment.

Columns 2-5 of Table 3.3 demonstrate the empirical importance of including both economic and case management controls. Column 2 shows that a naive analysis of the data without any economic or case management controls would find a precisely estimated positive effect on all three outcomes. Column 3 shows that the positive effects on employment and hours in the naive model are explained by local economic and seasonal trends collinear with the timing of relocation. Column 5 shows that non-parametric controls for the duration of case management are also able to explain

[15] To give further context to this estimate, the average difference in wages for pre-move CN residents with some college as compared to residents who did not finish high school is $1.08, suggesting the increase in wages due to relocation was of a meaningful magnitude.

the positive employment and hours effects, but not the wage effect. Here it is also apparent that the duration dummies are positive and increasing for most outcomes; implying that, on average, outcomes are improving over time. However, as this specification does not control for any local economic conditions, it does not differentiate a causal impact of case management from trends in the local labor market. In fact, even in columns 1 and 5, where both case management and labor market controls are included, one should be cautious when interpreting the case management fixed effects as causal, because the plurality of individuals begin case management in one of two months corresponding to the beginning of the case management programs the CN and non-CN sites. Therefore, case management is correlated with time trends in the local labor market, and there is limited variation to separately identify both effects. However, if the case management duration fixed effects, polynomial of local outcomes, and seasonal fixed effects are sufficient to account for any unobserved events correlated with the timing of relocation, then my estimate for the impact of relocation in column 1 of Table 3.3 will be identified, as discussed in Section 3.2.2.

Is it reasonable to attribute the relocation effects estimated in Table 3.3 to the sizable change in neighborhood characteristics correlated with relocation? Table 3.4 presents evidence that rules out several other competing channels which could plausibly cause positive relocation effects and suggests that the estimated wage effect was caused by a shift in neighborhood quality. One such competing explanation is anticipatory behavior from CN households. By construction, households at the CN site only enter the panel after interacting with a case manager, and so were certainly aware of the approaching relocation. If residents changed jobs in anticipation of relocation, for example by switching to lower paid part-time work in order to leave more time to search for replacement housing, the estimated impact of relocation on wages would be positive, even in the absence of a causal impact of neighborhoods.

To minimize the risk that estimates are driven by anticipatory behavior, all analysis excludes observations for CN residents during the relocation window defined as the period after I observe the first residents moving and before the relocation deadline. However, it is possible that residents began to adjust their employment decisions before the window opened, as they would have been aware of the coming relocation before the start of this window. Column 2 of Table 3.4 tests if anticipatory behavior can explain the increase in wages by further restricting the pre-move CN sample to include only observations more than 3 months before the relocation window opened. We see that the estimated coefficient is nearly identical to the baseline, suggesting that anticipatory behavior cannot explain the estimated relocation effects on wages.

Another potential issue with attributing the estimates in Table 3.3 to neighborhood shifts is the possibility that voucher receipt itself changed residents' outcomes. For example, voucher receipt allows households a degree of choice in the quality of housing they consume, and so they may be encouraged to find higher-paying jobs to afford higher-quality housing. To test if voucher take-up itself can explain the estimated relocation effects, Table 3.4 restricts the CN sample to the 42% of CN households who chose to move into other public housing rather than take a voucher. We see that the estimated relocation effect in this restricted sample is very close to the baseline estimate, implying that voucher take-up itself cannot explain the estimated wage effects.

We cannot directly measure the causal impact of specific measures of neighborhood quality on labor market outcomes, due to the possibility of selection into neighborhoods based on these measures. However, a necessary condition for relocation effects to be attributed to changes in neighborhood quality is that neighborhood observables are correlated with improved outcomes. For example, suppose that moving itself, regardless of location, prompted households to search for higher-paying jobs. Then, one would expect to see an increase in wages after moving, but not due to improved neighborhood quality. To test if this is the case, column 4 of table 3.4 adds five neighborhood characteristics to the baseline specification. We see that the estimated coefficient on relocation decreases by more than 40% compared to the baseline and is only marginally significant. Therefore, I can reject the hypothesis that neighborhood observables are unrelated to wage improvements post-relocation. Care should be taken not to interpret the coefficient estimates in column 4 of Table 3.4 as causal, given the endogeneity of the five neighborhood regressors. Further, this analysis does not rule out the possibility that some other factors changed at the same time as the CN moves for the subset of households that were predisposed to move into higher-quality neighborhoods as measured by the five observable characteristics.

Overall, estimates indicate that relocated residents experienced a relatively large wage increase after relocation to better neighborhoods. As I only observe wage changes due to switching jobs, this directly implies that some residents found higher paying jobs after relocating.[16] Evidence suggests that this wage increase was caused by the significant improvement in neighborhood quality, not other factors associated with the move.

3.3.3 Relocation Effects Persist Over Time

The baseline empirical strategy allows for the average effects of moving to be quantified in a panel where the average relocation duration is approximately 1 year. However, these estimates do not address the possibility that the impact of relocation may be time-sensitive, and either diminish with additional exposure or become realized only after a delay. Figure 3.4 presents estimates of the relocation effect in an

[16] Appendix Table B.6 displays estimated impacts on wages and hours worked assuming linear wage increases and finds similar results.

event study framework, where the coefficients of interest are dummies for six-month bins relative to the bin before relocation. For example, the coefficient for bin 0 can be interpreted as the average impact of relocation 0-5 months after relocation, relative to the base period 1-6 months before relocation.[17] Here I again find statistically significant increases in wages, but not in hours worked per week. Interestingly, the increase in wages appears to be realized within the first 6 months of relocation, with only small increases after that point. This implies that mechanisms that one may expect to work relatively slowly, such as better social networks, may not be of first-order importance in understanding the observed wage effects. Second, the constant or slight upward trend in the wage effect over time implies that these effects were not transitory up to two years after relocation.[18] Finally, in support of previous results suggesting minimal anticipatory behavior, I find no evidence of a significant pre-relocation trend in either hourly wages or hours worked per week.

Panel (a) of Figure 3.5 plots a similar event study estimation using employment as the outcome of interest. As with the baseline specification, I do not find statistically significant impacts on employment. However, the point estimates do show a downward trend in employment. Panels (b) and (c) break down employment into full-time (≥ 35 hours per week) and part-time employment. Here I see a statistically significant decline in full-time employment 1 year after moving. Over this same time, I see a smaller and statistically insignificant increase in part-time employment. Given that I see a decline in full-time employment and an increase in wages, it is possible that individuals transitioned out of full-time employment into part-time to prevent

[17]As not every individual in the panel begins case management or moves away from the site at the same time, the beginning and ending bins are extended further than 6 months.

[18]If I assume that the impact of case management duration is constant after two years, I can then extend the panel to estimate relocation effects up to three years after relocation. We find positive wage effects on the order of $1 per hour up to three years after relocation under this assumption.

their income from disqualifying their household from receiving need-based government transfers. This would imply the correlation between wage and full-time employment decreased after relocation. However, the opposite is true in the data, which shows an increase in the correlation coefficient from 0.012 to 0.147, implying that individuals did not use the intensive margin of employment to adjust their total earnings in response to a wage increase. Finally, panel (d) displays the impact of relocation on the proportion of households with at least one adult working full-time. While the standard errors for this outcome are large, the point estimates suggest the proportion of households with at least one adult working full-time remained unchanged after relocation, implying that any decline in full-time employment was experienced by households with two or more full-time workers. The estimated impacts of relocation on employment at the extensive and intensive margin presented in Figure 3.5 imply full-time employment decreased over time after moving and hint at potentially interesting substitution patterns within households. However, data limitations leave the analysis under powered, with large confidence intervals, especially at the household level.

3.3.4 Heterogeneous Effects

Given wages increased after moving and more ambiguous evidence suggests employment may have decreased, it is not clear if these effects were evenly distributed across the population. Labor market outcomes differ with education and over the life cycle, so one may expect different demographic groups to have significantly heterogeneous reactions to new neighborhoods. Table 3.5 displays the estimated relocation coefficients from an interaction model, which tests for heterogeneous reactions to relocation by education and age. The full empirical specification is:

$$Y_{i,t} = \theta + \beta_1 R_{i,t} demo_i + \beta_2 R_{i,t}(1 - demo_i) + \gamma_i + \psi_{CM_{i,t}} + \kappa_{s_{i,t}} + g(\boldsymbol{E_t}) + \phi_{i,t}$$

where $demo_i$ is a dummy variable for the specific demographic group in question. Notice that the level effect of $demo_i$ is subsumed by the individual fixed effects γ_i. Since these characteristics were measured before relocation, they are exogenous to the timing of relocation, and so one may interpret the estimated coefficients in the same way as in the baseline specification.[19]

The first panel of Table 3.5 demonstrates that individuals with at least a high school degree had statistically different employment effects as compared to their less educated peers. High school graduates are estimated to have experienced a near-zero effect on employment, while those without a high school degree experienced relatively large and significant declines after moving. While the difference in wage effects between the educational groups is not statistically significant, it is unlikely that more educated residents experienced a lower wage increase than less educated residents (p=0.112). The second panel of Table 3.5 shows that relocation effects differed significantly by age as well. Individuals 45 or younger at the start of the panel experienced increases in wages and hours worked after relocating while avoiding the negative impact on employment that older residents experienced. Individuals older than 45 are estimated to have had near-zero changes in wages or hours worked while experiencing large declines in employment. Together these results suggest that changes in neighborhood quality unambiguously improved the labor market outcomes for younger or more educated residents. Further, while the average impact of relocation on employment is somewhat ambiguous, the effects become much more clear after disaggregating by demographic group. Both less educated and older residents are estimated to have

[19] In order to attribute differences in relocation effects across demographic groups to different neighborhood effects, one must rule out the possibility that differential neighborhood sorting patterns may have caused any difference in outcomes. Table B.2 displays the average standard deviation shift in neighborhood characteristics for the demographic groups in question.

experienced large declines in employment, while younger or more educated residents were estimated to have faced a fairly precisely measured zero impact on employment.

This evidence suggests that targeting neighborhood mobility initiatives towards specific demographic groups (younger, or more educated individuals) living in low opportunity public housing neighborhoods may be able to improve neighborhood match quality for some households while minimizing the disruption to households with little to gain from moving. It may not be surprising that younger and more educated individuals experienced better outcomes post-relocation, as they potentially had the most to gain from employment. On average, individuals with some college earn \$1.12 more than those without a high school degree post-relocation. Further, younger adults have more time to benefit from the accumulation of work experience and the public housing population is disproportionately likely to work in manually intensive occupations such as landscaping, food service, and maintenance where younger individuals may have a comparative advantage. The data supports the claim that younger adults have a higher earning potential, with individuals under 45 years old (at baseline) earning \$0.27 more than older adults post-relocation.

3.3.5 Wage and Employment Transitions

Section 3.3.2 establishes that on average non-voluntary moves into better neighborhoods increased wages while finding more ambiguous results for employment. Section 3.3.3 establishes that full-time employment decreased significantly after moving. This section will expand upon these results by evaluating the wage and employment transitions CN residents experienced over the move. We see that the maximum wage at the CN site increases substantially after the move, as well as a thickening of the wage distribution in the \$12+ range. While the available data on income from government transfers is extremely limited, it is consistent with the conclusion that economic-self sufficiency increased after the move. Finally, while employment increased at both the CN and non-CN site, I estimate a negative employment effect because both fewer unemployed CN residents transitioned into employment and more employed CN residents transitioned out of employment compared to their non-CN counterparts. This suggests that the disruption cost of non-voluntary moves was similar in magnitude across pre-move employment status.

Figure 3.3 displays kernel density estimates for wages of CN residents with jobs both 4-9 months before relocation and 18-24 months after relocation. We see that post-relocation, the right tail of the distribution is significantly extended, implying the maximum wage increased quite substantially. A Kolmogorov Smirnov test rejects equality of the two distributions (p=0.018). While the extreme of the post-relocation right tail is very sparsely populated, it may still have policy relevance because a full-time $15+ per hour job would put a family of four comfortably above the 2018 poverty line in Tennessee, and thus represents a transition towards economic self-sufficiency.[20] Further, there is also a thicker distribution in the $12+ wage range after relocation, implying that the estimated impacts of wage increases were not due to a few outlier observations transitioning into $20+ per hour jobs, but rather a distributional shift experienced by a number of individuals.

If the wage increase due to relocation did, in fact, increase economic self sufficiency then one would expect the amount of government transfers individuals receive to decrease after relocation. Only limited data on the amount of government transfers individuals received is available. However, Table B.3 displays the average within-individual change in transfers over the panel, for the subset of individuals with multiple income records. For the 50 individuals from the CN site with multiple income reports, I find an average decline of $69.4 per month in the combined

[20]The 2018 poverty line in Tennessee for a family of four was $24,860

amount of government transfers from several common programs.[21] A similar within-individual decline for residents of the non-CN sites shows only a \$18.5 decline. While the small sample sizes of this analysis make it difficult to put too much weight on any of the point estimates, the available data does support the conclusion that relocation increased economic self-sufficiency, at least for some households.

The baseline estimates of the employment effect in Table 3.3 is negative but imprecisely measured. To further investigate the employment transitions individuals experienced after moving, Table B.4 presents the empirical employment transition matrices for both CN and non-CN residents.[22] Again, the estimates are not very precisely estimated, as evidenced by the weak statistical differences across the two populations. However, the relocated residents at the CN site were 10 percentage points less likely to transition into employment and 15 percentage points more likely to transition out of employment compared to residents at the non-CN site after an equivalent duration of case management. While these estimates do not control for economic trends or seasonal effects, they do suggest that relocation causes a significant disruption to employment for some residents. However, given the relative samples of baseline employed an unemployed, the employment rate significantly increased in both the CN and non-CN samples over time.

[21]The specific programs are Temporary Assistance for Needy Families (TANF), Unemployment Insurance, Special Supplemental Nutrition Program for Women, Infants, and Children (WIC), and Food Stamps

[22]The sample used to compute Table B.4 is restricted to CN individuals who began case management prior to moving and non-CN individuals who began case management early enough to have completed at least 18 months.

3.4 Mechanisms

A large number of factors could potentially explain why moving into better neighborhoods increased wages for the Memphis CN residents. In this section, I evaluate two separate mechanisms, increased proximity to job opportunities and case management services helping ease disruption costs. Despite being a compelling story to explain why neighborhoods may be linked to wages, I find no evidence to support the conclusion that proximity to job opportunities played in role in this setting. However, I do find evidence that case manager assistance may be key to understanding these results in a broader context.

3.4.1 Proximity to Job Opportunities Did Not Change After Moving

One way relocating to better neighborhoods has been hypothesized to benefit low-income households is through access to better or more abundant job opportunities. We may expect low-income households to be more geographically constrained when searching for job opportunities because they are much less likely to own a car, and thus must rely on alternative methods of transportation.[23] Relocating nearer to more economically prosperous areas of a city may allow individuals the opportunity to pursue jobs that were inaccessible prior to relocation. However, evidence suggests that after relocation, the CN residents did not find jobs that were inaccessible due to physical distance.

The first panel of Table 3.6 displays the mean linear distance between an individual's residence and employment location, where location is measured as the centroid of the census block group for the residence or employer. The unit of observation for this table is a job held by a resident of the CN site. The upper left quadrant

[23] Only 31% of the individuals in the panel own a car when they were first observed.

of the table states that the mean distance from the place of employment to the CN site, for jobs that began before relocation, was 6.94 miles, with a median distance of 7.74 miles. Comparing this quadrant to the upper right quadrant in the first panel, I find that the linear distance from the CN site to jobs that residents started after moving (and thus which offered higher wages) was 7.54 miles. This is not statistically different than the distance to jobs started before moving (p=0.166). The difference in median distance is even less dramatic, at 0.38 miles.

It is possible that while better jobs were not physically too far away from the old site, they were inaccessible due to commute times. The second panel of table 3.6 tests this by displaying mean commute times measured using Google Maps API for a Wednesday bus trip arriving to work at 8:00 am in mid-October 2019.[24] We see that the average job which began after relocation was no further away via bus from the CN site than the average job started pre-move at 60.9 minutes and 59.5 minutes respectively. I find that jobs started after relocation were on average 10 minutes further away from home via bus than jobs started pre-move (p=0.034).[25] This could be due to relocation increasing individuals' willingness to commute, or simply that the best available housing was in areas with less public transportation coverage.

Further, in both panels of Table 3.6, I find that residents who had jobs at the time of relocation did not choose to move significantly closer to their workplace, implying that the distance/commute to work did not impose a large cost on those with jobs at the time of relocation.[26]

[24]This approximation is required because Google does not allow for historic route calculations, and thus I assume bus routes have been relatively stable over time in Memphis.

[25]This can be seen by comparing the upper left quadrant to the lower right quadrant in the second panel of Table 3.6.

[26]This can be seen by comparing the upper left quadrant to the lower left quadrant in both panels of Table 3.6. Similarly, I do not observe differences for jobs quit before moving, implying the

If proximity to better jobs was of first-order importance in explaining the estimated wage effects, one would have expected to find noticeable differences in commute distances over relocation. However, it does not appear as if the higher wage post-move jobs were inaccessible from individuals' pre-move residences.

3.4.2 Case Management Prevent Job Loss over the Move

It is worth asking why I find sizable relocation effects in this setting, while previous studies have not consistently found similar results. One possibility is the magnitude of neighborhood shift experienced by CN residents, which was significantly larger than in other studies of exogenous shocks to neighborhood quality. The interim MTO report notes both baseline census tract poverty rate and treatment on the treated changes in neighborhood poverty rate of approximately 50% of the level and change experienced by CN residents.[27] Another possibility is that, in this setting, relocation happened over the course of an intensive case management regime. It could be in the absence of these services, residents would have faced larger disruption costs or not been able to capitalize on the potential gains in their new neighborhoods.

Throughout the analysis presented above, changes in outcomes due to relocation have been separately identified from the impact of case management through the use of case management duration fixed effects. However, it is possible that in the absence of case management, residents would have been unable to capitalize on the benefits of relocation, *i.e.* it's possible that case management had a positive interaction with relocation in my empirical model. This interaction could be directly estimated if a random subset of CN households had been left to relocate without receiving case management service. While this is not the case (all observable households received case

relocation did not cause some households to quit a job due to distance from new residence.

[27] See Orr *et al.* (2003) for details.

management services), as a second-best approach I can use quasi-random variation in the intensity of case management to investigate the relevance of such an interaction.

As discussed in section 3.1.2, the NPO calculated quarterly risk scores for each adult, in order to assist case managers in targeting high need households. Because scores below 100 were considered relatively safe, case managers were suggested to focus more of their efforts on residents with scores above 100. To test if this cutoff in the designation of "safe" scores generated exogenous variation in the amount of attention case managers gave to residents, Figure 3.6 plots the number of outside service providers case managers referred individuals to contact (service links) against the distance of the individual's cumulative maximum risk score to the cutoff of 100. The sample is restricted to a relatively narrow bandwidth around the cutoff so that 20% of the observations in the panel fall above and 20% below the cutoff. The 20% bandwidths were selected to keep observables and unobservables above and below the cutoff as similar as possible while maintaining a reasonable sample size, but the robustness of the results to other bandwidths is explored.[28] We see a distinct jump in the average number of service links at the cutoff, highlighted by the discontinuity in the polynomials estimated separately above and below the cutoff. Completed service links capture only a piece of the case management process, but a discontinuity in part of the overall package of case management suggests that individuals with scores over 100 may have received more case management attention in multiple domains of the case management strategy.

While the risk scores are in principle continuous, in practice the NPO's scoring algorithm resulted in a very lumpy distribution of scores. Therefore, making a standard McCrary test infeasible. However, for multiple reasons, manipulation is unlikely to play a significant role in this setting. First, residents were not made aware of

[28]See Figure 3.8 to see how the regression discontinuity analysis changes with different bandwidths

the scoring system, and so would not have had a reason to manipulate their score. Second, the risk scores were a relatively new tool rolled out by the NPO, and so case managers were not evaluated on the risk scores of the residents they were assigned. Therefore, there was little reason for case managers to manipulate the scores of the households in their portfolios. And finally, scores were constructed automatically from the regular reports case managers submitted, and thus the case managers had both limited understanding of the actual scoring algorithm and would have had to lie about a resident's responses to verbal questionnaires if they wanted to manipulate scores.

Even if risk scores were not manipulated around the cutoff, it may be the case that the distribution of scores around the cutoff changed after relocation. Figure 3.7 plots the distribution of risk scores before and after relocation. We see that the lumpy distribution of scores is persistent after relocation, suggesting that any selection on unobservables across the cutoff is consistent pre- and post-move, and so comparing the difference in cutoff effects pre- and post-move is consistent.

In order for this discontinuity to be exogenous, it must be that individuals just above and just below the cutoff have similar observables, and thus it can be assumed similar unobservables. Table 3.7 tests the balance of observables at the start of the panel for individuals on either side of the risk score cutoff of 100, within the same 20% bandwidths used in Figure 3.6. Here I find that individuals had similar observable characteristics at the beginning of the panel, supporting the assumption that there is little difference in unobservables within a relatively narrow band of the cutoff.

Exploiting the discontinuity in case management intensity around risk score 100, I can test the interaction of case management intensity and relocation by adapting my baseline empirical strategy to the standard regression discontinuity model. In practice, this means including separate risk score polynomials above and below the cutoff, restricting the sample to within a narrow band around the cutoff, and including a dummy for above the cutoff and an interaction of the relocation dummy and above the cutoff dummy. The full empirical specification is:

$$
\begin{aligned}
Y_{i,t} = \theta &+ \beta_1 R_{i,t} + \beta_2 R_{i,t} risk_{i,t}^{100} + \beta_3 risk_{i,t}^{100} + h(risk_{i,t}) + h'(risk_{i,t}) risk_{i,t}^{100} \\
&+ \gamma_i + \psi_{CM_{i,t}} + \kappa_{s_{i,t}} + g(\boldsymbol{E_t}) + \phi_{i,t}
\end{aligned}
$$

where $risk_{i,t}$ is the cumulative maximum risk score an individual has been assigned by month t, $h(\cdot)$ and $h'(\cdot)$ are second-degree polynomials, and $risk_{i,t}^{100}$ is a dummy equal to 1 if $risk_{i,t} \geq 100$. If the interaction of high-risk score and relocation is estimated as positive, that would be evidence of the hypothesized complementarity between relocation and receiving more intensive case management. The estimation is carried out on the sample with risk score between 43 and 241, in order to minimize the effect of unobservables correlated with risk score.[29] Figure B.2 provides a visual representation of the empirical strategy, without controls. Table 3.8 displays the relevant estimated coefficients from this regression discontinuity model applied to employment, wages, and hours worked. The positive coefficient on the interaction term in column 3 demonstrates that having a cumulative maximum risk score above 100 mitigated nearly 75% of the negative employment effect from relocation, consistent with my hypothesis. Figure 3.8 demonstrates that estimates of the interaction of high-risk score (above 100) and relocation are relatively stable for other bandwidths, and statistically significant for all bandwidths greater than 20%. The marginally significant coefficients in columns 1 and 2 suggest that, while case management helps prevent unemployment after relocation, the marginal job obtained due to extra case

[29] These are the same bandwidths used in Table 3.7 and Figure 3.6, and are selected such that 20% of the observations in the panel are above and 20% are below the cutoff.

management was slightly lower wage or more likely to be part-time relative to an individual's pre-relocation job(s). While the case management program implemented in this setting is well designed and has been refined over more than a decade, there is no reason to assume that it is the optimal strategy to prevent negative disruption effects. Future work to determine which aspects of case management support are crucial for households to thrive in new environments may suggest even more effective methods to dampen the disruption costs from involuntary moves for low-income households.

3.4.3 Robustness

This section will review several robustness checks for key estimates in the paper.

First, Table B.5 tests if the heterogeneity in household composition across the CN and non-CN sites as documented in Table 3.1 can explain the estimated relocation effects by restricting the sample to only households with 1 or 2 adults. In this restricted sample household composition is balanced across sites. We see that estimates are very close to those in Column 1 of Table 3.3, suggesting that the differences in household composition presumably caused by different housing unit sizes across sites cannot explain the estimated impact of moving to better neighborhoods.

Next, Table B.6 tests if estimated relocation effects are a mechanical result of the fact that only starting wages are known for all jobs. The second column demonstrates that moving made it more likely that the average individual started a new job in a given month. Then if on average, wages were increasing over time, the increase in job finding rates in the CN sample post-move may lead us to falsely infer that wages increased after moving. Columns 1 of Table B.6 tests this theory by applying the baseline empirical specification to a constructed wage variable that takes into account either the observed or average monthly wage growth within the panel. Specifically, if a job ended during the period of observation, the ending wage is known, and so a linear growth rate is applied to the interim months the individual was employed. If a job has not ended by the end of the observation period, the ending wage is not known, and so the average monthly wage growth rate within the set of ended jobs is applied to the observed starting wage. This adjusted wage may still mask within job wage growth if wage growth does not happen relatively smoothly over the duration of employment, or if jobs are ended based on the rate of wage growth. However, the fact that the point estimate of the relocation effect only decreases by $0.02 strongly suggests that the fact I do not observe with-job wage changes cannot explain the estimates.

3.5 Conclusion

This paper presents evidence that neighborhoods in adulthood can have an economically relevant impact on labor market outcomes. We find that adults required to relocate from a severely disadvantaged neighborhood in Memphis Tennessee earned more per hour after relocating. This relocation was found to substantially improve the neighborhood characteristics of relocated households, which provides causal evidence that neighborhoods have an economically relevant impact on adults. By comparing residents just above and below a cutoff used to evaluate a household's case management needs, I find that additional attention from a case manager helped mitigate some of the negative employment effects of relocation, suggesting a degree of complementarity exists between individual case management and relocating out of low-opportunity public housing.

The benefits of relocation were not evenly distributed across the relocated population. Individuals less than 45 years old or with at least a high school degree saw non-trivial wage increases after moving, without experiencing the negative employment shock that their peers faced. These two demographic groups also earned significantly more post-move, suggesting that moves out of low-opportunity public housing neighborhoods unambiguously benefited individuals with a higher earning potential.

Finally, after relocation residents did not find jobs that were further away, in distance or time, from their initial neighborhood than the average pre-relocation job, indicating that physical barriers to better job markets are unlikely to explain the estimated neighborhood effects.

3.6 Tables

Table 3.1: Sample Balance

	Non-CN	CN	P-Value
Number of Individuals	197	475	
Age	36.15	36.71	0.609
Mean Dependent Age	10.81	11.36	0.288
At Least HS Degree	0.68	0.65	0.493
College Degree	0.04	0.03	0.794
Female if HoH	0.96	0.94	0.543
Employmed	0.30	0.25	0.212
Hours Worked per Week if Employed	32.73	28.83	0.004
Hourly Wage if Employed	9.98	9.46	0.238
Bedrooms in Housing Unit†	2.93	2.25	
Number of Dependents in HH	2.12	1.43	0.000
Number of Adults in HH	1.72	1.49	0.001
Female	0.88	0.83	0.100

† Average taken over all housing units at a site, not at individual level. Number of rooms known for two of the three non-CN sites.

Notes: Table displays mean observables in the first month of case management for all individuals in the analysis sample. CN is short for Choice Neighborhoods, and designates the housing site which was redeveloped. All CN residents were required to move away from the site.

Table 3.2: Shifts in Neighborhood Characteristics Over Move

	Pre-Move Level	Post-Move Level	Diff in County SD
Proportion of Neighborhood:			
Below Poverty Line	0.76	0.35	-2.09
Racial Minority	0.93	0.92	-0.04
HS Degree or Less	0.72	0.57	-0.73
Renting Home	0.95	0.55	-1.52
Female HoH	0.41	0.32	-0.56

Notes: The first two columns report average pre- and post-relocation block group neighborhood characteristics for CN households. Characteristics data from ACS 5 year estimates, 2012-2016. Column three normalizes the difference between columns one and two by the county standard deviation across block groups.

Table 3.3: Relocation Effect Estimates

	(1)	(2)	(3)	(4)	(5)
Employment					
Relocation	-0.037	0.070***	0.000	-0.045***	-0.037
	(0.034)	(0.007)	(0.018)	(0.010)	(0.046)
CM Duration = 2mo	0.005			0.007	0.005
	(0.010)			(0.015)	(0.014)
CM Duration = 12mo	0.059*			0.098***	0.059
	(0.032)			(0.014)	(0.037)
CM Duration = 23mo	0.071			0.145***	0.071
	(0.048)			(0.017)	(0.058)
Pre-Move Mean	0.297	0.297	0.297	0.297	0.297
Avg. Months Moved	13.1	13.1	13.1	13.1	13.1
N	11,969	11,969	11,969	11,969	11,969
Hourly Wage if Employed					
Relocation	0.69***	0.46***	0.48***	0.44***	0.69**
	(0.26)	(0.07)	(0.08)	(0.08)	(0.21)
CM Duration = 2mo	0.02			0.13	0.02
	(0.10)			(0.13)	(0.13)
CM Duration = 12mo	-0.41*			0.10	-0.41**
	(0.24)			(0.10)	(0.09)
CM Duration = 23mo	-0.91*			0.13	-0.91*
	(0.49)			(0.13)	(0.37)
Pre-Move Mean	9.36	9.36	9.36	9.36	9.36
Avg. Months Moved	12.0	12.0	12.0	12.0	12.0
N	3,972	3,972	3,972	3,972	3,972
Hours Worked per Week if Employed					
Relocation	1.07	0.68***	-0.73	-0.09	1.07*
	(0.71)	(0.21)	(0.80)	(0.27)	(0.39)
CM Duration = 2mo	0.25			0.13	0.25
	(0.17)			(0.32)	(0.15)
CM Duration = 12mo	0.30			0.07	0.30
	(0.53)			(0.32)	(0.28)
CM Duration = 23mo	1.43*			1.19***	1.43**
	(0.74)			(0.40)	(0.25)
Pre-Move Mean	27.87	27.87	27.87	27.87	27.87
Avg. Months Moved	12.0	12.0	12.0	12.0	12.0
N	3,972	3,972	3,972	3,972	3,972
Indiviudal FE	X	X	X	X	X
Econ Poly + Season FE	X		X		X
Non-Parametric CM Effects	X			X	X
SE Clustered on Individual	X				
SE Clustered on Site and Date					X

Notes: Column 1 corresponds to the preferred specification. CM stands for case management. Standard errors reported in parentheses. If not clustered, heteroscedasticity consistent standard errors reported. Non-parametric case management controls include dummies for case management duration, only three case management coefficients are displayed for reference. Economic controls include polynomial for 12 month rolling average metro wage, metro hours worked, or metro unemployment. (*= p<0.10, **= p<0.05, ***= p<0.01)

Table 3.4: Relocation Wage Effects and Neighborhood Shifts

	(1) Baseline Specification	(2) Only Early Pre-Reloc. Obs.	(3) Only Public Housing Residents	(4) Include Nbhood. Controls
Relocation	0.69***	0.73***	0.75**	0.41*
	(0.26)	(0.27)	(0.29)	(0.22)
Tract Poverty				-0.80
				(2.63)
Tract Minority				2.67
				(2.16)
Tract No College				-9.61
				(5.92)
Tract Rent				3.09
				(3.08)
Tract Fem. HoH				9.99
				(6.39)
N	3,972	3,700	2,058	3,972

Notes: Dependent variable is hourly wage conditional on working. All columns include controls for case management duration, a polynomial of 12 month rolling average wages in metro, seasonal fixed effects, and individual fixed effects. Standard errors in parenthesis are clustered at the individual level. Column (1) represents the preferred specification. Column (2) drops observations from the CN site closer than 4 months before the first relocations away from the site, and suggests that the results in column (1) are not driven by anticipation of relocation. Column (3) restricts the CN sample to only individuals who chose to relocate to another public housing site, rather than use a housing voucher. These individuals should have experienced minimal changes to their housing budget or net transfers due to the relocation. Column (4) includes the indicated census tract observables as non-causal controls and demonstrates that the results in column (1) can be partially explained by changes in a selected set of neighborhood observables. (*= $p<0.10$, **= $p<0.05$, ***= $p<0.01$)

Table 3.5: Heterogeneous Relocation Effects

	If Employed		
	Hourly Wage	Hours Worked Per Week	Employed
Relocated X Education $\geq$ HS [60.5% of indiv.]	0.87** (0.38)	0.76 (0.81)	0.010 (0.040)
Relocated X Education $<$ HS [39.5% of indiv.]	0.41** (0.18)	1.58 (1.03)	-0.102*** (0.036)
Relocated X Age $\leq$ 45 [69.8% of indiv.]	0.94*** (0.33)	1.58* (0.88)	-0.002 (0.039)
Relocated X Age $>$ 45 [30.2% of indiv.]	0.09 (0.20)	-0.11 (0.69)	-0.094** (0.037)

Notes: Each panel displays estimated coefficients from an interaction of relocation and the indicated demographic group. Demographics measured at baseline observation. Proportion CN residents which belong to demographic group reported in brackets. Each regression includes individual fixed effects, dummies for duration of case management, seasonal fixed effects, and polynomial for corresponding average outcomes across metro (wage, hours worked, and unemployment respectively). Standard errors clustered at the individual level. (*= p<0.10, **= p<0.05, ***= p<0.01)

Table 3.6: Distance and Time to Work

		Jobs Started Before Move	Jobs Started After Move
Linear Distance (Miles)	Pre-Move Residence	6.94 (0.55) [7.74]	7.54 (0.33) [8.12]
	Post-Move Residence	6.96 (0.68) [7.22]	7.03 (0.37) [6.83]
Commute Time Via Bus (Minutes)	Pre-Move Residence	59.5 (3.6) [58.9]	60.9 (2.4) [62.7]
	Post-Move Residence	78.0 (14.3) [67.0]	71.7 (4.8) [75.6]

Notes: The first panel displays mean linear distance, in miles, from place of employment to residence pre/post-relocation. Means are calculated separately for jobs accepted pre/post-relocation. Standard errors are displayed in parentheses, medians in brackets. The second-panel displays mean commute times, via bus, from place of employment to residence pre/post-relocation. Times are calculated via google maps using google's standard traffic model to arrive at 8:00 am on a Wednesday in October 2019. All locations are to the centroid of the block group.

Table 3.7: Baseline Characteristics By Risk Score

	Ever $43 \leq \text{Risk} \leq 100$	Ever $100 < \text{Risk} \leq 241$	P-Value
Number of Indiviudals	80	96	
Age	40.01	41.74	0.328
Female	0.93	0.96	0.344
Number of Dependents in HH	1.34	1.44	0.667
Number of Adults in HH	1.34	1.40	0.563
Mean Dependent Age	10.40	10.31	0.916
At Least HS Degree	0.50	0.61	0.129
College Degree	0.04	0.03	0.821

Sample includes adults who are ever observed with a risk score within the indicated bandwidths. Bandwidths were selected so that 20% of the observations in the panel fall above and below the cutoff. N= 5 individuals were observed with both a high and a low-risk score. Only pre-move baseline observations included.

Table 3.8: Risk Score Interaction

	If Working		
	Hourly Wage	Hours Worked per Week	Employed
Relocated × Above Cutoff	-0.61*	-3.26*	0.17***
	(0.33)	(1.67)	(0.05)
Relocated	0.84**	3.62	-0.23**
	(0.41)	(2.26)	(0.12)
N	1,456	1,456	4,176

Notes: Each regression includes a second degree polynomial for cumulative maximum risk score both above and below the risk score cutoff, a dummy for cumulative maximum risk score above the cutoff, individual fixed effects, dummies for duration of case management, seasonal fixed effects, and polynomial for corresponding average outcomes across metro (unemployment, hours worked, and wage respectively). Standard errors clustered at the individual level. Sample restricted to observations with a cumulative maximum risk score between 43 and 241, so that 20% of observations are above and below the cutoff. (*= $p<0.10$, **= $p<0.05$, ***= $p<0.01$)

3.7 Figures

Figure 3.1: Distribution of Moving Dates

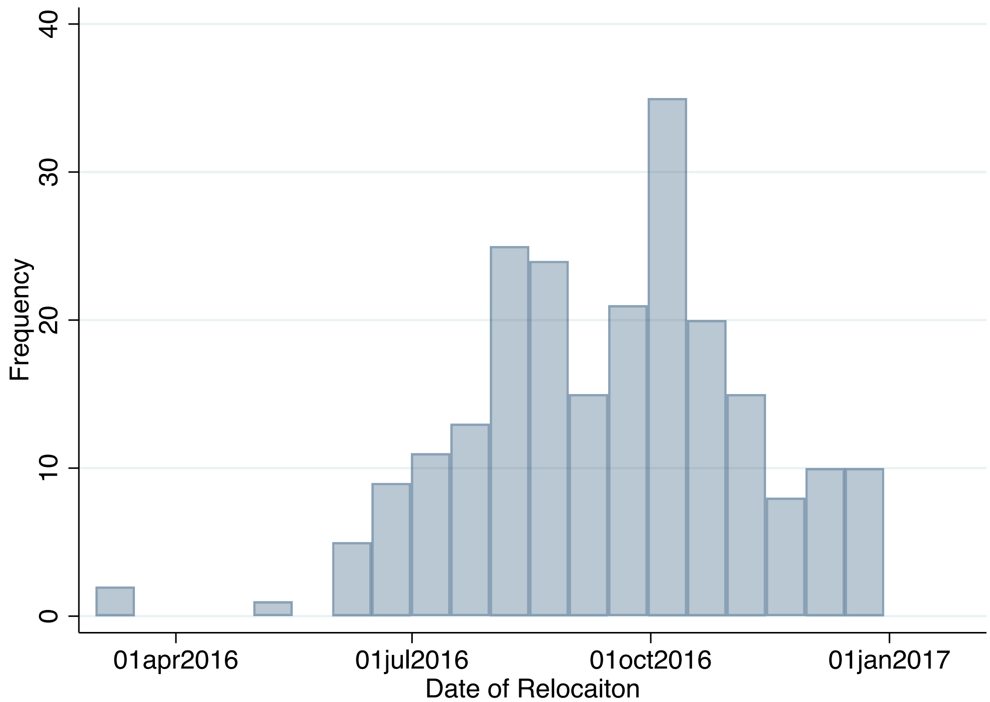

Notes: Figure displays the distribution of known moving date for households at the CN site. Exact relocation dates know for 320 individuals (224 households, 67% of the sample).

Figure 3.2: Neighborhood Poverty and Relocation

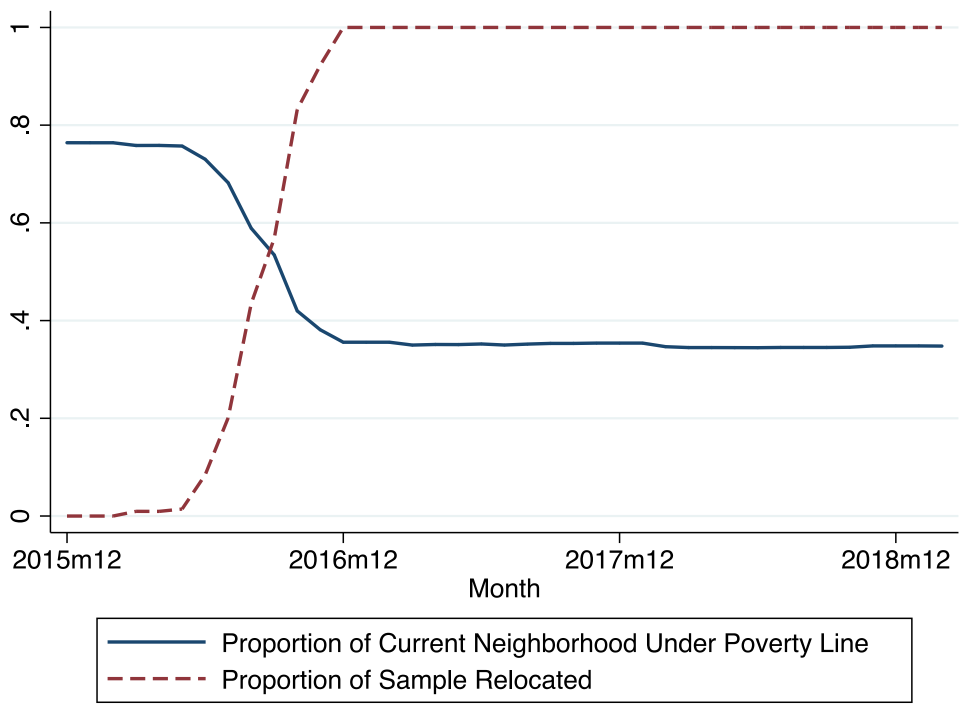

Notes: The solid blue line displays mean census block group poverty level within the sample of CN households with known relocation dates, by month. The proportion of the same sample which has moved away from the initial public housing site is shown as a dashed red line.

Figure 3.3: Wage Distribution Pre- and Post-Move

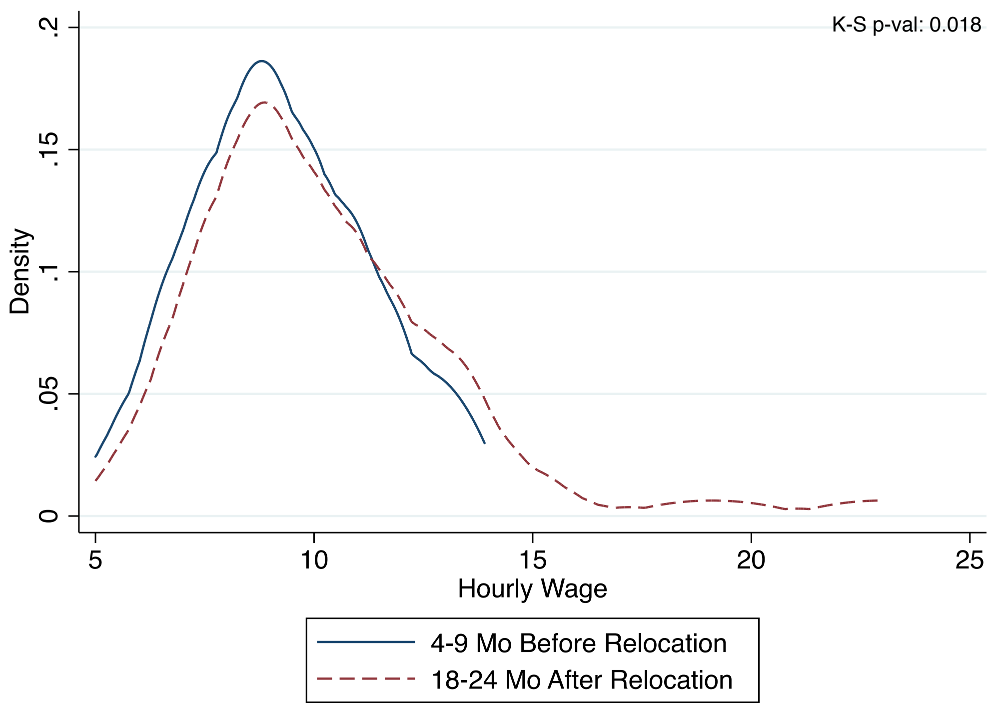

Notes: Figure plots kernel density estimates for hourly wages before and after relocation for a balanced panel of individuals at the CN site. P-value from Kolmogorov Smirnov test for equality of the two distributions displayed in the top right.

Figure 3.4: Wage and Hours Effects over Time

a: Hourly Wage if Working

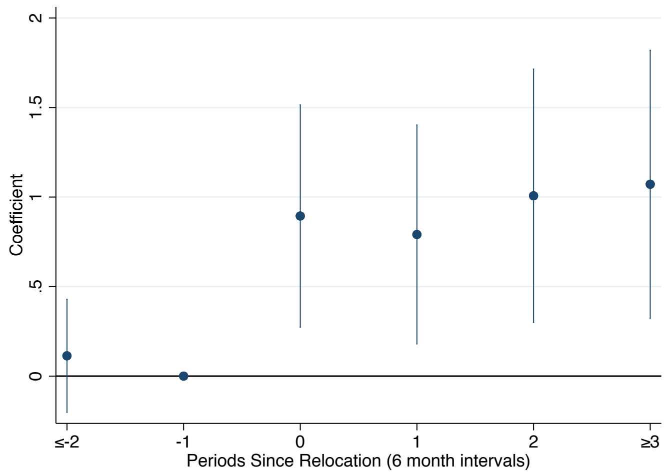

b: Hours Worked per Week if Working

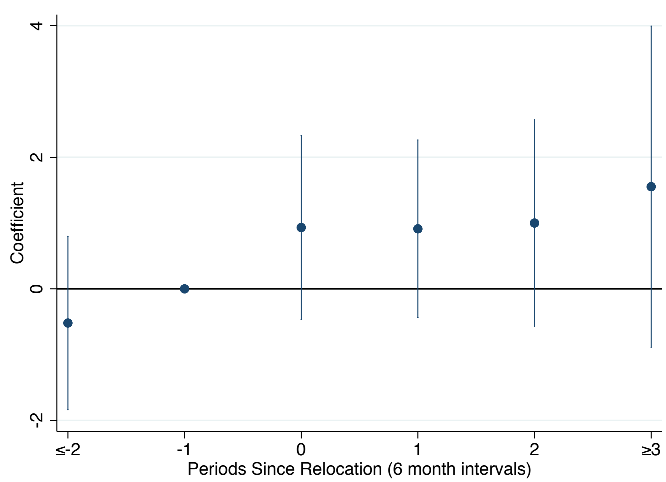

Notes: Each panel displays the coefficients from a regression of the indicated outcome onto a vector of dummy variables corresponding the number of quarters since (or until) an individual moves away from their initial housing site. Bars represent 95% confidence intervals. Each regression includes individual fixed effects, dummies for duration of case management, seasonal fixed effects, and polynomial for corresponding average outcomes across metro (wage and hours worked respectively). Standard errors clustered at the individual level.

Figure 3.5: Employment Effects Over Time

a: Employment

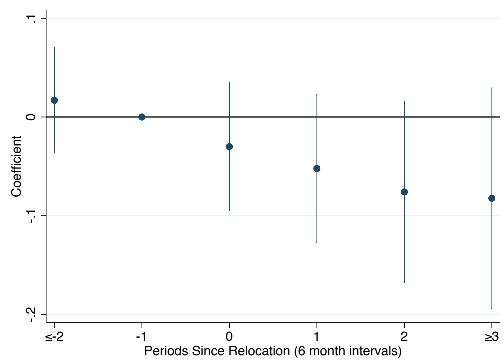

b: Part-Time Employment

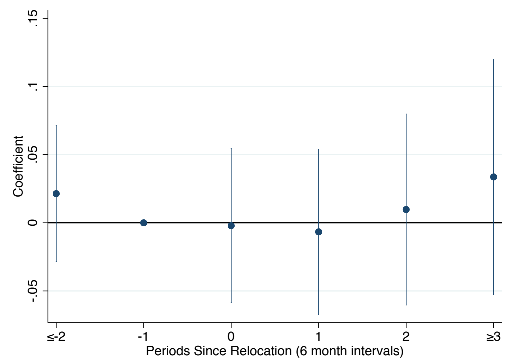

c: Full-Time Employment

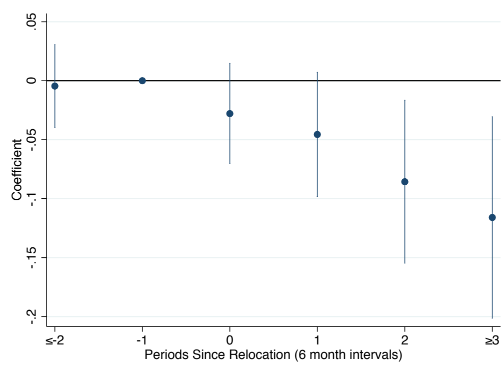

d: Any Family Member Full-Time Employed

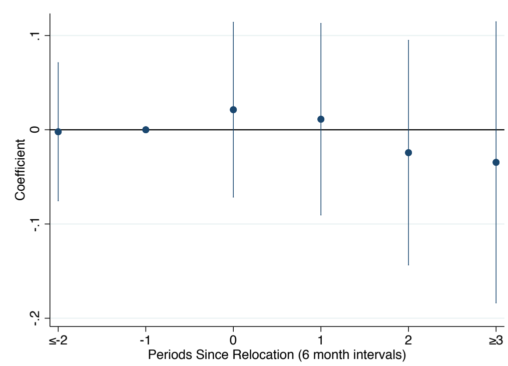

Notes: Each panel displays the coefficients from a regression of the indicated outcome onto a vector of dummy variables corresponding the number of quarters since (or until) an individual moves away from their initial housing site. Bars represent 95% confidence intervals. Each regression includes individual fixed effects, dummies for duration of case management, seasonal fixed effects, and polynomial for average unemployment across the metro. Standard errors clustered at the individual level.

Figure 3.6: Discontinuity in Case Manager Attention

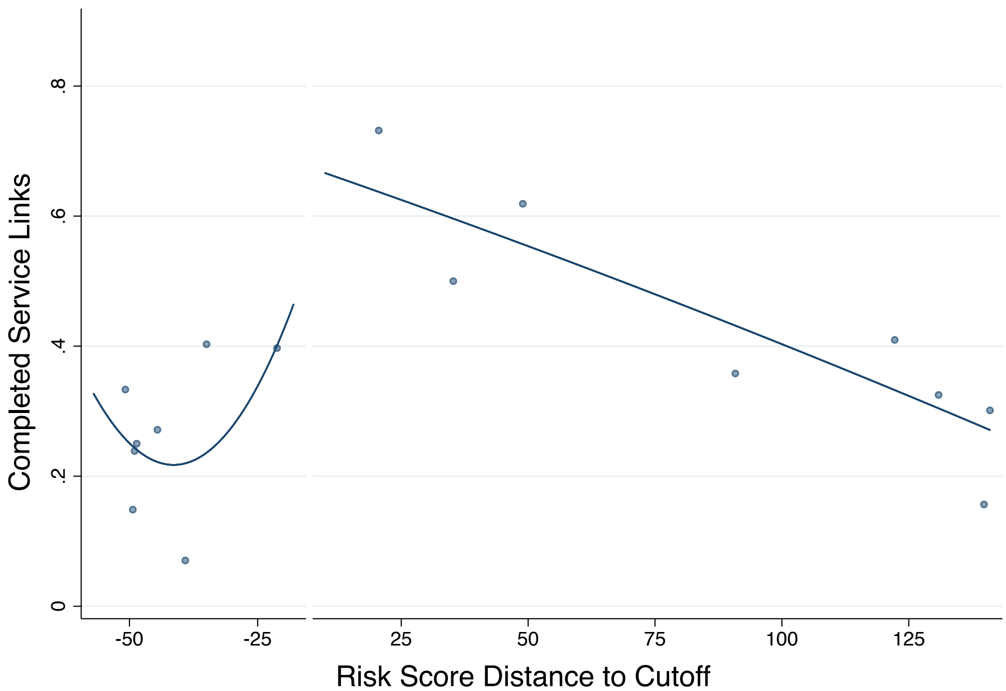

Notes: Figure plots the average number of service linkages individuals were referred to by their case manager against the distance of their cumulative maximum risk score to the cutoff of 100 in 16 evenly sized bins. Service links represent a part of the overall case management strategy, and so are an incomplete proxy for the amount of attention or effort a case manager spends on an individual. Also plotted are second-degree polynomials separately fit to the data above and below the cutoff. Score bandwidth selected so that 20% of the sample is above and below the cutoff.

Figure 3.7: Risk Score Distribution Pre- and Post-Relocation

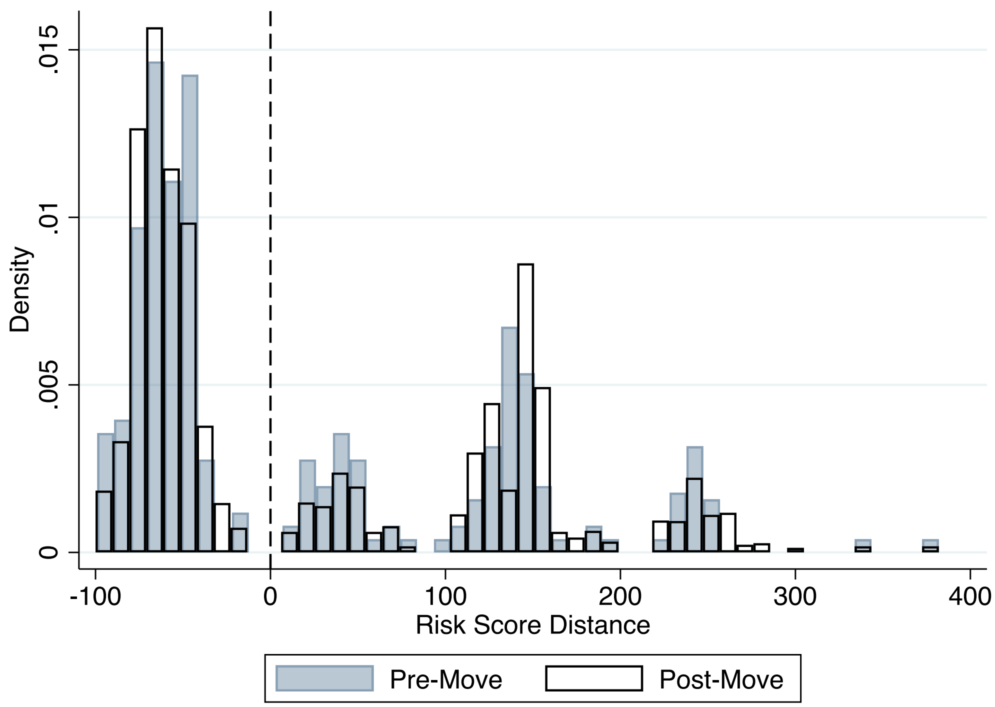

Notes: Figure plots the density of cumulative maximum risk score distance to the discontinuity cutoff of 100 within the CN sample by relocation status. The lumpy distribution is a function of the case management organization's scoring algorithm, and does not appear to differ significantly across relocation status.

Figure 3.8: Regression Discontinuity Bandwidths: Employment

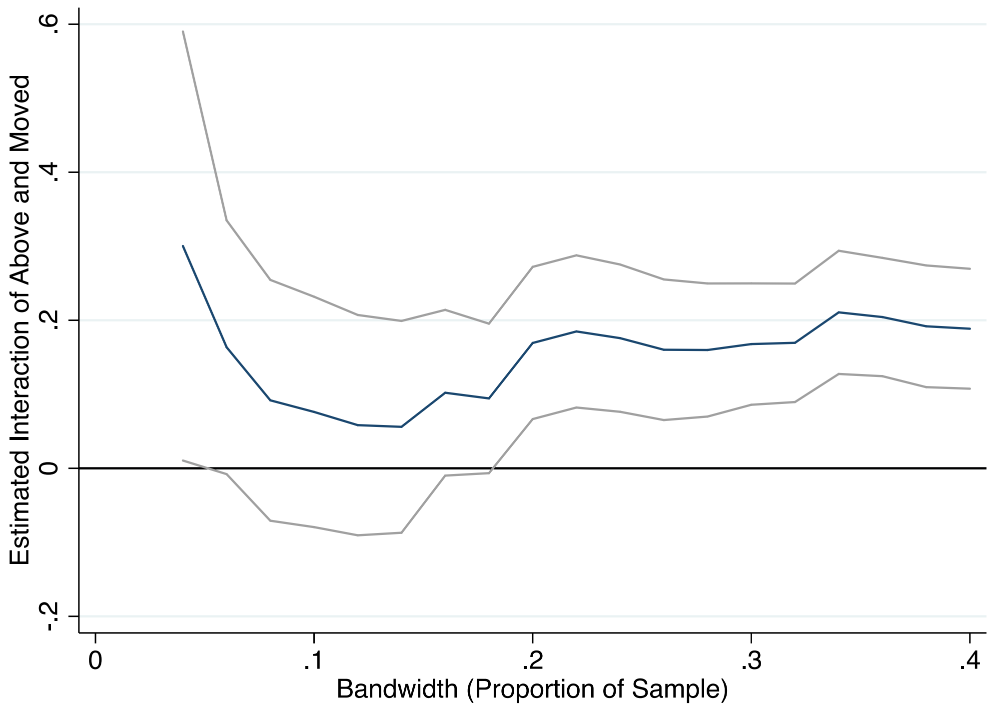

Notes: Figure plots the estimated coefficients on the interaction of above the risk score cutoff and relocated for various bandwidths where an employment dummy is the dependent variable. The grey lines represent a 95% confidence interval for each estimate. Each regression includes individual fixed effects, separate second-degree polynomial for risk scores above and below the cutoff, dummies for duration of case management, seasonal fixed effects, and polynomial for corresponding average outcomes across metro (wage and hours worked respectively). Standard errors clustered at the individual level.

REFERENCES

Acemoglu, D., "Technical Change, Inequality, and the Labor Market", Journal of Economic Literature p. 67 (2002).

Acemoglu, D. and P. Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor", Journal of Economic Perspectives **33**, 2, 3–30 (2019).

Andrews, I., J. H. Stock and L. Sun, "Weak Instruments in Instrumental Variables Regression: Theory and Practice", Annual Review of Economics **11**, 1, 727–753 (2019).

Angrist, J. D. and G. W. Imbens, "Two-stage least squares estimation of average causal effects in models with variable treatment intensity", Journal of the American Statistical Association **90**, 430, 431–442 (1995).

Anil, B., D. L. Sjoquist and S. Wallace, "The effect of a program-based housing move on employment: Hope vi in atlanta", Southern Economic Journal **77**, 1, 138–160 (2010).

Autor, D. H., "Work of the Past, Work of the Future", AEA Papers and Proceedings **109**, 1–32 (2019).

Bayer, P., S. Ross and G. Topa, "Place of work and place of residence: Informal hiring networks and labor market outcomes", Journal of Political Economy **116**, 6, 1150–1196 (2008).

Bergman, P., R. Chetty, S. DeLuca, N. Hendren, L. F. Katz and C. Palmer, "Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice", Working Paper 26164, National Bureau of Economic Research (2020).

Besbris, M., J. W. Faber, P. Rich and P. Sharkey, "Effect of neighborhood stigma on economic transactions", Proceedings of the National Academy of Sciences **112**, 16, 4994–4998 (2015).

Bessen, J. E., M. Goos, A. Salomons and W. Van den Berge, "Automatic Reaction - What Happens to Workers at Firms that Automate?", SSRN Electronic Journal (2019).

Binswanger, H., "Agricultural mechanization: A Comparative Historical Perspective", The World Bank Research Observer **1**, 1, 27–56 (1986).

Boone, C. D. A. and L. Wilse-Samson, "Structural Change and Internal Labor Migration: Evidence from the Great Depression", The Review of Economics and Statistics pp. 1–54 (2021).

Bryan, G., S. Chowdhury and A. M. Mobarak, "Underinvestment in a profitable technology: The case of seasonal migration in bangladesh", Econometrica **82**, 5, 1671–1748 (2014).

Butz, E. L. and O. G. Lloyd, "The cost of using farm machinery in Indiana", Bulletin 437, Purdue University Agricultural Experiment Station, Lafayette Indiana (1939).

Chetty, R., N. Hendren and L. F. Katz, "The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment", American Economic Review **106**, 4, 855–902 (2016).

Chyn, E., "Moved to opportunity: The long-run effects of public housing demolition on children", American Economic Review **108**, 10, 3028–56 (2018).

Clampet-Lundquist, S., "Moving over or moving up? short-term gains and losses for relocated hope vi families", Cityscape **7**, 1, 57–80 (2004).

Dvorkin, M. A. and A. Monge-Naranjo, "Occupation Mobility, Human Capital and the Aggregate Consequences of Task-Biased Innovations", Working Paper (2019).

Eisner, E., "Have Robots Taken All Our Jobs Before? The Case of Mechanization in US Agriculture", Working Paper (2021).

Feigenbaum, J. and D. Gross, "Automation and the Future of Young Workers: Evidence from Telephone Operation in the Early 20th Century", Tech. Rep. w28061, National Bureau of Economic Research, Cambridge, MA (2020).

Fiszbein, M., J. Lafortune, E. Lewis and J. Tessada, "New Technologies, Productivity, and Jobs: The (Heterogeneous) Effects of Electrification on US Manufacturing", Tech. Rep. w28076, National Bureau of Economic Research, Cambridge, MA (2020).

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

Garcıa-Couto, S., "Beyond Labor Market Polarization", Working Paper (2021).

Goetz, E. G., "Forced relocation vs. voluntary mobility: The effects of dispersal programmes on households", Housing Studies **17**, 1, 107–123 (2002).

Haines, M., P. Fishback and P. Rhode, "United States Agriculture Data, 1840 - 2012", (2018).

Hilger, N. G., "Parental Job Loss and Children's Long-Term Outcomes: Evidence from 7 Million Fathers' Layoffs", American Economic Journal: Applied Economics **8**, 3, 247–283 (2016).

Humlum, A., "Robot Adoption and Labor Market Dynamics", Working Paper (2019).

Jacob, B. A., "Public housing, housing vouchers, and student achievement: Evidence from public housing demolitions in chicago", The American Economic Review **94**, 1, 233–258 (2004).

Kain, J. F., “Housing segregation, negro employment, and metropolitan decentralization”, The Quarterly Journal of Economics **82**, 2, 175–197 (1968).

Katz, L. and R. Margo, “Technical Change and the Relative Demand for Skilled Labor: The United States in Historical Perspective”, Tech. Rep. w18752, National Bureau of Economic Research, Cambridge, MA (2013).

Lindert, P. H., “Long-run trends in american farmland values”, Agricultural History **62**, 3, 45–85 (1988).

McKibben, E. G. and R. A. Griffin, “Changes in farm power and equipment: Tractors, trucks, and automobiles.”, Tech. Rep. A-9, Federal Works Agency, Philadelphia, Pennsylvania (1938).

McKibben, E. G., J. A. Hopkins and R. A. Griffin, “Changes in farm power and equipment: Field implements.”, Tech. Rep. A-11, Federal Works Agency, Philadelphia, Pennsylvania (1939).

Mendenhall, R., S. DeLuca and G. Duncan, “Neighborhood resources, racial segregation, and economic mobility: Results from the gautreaux program”, Social Science Research **35**, 4, 892 – 923 (2006).

Nakamura, E., J. Sigurdsson and J. Steinsson, “The gift of moving: Intergenerational consequences of a mobility shock”, (2019).

Nunn, N. and D. Puga, “Ruggedness: The Blessing of Bad Geography in Africa”, The Review of Economics and Statistics **94**, 1, 20–36 (2012).

Olea, J. L. M. and C. Pflueger, “A Robust Test for Weak Instruments”, Journal of Business & Economic Statistics **31**, 3, 358–369 (2013).

Olmstead, A. L. and P. W. Rhode, “Reshaping the landscape: The impact and diffusion of the tractor in american agriculture, 1910–1960”, The Journal of Economic History **61**, 3, 663–698 (2001).

Orr, L., J. Feins, R. Jacob, E. Beecroft, L. Sanbonmatsu, L. F. Katz, J. B. Liebman and J. R. Kling, “Moving to opportunity: Interim impacts evaluation”, (2003).

Ran, A., B. Leah and R. Myera, “Census Linking Project: Version 1.0 [dataset]”, (2020).

Riley, S. J., S. D. DeGloria and R. Elliot, “Index that quantifies topographic heterogeneity”, intermountain Journal of sciences **5**, 1-4, 23–27 (1999).

Rosenbaum, J. E., “Changing the geography of opportunity by expanding residential choice: Lessons from the gautreaux program”, Housing Policy Debate **6**, 1, 231–269 (1995).

Ruggles, S., S. Flood, S. Foster, R. Goeken, J. Pacas, M. Schouweiler and M. Sobek, “Ipums usa: Version 11.0 [dataset]”, (2021).

Sanbonmatsu, L., J. Ludwig, L. F. Katz, L. A. Gennetian, G. J. Duncan, R. C. Kessler, E. Adam, T. W. McDade and S. T. Lindau, "Moving to opportunity for fair housing demonstration program – final impacts evaluation", (2011).

Schmutte, I. M., "Job referral networks and the determination of earnings in local labor markets", Journal of Labor Economics **33**, 1, 1–32 (2015).

Stock, J. and M. Yogo, *Testing for Weak Instruments in Linear IV Regression*, pp. 80–108 (Cambridge University Press, New York, 2005).

van Dijk, W., "The socio-economic consequences of housing assistance", (2019).

Weinberg, B. A., P. Reagan and J. J. Yankow, "Do neighborhoods affect hours worked? evidence from longitudinal data", Journal of Labor Economics **22**, 4, 891–924 (2004).

APPENDIX A

APPENDIX TO CHAPTER 2

Appendix Tables

Table A.1: Other Outcomes

	Δ County (1)	Δ State (2)	New child '10-'20[1] (3)	Married '10-'20[2] (4)	Divorced '10-'20[3] (5)
Panel A: 1910 Agricultural Non-wage Worker					
Pct. Δ equip. value per acre '10-'20 (SD)	0.004	0.012	-0.073**	-0.006	0.003
	(0.031)	(0.022)	(0.037)	(0.052)	(0.003)
Obs.	243,612	243,612	100,540	36,578	207,034
Clusters	2407	2407	2399	2374	2406
Y mean	0.356	0.169	0.636	0.508	0.005
Montiel Olea-Plueger F-stat	15.9	15.9	14.5	19.1	14.9
Controls					
State FEs	x	x	x	x	x
Standard soil and topo	x	x	x	x	x
Panel B: 1910 Agricultural Wage Worker					
Pct. Δ equip. value per acre '10-'20 (SD)	0.011	-0.004	-0.049	0.112^{+}	0.018^{+}
	(0.050)	(0.042)	(0.058)	(0.072)	(0.012)
Obs.	32,871	32,871	19,432	13,920	18,951
Clusters	2354	2354	2281	2168	2228
Y mean	0.525	0.255	0.531	0.533	0.010
Montiel Olea-Plueger F-stat	16.9	16.9	14.8	18.0	13.8
Controls					
State FEs	x	x	x	x	x
Standard soil and topo	x	x	x	x	x

***:p<0.01, **: p<0.05, *:p<0.10, +:p<0.15
[1]: Conditional on age ≤ 35
[2]: Conditional on not married
[3]: Conditional on married

Appendix Figures

Figure A.1: First Generation Wage Distribution

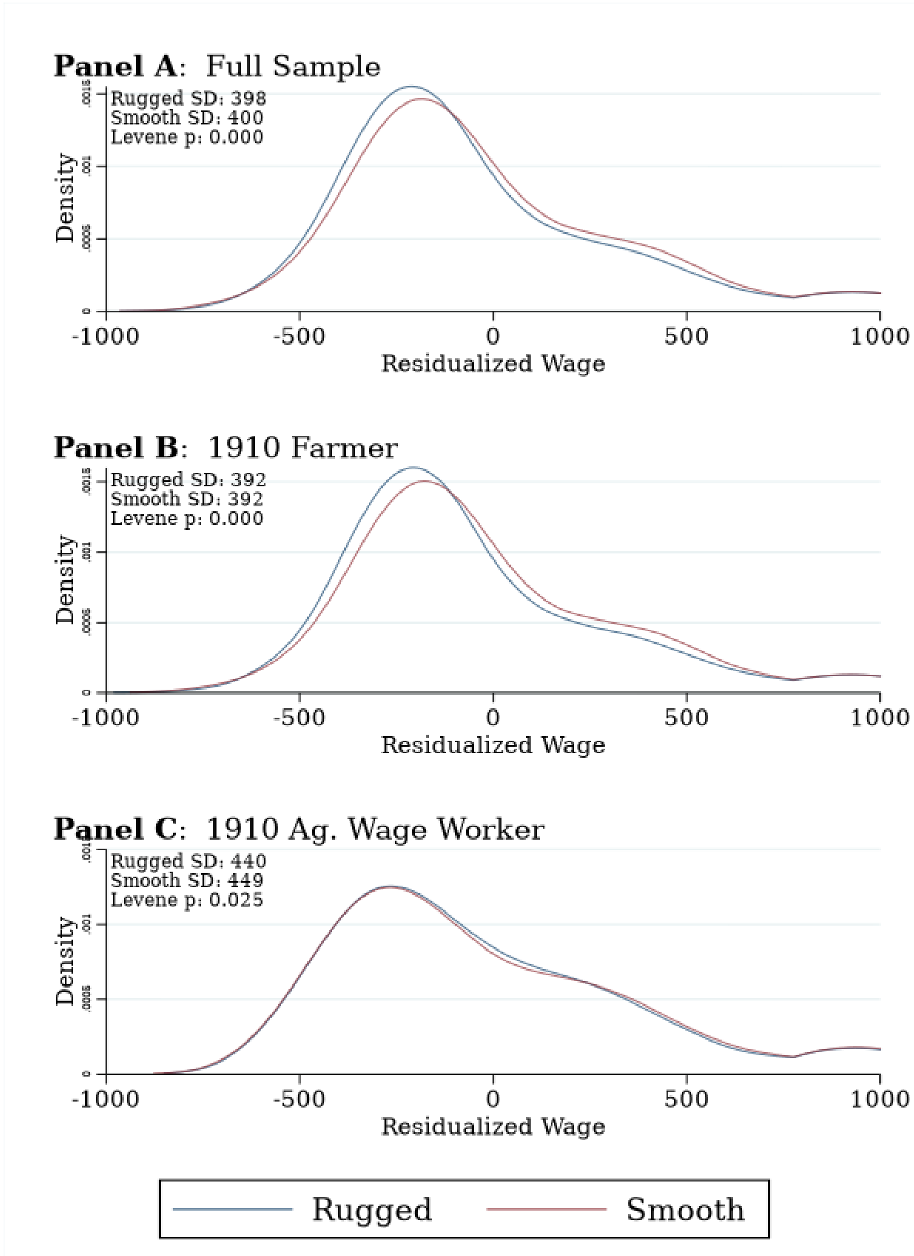

Notes: Figure displays kernel density plots for the residualized occupational wages of agricultural workers separately for the upper and lower terciles of terrain ruggedness in their baseline (1910) county. Plots winsorized above $1,000 for clarity, but standard deviations derived from underlying data. The sample in Panel A includes all matched workers between 25 and 55 in 1910, and Panels B and C present distributions for the indicated sub-samples. Wages are residualized against the same controls as in Table 2.6.

Table A.2: Displaced Worker Mechanisms

	Log National Occ. Wage 1920	Literate 1910 (0/1)	1920 Occ. Avg. Yrs. School	1920 Occ. Share $\geq$ 8th Grade (0/1)
	(1)	(2)	(3)	(4)
Pct. Δ equip. value	-0.0275	-0.0091	-0.0240	-0.0063
per acre '10-'20 (SD)	(0.0473)	(0.0387)	(0.1870)	(0.0169)
Y mean	1088.0[1]	0.882	8.251	0.658
Obs.	16,071	16,077	12,021	12,021
Clusters	2189	2195	2079	2079
Montiel Olea-Plueger F-stat	17.9	17.9	16.1	16.1
Controls				
State FEs	x	x	x	x
Standard soil and topo	x	x	x	x

***:p<0.01, **: p<0.05, *:p<0.10, +:p<0.15
1: mean in 1940 $

Table A.3: Effects on Second Generation Skills

	Years Schooling	≥ 8th Grade (0/1)	1920 Occ. Share ≥ 8th Grade (0/1)	Years Schooling	≥ 8th Grade (0/1)	1920 Occ. Share ≥ 8th Grade (0/1)
	(1)	(2)	(3)	(4)	(5)	(6)
	Father Left Ag. 1920			Father Ag. Non-Wage Worker		
Pct. Δ equip. value	-0.6727	-0.0928	0.0036	0.5439^{+}	0.0624^{+}	0.0312*
per acre '10-'20 (SD)	(0.6289)	(0.0726)	(0.0266)	(0.3378)	(0.0387)	(0.0162)
Y mean	9.440	0.796	0.744	9.127	0.750	0.733
Obs.	12,621	12,621	12,788	450,435	450,435	456,249
Clusters	2201	2201	2203	2418	2418	2418
Montiel Olea-Plueger F-stat	12.9	12.9	13.1	12.6	12.6	12.6
Controls						
State FEs	x	x	x	x	x	x
Standard soil and topo	x	x	x	x	x	x

***:p<0.01, **: p<0.05, *:p<0.10, +:p<0.15

Table A.4: Human Capital Transition Cost Parameterization

		Destination Occupation					
		h	l	Tf	Mf	Tw	Mw
Origin Occ.	h	-	τ_1	τ_2	τ_2	τ_2	τ_2
	l	τ_3	-	τ_4	τ_4	τ_4	τ_4
	Tf	τ_5	τ_6	-	τ_8	τ_7	τ_7
	Mf	τ_5	τ_6	τ_8	-	τ_7	τ_7
	Tw	τ_5	τ_6	τ_7	τ_7	-	τ_8
	Mw	τ_5	τ_6	τ_7	τ_7	τ_8	-

Notes: Table displays the parameterization of the human capital transition costs in the model. Each τ in the matrix represents a separately estimated parameter.

Figure A.2: Selection of the Investment Cutoff to Define Agricultural Regions in Model

Notes: Figure displays the change in average hours needed to work one acre for each of the 5 categories of farm equipment as defined in Figure 2.3 by the equipment investment level in the county of measurement. The sample construction is identical to Figure 2.4 Panel C. The red lines indicates the selected cutoff value used to discipline the difference between high-tech and low-tech regions in the model.

Figure A.3: Cumulative Distribution Functions of Changes in Equipment Value by Ruggedness

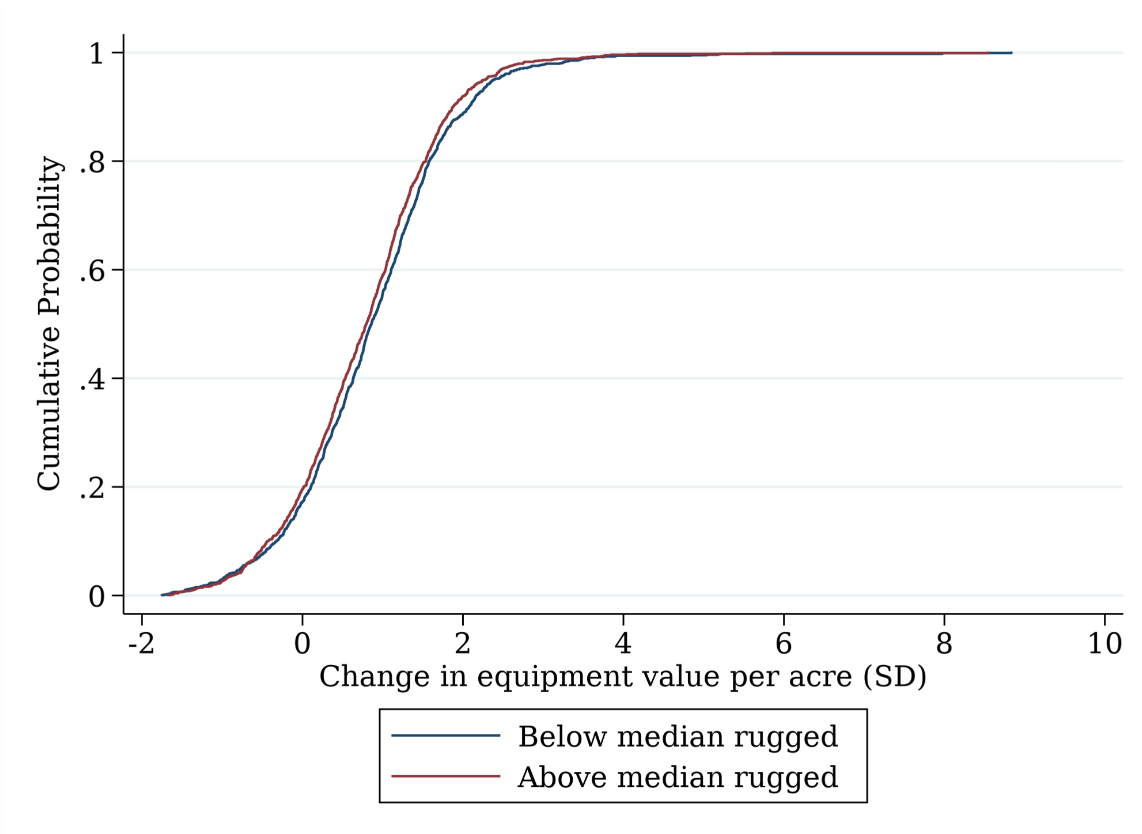

APPENDIX B

APPENDIX TO CHAPTER 3

Appendix Tables

Table B.1: Neighborhood Shift Over Relocation: Census Tract Observables

		Pre-Move Level	Post-Move Level	Diff in County SD
Proportion of Census Tract...	Below Poverty Line	0.66	0.45	-1.25
	Racial Minority	0.83	0.88	0.18
	HS Degree or Less	0.68	0.60	-0.45
	Renting Home	0.97	0.69	-1.27
	Female HoH	0.31	0.31	0.03

Notes: The first two columns report average pre- and post-relocation census tract neighborhood characteristics for all CN households. Characteristics data from ACS 5 year estimates, 2012-2016. Column three normalizes the difference between columns two and one by the Shelby County standard deviation of the variable across census tracts.

Appendix Figures

Table B.2: Shifts in Neighborhood Characteristics Over Move by Demograph Group

	High School Degree or More	Less than High School Degree	Younger than 45	Older than 45
Poverty Share	-2.14	-2.02	-2.15	-2.00
Minority Share	-0.03	-0.07	-0.02	-0.08
HS or Less Share	-0.71	-0.76	-0.77	-0.66
Renter Share	-1.56	-1.47	-1.68	-1.29
Female HoH Share	-0.48	-0.67	-0.53	-0.60

Notes: Each column displays the post-move change in the indicated neighborhood characteristics normalized by the Shelby County standard deviation of the characteristic across census block groups. Characteristic data from ACS 5 year estimates, 2012-2016.

Table B.3: Average Within Indiviudal Change in Transfer Receipt

	CN	non-CN	p-value
TANF	-12.1	0.0	0.339
UI	0.0	0.0	-
WIC	-14.4	0.0	0.019
Food Stamps	-42.9	-18.5	0.399
Total	-69.4	-18.5	0.115
CM Months First Obs.	0.0	2.3	
CM Months Second Obs.	19.0	11.2	
Relocated Months Second Obs.	9.4		
N	50	23	

Notes: Table displays the within indiviudal average change in transfers from several federal assistance programs. The sample is restricted to the subset of individuals with more than one income report.

Table B.4: Employment Transitions

		CN Residents		Non-CN Residents	
		Starting Employment Status		Starting Employment Status	
		* Unemployed [n=220]	* Working [n=88]	* Unemployed [n=77]	* Working [n=27]
Ending Employment Status	Unemployed	80.5	18.2	70.1	3.7
	Working	19.5	81.8	29.9	96.3
	Sum	100	100	100	100

Notes: Table displays the proportion of the column group which transitioned into the corresponding row groups. Employment measured as binary outcome, all individuals either employed or unemployed. Stars denote significant differences in column transition probabilities across housing sites based on Chi-Squared test. Starting employment measured at the first month of case management. Ending employment measured at 18 months of case management. The CN sample was further restricted to only include unmoved initial observations. All CN residents have been moved in the ending sample. Average relocation duration in this sample is 12.0 months. (*= p<0.10, **= p<0.05, ***= p<0.01)

Table B.5: Relocation Effects: 1-2 Adult Households

	If Employed		
	Hourly Wage	Hours Worked Per Week	Employed
Relocation Effect	0.64** (0.27)	1.13 (0.78)	-0.043 (0.036)
Pre-Move Mean	9.36	27.87	0.297
N	3,507	3,507	10,383
Avg. Months Moved	11.5	11.5	12.9

Notes: Table reports estimated relocation effects. Sample restricted to only households with 2 or fewer adults at the baseline. Each regression includes individual fixed effects, dummies for duration of case management, dummies for season of year, and polynomial for corresponding average outcomes across metro (wage, hours worked, and unemployment respectfully). Standard errors clustered at the individual level. (*= p<0.10, **= p<0.05, ***= p<0.01)

Table B.6: Effect of Relocation on Adjusted Labor Market Outcomes

	Adjusted Hourly Wage if Employed	First Month Observed in Job
Relocation Effect	0.67***	0.163***
	(0.26)	(0.022)
Pre-Move Mean	9.37	0.054
N	3,972	11,969
Avg. Months Moved	12.0	13.1

Notes: Table reports estimated coefficients on a dummy for relocation. In column 1, wages are adjusted to incorporate the observed wage increases as a monthly linear trend if a job has ended (i.e. ending wages are known). If ending wages are not know, the average linear trend across jobs within the panel is applied. Each regression includes individual fixed effects, dummies for duration of case management, dummies for season of year, and polynomial for corresponding average outcomes across metro (wage and unemployment respectfully). Standard errors clustered at the individual level. (*= p<0.10, **= p<0.05, ***= p<0.01)

Figure B.1: Neighborhood Characteristics over Time

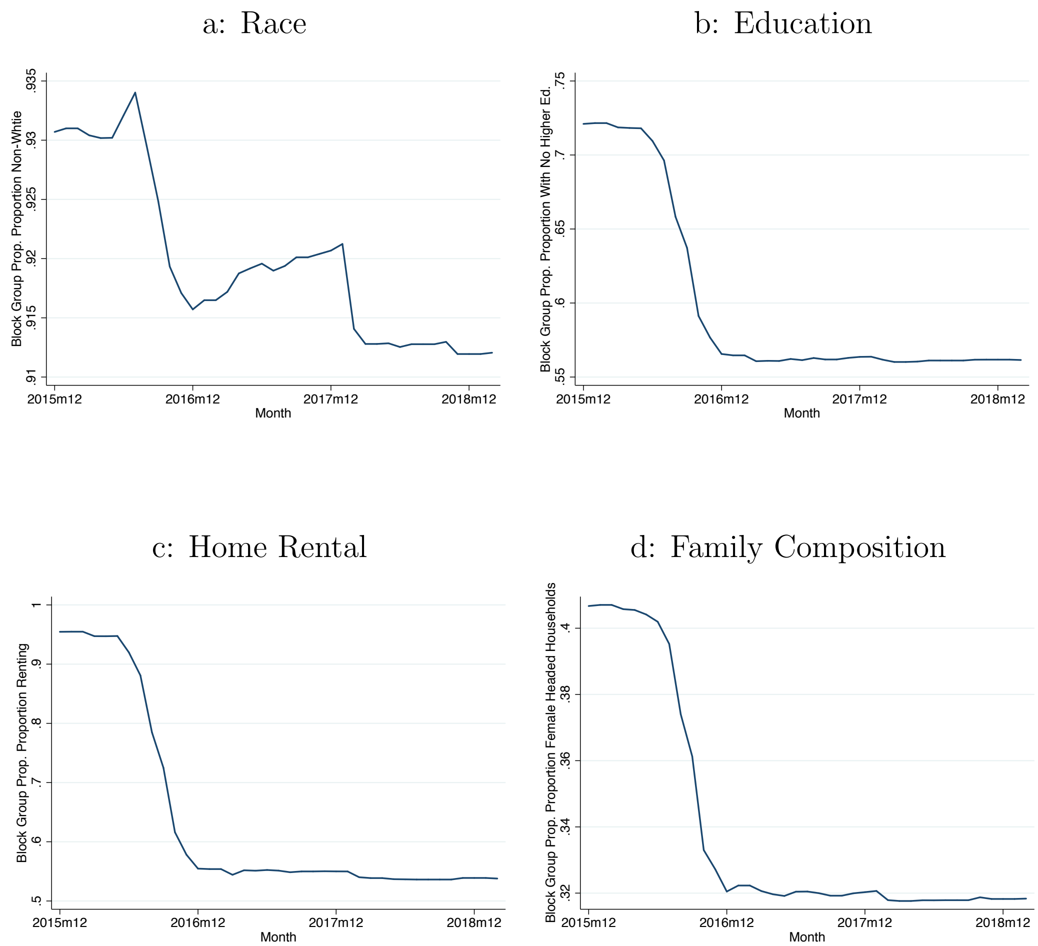

Notes: Each panel displays average census block group observables within the sample of CN residents with known relocation dates.

Figure B.2: Visual Regression Discontinuity Strategy

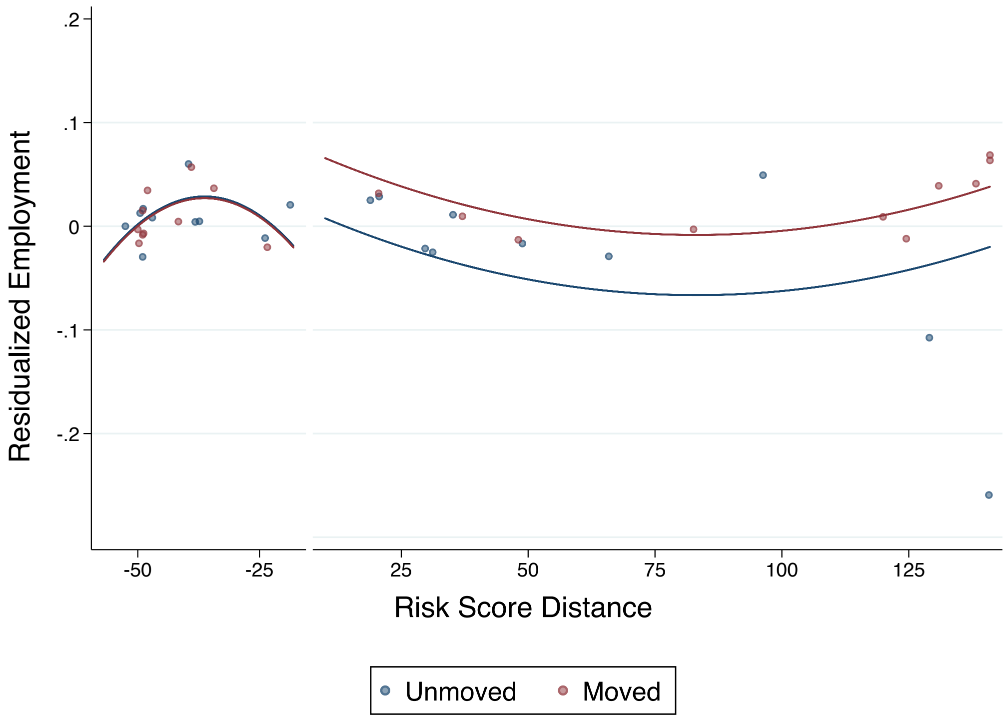

Notes: Figure demonstrates the regression discontinuity strategy which identifies the joint impact of additional attention from a case manager and relocating away from the initial housing site. The running variable is the difference between an individual's cumulative maximum risk score and the cutoff score of 100. The dependent variable is a dummy for employment status residualized against individual fixed effects. As the figure does not include any other controls, it does not fully represent the complete empirical specification displayed in Table 3.8.