OPPORTUNITY FOR WHOM?
SOURCES OF INTERGENERATIONAL MOBILITY IN THE U.S.

A Dissertation

Submitted to the Faculty

of

Purdue University

by

Timothy M. Smith

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

May 2019

Purdue University

West Lafayette, Indiana

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF DISSERTATION APPROVAL

Dr. Michael S. Delgado, Chair
Department of Agricultural Economics
Dr. Tom Hertz
United States Department of Agriculture, Economic Research Service (ERS)
Dr. Gerald Shively
Department of Agricultural Economics
Dr. Brigitte Waldorf
Department of Agricultural Economics

Approved by:

Dr. Nicole Olynk Widmar
Chair of the Graduate Committee

To Michael and Raymond.

ACKNOWLEDGMENTS

A lot of people did a lot of work to get this document to this point. First and foremost, my original co-chairs, Michael Delgado and Raymond Florax, put a lot of energy into helping me turn my ideas into something concrete and coherent, and in the process taught me a great deal about what it is that we're trying to do as economists. In the later years of my program, Michael has worked hard and carefully to keep me moving in the right direction, convincing me to abandon some ideas that weren't progressing, and to persist with some others that weren't progressing either, but in fact had promise.

The rest of my committee – Tom Hertz, Jerry Shively, and Brigitte Waldorf – also made important contributions to this project, and to my development as a scholar. Tom and Brigitte contributed both to the thinking and arguments in this dissertation itself, and to my thinking about mobility specifically and economics more generally, through their work on the "Poverty Project". Jerry's support throughout my program has helped to clarify my thinking about how this work fits into our broader understanding of investments in children, and without his willingness to mentor a master's student with an odd profile, I almost certainly wouldn't be finishing a Ph.D..

I learned a lot from Huan Li and Indraneel Kumar, both through the content of their work and feedback on the projects that became Chapters 3 and 5, and through valuable conversations about the project and scholarship in general. Their influence is particularly present in Chapter 5, which would not exist in its current form without their contributions, and the contributions of our PCRD colleagues who contributed to Kumar et al. (2019).

My many colleagues in the SHaPE group and in the graduate program more broadly have helped me a great deal by engaging with my ideas and, at least as importantly, giving me opportunities to clarify my thinking by engaging with theirs. I am particularly grateful to Eric Gerber, Stacy Prieto, Jeff Opgrand, and Jeff Young; while many of my friends heard a lot about my concerns as this dissertation developed, they heard the most.

Finally, I want to thank my parents, who gave me the opportunity to develop the preferences and bundle of cognitive, character, and social skills that produce the desire and, apparently, ability to complete an undertaking like a Ph.D..

TABLE OF CONTENTS

Page

Page

LIST OF TABLES

Table Page

LIST OF FIGURES

Figure Page

Figure Page

ABSTRACT

Smith, Timothy M. Ph.D., Purdue University, May 2019. Opportunity for Whom? Sources of Intergenerational Mobility in the U.S.. Major Professor: Michael S. Delgado.

Economists generally consider intergenerational economic mobility to be an important feature of market economies, as it allows people born into poverty to achieve a measure of prosperity in the presence of minimal government intervention or redistribution. The empirical literature on mobility in the U.S. has, however, found evidence that mobility is lower than previously thought, and scholars have responded by developing expansive literatures on many aspects of intergenerational mobility, including studies of its origins. In this dissertation, I contribute to this strand of the literature by reviewing recent trends in the literature, with a particular emphasis on studies aimed at explaining the sources of mobility, and then discussing three empirical studies into specific sources of mobility, using data organized at different geographic and temporal scales. These empirical chapters focus on the role of different aspects of childhood poverty in determining income rank in adulthood, modeling variation in racial mobility gaps across different kinds of communities and local economies, and measuring the relationship between trends in intergenerational mobility and the structural transformation of agriculture in the 20^{th} century U.S..

CHAPTER 1. INTRODUCTION

The claim that market economies allow a high degree of intergenerational mobility is a common rejoinder to criticisms resting on high levels of inequality, poverty, and other circumstances widely considered to be undesirable. Economists often argue that if, through hard work and talent, people born into unfavorable circumstances are able to attain a comfortable middle class standard of living, we ought not concern ourselves with the existence of poverty or inequality, because under high mobility, growth will continue to make everyone better off and the redistributive policies often recommended for limiting poverty and inequality will interfere with that process. This paradigm is based on sound economic reasoning, but in the U.S. context, it also relies crucially on the assumption that intergenerational mobility is particularly high, an assumption closely related to common perceptions about the accessibility of upward mobility in the U.S. economy, often referred to as the 'American Dream'. Empirical research on the level of intergenerational mobility in the U.S. has, however, revealed that not only is the U.S. not exceptionally mobile, it is, in fact, relatively immobile, in absolute and relative terms within the population and in a number of cross-country comparisons. This finding, first widely publicized in the early 1990's, has laid the foundation for a vibrant and increasingly broad literature on the levels and causes of intergenerational mobility, which has often overlapped productively with broader literatures on human capital development, economic history, and social externalities.

This research agenda has evolved considerably since its genesis in the early 1990's. Much of the energy has been devoted to producing more and more credible estimates of the degree of intergenerational mobility, in the U.S. and elsewhere. This effort has clarified the definitions and empirical conventions that act as the literature's *lingua franca*, in addition to producing a nuanced framework for measuring and comparing mobility statistics. With this common language developed, scholars have begun to pay more and more attention to understanding the sources of intergenerational mobility, an enterprise which relies crucially on exploiting variation in mobility. One notable feature of this "new mobility" literature is its creativity in exploiting variation occurring at different levels of data aggregation. The core of this literature must rely, of course, on relationships in incomes across generations within families, but the influence community characteristics and attributes of local economies exert on children and households leaves scholars to ignore units larger than the family at their own peril. Recent work by Chetty et al. (2014a) makes this clear, as their results show that mobility varies a great deal across space, and that the size of that variation can be explained by community covariates as well as more standard household-level information.

In this dissertation, I contribute to the intergenerational mobility literature by examining sources of intergenerational mobility in the U.S. at different geographic and temporal scales. Varying the units of analysis lets me consider the importance of qualitatively distinct kinds of factors, and considering different time periods of differing lengths lets me comment on the influence of structural changes, in addition to discussing the influence of family and community factors. In my first chapter, I review the evolution and key contributions of the literature I allude to in this introduction, both to provide context and to underscore the salience of the gaps I seek to contribute to. In my second chapter, I discuss the data I use in my empirical chapters. I focus on common features of these datasets and on the attractive features that they bring to the analyses I undertake, leaving more detailed discussion of the measurements and specific sampling restrictions I apply in each chapter for discussion within that chapter in question. In my final four chapters, I describe three empirical projects designed to explain specific aspects of intergenerational mobilities at different scales, and then provide a conclusion in which I summarize my results and their contribution, and discuss the advantages and drawbacks of the approach I have deployed in this dissertation.

These empirical chapters are the core of my dissertation, and each focuses on a different aspect of the sources of intergenerational mobility, with a different level of data aggregation. In my first empirical chapter, I examine the relationship between specific aspects of childhood poverty and adulthood rank in the income distribution, in an effort to better understand the specific ways in which poverty acts as a barrier to mobility. This is important because it contributes to our understanding of variation in mobility across the income distribution, as I focus on the lower end of the distribution, while engaging with results from the structural and quasi-experimental literatures on human capital accumulation, all in a policy-relevant framework. In my second empirical chapter, I model the substantial differences in income mobility between black and white Americans as a function of features of neighborhoods and local economies, an emphasis that is particularly salient because a great deal of black-white residential segregation persists, often creating very different social and economic outlooks for black communities than white ones. As Hertz (2005) has shown, many of the stylized facts about mobility in the U.S. are driven by the extremely low level of mobility black Americans face, so improving our understanding of these differences provides an exciting opportunity to learn about the sources of mobility in the U.S. more generally. In my final empirical chapter, I take a longer view, using census data to compute state-level mobility statistics going back to 1940, which I then use to examine the influence of structural changes in the agricultural sector on trends in intergenerational mobility, with a particular emphasis on the South. This project contributes to a growing economic history strand of the mobility literature, which takes advantage of intertemporal variation to capture trends that longitudinal datasets starting in the late 1960's may not reflect clearly.

Overall, this dissertation contributes to our understanding of the sources of intergenerational mobility in several distinct contexts. This is important for both scholarly and practical reasons: learning about where mobility comes from helps scholars to understand the relative merits of conceptual and theoretical predictions, which can in turn refine those predictions, and it can help policymakers and other stakeholders with preferences for more mobility to understand which interventions might be effective or ineffective in reaching that goal. My work in this dissertation is far from definitive, even within the three relatively narrow areas I study in my empirical chapters, but there remains a great deal to learn about the mechanics of intergenerational mobility, and I am confident that my work here provides both useful answers to some of the questions in this literature, as well as a valuable perspective on how the literature might proceed in this direction.

CHAPTER 2. MOBILITY IN ECONOMIC STATUS AMONG PEOPLE AND PLACES

2.1 Introduction

In a society with low intergenerational economic mobility, advantage endowed by an individual's parents drives an individual's economic stature. Someone born without the advantage of wealthy parents is unlikely to rise to the top of the income distribution solely based on individual talent and effort. For this reason, many policymakers and citizens tend to prefer a more mobile society, in which individual talent and effort serve to overcome disadvantages outside of one's control; in the United States, this possibility is central to the ideology termed the "American Dream." Economic mobility provides a foothold for policymakers seeking to provide opportunities for advancement to disadvantaged individuals, and social scientists have responded by studying individual economic mobility in a variety of ways.

Many scholars have chosen to focus on intergenerational economic mobility, or the link between the economic status of an individual and the economic status of his/her parents. The empirical literature on intergenerational economic mobility falls largely into two distinct segments. The first focuses solely on the individual and his/her parents, built on the Becker and Tomes (1979) theory of the family that emphasizes the role of parental human capital and income in determining an individual's income later in life. The second emphasizes the importance of an individual's neighborhood as a factor that facilitates the intergenerational transmission of economic status, and within economics, it draws heavily on theoretical work by Loury (1977). Loury's key insight is that an individual's community bounds the set of inputs he/she can receive while developing human capital as a child, playing a key role in intergenerational income dynamics and in opportunity.

We review some of the key economic and econometric insights underlying this vast literature, review important empirical findings, and discuss what we believe to be promising directions for policy development and future research. To develop a sense of the scope of this literature, Table 2.1 summarizes the main types of models that have been used to investigate intergenerational economic mobility; these models include cross-sectional models, panel data models (to study time trends), models of neighborhood effects, and experimental and quasi-experimental studies. We discuss the basic econometric model of individual mobility, as well as several extensions of that model in order to understand the extent to which common modifications of the basic model generate new empirical insights. We also summarize the key mechanisms through which neighborhood characteristics influence individual economic mobility. Our review of these literatures complements and synthesizes other excellent reviews of intergenerational income mobility (Solon, 1999; Black and Devereux, 2011), neighborhood effects (Durlauf, 2004; Leventhal and Brooks-Gunn, 2000), and international comparisons of intergenerational mobility estimates (Corak, 2006; Blanden, 2013).

Our review also provides a synthesis of recent research that focuses on the interactions between neighborhood effects and individual economic mobility. Much of this work has relied on field experiments, and has revitalized the field for at least two reasons. First, these field experiments provide exogenous assignment of treatment (i.e., moving from a poor neighborhood to a rich or middle income neighborhood), which removes many potential sources of bias that are thought to influence many estimates of treatment effects based on observational data. Second, these experiments naturally integrate two traditionally distinct sources of economic advantage and disadvantage: families and neighborhoods. Our goal with this discussion is to clarify what we know about the intersection of people and places in terms of economic mobility, and to characterize the current frontier of research.

At this frontier, a number of questions remain unanswered. Regression models focused on specific neighborhood mechanisms often lack a compelling identification strategy because it is difficult to rigorously link cause and effect in a context characterized by so many interacting choices and influences. On the other hand, more plausibly identified approaches, typically based on randomized housing relocation, produce a treatment effect that is likely causal but cannot easily measure specific mechanisms, as relocation changes multiple neighborhood characteristics simultaneously; these approaches struggle to explain what drives the effect. We discuss what these empirical studies can tell us, given both their strengths and limitations.

We make headway towards a common ground in these strands of the literature through recent theoretical work on poverty traps – defined by Azariadis and Stachurski (2005) as "self-reinforcing mechanisms that cause poverty to persist," and can exist at the individual or community level. The poverty trap literature is conceptually related to intergenerational mobility because it emphasizes human capital, access to investment, and dynamics that may produce multiple equilibria depending on individual and/or environmental factors. Each of these factors influence where an individual starts in the income distribution, and the extent to which he/she is able to move away from that position. This perspective serves as a means to unify much of the literature on individual economic opportunity, and clarifies new directions for both empirical research and for policy development. A relevant insight from this work is the consensus that multiple types of frictions must be present to generate a true poverty trap, and policy must simultaneously address multiple frictions in order to lead to real improvements in economic mobility.

The outline of our review is as follows. Section 2.2 describes the individual-parent model of intergenerational mobility, describes a variety of alternative model specifications, and summarizes general findings and implications of this research. Section 2.3 describes the early literature on the impact of neighborhood characteristics on individual economic status, and Section 2.4 describes recent work based on experimental and observational housing relocation that explores causal reduced form models of the relationship between neighborhoods and intergenerational mobility. Section 2.5 discusses lessons from recent theoretical work on poverty traps that relates to policy development for individual economic opportunity. We offer our final thoughts in Section 2.6.

2.2 Intergenerational Mobility

2.2.1 The Classic Intergenerational Mobility Model

Dating back at least as far as Sewell and Hauser (1975), social scientists have estimated a society's level of economic mobility using the regression model

$$Y_i = \alpha + \beta Y_i^{\pi} + \epsilon_i, \qquad i = 1, 2, \ldots, n \tag{2.1}$$

where Y_i is individual i's permanent income, Y_i^{π} is his/her parent's permanent income, ϵ_i is the disturbance, and the parameter β captures the relationship between Y_i^{π} and Y_i. We refer to Y as permanent income to be consistent with a majority of this literature, noting that it might be more appropriately called lifetime average income. If the variances in income are the same for both generations, the regression slope β is equal to the intergenerational correlation in income, $cor(Y, Y^{\pi})$, and is interpretable as a dollar-for-dollar coefficient if incomes are measured in levels, or as an elasticity if they are measured in logs (Solon, 1999; Lee and Solon, 2009; Black and Devereux, 2011). Otherwise, the regression coefficient β is a biased estimate of either intergenerational income relationship, and without knowledge of the change in income variances, both the size and sign of this bias is unknown (Zimmerman, 1992; Solon, 1992).[1] We summarize key results from the literature employing this model in Table 2.

In any case, β is a descriptive statistic that provides insight into how closely related an individual's permanent income is to his/her parent's permanent income (Hertz, 2005, 2007). A large value of β (i.e., β close to 1) implies that a society is

[1]This is because the regression slope is equal to $\rho \times (\sigma_Y/\sigma_{Y^{\pi}})$, where ρ denotes the intergenerational correlation in income, and σ is the standard deviation. If there is a difference between σ_Y and $\sigma_{Y^{\pi}}$, then the parameter β is not equal to the intergenerational correlation in income (or the intergenerational elasticity, if income is in logs).

not economically mobile across generations, meaning that an individual's economic status is closely tied to his/her parent's economic status – the rich stay rich, and the poor stay poor. A small value of β (i.e., β close to 0) means that a society is almost perfectly economically mobile across generations, as an individual's economic status is almost unrelated to his/her parent's economic status – all other things being equal, two people have the same chance of ending up at a certain point in the income distribution, regardless of who are their parents.

An estimate of β can also be used to compute the probability that an individual ends up in a particular part of the income distribution given his/her parents' location in the income distribution (e.g., Solon, 1992; Björklund and Jäntti, 1997). Assuming Y and Y^π are jointly normally distributed, and using μ and σ as the mean and standard deviation of the income distribution,

$$Y|(Y^\pi = y^\pi) \sim \mathcal{N}\left[\frac{\sigma_Y}{\sigma_{Y^\pi}}(\alpha + \beta y^\pi), \frac{1}{(1-\beta^2)}\sigma_Y^2\right] \tag{2.2}$$

can be used to compute the probability that an individual achieves a particular level of income (e.g., above the median) conditional on his/her parents' income level.

This classic model is subject to three potential criticisms. The first is that the model assumes a simple parametric functional form, and a distributional assumption when $\widehat{\beta}$ is used to compute transition probabilities (assuming the underlying equation of motion follows a Markov process) via equation (2.2). Scholars have challenged both the distributional assumption and the parametric structure of this model based on empirical evidence of asymmetry in terms of upward and downward mobility (e.g., Solon, 1992; Dearden et al., 1997); asymmetry implies that these assumptions do not hold. The second criticism relates to the reliability of empirical measures of permanent income, and the third relates to heterogeneity in the intergenerational elasticity (or correlation) in income. These latter two criticisms have been subject to a substantial amount of empirical work, and in the following subsections we discuss each in turn.

2.2.2 Correcting for Measurement Error in Permanent Income

Whether an estimate of β from the classic model is useful as a descriptive statistic related to intergenerational mobility depends on the extent to which the income measurements capture permanent income. The distinction between permanent and transitory income is important because the theory that motivates much of the literature focuses on the intergenerational persistence of earning ability, which is best captured by permanent income that reflect an individual's earnings, earning ability, and position in the income distribution. In a typical dataset, an individual is not observed over a long enough time period for which permanent income might be directly observed. Hence, there are two potential sources of bias: mismeasurement of annual income, and the timing of the income observations used. It is straightforward why mismeasurement in annual income is problematic, and may not accurately reflect permanent income. The timing of the income observations is important for comparability: if different individuals' income is measured at different age ranges, some individuals will be at the height of their earning potential while others are measured early in their careers. Yet, consistently measuring annual income across individuals but at an age that is not likely to be correlated with permanent income may lead to biased estimates; for instance, estimating intergenerational income relationships for individuals in their early 20's may not generate unbiased estimates (because income measured in one's 20's may be a biased measurement of permanent income).

These questions of measurement have generated wide-ranging scholarly debate in the intergenerational mobility literature. Atkinson (1980) and Solon (1992) point to measurement error in parental income and nonrandom sampling as a source of downward bias in estimates of β, and both Solon (1992) and Zimmerman (1992) describe how the use of short-run income as a measure of long-run income induces an errors-in-variables bias. The discrepancy between short-run and long-run income can be caused by idiosyncratic shocks to annual income, or by systematic differences in income by stage of life (a lifecycle bias). To the extent that the measurement error comes through annual fluctuations in income, averaging the income measures across a number of years will reduce the effects of annual fluctuations on the income measures; an alternative strategy would be to use instrumental variables to correct for measurement error in parental income. Also, Solon (1992) uses the Panel Study of Income Dynamics (PSID) data as a nationally representative survey to overcome previous difficulties with non-random sampling (e.g., Sewell and Hauser, 1975); however, Atkinson (1980) cautions that such a survey may not be immune to attrition or endogenous adjustment in individual behavior because these surveys operate through repeated measurements of respondents (continued participation).

The lifecycle bias is more challenging to address. Haider and Solon (2006) show that the correlation between yearly earnings and lifetime permanent income varies by age, and that this correlation is lower during early adulthood and higher around age forty. They also show that the use of temporary income as a proxy for permanent income attenuates estimates of β, and consistent with their evidence that the correlation between yearly and permanent income is highest around age forty, the size of the attenuation bias is smallest around age forty. Their evidence of a lifecycle bias contradicts the conventional belief that deviations between short-run and long-run income are caused only by transitory shocks. Additionally, Mazumder (2005) shows that if there is autocorrelation in the measurement error, which is consistent with the lifecycle bias of Haider and Solon (2006), that standard methods will substantially underestimate the intergenerational elasticity. As a result, both Mazumder (2005) and Haider and Solon (2006) recommend using averages over periods longer than five years for both individual and parental income, and the results from Haider and Solon (2006) specifically point towards measuring income around age forty.

More recently, Mitnik et al. (2015) have reignited debates over the details of measurement and estimation in mobility models, arguing that the elasticity of the conditional mean of income with respect to parental income is a better measure of mobility than the elasticity of the conditional geometric mean of income, though the latter is what is technically estimated in the log-log regression model.[2] However, in practice the difference between the standard approach and their preferred method is fairly small. In contrast, Chetty et al. (2014b) have argued that the combination of a rank-based measure of income and a large dataset alleviate concerns over attenuation and lifecycle biases, but Mazumder (2015) shows that the Chetty et al. (2014b) estimates do not necessarily eliminate these biases. Through sample selection criterion analogous to those in Chetty et al. (2014b), Mazumder (2015) shows that the income measures those authors use tend to produce a lower intergenerational elasticity estimate, which he attributes to these sources of bias. He notes, however, that these biases are relatively small.

An alternative strategy is to use instrumental variables to overcome multiple sources of measurement error (e.g., Solon, 1992; Zimmerman, 1992; Lucas and Kerr, 2013; Lefranc et al., 2014). By the summary in Table 2, the instrumental variables approach does not provide much insight beyond the traditional (non instrumental variables) estimates of β. Apart from these papers, however, little work has been done to use instrumental variables to resolve identification concerns, in part because finding relevant and valid instruments for parental income is difficult.

What do the debates over measurement error tell us about intergenerational mobility? The statistical analyses in Mazumder (2005) and Haider and Solon (2006) suggest that measurement error leads to a substantial downward bias in estimates of β. This conclusion bears important implications for measuring income and interpreting estimates of β, and also suggests that the intergenerational elasticity is likely to be higher than indicated by traditional estimates. For instance, this measurement error may have contributed to the earlier findings that the United States is relatively mobile (i.e., a relatively small β).

[2]This point is somewhat technical, and we direct the reader to Section 3 of Mitnik et al. (2015). The crux of the argument is that the log transformation of income makes the intergenerational elasticity estimate a geometric expectation instead of the arithmetic expectation that is consistent with economic theory of economic mobility.

2.2.3 Heterogeneity in Intergenerational Mobility Estimates

Another refinement of the basic model which frequently been considered is the question of whether the intergenerational elasticity is constant across groups and over time – in essence, the extent to which the intergenerational income relationship is heterogeneous. Related to this point is a literature that has evolved by trying to understand factors or mechanisms that are reflected in β.

Trends in Mobility

To what extent is the intergenerational elasticity (or correlation) in income changing over time? It is conceivable that structural changes in an economy over time lead to changes in the intergenerational elasticity (or correlation). One common question, for instance, is whether the United States is becoming more intergenerationally mobile.

Levine and Mazumder (2002), Hertz (2007), and Lee and Solon (2009) use repeated cross-sectional models to test for trends in the intergenerational elasticity, but they find no evidence of a trend. Aaronson and Mazumder (2008) offer a dissenting perspective: they create synthetic parental generations from average values of the census data, and use these synthetic parents to estimate the individual-parent correlation. They find that mobility is decreasing over time. Lee and Solon (2009) have argued against this result, noting that the validity of the intergenerational elasticity estimates computed using the synthetic data rests on the assumption that the unobserved families are relatively close to the average of census data observations. Using the population of intergenerationally-linked tax data for people born over several years in the 1980's, Chetty et al. (2014a) confirm earlier findings of stable mobility in the late 20th century United States.

International Comparisons

Scholars of intergenerational mobility have also developed an interest in comparing mobility between countries, as the relative ease of obtaining data for OECD countries makes potentially illuminating comparisons a natural extension of the literature. Table 2.3 summarizes these results, which generally suggest that the United States is less mobile than many European countries, particularly Nordic countries, and that the intergenerational elasticity falls between approximately 0.15 and 0.5 in the developed world.

Solon (2002) cautions that the methodological challenges associated with intergenerational mobility estimates may be compounded by cross-national comparisons, as the patterns of bias may vary between countries. He cautions that cross-country differences may have more to do with data preparation and modeling choices than with true cross-national differences. In general, researchers making these international comparisons have been circumspect in making claims about what drives these differences, because data limitations combined with statistical and conceptual difficulties makes testing hypotheses about these differences difficult. Corak (2006) highlights returns and/or access to education as an important component of intergenerational mobility, noting that both rose dramatically through the 20th century in the developed world. Blanden (2013) finds a positive association between mobility and educational spending, as well as a negative association between mobility and both returns to education and inequality.

Heterogeneity in Mobility Across Groups

When considering heterogeneity in β, it becomes clear that many different possible mechanisms are built into $\widehat{\beta}$; parents affect the economic status of their children in many ways. For instance, Becker and Tomes (1979) describe the intergenerational transfer of family background and resources; Solon (1999) describes the division of family resources among different children in the household; and Durlauf

(1996) develops a model in which parents sort into neighborhoods that have different characteristics that influence the intergenerational transmission of economic well-being. Hence, the mechanisms by which parents affect the economic status of their children are both numerous and difficult to observe (or measure). The homogeneous β is some agglomeration of these forces, but may not be representative of the intergenerational link for any particular individual-parent pair. Keeping in mind that the intergenerational elasticity parameter is fundamentally a descriptive statistic, allowing for heterogeneity within the econometric model does not allow one to draw causal inferences about how family background influences adulthood outcomes. Nevertheless, econometric heterogeneity moves beyond the baseline intergenerational elasticity model by generating subpopulation parameters that may lead to the generation of new hypotheses related to intergenerational mobility for these subgroups of individuals, although intergroup heterogeneity in this descriptive statistic cannot answer these questions itself.

In the simplest sense, modeling heterogenity in β amounts to estimating group-specific mobility parameters, as in Hertz (2005), where an intercept offset for black individuals implies differences in transition probabilities for this population, even in the presence of a relatively small difference in the level of black and white intergenerational elasticities. A natural way to model these differences is by specifying interactions between socio-demographic characteristics and parental income in the intergenerational regression model. Research in this area has indeed found that different groups experience different values of the intergenerational elasticity. A more ambitious goal is to try to capture causal effects of neighborhood characteristics as mechanisms through which intergenerational income links are forged. We return to this issue after presenting our review of the neighborhood effects literature in the following sections.

2.2.4 Contributions from Economic History

The need for many generations of data plays to a comparative advantage of economic historians, who have made substantial advances in matching individuals to their forebears, sometimes tracing families as far back as 1850. Like the other studies we discuss, historical economic research on intergenerational economic mobility focuses on links between individuals; yet, this research is unique in providing inventive ways to learn about intergenerational economic mobility, marking recent advances in the economic history of mobility as an exciting frontier in the mobility literature. We summarize this work in Table ??.

This literature makes three contributions. First, using further-reaching datasets provides data on additional generations, allowing researchers to estimate grandparent and great-grandparent effects that are not feasible to study using PSID, NLSY, or tax return data (Chetty et al., 2014a). Second, the focus on change over long periods of time allows researchers to study sources of mobility in ways that complement the traditional mobility literature. For example, asking whether or not the Great Depression affected mobility (Feigenbaum, 2015) or measuring mobility among the descendants of former slaves in the early 20th century (Collins and Wanamaker, 2017), is impossible with datasets such as the NLSY and PSID. Finally, economic historians have also developed a unique instrumental variables strategy based on similarities among names that provides alternative estimates of intergenerational mobility that are much higher (intergenerational elasticities of around 0.75 or higher) than standard estimates.

Scholars are interested in the role of grandparents on intergenerational economic mobility because grandparents often make monetary investments in their grandchildren, because of genetic bequests across generations, and because there may be interactive effects between parents' and grandparents' incomes (Solon, 2014, 2015). This process matters for the rate of income convergence across generations, as a positive grandparent effect implies slower convergence.[3] The existence of a grandparent correlation conditional on a parental correlation also implies, contrary to the canonical Becker and Tomes (1979) model, that intergenerational income relationships last longer than a single generation. This literature is relatively new, however; a broad consensus has not emerged on the size and interpretation of the grandparent coefficient. Ferrie et al. (2016) find positive and significant elasticities, but acknowledge measurement error concerns, while Olivetti et al. (2018) and Modalsli (2016) find positive grandparent elasticities. These results suggest that the typical parent-child model is missing the child-grandparent link.

The economic history literature also focuses on long term trends in mobility, based on datasets that stretch back farther than the PSID and NLSY surveys. In addition to estimating and explaining long term trends, this literature provides insight into how major events in U.S. history, such as the Great Depression and the Civil Rights Movement, affected economic mobility. Evidence suggests that mobility did decrease between the early 20th century and the 1970's in the U.S.: Feigenbaum (2018) estimates the intergenerational elasticity to be about 0.2 during this period, while Hilger (2017a) finds a gradual decrease in the intergenerational elasticity between 1930 and 1970, and an increase after 1970. This research has also uncovered new directions for future work. Hilger (2017a), for example, finds that many changes in U.S. law and society – such as the G.I. bill, the Civil Rights movement, and the integration of the education system – had a limited impact on mobility rates, but that rising incomes and equality did have an impact. Emphasizing the role of racial disparities, Collins and Wanamaker (2017) find that since 1880 African-Americans have experienced lower rates of upward mobility than white Americans across the income distribution, and attribute this finding to persistent differences in human capital attainment that is related, in part, to neighborhood differences.

[3]As Solon (2014) shows, following Becker and Tomes (1979) in assuming mobility follows an AR(1) process is necessary for the IGE to be estimated with β alone; under an AR(2) process, the IGE is a function of both parent and grandparent parameters, and it increases in the grandparent coefficient, implying slower convergence.

Analyses using long-term data is important because through these data scholars directly observe the evolution in intergenerational economic mobility over a long time horizon, and also allow researchers to link changes in this evolution to discrete, historical events. Hence, this kind of analysis can more credibly relate change in the intergenerational elasticity to structural changes in the economy, compared to cross-sectional or cohort analyses of short term trends in mobility. There are certainly many unanswered questions in this segment of the literature, but these papers, as well as those summarized in Table ??, provide a foundation for a complementary perspective on these different approaches.

The final strand of economic history research that we discuss here approaches intergenerational mobility in a totally different way, capturing persistence in economic status between individuals in different generations matched using rare surnames. Clark (2014) – the earliest application of this approach – studies people and groups with distinctive surnames, often arising from specific circumstances in their countries' histories, and finds that individuals with distinctive surnames associated with high income and wealth are massively overrepresented in high prestige occupations and in upper wealth and income quantiles. He shows that his results are consistent with an intergenerational elasticity of approximately 0.8 across several countries, suggesting a much lower level of short-term mobility than the mobility literature building on labor economics; he notes that this value is consistent over time, implying long-run reversion to the mean.

Yet, the surname approach is subject to methodological criticisms based on the distinction between group-based and person-based mobility measures. Torche and Corvalan (2016) argue that comparing the surname-based aggregate measures with intergenerational elasticities estimated using individual data is invalid, because Clark's use of surname averages conflates group-level and individual mobility measures. Standard intergenerational elasticity models average incomes within individuals to reduce the influence of measurement error; Clark argues that he averages within surnames for the same reason, but Torche and Corvalan argue that this shifts the unit of analysis to the surname groups. It is important to account for within-individual variation, but partialing out within-surname variation results in a regression that finds within-group persistence by construction. The implication is that Clark's claim that standard intergenerational elasticity estimates overstate mobility are incorrect, as his measure does not capture the same relationship as the traditional intergenerational elasticity measure. Finally, between-group estimates like Clark's are typically much higher than individual analogues (Torche and Corvalan, 2016).

2.2.5 Alternative Model Specifications

Capturing Non-Monetary Influences

The mobility literature, going back to Becker and Tomes (1979, 1986) relies on the intergenerational elasticity as its primary summary statistic not because scholars expect parents' money to define adulthood outcomes per se, but because microeconomic theory suggests that an individual's permanent income is related to his/her parents' permanent income through both direct investment and non-monetary inputs. Several authors have examined the distinction between these categories of mobility-relevant inputs. Mayer (1997) uses data from a variety of sources – including Great Society era anti-poverty interventions – to study the intrinsic role of money. She finds that doubling parental income is associated with only a modest improvement in a child's educational outcome, but that non-monetary factors associated with income (e.g., residence in a single parent household, parental psychological help, or parental time investment) are closely associated with a range of outcomes, such as employment, educational attainment and achievement, or the avoidance of teenage pregnancy. In a similar vein, Shea (2000) exploits random shocks to family income ('luck') during an individual's childhood (specifically, temporary unemployment, changes in pay, or industry of employment) to identify the effect of parental money income on individual income later in life. He finds that the relationship between instrumented parental income and individual income is not significant, indicating that statistically significant intergenerational correlations reported in other studies may be driven by factors other than income – for instance, non-monetary bequests. Lefgren et al. (2012) take a different, but related approach in estimating a three-factor intergenerational elasticity, in which the factors are 'luck' income and human capital factors. They evaluate the bounds of each factor's contribution to the overall elasticity estimate, and find that, at most, income itself accounts for 37 percent of the intergenerational elasticity, with the rest coming from human capital transmission. Their results are consistent with Shea (2000) and Mayer (1997).

The Sibling Correlation Model

One interesting alternative to estimating β in the benchmark model is to measure the correlation in income among sibling pairs. A high sibling correlation implies a high level of family determinism and low individual mobility, analogous to a high intergenerational elasticity measured with a high β. Yet, unlike the benchmark (individual-parent) model, the sibling model does not require as long of a time series, and is relatively robust given that siblings typically have a common childhood background. Further, the sibling model accounts for sources of family correlation beyond parental earnings or income (Solon, 1999), though decomposing a sibling correlation into different potential family mechanisms can be difficult. Solon (1999) shows how the sibling correlation model can be written as a function of the individual-parent correlation and a second term that represents additional sources of within-family correlation. The sibling correlation model is a complement to the individual-parent model, because an estimate of the individual-parent correlation makes it possible to derive the second component of the sibling correlation model, which captures the impact of common influences not related to family income.[4]

[4]Specifically, Solon shows that $\sigma_a^2 = \rho^2\sigma_x^2 + \sigma_z^2$, where σ_k^2 refers to the variance for variable k, a denotes the family component of a sibling correlation, ρ is a coefficient, x is permanent income, and z is a set of common family factors unrelated to income. Dividing by the variance of individual income, σ_y^2, yields $\frac{\sigma_a^2}{\sigma_y^2} = Corr(y_{ij}, y_{ij'}) = (\rho^2\frac{\sigma_x^2}{\sigma_y^2}) + (\frac{\sigma_z^2}{\sigma_y^2})$, where y_{ij} and $y_{ij'}$ denote incomes for each of a pair of siblings, which in turn can be written as $Corr(y_{ij}, y_{ij'}) \approx \rho^2 + (\frac{\sigma_z^2}{\sigma_y^2})$. This final expression

The sibling pair approach has produced substantially different results than the intergenerational approach. As early as the 1970's, empirical work employing this framework estimated sibling correlations to be as high as 0.34 (Solon et al., 1991) and 0.44 (Brittain, 1977), suggesting less mobility than contemporaneous intergenerational elasticity estimates. Since the sibling correlation model captures sources of within-family correlations beyond income, the interpretation of these estimates is not analogous to the interpretation of the intergenerational elasticity estimate from the individual-parent model; nevertheless, these estimates suggest a markedly lower level of mobility than the early intergenerational elasticity estimates of 0.1 to 0.2. This lower level of mobility may be driven by both income and non-income sources.

A Role for Genetic Bequests?

Another interesting angle that has been studied is the effect of genetics on economic outcomes. Some scholars have studied the role of genetics in economic mobility through samples of adopted and biological children, in order to decouple prenatal and genetic factors from other factors that drive economic outcomes. Björklund et al. (2006) find a significantly positive correlation between an individual's educational attainment and the educational attainment of his/her parents, regardless of adoption/biological status, and some evidence of a correlation with his/her biological father's income and not the income of an adopted father. Sacerdote (2007) finds the opposite, concluding that genetics and prenatal health are more than twice as important as family environment for explaining variance in education and income. Similarly, Liu and Zeng (2009) estimate a significantly higher intergenerational elasticity for children living with their biological parents (0.37) than for children who were adopted (0.10), which confirms that there may be a large genetic/perinatal component, if we assume that parents invest equally in adopted and biological children.

states that the sibling correlation is the sum of the intergenerational correlation and the ratio of the variances of income and non-income inputs during childhood. This second term is what we discuss in the text.

These results are fairly striking, and suggest that there is a strong biological component to economic mobility and potentially only a limited role for environmental factors. It is worth bearing in mind that it is difficult to differentiate between perinatal health and genetic bequests; these results suggest that policymakers interested in economic mobility and opportunity may be wise to consider mother and child health soon before and after birth. The implications of the empirical results regarding genetics per se are harder to parse; one implication is that genetics cannot explain all the variation in the economic outcomes of a child. It is also worth noting that while it is difficult to differentiate between genetics and perinatal environment, only the latter can be influenced by policy.

2.3 Early Investigation of Neighborhood Impacts

2.3.1 The Scope of the Literature

A parallel literature is based on the notion that an individual's neighborhood influences individual economic outcomes by affecting this process is not controversial, and a vast social science literature points towards "reduced buying power, increased welfare dependence, high rates of family disruption, elevated crime rates, housing deterioration, elevated infant mortality rates, and decreased educational quality" (Massey, 1990, p. 342) as important influences on these human capital inputs. Seminal papers – summarized here and in Table 2.5 – have explored racial and class segregation, social externalities, and local structural economic change as drivers of heterogeneity in neighborhood characteristics across the United States.[5] Yet, the early empirical literature investigating the role of neighborhood characteristics on the economic status of individuals and intergenerational economic mobility struggled to surmount the econometric challenge of distinguishing between correlations describing equilibria and causal effects. Many empirical analyses instead describe a particular

[5] Social externalities refers to a self-reinforcing dearth of community engagement and weakens community cohesion.

mechanism through which neighborhood characteristics might affect economic mobility, and provide statistical evidence that is consistent with that mechanism. To proceed, we withhold discussion of the experimental and observational relocation studies until Section 2.4, to first describe several important mechanisms through which neighborhoods affect individual opportunity and intergenerational economic mobility.

2.3.2 Key Empirical Findings

In a seminal theoretical contribution, Loury (1977) describes the neighborhood as a factor that influences childhood absorption of human capital both through the inputs the neighborhood provides, and through its role in forming an individual's expectations about the returns to certain activities, such as schooling. In this model, neighborhood characteristics can have far-reaching effects on the lifelong economic status of an individual, even holding parental attributes constant. In an effort to understand these mechanisms, scholars have found evidence that neighborhood characteristics significantly correlate with individual economic outcomes such as employment, income, single-parenthood, and labor market participation. Table 2.5 summarizes some important contributions; we conclude that neighborhood characteristics are at least associated with long-term differences in individuals' economic status.

Neighborhoods and Opportunity

Some scholars have noted the relevance of community characteristics for attempts to understand intergenerational mobility. Building on sibling correlation models common in this literature, Solon et al. (2000) and Page and Solon (2003) estimate neighborhood correlations in both educational attainment and men's earnings. They find robust associations that suggest that neighborhood characteristics have a significant impact on long-run educational attainment and earnings, linking individual economic mobility and the community, although the neighborhood correlation is much weaker than the family correlation. Brooks-Gunn et al. (1993) and Leventhal and Brooks-Gunn (2000) find relatively strong neighborhood effects. One important difference between these sets of papers is that Solon et al. (2000) and Page and Solon (2003) are agnostic about the underlying neighborhood effects mechanism. Page and Solon (2003) find evidence that growing up in an urban area is associated with higher earnings in adulthood, but they are unable to determine where this relationship comes from, and so they do not point to any specific mechanism through which neighborhoods influence mobility. In contrast, Brooks-Gunn et al. (1993), Leventhal and Brooks-Gunn (2000), and others, emphasize the role of peer influences stemming from neighbors' characteristics (e.g., income, education, single parent status, or drug use). Nevertheless the empirical findings indicate that neighborhood characteristics affect individual economic outcomes, and therefore play a role in intergenerational mobility.

Racial Discrimination and Segregation

Other scholars have approached neighborhood effects by studying statistical differences between racial and ethnic groups, largely because of the degree of racial segregation common in American cities prior to and during the 1970's. This research has wide-ranging implications for the mobility literature, both because of its implications for neighborhood effects, and because a substantial racial disparity in income and socioeconomic opportunity bears strong intergenerational implications. Wilson (1987) and Massey (1990) provide seminal contributions in the development of this race-place-opportunity literature. Wilson (1987) argues that economic structural change, particularly the combined decline of manufacturing and the out-migration of middle-class blacks from the distressed inner city neighborhoods, led to concentrated black poverty in inner cities. Massey (1990) argues that while structural economic change provided the impetus for black poverty, pre-existing racial and class segregation was the necessary condition for the focus of those changes to be on the inner-city black population. It is both interesting and important to recognize that segregation need not coincide with racial discrimination for it to foster localized poverty. Durlauf (2004) notes that segregation arises when "individuals have a preference [for] others of similar ethnicity" (page 2197), which need not develop out of racial discrimination; but despite its origin, "segregation creates a structural niche within which a self-perpetuating cycle of minority poverty and deprivation can survive and flourish" (Massey, 1990, p. 350).

A number of papers have empirically studied racial disparities, often focusing explicitly on segregation as a measure of disadvantages faced by minority groups. Cutler and Glaeser (1997) evaluate the effects of segregation on city-level black employment, income, and family structure, and find evidence that blacks in segregated areas suffer worse outcomes than blacks in less-segregated areas. They conjecture that racial segregation coincides with income segregation, negative social externalities (e.g., lack of positive role-models), as well as structural economic disadvantage (e.g., workers live farther from their jobs); however, the negative association between segregation and economic outcomes seems to pertain only to data collected from the 1990 census (Collins and Margo, 2000). Ananat (2011) finds that segregation increases poverty and inequality for the black population (relative to the white population), but has the opposite effect on whites. The positive effect on whites is consistent with intuition in Massey (1990): while segregation magnifies the harmful effects of economic shocks on vulnerable groups in the population, especially minorities, segregation also insulates relatively well-off groups. Massey et al. (1991) show that joblessness, teenage pregnancy, and single motherhood are associated with segregation, which is in turn influenced by the economic structure of the city, and Massey et al. (1994) find evidence that the geographic concentration of poverty in inner cities is, at least in large part, the result of racial segregation in housing.

Borjas (1992, 1995), and Vigdor (2002) provide a different perspective, focusing on social externalities that result from ethnic clustering. Borjas (1992) focuses on the intergenerational effects of ethnic community characteristics as suggested by

Loury (1981), which he terms "ethnic capital." Ethnic capital is, essentially, a human capital externality, arising from the influence of members of a particular ethnic group of younger members, through their roles as role models, vectors of information, etc., which affects outcomes for the next generation within that group. Borjas (1995) expands this work by using neighborhood-level measurements of ethnic capital as well as neighborhood fixed effects, and concludes that much of the effect of ethnic capital estimated in Borjas (1992) was, in fact, driven by ethnic sorting into neighborhoods, rather than by the human capital externalities. Vigdor (2002) uses the concept of ethnic capital as a theoretical foundation on which to reconcile inconsistent findings regarding the relationship between segregation and economic outcomes for blacks over time (Cutler and Glaeser, 1997; Collins and Margo, 2000). He argues that the self-selection of skilled workers away from segregated cities induces changes in the level of ethnic capital, which affects intergenerational mobility.

The literature on race and segregation is, also, somewhat limited in its discussion of the mechanisms by which the observed effects occur, but it still has valuable implications for our understanding of race and opportunity. First, racial disparities exist, and working to understand those disparities within the broader context of mobility and space may be helpful in clarifying which mechanisms may influence both intergenerational mobility and racial disparities in economic outcomes. An important policy implication arising from this literature is that understanding the role of racial (or class) segregation must be an integral part of efforts to develop policy prescriptions designed to overcome neighborhood poverty in the context of the United States. Second, some papers seem to side with either a structural or social explanation; that is, focusing on a mechanism rooted in interactions between neighbors, either in the 'ethnic capital' sense (Borjas, 1992) or in the more traditional peer effects terms (Brooks-Gunn et al., 1993). Yet, one need not accept an either-or framing: it is likely the case that both structural and social effects are active determinants of individual economic mobility.

2.3.3 Focusing on Rural Poverty and Regional Economics

The micro-level literature on poverty and space – particularly as it pertains to opportunity and mobility – has focused primarily on urban areas, as much of the earlier literature is motivated by urban racial segregation. A parallel strand of the literature focuses on the economic status of places rather than people, and this literature is more engaged with rural poverty. Paying close attention to rural areas is important: Farrigan and Parker (2012) note that the persistent poverty of the same geographic groupings of U.S. counties (e.g., in the Mississippi Delta, Appalachia, the 'Cotton Belt' of the Southeast) bears some conceptual similarity to the concentration of urban poverty that motivated Wilson (1987). Beyond this shift in focus, engaging with the regional economics literature adds additional context in the discussion of local economies. Studying poverty of places, instead of individuals per se, makes clear that problems arising from region-based frictions may be solvable through a different set of levers than problems associated with individual or neighborhood poverty.

We summarize some results from this literature in Table 2.6. Several themes emerge from the regional poverty literature: an emphasis on labor market and spatial equilibria, a concern over whether persistently or extremely poor regions respond differently to economic change than areas that are not so poor, and an interest in place-based – in contrast to people-based – policy solutions. The equilibrium emphasis manifests itself in an emphasis on the role of remoteness as a mechanism for regional poverty by influencing the equilibrium distribution of worker types and firms, and in the presence of structural equations models that model markets for labor, goods, and housing simultaneously (e.g., Wu and Gopinath, 2008; Gebremariam et al., 2011, for recent examples). Whether regional poverty responds differently to changes expected to foster growth, such as an increase in employment, is related to the idea of a poverty trap: in the presence of a poverty trap, benefits from positive developments may be blunted by the frictions creating the trap. Both of these themes naturally lead to an interest in place-based policy, because if regional poverty is the outcome of a spatial equilibrium, and if some economic changes are associated with stronger or weaker reductions in poverty rates, a policy aimed at shifting this equilibrium directly might be more effective than a policy that targets individuals. Since the community variables determined in equilibrium seem to contribute to mobility and opportunity, this regional literature is germane to these policy discussions.

As Table 2.6 shows, spatial remoteness is correlated with higher poverty rates, though persistently poor and remote counties do not respond differently to economic growth than other counties, suggesting that place-based policies may be effective. On the other hand, structural economic models suggest that amenities (e.g., access to attractive natural features, a more pleasant climate, and/or urban cultural amenities) influence the spatial distribution of workers, and so place-based interventions may be less effective if they are unable to alter the spatial distribution of amenities.

Studies in this literature typically employ national census data that is both large and geographically representative. An alternative that has become popular in recent years is to study the effects of specific place-based policies. A literature has developed around "Empowerment Zones," a policy intervention by which a state or federal government offers tax credits to firms in a particular distressed region, though evidence suggests that these interventions are not effective in reducing poverty (Hanson and Rohlin, 2013; Lockwood-Reynolds and Rohlin, 2015). Results suggest that at a higher level of aggregation (i.e., the county level) the structural perspective is most important, not social externalities, as the factors that seem to matter have to do with production structure and migration rather than compounding social factors, such as poverty rates or demographics. Like Wilson (1987), authors in this literature emphasize the role of structural economic forces rather than social factors.

A Brief Conclusion

As a whole, the early neighborhood effects literature provides a detailed descriptive analysis of multiple factors that connect an individual's childhood environment to his/her socioeconomic outcome in adulthood. As Table 2.5 shows, there is strong empirical evidence that neighborhoods (or communities) play a role, even if the mechanisms driving these links are difficult to disentangle from each other. We also notice a particular emphasis on race as a factor that has complex links to neighborhood characteristics, and questions surrounding this "race-place" intersection are far from resolved. We believe that future research continuing to explore the race-place-mobility domain may yield valuable insight.

2.4 New Directions in the Mobility Literature

Recently, the conceptual and methodological orientation to studying intergenerational mobility shifted. Scholars increasingly approached questions of intergenerational mobility in terms of the causal effects of childhood circumstances on adulthood outcomes. For example, Chetty and Hendren (2015) focus on the (hypothetical) effect of growing up in a different county or commuting zone on adult income, obtaining precise estimates of neighborhood effects. This approach involves estimating local intergenerational elasticities, and uses geographic-specific intercept shifts to capture differences in absolute mobility; hence, emphasis is placed on geographic variation in intergenerational mobility, rather than on a single parameter. Such an analysis has implications for mobility and opportunity because it connects circumstances determined by previous generations (such as neighborhood choice) to individuals' outcomes.

2.4.1 Relocation Experiments

In the early 1990's scholars began using randomized relocation experiments to identify the effects of neighborhood characteristics on long-term economic outcomes such as earnings or educational attainment. Exploiting randomized relocation experiments is attractive because the randomization generates exogenous variation in neighborhood characteristics across individuals, allowing researchers to circumvent the endogeneity problem that motivated the equilibrium models of social interactions. Popkin et al. (1993) were the first to exploit such randomization as a source of identification via the Gautreaux Program, a relocation program for individuals living in Chicago's public housing. They found that treatment resulted in improved labor market participation, but had limited effects on wages. These findings are consistent with the spatial mismatch hypothesis – that poor workers in inner cities often struggle to find employment because of high transaction costs which foster continued unemployment (e.g., monetary costs of transportation, time costs of searching for work, and working in an unfamiliar place) that arise because of their location (Kain, 1968).

More recently, a series of studies have analyzed data from the Moving to Opportunity program (MTO), another randomized relocation program similar to the Gautreaux Program.[6] Follow-up interviews with members of treated households lead to the construction of a longitudinal data set, making the data increasingly useful for identifying the causal effects of neighborhood exposure on long term outcomes in an experimental setting. In Table 2.7, we summarize several studies that have exploited randomized relocation experiments to assess the impact of neighborhood characteristics on economic outcomes. This literature has investigated whether there are long-run neighborhood impacts on adults and also whether there are medium-run

[6]The MTO sample includes approximately 4,600 low income families living in poor census tracts in Baltimore, Chicago, Boston, Los Angeles, and New York City. In the mid 1990's, these households were offered housing vouchers that would allow them to move to a better neighborhood, and some households offered vouchers faced the restriction of having to move to low poverty tracts. See, for example, Katz et al. (2001) for further details.

impacts on teenagers; in general, studies have not found much evidence that housing relocation led to improved economic status.

In contrast to many papers in this area, Chetty et al. (2016) find evidence in support of the neighborhood effects hypothesis, using data that has only recently become available as individuals exposed to MTO as young children reach adulthood. They study the long-term effects on a group of individuals not studied in the previous papers – children who were younger than thirteen at the time of relocation – and find significant effects of relocation on a number of economic outcomes. Their findings are consistent with previous conclusions that MTO did not significantly affect individuals who were older teenagers at the time of relocation. Yet, the authors do not find an age-threshold at which relocation becomes effective; the relocation effect is decreasing in age at the time of relocation. From this evidence, the authors surmise that there are two opposing effects of relocation: a fixed developmental 'cost' of relocation arising from the disruption caused by moving (e.g., stress from moving and/or having to acclimate to a new social environment), and the positive effect of exposure to a better neighborhood. Since younger children ultimately spend more time in the better neighborhood, they can better overcome the disruption costs, even if the benefit of relocation is positive and constant across the age distribution. By contrast, an individual who relocates at an older age may show a moving effect that is close to zero (or even negative) as the cumulative positive effects of relocation may not be sufficient to overcome the disruption costs.[7] Chyn (2016) finds a similar result using data from randomized demolitions of housing projects in poor Chicago neighborhoods, and shows that relocation improves economic self-sufficiency and reduces the probability

[7] Chetty et al. (2016) also discuss several new mechanisms that might drive the effects of MTO. First, they show that MTO children live in neighborhoods that are slightly richer and slightly less black than their previous neighborhoods, and interpret this evidence as evidence of segregation leading to poverty in the original neighborhood. They also revisit the gender heterogeneity observed in Kling et al. (2007) and Ludwig et al. (2013) and find effects are generally similar between genders, though one important area in which boys fare worse than girls is in the incidence of risky behavior. Chetty et al. (2016) speculate that the richer, safer environments to which MTO boys relocated punish risky behavior less severely.

of being arrested among individuals treated as children or, in contrast with Chetty et al. (2016), as teenagers.

Yet, Durlauf (2004) cautions against drawing overzealous policy prescriptions from these relocation studies, because they ignore the possibility of general equilibrium effects that arise from the relocation of a large number of poor people to non-poor neighborhoods. The fact that MTO may have positive effects does not imply that poverty could be ended by moving everyone in poor neighborhoods to rich ones. Kling et al. (2007) argue that because the relocated individuals did not all relocate to the same area, the potential for general equilibrium effects is small; however, Chetty and Hendren (2015) note high frequencies of individuals relocating to certain popular areas. Nevertheless, it is clear that scaling up of relocation programs on the basis of certain successes of existing experiments may lead to general equilibrium effects, such as a change in behavior from relatively wealthy individuals (Durlauf, 2006).

Additionally, one might argue that the findings that relocation experiments do not induce positive effects for relocated individuals is inconsistent with the foundational theoretical perspectives of Wilson (1987) and Massey (1990), and of related empirical work (e.g., Borjas, 1992; Cutler and Glaeser, 1997). If the barriers to economic empowerment were caused by structural economic changes situated specifically in space (racial or class segregation, or ethnic capital) as previous research suggests, one would expect to see substantial effects of relocation; yet, the only effects are those for young children identified by Chetty et al. (2016). These findings are consistent with the neighborhood effects literature in sociology and psychology that finds significant neighborhood effects on health and behavior, particularly for women (e.g., Jencks and Mayer, 1990; Brooks-Gunn et al., 1993). Based on this empirical evidence, our perspective is that neighborhoods are important, as early work indicates, but the mechanics of these relationships seem to work through a more subtle channel than a direct correspondence between exposure to certain neighborhood characteristics and income.

2.4.2 Mobility and Neighborhood Inputs

Chetty et al. (2014a) study variation in opportunity in the United States by estimating mobility measures for specific geographic units.[8] They show that the chances of achieving upward mobility vary widely across space: in Charlotte, North Carolina, the probability of moving from the 25th percentile to the top half of the income distribution is 4.4 percent, while in San Jose, California, it is 12.9 percent. This difference is similar in magnitude to the difference between an intergenerational elasticity of 0.2 and an intergenerational elasticity of 0.4. This finding is particularly noteworthy because this variation suggests that there is a direct role for geographic mechanisms to influence economic mobility. Chetty and Hendren (2015) expand on this work, and find that moving to a more mobile place improves an individual's income later in life, and that at least 49 percent of the variation in adult income across space is based on differences in place of residence during childhood. Further, their model predicts, for example, that a year spent in Salt Lake City increases income at age 26 by \$135.90, while a year spent in New Orleans reduces it by \$175.30. Using these place-specific estimates, they correlate segregation, poor school quality, and a dearth of social capital as factors that suppress long-run incomes.

These papers emphasize geographic variation, and engage less with explaining the underlying mechanisms. Studies by Knapp and White (2016) and Bosquet and Overman (2019) take the opposite approach, emphasizing relationships between urbanity and rurality and mobility, and taking steps towards explaining these relationships. Knapp and White (2016) find that a negative association between childhood-county poverty rates and wages, and that this association is much stronger in rural counties relative to urban counties. Yet, the mechanism through which the youth

[8] Chetty et al. (2014a) use a rank-rank income regression that is conceptually similar to an intergenerational elasticity but measures movement across the income distribution rather than in dollar terms, to estimate the degree of mobility across U.S. counties and commuting zones. Their emphasis is on transition probabilities, particularly the probability of moving from the bottom quartile to the top, which they refer to as 'absolute mobility', in contrast to the 'relative mobility' denoted by the intergenerational elasticity. They argue that the intergenerational elasticity itself is a relative measure because all other things being equal, a given intergenerational elasticity characterizes how well a person born to poor parents will do compared to one born to rich parents.

poverty/rurality-earnings correlation arises is unclear; the authors argue that it is a matter of opportunity, but the same correlation could also arise from differences in service quality, peer effects, or from other household level unobservable forces.

Bosquet and Overman (2019) estimate adult wage elasticities with regard to the population size of individuals' birthplaces. They distinguish between three different mechanisms through which such an elasticity could emerge: *sorting*, whereby types of parents who choose to live in certain places produce a correlation; *learning*, whereby population size affects human capital accumulation; and *geography*, whereby birthplace affects migration decisions later in life, which contribute to location decisions, and thus the labor market the individual faces. They find a raw wage elasticity of birthplace size of 6.8 percent, and argue that this is driven primarily by sorting, because parents with professional occupations are substantially more likely to live in large cities. Since controlling for parental status removes the association between size and education, the authors argue that the elasticity does not operate through learning. Finally, the evidence on the geography hypothesis is inconclusive, although Bosquet and Overman (2019) do observe differences between movers and stayers (although this choice is, of course, endogenous to many other labor market outcomes). This outcome affirms the concerns continuously expressed throughout the neighborhood effects literature that apparent neighborhood effects can arise from sorting.

This segment of the literature remains small, but it has drawn some conclusions that help to reconcile conflicting findings from the neighborhood effects and mobility literatures. These papers suggest that places do play a substantial role in mobility, but that a large part of this story is heterogeneity between places, which may not be apparent in the earlier studies that suggested small to null neighborhood effects. Like the experimental papers, they estimate long-term outcomes rather than mobility parameters, and while this departure from mobility models provides valuable flexibility, it means that they answer questions of mobility only indirectly. Taken together, however, this body of work lays a foundation for further investigation into the sources of mobility, as opposed to its level, particularly with regard to the role of community inputs in determining individual mobility.

2.4.3 Some Open Questions

The work on relocation experiments has established that childhood exposure to certain neighborhood characteristics has long-term effects on individual socioeconomic outcomes, establishing a causal link between exposure to poor neighborhoods and economically important outcomes such as wage, education, and employment. Recent extensions of the traditional mobility literature – relying on descriptive statistics about mobility itself, rather than causal effects of childhood experiences – have also answered important questions, most notably by establishing that there are substantial spatial differences in economic mobility, along with provisional evidence about the sources of these differences. The shrinking gap between studies that estimate intergenerational elasticity parameters and those that focus on causal effects raises important conceptual and empirical questions. Research has established a relationship between place and mobility, but additional work is needed to answer questions about why mobility is higher in San Jose than in Baltimore or the Mississippi Delta as Chetty et al. (2014b) and Chetty and Hendren (2015) suggest. As we have discussed, it is difficult to answer such questions within the traditional mobility framework because the intergenerational elasticity parameters are descriptive statistics and not causal, while at the same time, the field experiment models have so far not been able to disentangle particular mechanisms.

2.5 Lessons for Research and Policy

So far in this review we have touched on three strands of research related to intergenerational mobility: studies that estimate individual level intergenerational mobility parameters, studies that explore the impact of neighborhood characteristics on economic outcomes, and studies that investigate mechanisms driving rates of mobility. The evidence generally indicates that there are significant individual and neighborhood factors that influence long term economic outcomes, although only a few studies can conclusively identify the mechanisms through which these factors affect outcomes. In this section, we focus on the road ahead, asking two questions: "What is next along the research frontier?" and "What policy lessons can we draw from this vast body of research?"

2.5.1 Insights from Theoretical Models of Poverty Traps

In addressing these two questions, we believe it is useful to consider the theoretical foundations from the of poverty trap literature. Thinking about mechanisms through which poverty traps may influence the process of intergenerational mobility is one way to gain some clarity as to the different types of mechanisms and policies that are related to intergenerational mobility. More importantly, whether or not economic outcomes are subject to true poverty traps – defined in Azariadis and Stachurski (2005) as "self-reinforcing mechanisms that cause poverty to persist" – bears important implications for the types of policies that might improve economic outcomes. This point deserves attention because it has serious implications for our understanding of the extent to which history influences long-run steady-state outcomes.

An overarching conclusion from simple models of poverty traps is that multiple frictions need to exist simultaneously in order to generate a poverty trap (i.e., a stable steady state at which individuals remain in poverty; Azariadis, 2006; Ghatak, 2015). While this conclusion may seem counterintuitive, it is important to understand that frictions such as capital market imperfections or nonconvexities in returns to labor do not by themselves prevent two individuals from converging to the same long run steady state (Ghatak, 2015); in other words, these factors need not lead to a poverty trap. This distinction is important and separates the idea of poverty persistence from a poverty trap: a person, place, or even country could be consistently poor over a long period of time, but not necessarily be afflicted by a true poverty trap. In such a case, a relatively simple shock such as a cash transfer, an increase in local employment, or improved technology, may be sufficient to put the individual back on the path toward the long-run steady state (and away from poverty).

Yet, many types of frictions might lead to a poverty trap, and Ghatak (2015) classifies these mechanisms into external and preference-based frictions.[9] From this perspective, understanding interactions between individuals, families, and communities is key because in the case of a 'social' poverty trap the equilibrium depends on the relative strength of social and individual behavioral incentives. The predictions of these abstract models appear to be consistent with the empirical literature on individual mobility and neighborhood effects: scholars have found a wide variety of neighborhood associations that correspond to both external and preference-based explanations. In addition, the presence of imperfect markets for noncapital goods (such as human capital) or limitations on intergenerational transfers, the perfect capital markets assumed in the standard Becker and Tomes (1979) model will not guarantee income convergence. A main implication is that preference or scarcity-based frictions can generate poverty traps even in the absence of external frictions, just as other market imperfections make poverty traps possible in the presence of perfect capital markets.

What do these models tell us about how to think about mobility, neighborhoods, and future empirical research? One answer is that we need a better understanding of *how* both families and neighborhoods affect human capital accumulation, combined with an understanding of how nonconvexities in returns to human capital affect permanent income. The findings in Chetty et al. (2016) are illustrative, as a large moving effect for children is consistent with the poverty trap hypothesis, because this treatment changes many factors at the same time. The paper shows that MTO had a substantial effect on educational attainment and income of relocated children, but whether this effect arises from differences in the allocation of educational quality,

[9] Specifically, the ideas that the poor are rational but subjected to adverse constraints (external) vs. the notion that the poor make different choices because they are poor (preference-based).

differences in one of a variety of possible peer effects, or from another unobservable factor (such as parental valuation of education that may influence both participation in MTO and these outcomes) remains unclear. This ambiguity emerges precisely because treatment changed many factors simultaneously, and so an effect driven by a change in a single factor and an effect driven by breaking a multifaceted poverty trap are observationally equivalent.

The poverty trap perspective also suggests that poverty traps may be a source of heterogeneity in intergenerational mobility parameter estimates, which further implies that the classic homogeneous parameter models do not adequately measure mobility (e.g., the models from Becker and Tomes, 1979; Solon, 2002). In a simple model with two steady states – a low income steady state that corresponds to a poverty trap, and a higher income steady state that corresponds to improved economic outcomes – it is clear that there will be a high degree of intergenerational income persistence for individuals in the poverty trap, and a lower degree of intergenerational persistence for individuals outside the trap as these individuals are moving towards the steady state.[10] Yet, these dynamics are not consistent with the assumption of a single intergenerational mobility parameter for all individuals. In other words the poverty trap perspective contradicts the implicit assumption of (at least conditional) convergence to a single steady state level of economic well-being. Allowing for heterogeneity in the estimates of intergenerational parameters, perhaps across the income distribution and across different groups (e.g., ethnic groups, those living in urban vs. rural areas), might provide a first step in assessing the extent to which the magnitude of the mobility estimates are impacted by these sources of heterogeneity. One way to capture this type of heterogeneity is to use a hierarchical model, as in Sampson and Morenoff (2006). Similarly, Hertz (2005) shows that black Americans experience lower mobility overall, and substantially higher rates of down-

[10]Implicit in this argument is the belief that the economy is not already in a steady state, and so individuals in the poverty trap remain fixed with low mobility and individuals who are not in the trap are free to travel towards the steady state. For our discussion, we maintain this implicit assumption.

ward mobility, compared to white Americans; this evidence is consistent with the membership-based poverty trap theory. Another alternative is to think about how neighborhood effects might have heterogeneous impacts on individual mobility across different age cohorts of individuals (Wodtke et al., 2016), or that IGE-like parameters can be decomposed across different groups or factors (Sacerdote, 2007; Lefgren et al., 2012).[11]

2.5.2 How Might We Think About Policy?

The insights from the theory of poverty traps is also important for understanding continued policy development. If, on the one hand, empirical evidence shows that economic outcomes are not subjected to the forces of a poverty trap (though poverty may be persistent), then history does not matter and certain straightforward (e.g., lump sum) policies can aid individuals in upward economic mobility (though these investments may need to be large; Azariadis, 2006). On the other hand, if empirical evidence indicates a poverty trap, then policymakers must proceed in a fashion that simultaneously targets multiple frictions.

One implication is that given nonconvexities in labor markets and preference-based traps, many antipoverty policies – e.g., Empowerment Zones, expanded early childhood education, training aimed at the development of character skills, education reform in impoverished areas – are unlikely to substantially increase mobility. An optimistic way to view this insight is to suggest a rehabilitation of policies that may have already been deemed ineffective, if they could be combined with other types of interventions aimed at other frictions. Policy aimed at breaking poverty traps likely needs to target both people and places, accounting for both structural (e.g., economic and institutional) and preference-based frictions. To combat social effects, Durlauf (2006) suggests policies of associational redistribution – that is, a redistribution of access to certain services and opportunities, such as in the case of school desegregation

[11] Wodtke et al. (2016) focus on heterogeneity in high-school graduation induced by variation in both age cohorts and neighborhood effects, and not specifically on income or earnings.

during the 1950's and 1960's. These policies aim to alter group composition, such as busing children to different schools, or affirmative action. In cases in which associational redistribution policies may not be politically feasible, Durlauf (2006) suggests considering supply-side associational redistribution policies, to increase the ability of individuals of disadvantaged groups so that they can compete under meritocracy and benefit from existing economic structures.

2.6 Conclusion

The literature on intergenerational mobility is vast, and rapidly developing. In this article, we weave together several strands of the literature, touching on key themes at the core of intergenerational economic mobility. The integration of individual and neighborhood effects in efforts to understand mobility and inform policy represents the forefront of this literature, and while the literature we summarize provides a solid foundation, there is still much to learn. Recent work that uses relocation experiments for identification has made substantial progress towards understanding the role of neighborhood context in economic mobility, as research has convincingly established a strong and persistent relationship between the two. Yet, as we have discussed, much remains to be learned about how this relationship works. It is, however, our hope that the relevant insights in the older literature will be retained and used to ask and answer deeper questions. The primary goal must be to understand the microeconomics underlying the most credible results, solidifying strong links between theory and empirics, to support stronger policy development.

This review highlights opportunities to contribute to the mobility literature in at least two concrete ways: by integrating community characteristics more fully into models of individual mobility, and by focusing on explanations of mobility, rather than adding more discussion to the important literature on estimating mobility parameters. These opportunities serve as the motivation for my empirical chapters, with all three emphasizing explanations for patterns of intergenerational mobility, and

Chapters 4 and 6 emphasizing the role of factors which display large geographic variation. Furthermore, these empirical chapters each engage with more specific themes that appear in this review. I devote a substantial portion of this chapter to literatures on childhood development and racial differences in income and human capital accumulation, and Chapters 4 and 5 both focus directly on drawing insights from the intersections between those literatures and the mobility literature. I also discuss the value of the perspective economic history brings to the study of intergenerational mobility, emphasizing its capacity to capture longer trends and important changes that occurred before the workhorse longitudinal surveys, and my analysis in Chapter 6 is designed to explain historical sources of mobility by leveraging precisely these advantages. The perspective on the mobility literature that I provide in this chapter thus provides not only the context that informs my approach, but also the motivation for my choice of research questions in Chapters 4 through 6.

Table 2.1.: Summary of the scope of the intergenerational mobility literature and a list of key citations

Model	Origin	Key citations
Cross-sectional	Sewell and Hauser (1975)	Becker and Tomes (1986), Solon (1992), Zimmerman (1992), Mazumder (2005), Mitnik et al. (2015)
Cross-sectional IV	Behrman and Taubman (1985)	Solon (1992), Shea (2000), Zimmerman (1992), Björklund and Jäntti (1997)
Panel	Levine and Mazumder (2002)	Lee and Solon (2009), Hertz (2007), Aaronson and Mazumder (2008)
Sibling correlations	Gorseline (1932), Chamberlain and Griliches (1975)	Solon et al. (1991), Björklund et al. (2006), Björklund et al. (2009), Levine and Mazumder (2007), Mazumder (2008)
Neighborhood effects	Loury (1977), Wilson (1987)	Case and Katz (1991), Massey and Denton (1993), Page and Solon (2003), Chetty and Hendren (2015)
Experimental studies	Popkin et al. (1993)	Oreopoulos (2003), Ludwig et al. (2013), Jacob et al. (2015), Chetty et al. (2016)
Quasi-experimental place studies	Chetty and Hendren (2015)	Chetty and Hendren (2015), Bosquet and Overman (2019)

Table 2.2.: Summary of intergenerational correlation/elasticity literature that focuses on the individual-parent model

Citation	Time frame	Model	Methodological Summary	Intergenerational elasticity
Sewell and Hauser (1975)	1957-1967	Cross-sectional	Average parental income over four year period	0.18
Solon et al. (1991)	1968-1982	Cross-sectional sibling	Corrects for variance components; measures correlation in family (sibling) component rather than transitory individual shocks	0.45
Zimmerman (1992)	1966-1981	Cross-sectional	Instruments parental earnings with Duncan index	At least 0.4
Solon (1992)	1968-1984	Cross-sectional	Correct for transitory bias; selection of homogeneous sample	0.413
Shea (2000)	1968-1989	Cross-sectional	Uses instrumented 'luck income' to evaluate direct effect of money, as opposed to household characteristics associated with income	Insignificant
Levine and Mazumder (2002)	1968-1994	Panel	Decompose intergenerational elasticity into direct income effect, and indirect education/human capital effect	0.41 (NLS), 0.29 (PSID)
Mazumder (2005)	1968-1998	Panel	Averages many years; use of multiple data sets (NLSY, SIPP, Social Security Earnings data); corrects for autocorrelated transitory shocks and lifecycle biases	0.6
Aaronson and Mazumder (2008)	1950-2000	Cross-sectional sibling	Relies on synthetic parental generation means to produce sufficient data, but finds significant upward trend in correlation between individual incomes and this average	0.40 (1950), 0.58 (2000)
Lee and Solon (2009)	1977-2000	Panel	Controls for parental age; observes individuals in multiple years; uses robust covariance estimator	0.44
Mitnik et al. (2015)	1972-2010	Nonparametric cross-sectional	Modifies functional form to relax standard intergenerational elasticity model; applies corrections following Mazumder (2005)	0.52 (men), 0.47 (women)

Table 2.3.: Summary of international intergenerational elasticity estimates for selected OECD countries

Country	Estimate	Years Averaged	Outcome Year	Source
Sweden	0.28	10	1991	Björklund and Jäntti (1997)
United States	0.47	5	1993	Grawe (2004)
Canada	0.19	5	1996	Grawe (2004)
United Kingdom	0.50	-	1991	Grawe (2004)
Norway	0.17	3	1992	Corak (2006)
Finland	0.18	-	1990	Corak (2006)
Germany	0.32	-	1997	Corak (2006)

These estimates from Corak (2006) are averages of multiple intergenerational elasticities estimated for those countries. For the Finland study, heterogeneous age ranges are used: at measurement, the son's age was 39.7 on average, and the father's age was 45.7 on average.

Table 2.4.: Summary of key papers from the economic history literature

Citation	Topic	Main Finding
Long and Ferrie (2015)	Occupational mobility in U.S. and Britain, 1850-2000	U.S. mobility is high from 1850-1880, driven by returns to migration; as returns dropped, these countries became more similar
Ferrie et al. (2016)	Role of grandparents in mobility	Grandparents are important for the long-term outcomes of individuals
Feigenbaum (2018)	Examining mobility among early 20th century cohorts, 1915-1940	Compares traditional IGE with occupational imputations and finds similar results; the IGE is lower for this period than later in the century
Feigenbaum (2015)	How did the Great Depression affect mobility?	The Great Depression reduced mobility via selective migration by richer sons
Collins and Wanamaker (2017)	Long run estimates of racial mobility differences	Black mobility depressed from the late 19th century to present because of persistent inter-group human capital differences relative to whites
Hilger (2017a)	Examines sources of 20th century mobility trends	Limited effects, though there are increases in persistent growth and access to higher education
Hilger (2017b)	Long-run Asian-American mobility – high earnings despite discrimination	Interaction between positive selection of Asian immigrants and rapid change in anti-Asian discrimination in the postwar period led to high mobility, rather than higher educational attainment
Modalsli (2016)	Measuring grandparents' role in Norwegian mobility	Temporally varied grandparent effects that are sometimes large; the individual-parent model is inadequate
Modalsli (2017)	Long-run trends in Norwegian mobility	In the 19th century, Norwegian mobility was relatively low, but rose through educational access and movement away from agriculture
Clark and Cummins (2015)	Wealth mobility in England using surnames	Persistence of wealth among surname groups is consistent with an intergenerational wealth elasticity of 0.7-0.75
Olivetti and Paserman (2015)	Use surname approach to study U.S. mobility from 1850-1940	Consistent mobility in the 1800's with a decrease in the early 20th century; an overall decrease in mobility; trends are a function of expanded educational access and migration
Olivetti et al. (2018)	Grandparents' role in mobility in the U.S. from 1850-1940	Relatively large, robust, and gender specific grandparent effects that are explained using a model of trait matching

Table 2.5.: Summary of the empirical literature investigating the effect of neighborhoods on economic outcomes

Citation	Mechanism	Outcomes	Result
Case and Katz (1991)	Neighborhood social externalities	Crime, drug and alcohol use, single parenthood, labor market 'idleness'	Strong association between neighborhood poverty and the outcomes; positive and significant Moran's I statistics
Massey et al. (1991)	Residential racial segregation	Single motherhood, labor market nonparticipation	Individuals in a high poverty neighborhood have a higher probability of being a single mother or not working
Borjas (1992)	Ethnic capital	Individual incomes, occupational prestige, educational attainment	Ethnic capital intergenerational elasticities are comparable to family intergenerational elasticities; approx. 0.23
Brooks-Gunn et al. (1993)	Social externalities, measured with neighborhood means of socioeconomic and behavioral variables	IQ, behavioral problems, out of wedlock birth, dropping out of high school	Significant association between these outcomes and neighborhood poverty rates; authors argue this arises from lack of high-income peers rather than "contagion" among low-income peers.
Borjas (1995)	Ethnic capital	Income, education	Insignificant ethnic capital intergenerational elasticity after adding fixed effects and adjusting wages for education
Cutler and Glaeser (1997)	Racial segregation	Employment, income, single parenthood	Segregation is associated with a decrease in income and probability of employment; a 10% increase in the dissimilarity index is associated with a 1.83% reduction in the probability of graduating high school and a 2.1% increase in probability of being unemployed
Solon et al. (2000)	Investigates neighborhood and sibling correlations	Educational attainment	Significant neighborhood-level correlations between .08 and .12, but smaller than family correlations of about 0.4
Page and Solon (2003)	Investigates the relative sizes of family vs. neighborhood correlations	Adulthood earnings for men	Neighbor correlations of about 0.15 are smaller than brother correlations of about 0.3; suggests limited long-term neighborhood effects on earnings
Ananat (2011)	Racial segregation	Poverty, black-white inequality	Segregation increases poverty and inequality for blacks; a 10% increase in the dissimilarity index causes a 2.58% increase in the black poverty rate

Table 2.6.: Summary of the empirical literature investigating the effect of regional characteristics on regional poverty

Citation	Mechanism	Outcome	Result
Crandall and Weber (2004)	Job growth	Census tract poverty	Job growth associated with a reduction in poverty, especially in high poverty tracts and enhanced by increased social capital; spatial autocorrelation in poverty rates
Partridge and Rickman (2005)	Employment growth	County poverty	Employment growth and human capital are strongly associated with a reduction in county-level poverty; evidence in favor of place-based policies
Partridge and Rickman (2007a)	Geographic remoteness	County poverty	Local job growth is strongly associated with a reduction in poverty; evidence against poverty traps
Partridge and Rickman (2007b)	Geographic remoteness	County poverty	Poverty reductions from job growth in the nearest MSA are attenuated by distance from the MSA
Partridge and Rickman (2008b)	Geographic remoteness	County poverty	Association between job growth and poverty reduction differs by MSA size, with county poverty reductions coming from job growth in the central county; evidence of metropolitan spatial mismatch as peripheral counties do not benefit from growth
Partridge and Rickman (2008a)	Geographic remoteness, spatial mismatch	County poverty	Remoteness is correlated with poverty; local job growth strongly associated with poverty reductions in remote counties
Wu and Gopinath (2008)	Geographic remoteness, labor market, natural amenities	County wage rates, employment density, housing prices	Remoteness accounts for the largest proportion of wages and employment density; natural amenities play a substantial role in housing prices which affects aspects of the regional equilibrium
Gebremariam et al. (2011)	Employment, migration	County-level median income	Substantial interdependence between migration, employment, and incomes, as well as spatial dependence between counties; evidence of regional interdependence in outcomes, and argue for an increased role for place-based policy

Table 2.7.: Summary of the empirical literature investigating the impact of relocation experiments on economic well-being

Citation	Intervention	Outcomes	Results
Popkin et al. (1993)	Gautreaux Program	Employment, wages, hours worked	No effect of movement on outcomes
Katz et al. (2001)	Moving to Opportunity	Employment, earnings, welfare recipiency, health, safety, problem behavior incidence	No effect on economic outcomes, but positive effects on safety, behavior, and health
Oreopoulos (2003)	Random subsidized housing selection in Montreal	Adulthood (age 30+) earnings, employment, and welfare recipiency for treated children	No effect on these outcomes, and neighborhood variation explains little; family variation makes a large difference
Kling et al. (2007)	Moving to Opportunity	Adult economic status, physical health, mental health 4-7 years after treatment	No long-term effect on economic status or physical health, substantial positive effects on mental health
Ludwig et al. (2013)	Moving to Opportunity	Long-term (10-15 years after treatment) effects on adult economic status, health, and children's education	Improvement in adults' mental and physical health, no effect on economic status or children's education
Jacob et al. (2015)	Housing lottery in Chicago	Long-term (14 years after treatment) effects on treated children's test scores, health, high school graduation, arrests, earnings, and welfare recipiency	Small to null effects on all outcomes
Chetty et al. (2016)	Moving to Opportunity	Long-term effects (15+ years after treatment) on college attendance, earnings, problem behavior, single parenthood, and health	Substantial increases in earnings and college attendance, reduction in single parenthood probability
Chyn (2016)	Randomized moves due to housing project demolition	Earnings, employment, education, criminal behavior	Long-term positive effects on earnings, employment, and education; drop in arrests.

CHAPTER 3. DATA SOURCES AND CONSIDERATIONS

3.1 Introduction

Empirical research into intergenerational mobility relies on datasets which include large numbers of parent-child pairs, as I note in the previous chapter. This means that the literature generally relies on longitudinal household surveys, so that linking people to other people, rather than group statistics in one period to group statistics in a previous period, is possible. In addition to this basic requirement, the possibility of linking to additional data sources is often necessary for explaining variation in mobility when relevant variables are not observed in the survey that provides the linked income data. Like others, I rely heavily on the Panel Study of Income Dynamics (PSID), but I emphasize different features of that dataset in Chapters 4 and 5. I exploit the rich set of household covariates the survey provides in Chapter 4, and in Chapter 5, I supplement these data with county and Census tract data from other sources by using confidential geographic crosswalks. In Chapter 6, I use data from the U.S. Census' Integrated Public Use Microsample (IPUMS) Ruggles et al. (2018), computing mobility statistics using a procedure proposed and described by Hilger (2017a), along with aggregate data from the decennial census, merged at the state level.

Individual chapters discuss a number of specifics, such as the definitions of key variables and discussion of sampling restrictions. This chapter provides the reasoning behind my choice of datasets, and on explanations of general features of these data that do not necessarily need to be discussed in the empirical chapters. This means that detailed discussion of the datasets I use in each chapter remain within each chapter. Including this chapter lets those specific sections be more precise and focused, while allowing me to explain my thinking regarding the choice of data in broader terms here.

3.2 The Panel Study of Income Dynamics

The PSID database is a longitudinal survey of a representative sample of individuals in the United States, and their families. The survey began with a sample of about 18,000 individuals from 5,000 families in 1968, and has followed these individuals and their descendants annually until 1997, at which point the sample began to be tracked biennially. The PSID originally emphasized income, demographics, and family composition, but these factors naturally overlap with other economic factors, such as education, employment, job type, and health, so the dataset has become a rich source for applied social scientists interested in a broad array of topics. In addition, PSID contains geospatial identification codes for each family, reporting census block, census block group, census tract, and county residence.[1] The geospatial identifiers allow us to link the individuals in our sample to the spatial measurements from the United States Census data, which is necessary for the measurement of community characteristics, as far too few PSID respondents live in any one place to produce credible local aggregates from PSID responses alone.

These data have noteworthy advantages for both Chapters 4 and 5. In both chapters, the number of parent child links is the primary advantage, however. This advantage is particularly pronounced in Chapter 4, however, because for reasons I discuss in detail there, I impose a number of fairly stringent sampling restrictions. My preferred sample ultimately includes just under 1000 parent-child pairs, so the fact that a ‘naive’ initial match provides several times that, approximately 5200, is extremely valuable, as starting from a smaller sample could have led to a substantially smaller final dataset. The ease of tracking families across years is very valuable as well, as I collapse childhood into a single cross-section in both chapters, which

[1] These identifiers are not available in the public use sample, but may be made available upon request.

requires confidence that I have captured all of the available data, that it is, in fact, comparable across years. The fact that the data are explicitly organized by family makes executing these transformations straightforward. Finally, the relatively large number of observations available for many respondents is very valuable in Chapter 4, where I define poverty as a multivariate treatment varying in different aspects of the timing of poverty. Capturing this variation requires many observations at widely spaced points, as opposed to a less frequent survey or a retrospective cross-section.

3.2.1 Harmonizing Spatial Data

Answering questions about the role of community characteristics and local economic conditions requires credible measures of these factors. To answer the questions I raise in Chapter 5, I need a dataset that connects parent-child pairs to specific locations, ideally at several points during childhood. Fortunately, census tract and county data are available through the University of Minnesota's National Historical Geographic Information System (NHGIS) project. The NHGIS is a publicly available database including decennial census and American Community survey data along with relevant Geographic Information System (GIS) files. Using NHGIS resources, I construct tract and county datasets spanning the years 1970 to 2010 for the continental United States. One challenge of using these data is that, unlike the PSID families, a specific location's membership in a tract or county can change over time, because the boundaries of both units change over time. This variation comes primarily from changes in census tract geographies, as the number of tracts has expanded dramatically during this period, rising from approximately 61,000 to approximately 74,000 between 1990 and 2010 alone. This rapid change is driven by the fact that tracts are defined to always include approximately 4,000 people, as opposed to representing a specific political unit, as counties do. This means that population changes can cause dramatic changes to tract boundaries. To address this issue, I use a spatially harmonized version of the United States Census database that is consistent at 2010 geographies, developed by Kumar et al. (2019). The authors developed the census tract data using the Brown University Longitudinal Tract Database (LTDB), which contains appropriate weights to adjust geographies from 1970 to 2000 to the 2010 geographies. Unlike census tract geographies, the boundaries of most counties do not change over time, and there does not exist any set of weights to adjust county boundaries for the few that have changed. We therefore developed our own weights to spatially harmonize the 130 noticeable changes in county boundaries across the five decades.[2]

Kumar et al. (2019) go on to interpolate the census data between decadal measurements to obtain a sample of annual, spatially harmonized census measurements. For our purposes, this is an extremely helpful feature of their approach to harmonization procedure, because it means that the database includes measurements at each year, rather than only decennially. These interpolated data are clearly measured with error, but I strongly prefer them to simply using the closest year's value or using only individuals observable in census years, as this latter solution would make this study all but impossible. For variables that have annual data at county level but not at census tract level, they use a ratio method of interpolation to adjust for intercensal years at the tract levels. Specifically, they interpolate the ratio of a variable at census tract to county (τ) for two consecutive census years, for example years 1980 and 1990. They then calculate the difference in the ratio between the two years by dividing by 10 to develop a value (ψ), which is used as the interpolated change rate yearly. This means that the intercensal census tract values at year $(s + n)$ can be interpolated by $\bar{o}_{s+n} = (\tau + \psi * n) \times \bar{\bar{o}}_{s+n}$. Variables that are interpolated by this ratio method include the unemployment rate, black population share, Hispanic population share, population level, and poverty rate. For other interpolated variables, including share of census tract single mother households, census tract ethnic capital, and county rurality, that do not have annual county measurement, they use a linear interpolation to deriving their intercensal years data.

[2]Additional details regarding this spatial harmonization project are available at https://dev.www.pcrd.purdue.edu/signature-programs/poverty-project.php.

For a study in which tract-level variables play a key role in the empirical strategy, this harmonization is extremely important, because without it, it becomes very easy to assign observations to the wrong tract, or to unnecessarily remove observations from the sample because they cannot be matched to a location. Relatively few scholars contributing to the mobility literature have exploited tract-level data, for a variety of reasons including changing boundaries. This database makes it much easier to take advantage of these data, which in turn facilitates incorporating community factors – long understood to be important in this context – into a mobility study.

3.3 Imputing Mobility Statistics from the Integrated Public Use Microsample (IPUMS)

In Chapter 6, I rely very heavily on data imputation methods developed in Hilger (2017a)[3], which provide credible estimates of educational mobility parameters[4] at the state level starting in 1940, based on readily-available Census data. This imputation procedure relies on the fact that some individuals in their mid to late 20's live with their parents in the IPUMS data, so for this subset of the population, the obstacle to using the IPUMS to calculate mobility parameters (i.e., the fact that I cannot link parents and children) does not exist. If I can mitigate sources of bias or mismeasurement that come from differences between the dependent and independent populations[5] populations, I can produce credible mobility statistics. Much of Hilger's contribution is demonstrating that his imputation method does, in fact, pro-

[3] I replicate a number of his equations, and paraphrase, in condensed form, much of his explanation of the validity of this approach. The interested reader is encouraged to read his explanation of the procedure.

[4] In this case, intergenerational education correlations, captured by the coefficient on parental education estimated with a regression of educational attainment on parental educational attainment. This measure captures mobility similarly to the traditional intergenerational elasticity of income, but education is measured much more clearly and consistently than income in the IPUMS data, due largely to changes in the tracking of income in the census, so the education-based measure is much more stable and reliable.

[5] That is, people in their 20's who live with their parents or on their own, respectively.

vide credible estimates of differences between dependent and independent outcomes, which bridges the gap in the base data.

Hilger explains the procedure in great detail in his paper, and includes a number of empirical exercises that validate the assumptions which it requires, but at least a brief explanation of the procedure is necessary here. First, note that mobility statistics are generally summaries of a conditional expectation function (CEF) $E[h_{i,t}|y_{i,t-1}]$, where h denotes some outcome, y denotes some parental characteristic (often, but not necessarily always, the same variable as h), and i and t index individuals and time periods, respectively. Furthermore, I can write the overall value of an outcome $h_{i,t}$ as

$$h_{a,y} = d_{a,y}h^D_{a,y} + (1 - d_{a,y}h^I_{a,y}), \tag{3.1}$$

where a indicates an age range, y denotes membership in a particular bin of parental outcome values (e.g. parental educational attainment of 12 years, 9th income decile), $d_{a,y}$ represents the proportion of dependent children in this age group and parental bin, and I and D indicate independence or dependence. Because I know the individual and parental values for dependent children, I know $h^D_{a,y}$, but not the other values. However, if I can impute them, I can compute a CEF, and thus the mobility statistics I want to study.

We also know that $d_{a,y} = \frac{N^D_{a,y}}{N^D_{a,y}+N^I_{a,y}}$, where N denotes number. However, by definition, the proportion of dependent children is the number of children in that bin divided by the total number of children in it. If I assume that proportions of children in each parental value group are the same in older cohorts as in younger ones within a census (i.e., the proportion of 17 year olds from the 8th decile of parental income is the same as the proportion of 26 year olds from the 8th decile), I can fill in the missing values from this proportion, and compute the CEF, however. Hilger refers to this assumption as a “smooth cohort assumption”, as it implies that compositions of subgroups across age cohorts within a single census are related by some smooth function, rather than varying in large discrete jumps. Because the overwhelming majority of children under 18 live with their parents, this population provides a convenient means of calculating proportion in each parental bin under this smooth cohort assumption: if $N^D_{a,y} \approx N_{a,y}$, I can use the observed proportions of 17 year olds in each parental attribute bin as proportions for older individuals, which lets us fill in the values of the missing terms in the definition of $d_{a,y}$ above.

With an observable $d_{a,y}$, all that remains to define $h_{a,y}$ is to define $h^I_{a,y}$, the mean outcome for independent children with parents from each parental outcome bin. Under the assumption that the functions relating parents to children follows a parallel path for dependent and independent children, given parents position in the education or income distribution, this is fairly straightforward. Formally, this assumption states that

$$f(h^D_y, h^I_y = \rho \forall y), \tag{3.2}$$

so that the relationship between an outcome for dependent and independent children is governed by a constant parameter. This appears to be a fairly strong assumption, but it is common, if rarely made explicit, in this literature, as researchers routinely summarize parent-child relationships in permanent income with a single parameter in datasets pooling many years, places, and characteristics in one regression model. Under this assumption, and rearranging the assumed identity $h^D_{a,y} = \rho + h^I_{a,y}$ I can estimate ρ as

$$\hat{\rho} = \sum_{j=1}^{J} \frac{\hat{N}^I_{a,j}}{\hat{N}^I_a} h^D_{a,j} - h^I. \tag{3.3}$$

With this estimate of ρ, it becomes possible to compute an estimate of independent children's outcome informed by the demographic structure of the young population in the census and the observed characteristics of parents and children who can be linked to one another based on their shared residencies.

With these estimates of individual education by parental education or parental income, I can estimate mobility parameters that are comparable to intergenerational educational coefficients. We follow Hilger in computing these at both the national and state level, although our analysis, like his, relies more heavily on the state analysis, because conducting the analysis at this level adds cross-sectional variation in both the degree of intergenerational mobility itself and in important explanatory variables.The coefficients on parental education and income in these regressions are our outcomes of interest, rather than the imputed education values. These imputed values are thus necessary for our analysis, but do not enter the regressions I discuss in the next section, except through their inclusion in the regressions which capture the degree of mobility by state in each census year.

The validity of this procedure depends crucially on the smooth cohorts and parallel trends assumptions that I have discussed only very briefly here. Hilger (2017a) goes into much greater detail on the derivation, explanation, and justification (based on validation through analysis of additional datasets) of these assumptions, however, and finds that they are credible based on these exercises. Clearly, direct household surveys, in which one could observe parents and children together in a harmonized enumeration framework, similar to, for example, the PSID, would be preferable, but since this is not an option for developing a dataset that captures the mid-20th century, this imputation approach is necessary to study long-term mobility trends.

There are reasons to have some reservations about this procedure, but these data are using the IPUMS data in this way is necessary in this paper, and it provides an exciting set of opportunities for similar studies. The possibility of considering a long time period presents opportunities that the PSID, for all of its many benefits, cannot replicate, because it does not begin until 1968. These longitudinal datasets, as well as the administrative datasets that some scholars have brought to bear on questions of mobility, can do a great deal to help us answer important questions about sources of mobility, but widening the lens of our scholarship to include older sources of variation can only help to enrich this literature. This is particularly helpful in the literature on mobility trends, which often emphasize trends during the period included in the longitudinal surveys. As these studies, discussed in detail in Chapter 2, note, the shape of these trends provides an important baseline for policy discussions, so adding additional context to this literature by both expanding the timeframe and focusing on variation, as I do, could contribute a great deal to the scholarly and policy discourse.

3.4 Summary

In my empirical chapters, I rely on several data sources, each of which have specific properties that make them appropriate for answering questions about the sources of intergenerational mobility. In Chapters 4 and 5, I rely on data, particularly data on incomes and family linkages, from the PSID. In addition to the generally attractive features of this dataset, the fact that it includes many childhood observations makes it particularly attractive. In Chapter 6, I use a combination of Census microdata from the IPUMS and state-level census data to investigate long-term mobility trends. Such an investigation would be impossible if I were to rely on the more standard longitudinal datasets, despite their many advantages in other types of study, including my own. None of these sources are new additions to the mobility literature, but recent advances in this literature have created opportunities to leverage these standard, but very rich, data sources to answer new questions.

CHAPTER 4. CHILD POVERTY AND INCOME MOBILITY IN THE UNITED STATES

4.1 Introduction

Compared to the rest of the OECD, in the United States rates of child poverty remain high while intergenerational economic mobility remains low. According to OECD child poverty statistics, 20.2 percent of children in the United States were poor, compared to an average rate of 13.6 percent across the OECD; out of the other North American or Western European OECD countries, the only one with a higher rate of child poverty than the United States is Spain, with a rate of 22.7 percent (OECD, 2018).[1] Corak (2006) and Blanden (2013) find that the United States is substantially less mobile than other Western countries when mobility is measured with an intergenerational income elasticity that measures income similarity across generations; in both studies, the United States ranks second lowest in mobility among the countries studied. We explore these issues through an empirical model of individual income rank (as an adult) that is a function of a multi-dimensional index of childhood poverty that captures the age at which the child entered poverty, the intensity and duration of that poverty, and the concentration of the individual's time in poverty.[2] Our goal is to better understand the way in which childhood poverty affects adulthood outcomes: accounting for both poverty intensity and duration allows us to determine the extent to which adulthood outcomes are affected differently by poverty spells

[1]The OECD child poverty rates for the United States are consistent with other recent United States sources. For instance the Annie E. Casey Foundation and the National Center for Children in Poverty both report a rate of 19 percent in 2018 based on American Community Survey data (Koball and Jiang, 2018; Foundation, 2018).

[2]By "concentration" we are referring to the degree to which an individual's time in poverty occurred as a single incidence of poverty of some length of time, or instead in multiple incidences spread throughout childhood. Our goal is to capture potential differences between a child who experienced poverty at different points in childhood compared to a child who experienced the same overall intensity and duration of poverty, but across a contiguous stretch of years. We describe our measurements in detail in Section 3.2.

of different natures, accounting for age at poverty exposure allows us to investigate whether younger children are disproportionately affected (as adults) relative to older children, and incorporating the concentration of poverty lets us consider how these factors vary depending on the spacing of periods of time in poverty. Our analysis is policy-relevant as our analysis informs policies that target childhood poverty by clarifying the context (nature and age of exposure) through which poverty affects individual socioeconomic outcomes; indeed, efficient use of social funds requires detailed knowledge of the processes that connect childhood inputs to adulthood outcomes.

Our work is motivated by the large literature that has investigated how different types of childhood interventions lead to positive socioeconomic outcomes for the individual as an adult. Several papers, for instance Cunha et al. (2010) and Agostinelli and Wiswall (2016), develop micro-theoretic foundations that link early-childhood human capital investments to socioeconomic outcomes as an adult, and provide supporting empirical evidence that parental investments made at an early childhood age compound over the remaining years of childhood. These papers build on the classic papers by Becker and Tomes (1979, 1986), focusing on the technology by which parents might influence an individual's economic outcomes. One alternative approach to focusing on the shape of a production function is to instead look for reduced-form (causal) evidence that children who experience a major socioeconomic adjustment (e.g., geographic relocation) subsequently experience better socioeconomic opportunities as adults (Chetty et al., 2016; Chyn, 2016), or to estimate intergenerational income elasticities or rank correlations (Chetty et al., 2014a). The structural investment models predict large differences in the effect of socioeconomic interventions by age at which the intervention occurs, while much of the (causal) reduced form evidence suggests that intervention effects are largely homogeneous across ages, although some recent work finds that age thirteen is a kind of threshold, before which certain interventions are more effective (Chetty et al., 2016). These findings lead to a policy-relevant empirical puzzle in the sense that different empirical findings suggest different interventions: if the effects of childhood experiences on adulthood outcomes are (largely) constant across age at which the experience occurs, policymakers can rely, for instance, on relatively broad policies such as housing vouchers, school desegregation (on income or racial lines), or improved primary education in poor neighborhoods to improve an individual's lifetime opportunities. On the other hand, strongly age-differential effects suggest that policymakers should take a much narrower approach, and focus on, for instance, neonatal care and accessible pre-Kindergarten childcare and education.

Approaching childhood poverty as a treatment with several distinct features helps us to shed further light onto this policy-dilemma. In measuring poverty as a continuous, multivariate process that accounts for duration, intensity, timing, and concentration components, we posit that differences in childhood poverty spells (along these dimensions) produce different effects on an individual's position in the income distribution as an adult. By considering these distinct components of poverty, we are able to detect nonlinearities in the poverty-adult outcomes relationship that the child development literature suggests, while maintaining a broad focus on childhood poverty, rather than focusing directly on production functions for skills. Our intuition is that each of these measures correspond to a distinct feature of an individual's exposure to poverty: duration – defined as the number of years an individual spent in poverty as a child, captures how long an individual spent in poverty; intensity – defined as the average difference between one's household income and the poverty line when in poverty, in percentage points, reflects how serious that poverty was in income terms; and timing – the age at which each individual first experienced poverty, accounts for the developmental stage at which a person experienced poverty. We measure the concentration of childhood poverty with several measures, but the most important are the standard deviation of ages at which each individual was poor and the number of contiguous poverty spells each individual experienced as a child, both of which reflect the extent to which an individual's experience of poverty was spread across childhood, rather than being concentrated in one period.

The economic foundation supporting our work comes from a subset of the human capital literature concerned with estimating elements of the technology that transforms inputs individuals receive in childhood into skills (e.g., Cunha and Heckman, 2010; Almond and Currie, 2010; Heckman and Mosso, 2014).[3] These skills, in turn, affect adulthood outcomes, such as income and educational attainment, but the magnitude of this effect decreases in the age at which the individual receives the investment. This age-decreasing relationship is driven by a phenomenon Cunha et al. (2010) refer to as 'dynamic complementarity': the extent to which investments in skill development in later childhood influences the individual's outcomes depends on previously acquired skills, which means that being poor for an additional year will not only lead to lost skills for the year of poverty but will depress the absorption of skills acquired in subsequent years. Translated into our view of poverty – i.e., distinguishing between age, duration, intensity, and concentration – we expect that poverty duration and intensity have distinct effects on individual outcomes, and that these distinct effects differ by age of poverty exposure and the degree of poverty concentration. Duration corresponds to a sustained reduction in parental investments, while intensity corresponds to a dramatic reduction in parental investments, the interaction between duration and intensity corresponds to a sharp and sustained reduction in parental investments. If a sharp, sustained reduction in investment occurs during early childhood, particularly if the reduction is concentrated in this period rather than spread into later childhood, then the micro-theoretic literature would suggest that we would find a significantly negative effect on socioeconomic outcomes as an adult, and that holding the size of the reduction in parental investment constant, the effect would be disproportionately smaller if it occurred later in childhood.[4] Yet,

[3] Almond and Currie (2010) review this literature and note in their abstract that "Child and family characteristics measured at school entry do as much to explain future outcomes as factors that labor economists have more traditionally focused on...". Cunha and Heckman (2010) provide a complementary overview, also emphasizing the disproportionate benefits of early childhood investment.

[4] Similarly, Agostinelli and Wiswall (2016) find several sources of nonlinearity in the relationship between childhood investments and adulthood outcomes: the average treatment effect of a transfer of $1000 drops substantially (60 percent or more) by the time a child is 11 or 12, compared to when she is 5 or 6, and that much of this drop occurs between the first period and the period when the child is 7 or 8. Further, the effects of the transfer on 5 and 6 year old children are approximately

the most striking estimates from the Moving to Opportunity experiment indicate that there was a strong positive effect of improved neighborhood quality for individuals who were younger than thirteen when they moved, but not for older children (e.g., Chetty et al., 2016).[5] Within this group of pre-thirteen year olds, however, Chetty et al. (2016) find that the effect of an additional year in a better neighborhood is linear, arguing that this piecewise linear relationship explains the difference in responses between younger and older children. Moving to a better neighborhood increases earnings (or other socioeconomic outcomes) linearly, but a fixed 'disruption cost' of moving decreases the effect of total exposure to the improved environment and so the treatment effect drops to zero for children who move later and receive fewer years in the better neighborhood. Translated into our model of poverty, these results imply that poverty spells do not have heterogeneous, age-differential effects on individual outcomes, and that instead entrance into poverty at an early age corresponds to a constant (by year) drop in individual income rank with no significant effect if the poverty spell was to begin after age thirteen.

We proceed using data from the PSID to define our poverty measures. These measures let us distinguish between different aspects of poverty which allow us to consider the results of relocation studies and the structural work on the technology of human capital formation in a common framework. This common framework is necessary for understanding what features of childhood poverty might produce these parallel bodies of results, which is the key contribution of this chapter. Our use of

---

twice as large as the average treatment effect at the tenth percentile of family income, and that they decline convexly as income rises, consistent with both nonlinearity in intensity, and with an interaction between intensity and timing.

[5]The Moving to Opportunity experiment is an intervention funded and implemented by the U.S. Department of Housing and Urban Development (HUD) that involved 4,604 families living in high poverty census tracts between 1994 and 1998 in Baltimore, Boston, Chicago, Los Angeles and New York City. Each family was randomly assigned to a group that received either a voucher that subsidized private rents in tracts with poverty rates below 10 percent, a voucher that allowed them to move into subsidized housing without a location constraint, or into the control group. The study has followed individuals who moved since initial assignment, and early studies (e.g. Katz et al. (2001), Kling et al. (2007), Ludwig et al. (2013)), found no evidence for long-term benefits on earnings or employment, although they did find positive effects on mental health and schooling for females. People treated in early childhood reached adulthood relatively recently, however, so these studies are not able to model the effects Chetty et al. (2016) detect for this subgroup.

these observational data in a reduced form setting places our contribution between the studies relying on field experiments and the studies developing structural models. We start by benchmarking our PSID data to the results from Chetty et al. (2016), and then expand our focus to the nonlinear, multivariate framework introduced in Section 2; yet, we do not consider treatment as at fine a scale as the structural literature, opting instead to bridge their results to our coarser, but somewhat more tangible, set of treatment and outcome measures. The neighborhood treatment that Chetty et al. observe differs from the household treatment we observe in important ways, but we are confident that the comparison we make is relevant, because the theory of child development engages with the role of investments of all kinds in determining the trajectory of skill development. We focus on comparisons with their paper because of its recency, its impact on the literature, and the strength of its identification strategy, but it is hardly the only robust study to suggest a linear effect of investment or deprivation; Milligan and Stabile (2011), Dahl and Lochner (2012), and Black et al. (2014) all exploit exogenous shifts in poor families' incomes, and find evidence consistent with linear returns to the long-term effects of these shocks.

We find a strong and statistically significant association between adulthood income rank and poverty duration, with much larger marginal effects at lower values (between one and six years of childhood poverty) than at relatively large values (between eleven and seventeen years of childhood poverty). This relationship is not mediated by the other treatments, and none of them are significantly associated with adulthood income rank in our preferred specifications. These results reinforce the conclusions Cunha and Heckman (2010) and Agostinelli and Wiswall (2016) reach regarding the importance of early interventions, because they suggest that interventions targeting individuals with a small to moderate dose of poverty will have larger effects than targeting someone who has already experienced a long poverty spell. Our results are also consistent with the argument that differences between Chetty et al. (2016) results and structural results are driven by complementarity between household poverty model and neighborhood inputs, which would imply that prolonged household poverty would reduce the positive effects of moving to a richer neighborhood. If these sets of inputs are complements, as the Agostinelli (2018) results suggest, we would expect to find the steep decreases in adulthood rank and income as childhood poverty rises that we find, as well as the linear and relatively small marginal decreases in income for additional years in poor neighborhoods that Chetty et al. (2016) find.

This chapter contributes to the mobility literature in general, and to my project specifically, by integrating aspects of the mobility and child development literatures. The mobility literature is increasingly interested in comparing expected mobility given one's childhood position in the income distribution (see, for example, Hertz (2005) and Chetty et al. (2014a)), and modelling expected rank as a function of poverty contributes to this effort by restating some of the concerns of the mobility literature in terms that directly overlap with the child development literature. This means that, while I do not model variation in mobility directly, this study has important implications for our thinking about mobility, because if expected rank varies widely over poverty profiles, the distribution of mobility may be less uniform than many studies have assumed. If this is the case, policies targetted at specific subsets of the left tail of the income distribution, as opposed to universal policies or broad economic changes, will be more likely to increase absolute mobility. This chapter thus helps to bridge the gap not only between the two strands of the childhood development literature I discuss in this introduction, but also between the mobility literature and the child development literature, and it does so in a policy-relevant manner. Taken together, this represents a contribution to the discipline's approach to considering interventions that could enhance mobility by targetting vulnerable children.

4.2 Hypotheses and Modeling Strategy

4.2.1 Null Hypotheses

We formulate and test two null hypotheses related to our multi-dimensioned measure of childhood poverty:

Null Hypothesis 1 The intensity and duration of childhood poverty do not have distinct impacts on the economic status of an individual during adulthood.

Null Hypothesis 2 The effect of childhood poverty on economic outcomes at adulthood is linear in the age at which the individual was exposed to poverty, and in the concentration of poverty.

The intuition we have presented in the introduction largely leads us to believe that empirical evidence leads to a rejection of these null hypotheses. Specifically, we expect that different aspects of poverty – captured through intensity and duration – have different effects on individual outcomes because they correspond to different natures of parental investments in the child. Likewise, a large majority of empirical evidence leads us to expect that the relationship between childhood poverty and individual economic outcomes are nonlinear in age of exposure. Furthermore, we expect that these effects in turn vary in the concentration of childhood poverty, with greater concentration in early childhood corresponding to more negative and nonlinear effects than poverty that is otherwise equivalent, but spread across childhood. This is because dynamic complementarity implies that a larger developmental deficit earlier will be harder to overcome, and that this difficulty increases convexly in the size of the deficit. Holding duration, intensity, and timing equal, a more concentrated poverty spell will imply a larger deficit because an otherwise-equivalent 'dose' of poverty occurs in a shorter period.

4.2.2 Empirical Models

Our empirical model of income rank is

$$rank_i = g(poverty_i) + X_i\beta + \epsilon_i, \tag{4.1}$$

where $i = 1, 2, \ldots, n$ indexes individuals, $rank_i$ is the rank of individual i's income in the national income distribution during the time of adulthood, $poverty_i$ is a measure or index of poverty (defined below) during the time the individual was a child, $g(\cdot)$ is a smooth (differentiable) but otherwise unspecified function, X_i represents a vector of control variables with coefficient vector β, and ϵ is a mean-zero error term. This model assumes cross-sectional data because we average across years as a child and adult separately for each individual to generate a dataset that links each individual's economic status as an adult to his/her childhood economic status. At the same time, this model also captures alternative models from the literature; for example, in the relocation studies the measure of poverty exposure is implicitly 17 minus the age at which the individual moved out of poverty. The relationship between this model and models of childhood investment technology is less clear because childhood investment levels do not easily translate into poverty; and yet, that literature implies that $g(\cdot)$ in our model should be convex and decreasing in continuous poverty measures for a fixed age of exposure. In our case, we define poverty as a function of age of exposure, duration, and intensity, so that Equation 4.1 becomes

$$Rank_i = g(Age_i, Duration_i, Intensity_i, Concentration_i) + X_i\beta + \epsilon_i. \tag{4.2}$$

We estimate our regression model using several different estimation strategies that allow for both parametric (i.e., relatively restrictive) and nonparametric (i.e., relatively flexible) assumptions on $g(\cdot)$. The simplest model allows only for linear interactions between age, duration, intensity, and concentration, and can be estimated using ordinary least squares. This type of model is potentially quite restrictive, because it does not allow for the general types of nonlinearity that are shown to exist in the early-development literature. The most general specification allows for general nonlinearities and interactions and requires the use of advanced estimation methods. We present technical estimation details for these models in the appendix; see, also, Li and Racine (2007). We consider the range of structure for $g(\cdot)$ so that we can statistically test the more stringent specifications against the more relaxed specifications, thereby allowing us to understand the extent to which there are nonlinearities (and possibly what type of nonlinearities) in the childhood poverty-economic status relationship.

This scope of models allows us to detect nonlinearities and interactions without requiring distributional or functional form assumptions. Our approach is different from the structural models, which are built on microeconomic models with two or more discrete periods during childhood, and which relate investments to skills, and skills to outcomes. This difference is intentional: though we are studying the long-term effects of poverty, as do the relocation studies, we also want our models to be able to detect nonlinearities that the theoretical and empirical results of the structural literature predict, albeit indirectly, in the context of a model of poverty. In other words, our empirical model has the potential to provide a bridge between the relocation studies and the structural model literature.

4.2.3 Benchmark Replication

In Section 4.4.1, I present a replication exercise in which we estimate models analogous to the age models in Chetty et al. (2016). We present these models to demonstrate that our data and measurements are capable of producing results similar to estimates found in the relocation literature. As we discuss in more detail in Section 4.4.1, our estimates of the Chetty et al. (2016) duration effect fall within the 95 percent confidence interval they compute. Our motivation is simple: if we cannot detect effects similar to the ones they detect using a version of their empirical model, it would be difficult to argue that we reach an explanation for some differences between linear and nonlinear poverty exposure effects in the context of our model. The evidence we present in Section 4.4.1 shows that our data are acceptable by this standard; through these regressions, we show that the baseline conclusions from Chetty et al. (2016) are borne out in our PSID sample.

4.3 Data and Measurements

4.3.1 Sample Description and Selection

Estimating intergenerational relationships in income ranks (or income) is plagued by various technical difficulties, many of which place constraints on the timing of observations during both childhood and adulthood, and consequently limit the sample. In our case, these issues are exacerbated by the fact that we need accurate measurements of our four primary treatment variables, all of which depend on different sample selection criteria, further limiting the size of our preferred sample. We impose requirements on the number of observations in both childhood and adulthood. In childhood, we require that each child be observed in each possible childhood year. In adulthood, we require that each individual is observed at least three times between the ages of 25 and 35. The adulthood requirement is necessary to reduce the influence of lifecycle bias, both from measuring income at inconsistent points in individuals' lives within the cross-section we use, and from measuring incomes in years in which it correlated less well with permanent income.[6] Our final sample includes observations across many birth year cohorts; we summarize the distribution of our two samples (described below) via histograms in Figure 4.1. The mode, 1968, is substantially more frequent than the second most frequent year; representation over the other cohorts is reasonably even, however. The frequencies in Figure 4.1 have standard deviations of about 33 and 27, respectively, which explains why the time period begins in 1963 when the PSID started in 1968.

Our second restriction – that we completely observe all years in childhood – is important. Relaxing this restriction, so that we do not observe every year of each

[6]See Haider and Solon (2006) and Hertz (2007) for further discussion of these lifecycle bias issues in the context of intergenerational mobility measures.

child's childhood, would substantially increase the risk of mismeasuring duration. If, for example, we relax the restriction and include anyone who is in the sample for ten or more years during childhood, and then measure duration as the percentage of observed years in which the individual's household was poor, duration would be measured with error, because the distribution of poverty in the unobserved years is unknown.

One restriction that we do not impose in our main sample merits discussion as well: we do not restrict the sample to individuals who experience only a single contiguous poverty spell in our main analysis. This restriction would ensure that each individual has experienced an equally concentrated period of poverty exposure, given their duration, which would mean that our estimates of g(·) are not affected by unobserved differences in concentration. This is helpful because it bolsters the credibility of our timing measure by ensuring that each child's time in poverty is spaced in the same way: no one leaves and re-enters, which could happen at different intervals that are difficult to hold constant. Under this restriction, we know that each person's entire duration is contiguous, that every person is observed at every possible PSID interview in childhood, and we know each person's age when she first entered poverty, we can map each timing and duration value to a specific age interval. In that case, age at first exposure would fully characterize the timing of total poverty exposure. By contrast, without this restriction age at first exposure reflects when poverty began, then rapidly loses precision as we move from that point in time. We do not use this sample in our primary results because of the cost in observations: we lose approximately 50 percent of our preferred sample if we impose this restriction. We do, however, repeat our analysis using this sample in Section 4.4.3. We do this because apply the restriction serves as one means of determining the impact of omitting concentration in our primary results, because it holds the level of concentration constant across individuals.

4.3.2 Key Measurements

Our outcome variable is family income rank measured during adulthood, as in Chetty et al. (2014a). In the PSID, this variable is defined as the sum across all sources of income at the household level, and consequently includes transfers.[7] We deflate total household income to real 2012 dollars before computing income ranks, and we report descriptive statistics of family income during both childhood years (ages 1-17) and adulthood years (ages 25-35), in Table 4.1. To provide a clearer impression of the distribution of these key variables, we also present scatter plots of age, duration, and intensity combinations in Figure 4.2, and of rank by age, duration, and intensity in Figure 4.3. Since the PSID is a sample of households and not the population, we approximate rank by computing each individual's percentile in the national distribution of family income rank, available from the Integrated Public Use Micro Sample (IPUMS) database (Ruggles et al., 2018). Since the IPUMS data are only available decennially, we use the distribution that corresponds to the nearest census year.[8]

Our approach to measuring poverty is perhaps the most important aspect of our data preparation. From the PSID, we are able to capture three key features of childhood poverty exposure: the duration of exposure, the intensity of exposure, and the timing of exposure. We measure duration as the number of years an individual spent in a household with a total family income below the poverty threshold, both of which are recorded as survey responses in the PSID at the family level, based on the relevant poverty line for the family, defined by the U.S. government. This removes the need for assigning or imputing a poverty line measure from other variables and/or data sources, as the poverty line for each family, defined by their characteristics, is already recorded. Duration values can thus take any whole-number value between 0 and 17,

[7] The correlation between total household income including transfers and household income without transfers is 0.98 among childhood observations and 0.996 among adulthood observations; thus, our results do not depend on whether or not we include transfers. These correlations are similar in the smaller subsample we discuss later in this section: 0.985 and 0.996, respectively.

[8] For example, we evaluate an individual's income rank in 1974 using the 1970 income distribution, while we use the 1980 income distribution for 1976.

because an individual is no longer a child at age eighteen. We define the intensity of poverty as the ratio of family income to the poverty line in years in which the family fell below the line, averaged over all the years of each individual's childhood. This means that intensity is continuous on $[0, 1]$, as averaging only over years in which the household was below the poverty line treats all years which would have ratios greater than one as zeroes in the summation. Finally, we capture the timing of poverty spells with the individual's age at his or her first exposure to poverty. Conditional on intensity and duration, this timing measure provides a straightforward measurement of when someone experienced poverty, which we then use to understand how the timing of poverty affects outcomes.

In addition to bounding the influence of concentration by holding it constant through sampling restrictions, we model the impact of concentration on rank and the relationship between rank and other treatments by including several measures of concentration in models of $g(\cdot)$. We expect that this factor only matters in the context of the other three treatments, and so we treat it as somewhat secondary. Our results, particularly regarding age at first exposure to poverty, may be biased by error in measurement of the timing of poverty. It is helpful to think of this in concrete terms: a child with three years of poverty that occur contiguously has experienced more concentrated poverty than a child who experienced the same amount of poverty in separate periods. This concentration is separate from timing, as two children who entered poverty at the same age and stayed equally poor for the same number of years could still have different levels of poverty concentration. It is, however, related to timing because greater concentration along with an early age implies, through dynamic complementarity, a more significant early source of deprivation for given levels of duration and intensity. Consequently, we expect that the relationship between rank and age (and perhaps the other relationships as well, particularly in models with interactions) depends on this measurement.

Our primary measure of concentration is the standard deviation of the ages at which each individual was poor, divided by duration. Individuals with contiguous poverty spells will have the minimum standard deviation for duration, which will increase as poverty ages spread out further, but the standard deviation will also increase as duration increases, even if contiguity is preserved.[9] To illustrate this measure, consider the example of two children with contiguous poverty spells, one of 2 years and the other of 17 years, each entering poverty at age 1. The standard deviation of (1,2), 0.71, is much lower than the standard deviation of (1,2,...,17), 5.05, despite the fact that both children experienced a concentrated period of poverty in early childhood. It does not make sense to say the higher-duration child experienced less concentrated poverty because that concentrated spell continued. Dividing the standard deviation by duration removes the majority of this artificial relationship between duration and concentration. Under this normalization, concentration values for possible values of duration, given a starting age of 1 with contiguous poverty years, range from 0.297 to 0.354, compared to 0.71 to 5.05 without it.

As an alternative measure of the concentration of poverty, we consider the number of poverty spells each individual experienced. To a certain extent, this measure reflects a feature of poverty similar to our standard deviation measure, but it differs in that it focuses narrowly on moves in and out of poverty. This variable has the advantage of being more easily interpretable, and, of potentially capturing the impact of disruption associated with repeated moves in and out of poverty, however. Figure ?? shows the distribution of spells in our sample; the coverage is reasonable across different values: we have a fairly large number of individuals with two or three spells, and enough with four or five spells to conduct statistical inference about these groups. We also consider the variance of the intensity measures in children's poverty years and the variance in children's poverty status, i.e. a binary variable taking a value of one in years when a child was poor, and zero otherwise. These do not affect our results, however, so for the sake of space and focus, we do not discuss them further here or in our results and discussion.

[9]We use this measure because we can compute this value for our full sample.

The distinction we impose on this conceptual model of poverty exposure – considering duration, intensity, timing, and concentration as distinct elements that interact to determine outcomes – enjoys support in the literature, and the results and arguments found in the papers we build on inform our hypotheses regarding the treatment surface. Wagmiller et al. (2006) identify several distinct strands of literature on child poverty and life chances: using proxies for permanent status, as in the case of intergenerational mobility models; measuring the effect of cumulative exposure to some circumstance; and capturing differences arising from the timing of a homogeneous experience – each of which include a diverse selection of papers going back decades. Duncan et al. (2010) build on this work, focusing on timing, find a different result based on a finer-grained dataset, consistent with strong timing effects. Both papers approach their measurement of poverty duration as a kind of 'dose' of treatment, in contrast to the more common approach of measuring average income or using a measurement at one time as a proxy for long-term income or deprivation, and find that long durations that occur early in life have a much larger effect than shorter ones. Brooks-Gunn and Duncan (1997) and Duncan et al. (1998) consider the intensity of poverty as well, distinguishing between different poverty line shortfalls, as well as making the duration and timing distinctions. Both studies find that more intense poverty is more closely associated with lower cognitive ability and academic achievement, as is earlier poverty and longer durations of poverty. Taken together, they establish a conceptual foundation consistent with the division of poverty attributes we have proposed. None of these papers, however, provide empirical evidence sufficient to identify the shape of a treatment surface for any of these components. Yet, in the context of Cunha et al. (2010), intensity and duration capture essentially the same thing, i.e., the degree of foregone investment that affects a given child due to a lack of financial resources. Our treatment of intensity and duration as distinct is, therefore, an important aspect of our analysis.

Our multidimensional approach is fairly straightforward, but even in this relatively simple framework, the breadth of our three measures lets us consider multiple aspects of poverty that may not be easily tested using alternative data sources. For example, the MTO design suppresses the distinction between duration and timing, because the only information available is age at the time of treatment assignment. This means that, for someone treated at 13, the duration of added time in a less poor neighborhood will always be 4, which need not be the case more generally. In the PSID, however, we observe many different duration values for individuals entering poverty at each possible age, and the inverse, which lets us disentangle the relationships between adulthood rank and each of these treatments, and to allow for heterogeneity in each of these relationships as the other value changes.

We summarize our key variables in Table 4.1. These summary statistics show that a large majority of the children in our sample (discussed in more detail in the next subsection) experience poverty starting fairly early in childhood, and experience short to medium-length poverty spells, both with a fairly large amount of variation based on the variables' respective standard deviations. Intensity is substantially less dispersed, however, with an interquartile range that extends over only about 20 percent of the support, compared to about 30 or 45 percent for age and duration respectively.

4.3.3 Control Variables

In most of our models, we control for an array of childhood factors which could otherwise bias the relationship between adulthood rank and childhood poverty. Given that we collapse the data into two time periods, childhood and adulthood, in most cases the control variables (summarized in Table 4.2) are averages of categorical variables describing the head of the individual's household during childhood; for example, 'Married' captures the percentage of the individual's childhood in which the head of her household was married. The exceptions to this definition are our education variables: we define the household head's education as the highest level of education of the household head observed during childhood, which we then collapse into high school, high school and some additional education, and a bachelor's degree or more (labeled 'college').

Our control variables all pertain to the individual's household during childhood, and fall into several categories: demographic variables, education variables, occupation variables, industry variables, employment, and residence. Head of household (HOH) race, the individual's sex, and parental marital status comprise our demographic variable group, and with the exception of the individual's sex, all are averages across each individual's childhood observations. In the case of parents' marital status, the interpretation is straightforward because it captures the percentage of childhood each individual lived in a household with married parents. In the case of parental race/ethnicity, however, the measure is somewhat less intuitive. We measure race with a mean value of categorical variables defined based on the value for the head of household, so while most observations have values of zero or one, a decimal value is possible. We maintain this measure because the head of household can change, and we expect that important mechanisms through which head race *per se* could affect outcomes conditional on other observables we condition on include residential sorting, labor market discrimination, and social networks, and ignoring variation in the head of household by using only the modal category may omit this information. Our education variables capture the maximum attainment the head of household achieved at any point during each individual's childhood, so that if the head of household had completed high school but nothing more at a child's birth, but finished college when the child was older, our 'college' binary variable would take a value of one, and the other attainment variables would take values of zero. We code education in this way because we believe it reflects underlying parental human capital better than percentage of time in each category. While education clearly has economic consequences, those occur partly through labor supply and job type, for which we also control with our employment, industry, and occupation variables.

We also measure occupation and industry variables as averages, and for completeness, we include all of the top-level categories for both, despite the fact that several categories (e.g. farm laborers) are very small both nationally and in our sample. These variables capture the kind of work each individual's head of household did across the individual's childhood. We expect job type to affect both economic security and the kinds of skills and norms that parents transmit to children. We measure employment and location in the same way, capturing parental labor supply with the percentage of the individual's childhood in which her head of household was employed, and location of residence with the percentage of childhood each individual spent in one of the four regions the PSID defines: South, Northeast, North-Central, or West. The majority of individuals in our sample never move, and the majority of nonmovers spent their childhoods in the south, but enough individuals moved between regions during childhood that we believe this measure of residence captures meaningful information that a binary measure would not capture.

4.4 Results

4.4.1 Replication Results

We summarize the results in Tables 4.3 and 4.4, where we present the standard intergenerational elasticity model applied to our data and a rank-rank mobility model following Chetty et al. (2014a), and models following the formula in Equation 4, which produces a duration coefficient comparable to the age effect in Chetty et al. (2016). We report models with and without controls and age-duration interactions. The results in Table 4.3 suggest a much lower degree of mobility than standard estimates – between 0.4 and 0.5 for the intergenerational elasticity, and approximately 0.35 for the rank-rank correlation – but neither is so large as to suggest that our data are incomparable with other samples. This sample includes only children who experienced some poverty during childhood, which likely tightens the income distribution in these regressions so that mobility is lower in this sample than in other samples. In Table 4.4, our estimates of the marginal effect of poverty duration are also higher than the Chetty et al. (2016) estimates of the effect of an additional year in a poor neighborhood (~ 724), but only in the first two models, which do not include controls.[10] The marginal effect of an additional year of duration includes the Chetty et al. (2016) estimate in its 95 percent confidence interval when we include controls, and it includes their estimate at the 90 percent confidence interval when we include both controls and interactions, so while our point estimates are somewhat larger in absolute value than theirs, they are qualitatively comparable.β_2 captures the marginal effect of an additional year of poverty on adulthood family income. We include age at first exposure as well, because, as we discuss above, the Chetty et al. (2016) model measures age and duration simultaneously. To approximate their regression as faithfully as possible, we must control for this aspect of poverty exposure. Our objective in fitting these models is not to compete with the Chetty et al. (2016) or Chyn (2016) estimates, but rather to establish the extent to which our data are comparable to the data they use. We see that our data leads us to similar conclusions, particularly in Models 3 and 4 in Table 4.4, where despite using different data sources, we obtain similar estimates of the effect of additional poverty exposure.

4.4.2 Poverty Duration, Intensity, and Age of Exposure

In presenting our results, we begin with a discussion of results from a model of rank as a function of duration, intensity, age, and controls only, and then expand that discussion by considering poverty concentration as well. We begin by fitting a series of parametric, linear in parameters models, which let us estimate $g(\cdot)$ from Equation 2 in a somewhat restrictive, but easily presented and interpreted, manner. In Table 4.5, we fit a simple additive model of $g(\cdot)$, a model with linear interactions, and an additive model with quadratic terms. In each of these models, duration is consistently significant, and the goodness of fit changes very little across specifications: based on the results of an F-test, we cannot reject the null hypothesis of no difference

[10] In the relocation studies, researchers obtain treatment exogeneity through the experimental (program) design. While we do not have comparable bias reductions, controlling for a range of household covariates is the closest approximation available to us, because in those studies, household covariates are the confounders that randomization circumvents.

between Model 1 and each of Models 2 and 3. Across specifications, the marginal effect of duration is economically meaningful: an additional year of poverty costs approximately two percentiles of adulthood rank in the linear specifications.[11] In the quadratic specification (Model 3), an additional year is associated with a decrease of approximately three rank percentiles at one year of duration, and approximately half a percentile at sixteen years of duration. When we add controls in Table 4.6, this pattern persists; duration coefficients are smaller in absolute value, and in the interaction specification (Model 2), the coefficient is only significant at the 10 percent level; nevertheless, these estimates indicate the loss of one percentile per year. The persistence of this association across specifications, and its robustness to the inclusion of a wide variety of controls, suggests that poverty matters primarily through duration. Across specifications, the rank-duration relationship does not vary in the other treatments as we hypothesize, and the other treatments are not significantly associated with income rank on their own. Overall, the impression that emerges from these parametric models is that duration matters, but that its effects are linear and relatively small, although certainly not negligible; they are consistent with the Chetty et al. (2016) linearity result, and not consistent with results from the structural models that predict large and nonlinear effects.

We are primarily interested in the shape of $g(\cdot)$, however, and these models only go so far in letting us model that shape. When we plot a surface of fitted values across the support of our treatment variables, we see many predictions outside the permissible range of adulthood income rank (i.e., $[0, 1]$), primarily in regions with limited support, such as high intensity and low duration. In these regions, we are effectively extending a nonlinear function out of sample assuming that the shape of the function is constant, which, based on these theoretical violations regarding the fitted values, we believe to be too strong an assumption.

[11] We report marginal effects of each treatment, evaluated at medians of the relevant data in these tables as well, although the qualitative interpretation of these marginal changes very little. We compute standard errors analytically in most cases, but the compound variance terms on age have very small magnitudes and are often negative, so we bootstrap them.

To produce estimates which respect these constraints on fitted values while allowing an unspecified form of nonlinearities and interactions, we fit a fully nonparametric model, regressing adulthood rank on an unspecified function of age, duration, and intensity.[12] When we fit this model, presented in Figure ??, both age and intensity are found to not influence income rank, which varies only in duration. In other words, the slope of duration does not vary in intensity, which we can see in Figure ?? because the duration curve is constant across the intensity axis, and the same pattern repeats as we evaluate the surface at different values of age. These estimates are significant at the 95 percent confidence level, computed by bootstrapping the mean of the nonparametric function.

In technical terms, this occurs because the optimized value of the bandwidth term, which governs the size of the neighborhoods for which the estimator fits local regressions, is very large, so the estimated conditional mean of rank is constant across these 'smoothed out' variables.[13] In other words, the optimal smoothing parameter for a certain variable is so high that in the local regression that relies on this parameter, the outcome does not vary across that variable, so it can be said to be irrelevant, and the smoothed variable does not affect the regression results. This result is not equivalent to a formal specification test, however, so we test the hypothesis that a nonparametric model of rank as a function of duration only is equivalent to a nonparametric model of rank as a function of age, duration, and intensity using standard nonparametric specification tests. We use two tests: a goodness-of-fit test comparable to a nonparametric F-test, in which the restricted model is the model with only duration and the unrestricted model includes age, duration, and intensity, and a direct test for the irrelevance of a subset of variables in a nonparametric regression, originally developed by Lavergne and Vuong (2000).[14] Both tests lead us to fail to

[12]We fit this model using bandwidths estimated by least squares cross validation, chosen for local constant least squares (LCLS) regression, which constrains the fitted values to fall on $[0, 1]$.

[13]See Hall et al. (2004) and Hall et al. (2007) for a more detailed treatment of bandwidth estimators smoothing out irrelevant variables.

[14]For a detailed explanation of these tests, and of the nonparametric testing paradigm in general, see Chapter 6 of Henderson and Parmeter (2015).

reject the null hypothesis of no difference between specifications at the 10 percent confidence level, confirming the validity of smoothing these variables out in these regressions.[15]

These results are consistent with the parametric results in suggesting that age and intensity do not matter apart from duration, and further indicate substantial nonlinearity in the rank-duration relationship, as well as a stronger relationship overall. The nonparametric models predict that a child who is poor for six years, the average value of duration in our sample, will reach an adulthood position in the income distribution approximately thirteen percentiles lower than if she had spent only one year in poverty, in contrast to values between six and nine percentiles in the parametric models. At higher levels of duration, the slope flattens substantially, however. They are also more consistent with the structural literature than with the Chetty et al. (2016) linearity result, as they suggest that a disproportionate amount of the effects of poverty comes from the first few years. This makes sense in the context of dynamic complementarity, because the structural models (e.g., Cunha et al., 2010; Agostinelli and Wiswall, 2016) predict that the benefits of additional family resources will be substantially lower after having spent several years in poverty than after having received a small dose. The fact that the slope levels off above the mean, and especially at high values, may seem to contradict the structural results, but the fact that those papers consistently find weaker effects for investments in older poor children reflects a similar leveling off process. They make the argument in terms of timing, but intervening early in a poor child's life also implies intervening at a relatively low value of duration, and that is consistently what the structural literature recommends.

The nonparametric results let us detect nonlinearities without imposing functional form assumptions, while also constraining the fitted values to fall on $[0, 1]$, but at the same time the results presented thus far are not conditioned on control

[15] The p-value in the goodness of fit test is about 0.95, and about 0.12 in the Lavergne and Vuong (2000) irrelevance test, both of which constitute evidence that we cannot reject the null hypothesis of no difference between the models.

variables, which seem to be important based on the parametric results.[16] We build on the fully nonparametric results by fitting a series of partially linear models, in which $g(\cdot)$ is allowed to take an arbitrary form as in the nonparametric model, but it is conditioned on a parametric, additive, linear in parameters function of our control variables. We summarize this model in Figure 4.7, where we plot the fitted values against duration only, because as in the fully nonparametric models, the other features are not relevant. The same rank-duration pattern that we see in the fully nonparametric models emerges here, and while the parametric models produce similar results – the 95 percent confidence intervals, not pictured for readability, overlap – the slopes do vary locally, particularly at low to medium durations. These are our preferred estimates, because they retain attractive features of the nonparametric model while allowing us to control for the parental and household characteristics summarized in Table 4.2. This requires the assumption that duration, intensity, and age are additively separable from the control variables. Making this assumption is warranted because we are interested in the marginal effect of deprivation imposed by different features of poverty, above and beyond whatever the inputs provided by the parental characteristics we hold constant.

We provide an alternative visualization in Table 4.8, in which we plot the mean values of a nonparametric regression of rank on duration and a partially-linear regression of rank on age, duration, intensity, and controls by years of duration, with bootstrapped confidence intervals. The shape of the rank-duration surface is similar between the two models: in both panels, we see sharp nonlinearity at low values of duration, and less nonlinearities at high values, although in the conditional model, the tighter confidence interval suggests substantial nonlinearity even at high levels of duration. Taken together, these results indicate that the rank-duration relationship

[16]We do not include controls in the nonparametric models because fully nonparametric estimators effectively allow arbitrary interactions and nonlinearities for all variables, which makes isolating and interpreting $g(\cdot)$ difficult. Furthermore, adding additional variables to a nonparametric model dramatically increases the amount of data needed for any degree of precision, and with only 984 observations in our preferred sample, precisely estimating a nonparametric function of between thirty and forty variables is not feasible.

is nonlinear, and that the shape of this relationship is robust to the inclusion of our controls.

4.4.3 Poverty Concentration

We repeat much of this analysis accounting for poverty concentration as well as duration, intensity, and age. First, we hold concentration constant by using only individuals who were poor in contiguous years only, so that no one in the sample moves in and out of poverty. Using this sample substantially reduces our sample size and the support of our treatment variables, particularly duration, so we also fit models using our preferred sample, in which we use the concentration measures discussed in Section 3.2 to control for poverty concentration.

In the contiguous poverty sample, each individual has the lowest possible concentration of poverty years for their duration and age at first exposure, so unobserved differences in concentration cannot bias our estimates. The most substantial difference between the partially linear model results in this sample and the results from the main sample is that intensity is not smoothed out by the kernel selection procedure.[17] This result suggests that leaving concentration unobserved suppresses the relationship between adulthood rank and intensity; individuals with the most concentrated poverty appear to be affected by the intensity of that poverty more than the average individual in the main sample. In economic terms, this is consistent with the predictions of the structural literature: dynamic complementarity (Cunha et al., 2010) suggests that a deeper shock followed closely by additional shocks will have a larger long-term effect, as the deprivation in this period will compound more than it would had the child experienced the same duration, intensity, and starting age spread over several years. This interpretation is consistent with the fact that our estimated duration curves look similar to those from the main sample, because they

[17]We fail to reject the null hypothesis of no difference between the age, duration, and intensity model and the duration only model using the Lavergne and Vuong (2000) irrelevance test (p = 0.516) at a 5 percent confidence level, but reject it at the 95 percent confidence level (p=0.001) based on the goodness of fit test. This is inconclusive, so we proceed with these intensity results.

are conditioned on intensity, and while it would make sense that a more concentrated poverty duration would also produce a compounding effect, that effect is mediated by intensity. Alternatively, this result may reflect the fact that intensity is better measured in this sample, because there is only one poverty spell. It is worth noting that the rank-duration curve remains nonlinear in this subsample, although less so than in the main sample.

We adjust our primary models by adding the standard-deviation based measure of concentration, discussed in Section 3.2, which increases as concentration decreases. Table 4.7 shows estimates from parametric models including this concentration measure. Adding concentration to these models does not affect the marginal effects of other treatment variables, and in many cases, it does not significantly change the model's fit, according to F-tests comparing the model with concentration to the base I-D-A models. The same result holds in nonparametric and semiparametric models. Overall, these results suggest that, while the contiguous sample does display different behavior, the omission of concentration in our main results does not affect our findings.

We summarize the fitted values by years of poverty in Figure **??**, by spell and years of duration in Figure 4.12, and by duration and intensity across different numbers of spells in Figure 4.13. The averages in Figure **??** follow a more locally linear pattern than the averages in Figure 4.8, but they are estimated much less precisely, in part because this figure ignores the intensity and spells dimension. Our original result falls well within the 95 percent confidence interval, so while this figure suggests that including poverty spells makes a difference, it does not undermine the original finding. In Figure 4.12, it is clear that the shape of the duration curve does vary across different numbers of spells, with children with more spells seeing a steeper drop before ten years of duration, but a similar flattening at high values. The lack of support in some duration-spell number bins, along with the width of the 95 percent confidence interval, makes this comparison tenuous, however. We reach a similarly ambiguous conclusion based on the semiparametric regressions presented in Figure

4.13, in which no clear pattern emerges across numbers of spells, although the general shape of the rank-duration curve persists, albeit with many deviations in the (wavy) surfaces. Taken together, these figures suggest that the number of spells matters in a way concentration does not, but their inclusion does not substantially alter the rank-duration curve.

4.5 Robustness Checks

4.5.1 Adjusting for Selection on Observables

Our first robustness assessment is to reduce, if not eliminate, sample selection bias using propensity score methods developed for continuous treatments through the use of a Covariate Balancing Generalized Propensity Score (CBGPS) approach, developed by Fong et al. (2018). In brief, this method is designed to reduce the bias in a regression of a continuous outcome on a continuous treatment and controls by minimizing the association between the treatments and the controls, and by reducing sensitivity to model misspecification, without making strong distributional assumptions. In our setting, this is helpful because many of the relatively large number of covariates we include as controls are correlated with our treatments. Furthermore, with multiple continuous treatments and many controls, the number of potentially credible specifications is high, so reducing bias due to misspecification also becomes particularly important. We are particularly concerned with the degree to which our estimates of the rank-duration slope are biased by covariate imbalance, as duration consistently has a strong relationship with rank, while the other two treatments do not. To this end, we apply the CBGPS procedure to balance covariates with regard to duration, and not the other treatments.[18]

We re-estimate models of rank as a function of duration adjusting for CBGPS-based weights, designed to consistently estimate the model after adjusting for the

[18]To our knowledge, the procedure cannot be applied in a way that adjusts for balance over three covariates.

CBGPS itself, which we summarize in Figure ?? and Table 4.8. In Figure ??, we summarize the differences between correlations between each control variables and duration, before and after CBGPS weighting. The height of the second bin indicates that for eighteen of our control variables variables, weighting reduced their correlation with duration by between 0 and 0.1. In almost all cases, the decrease is positive in absolute value, indicating that the adjustment did, in fact, help to balance the covariates with respect to treatment. The relatively large number of covariates showing only a small difference may suggest the effect was limited, but in absolute terms, the decreases we see are much greater than those in Fong et al. (2018), whose decreases were all between 0 and 0.1 in absolute value, so that they would have no frequency in bins beyond the second, corresponding to larger reductions in imbalance. This occurs primarily because our baseline imbalance was much larger than the imbalance in Fong et al. (2018).

In Table 4.8, we summarize parametric models equivalent to the parametric models with controls in Table 4.6, but adjusted for CBGPS weights. The duration coefficients become larger across all specifications, and the fit improves, suggesting that bias from covariate imbalance was, if anything, limiting the size of the rank-duration association. At the same time, however, these models produce implausibly large positive coefficients on intensity, and while this likely has to do with the fact that we can only adjust for imbalance with respect to duration, it does limit the robustness of the CBGPS specifications. Our semiparametric results thus remain our preferred specification, but the CBGPS results do reinforce the semiparametric results by showing that our duration results, if not necessarily the null results for the other treatments, are robust to adjustment for selection on observables.

Adjusting for the CBGPS changes the qualitative interpretation of our results only slightly. This exercise, does, however, serve to enhance the credibility of our baseline estimates of the rank-duration curve, because it establishes that our initial estimates were affected by bias arising from misspecification and imbalance in covariates given treatment only to a limited degree.

4.6 Discussion

We make two related contributions: we show that childhood poverty is closely related to adulthood income rank through cumulative disadvantage imposed through duration, and we find that, contrary to our expectations, the timing and intensity of poverty have very little effect on the model given duration, either on their own or by shifting the duration slope. Our substantive finding is that poverty duration is, conditioning on a wide array of observables and using a sample designed to minimize bias from various sources of measurement error, strongly associated with a lower adulthood rank, and the rank cost of more duration grows very quickly at low values but then tapers off. Our model predicts that, all other things being equal, someone who spent her whole childhood in poverty would achieve an adulthood income rank approximately 23 percentile ranks lower than she would have otherwise. The majority of this rank reduction occurs in the first six years of poverty (which need not occur in the first six years of life). After six years of poverty, there is a relatively flat period through eleven years of poverty, then a more or less linear reduction from eleven to seventeen years. To a certain extent, this shape is consistent with predictions from the structural literature (Heckman, 2008; Cunha et al., 2010; Agostinelli and Wiswall, 2016) regarding the timing of shocks, because absorbing a medium to long spell of childhood poverty – which would imply a low adulthood rank based on this curve – is impossible if one enters poverty for the first time at an older age. This means that our results do not directly contradict the established age results, because while we do not find a direct age-rank relationship, children who enter poverty early remain at greater risk because they can be exposed to a longer duration. Similarly, the relatively flat slope for children with a relatively long poverty duration suggests that a non-marginal reduction in poverty is necessary to improve their adulthood outcomes.

How, then, do our results help to bridge the structural, skill-based literature with the experimental treatment effect literature? Our results reinforce the policy implications of research that finds a particularly large effect of family resources on low-resource households. This research suggests that interventions, such as augmenting parental resources through the earned income tax credit (Dahl and Lochner 2012) or providing compensatory investments in skills (Cunha and Heckman, 2010; Agostinelli and Wiswall, 2016), would be most effective for children who are likely to experience, a relatively low to moderate degree of poverty duration, because as our results show, this initial period is when the negative effects of poverty are most severe. Our result is weakly age-neutral, because a child in her fourth or fifth year of poverty could be any age, but the general principle of targeting relatively young and disadvantaged children, prior to any sustained exposure to poverty, holds.

On the other hand, our model provides some insight into why Chetty et al. (2016) find a linear age effect despite the early childhood literature. Treatment in the Chetty et al. (2016) case is movement away from poverty, which roughly corresponds to an end of "treatment" in our model – i.e., preventing a child from reaching poverty duration of 13 or more years. In this region, however, the rank-duration slope is more or less linear and is remarkably shallow: the confidence interval includes a straight line with a slope of approximately -0.0675 ranks per year, lower than the linear slope we estimate in our parametric model, and much lower than the slopes we estimate elsewhere in the semiparametric and nonparametric models. Though our estimated rank-duration relationship is nonlinear, we see a relatively flat relationship at longer poverty spells. In this region, however, the rank-duration slope is more or less linear and is remarkably shallow: the confidence interval includes a straight line with a slope of approximately -0.0675 ranks per year of poverty, lower than the linear slope we estimate in our parametric model, and much lower than the slopes we estimate elsewhere in the semiparametric and nonparametric models. The Chetty et al. (2016) interpretation – that a fixed disruption cost of moving overwhelms the sum of the benefits for these children – also fits our estimates. At the same time, we believe that the overall nonlinearity in the rank-duration relationship stems from the differences between neighborhood and household poverty. Page and Solon (2003) show that within-household income covariances for children are about twice as high as within-neighborhood covariances, suggesting that while neighborhoods matter, families matter more. If MTO treatment had a small to null effect on children's parents – which earlier results (e.g., Kling et al., 2007; Ludwig et al., 2013; Jacob et al., 2015) find is the case – treatment changes neighborhood inputs while leaving the household characteristics, which include the propensity to enter poverty, constant. The structural literature does not generally draw a distinction between neighborhood and household inputs, instead specifying general investment, which may occur through positive neighborhood inputs; yet this literature does predict that children who incur significant disadvantage will gain less from new investments.[19] Thus, in the Chetty et al. (2016) context, treatment provides improved neighborhood inputs, but does not change the more important household inputs, and so the improved investment is not sufficient to take advantage of nonlinearities in returns to overall investment. Furthermore, our finding that the duration slope is constant across age and poverty intensity removes the objection that the MTO results conflate age and duration, as our results suggest that this conflation does not matter (although this is not equivalent to claim that the timing of poverty does not matter).

We control for a wide array of factors, but there are some parent and child characteristics we cannot control for well. To the extent that our controls account for the channels through which these factors affect adulthood rank, we can manage their influence partially. One example of such a variable is parental health. Our controls limit threat of bias from parental health variables, in part because of the breadth of poverty: in many cases, unobserved parental health variables will affect adult outcomes through our poverty variables (e.g., chronic conditions suppressing wages). In other cases, these shocks, e.g., a serious acute illness in a parent, could be realized in employment, which we capture with our employment and job type (i.e.,

[19]The notable exception to this characterization is Agostinelli (2018), who incorporates both household and neighborhood effects, and finds that parental investments and peer effects are substitutes within a period, but that parental investments in one period improve both the quality of peers and the effect of those peers.

industry and occupation) variables. Factors which reduce important parental factors such as income or time availability, but that do not affect one of our poverty measures, are a more serious concern. The persistence of the rank-duration relationship, across many specifications and conditional on many controls which should affect household resources, helps to mitigate it, however. This does not eliminate the possibility of bias from unobserved, sub-poverty disadvantage, but it would have to be the case that the poor households were consistently near poverty in non-poverty years for this to threaten our results. Doubling the poverty threshold and recalculating our poverty measures only slightly increases the average duration of poverty, however – shifting it from about 6 to about 7 years – so this does not appear to be the case in our sample. At the child level, the most notable omission is some measure of baseline ability: microeconomic models of skill formation, from Becker and Tomes (1979) to more recent work by Cunha and Heckman (2010) and Agostinelli and Wiswall (2016) emphasize the importance of initial ability stocks, affected by genetics, prenatal health, and random chance. For this to cause large bias, however, it would have to be the case that children more likely to experience adverse peri-natal conditions or poor genetic endowments are more likely to be poor conditional on our observables, and while this in fact seems reasonable, we also expect that our controls – particularly education, marital status, and employment – mitigate this source of bias.

4.7 Conclusion

In this chapter, we have explored the relationship between poverty and adulthood rank in the income distribution, by focusing on poverty duration, intensity, timing, and concentration. The literature on childhood skill suggests that these features are complementary, conferring increasingly large skill penalties as poverty increases. Consequently, we expect that separating these factors allows us to gain insight into the ways in which different facets of poverty affect individual economic outcomes. We find that the duration of poverty is the only component of poverty that matters in the long term, and that the slope of the rank-duration model is nonlinear in the manner suggested by the structural literature. This result reinforces the policy recommendations of the skill-development literature, i.e. that investments in children who have experienced relatively low amounts of poverty, but are at risk of experiencing more, are the most efficient. Our results also help to contextualize differences between Chetty et al. (2016) and this structural literature, as we find a nonlinear relationship between poverty duration and income rank for children with short to medium poverty spells, but an approximately linear and much flatter relationship for children with long poverty spells.

This suggests that people on the far left tail of the income distribution will have particularly low expected ranks, as those are the people whose average childhood incomes are particularly low and who experience high levels of duration, but that forgoing additional poverty will produce monotonic, if nonlinear, gains in absolute mobility. The fact that factors other than poverty duration appear to matter very little is very important here, because if the other treatments mattered on their own or by changing the rank-duration relationship, it could easily be the case that children with similar childhood average incomes – the standard explanatory variable in mobility models – would have very different prospects.These gains are particularly large for children who experience small doses of poverty duration rather than moderate doses, as opposed to children who experience medium doses rather than large doses, so they suggest that interventions aimed at children who experience moderate doses would be the most efficient means of improving upward mobility for poor children.

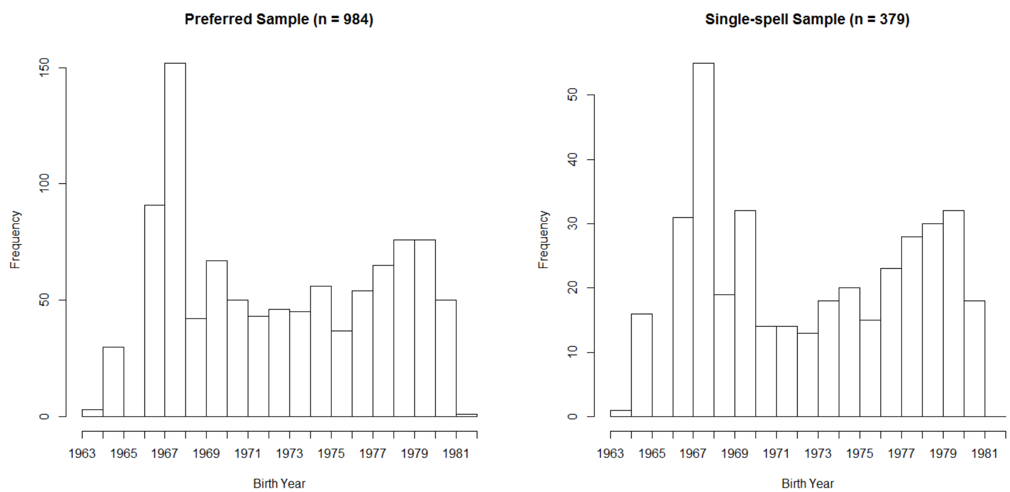

Figure 4.1.: Frequencies of observations by birth year and sample

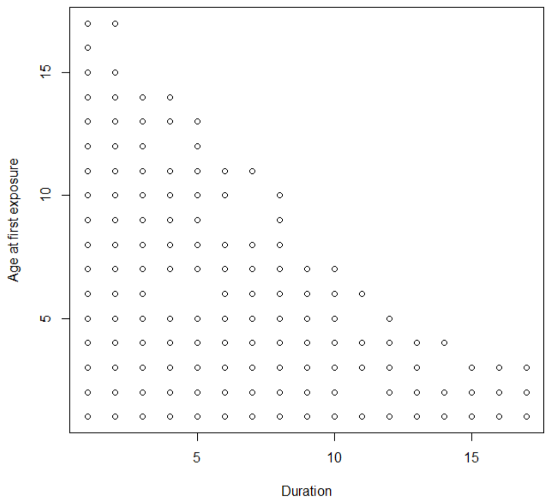

(a) Duration and age

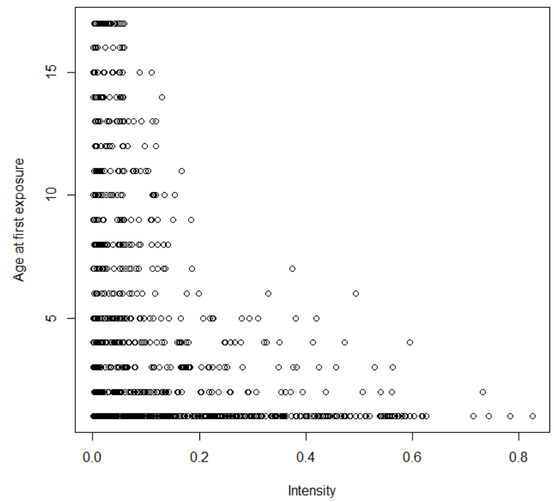

(b) Intensity and age

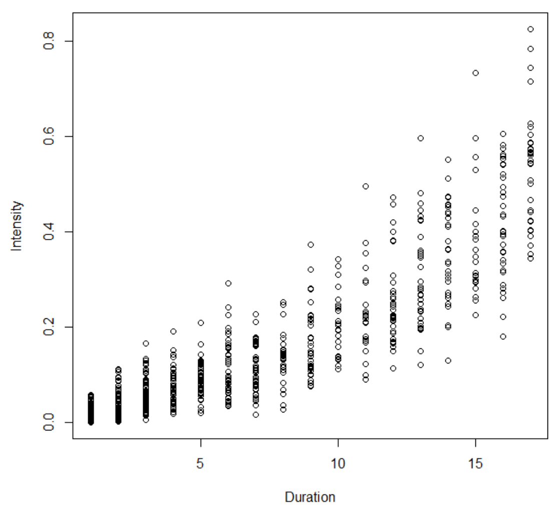

(c) Duration and intensity

Figure 4.2.: Bivariate Scatterplots of the three treatment variables

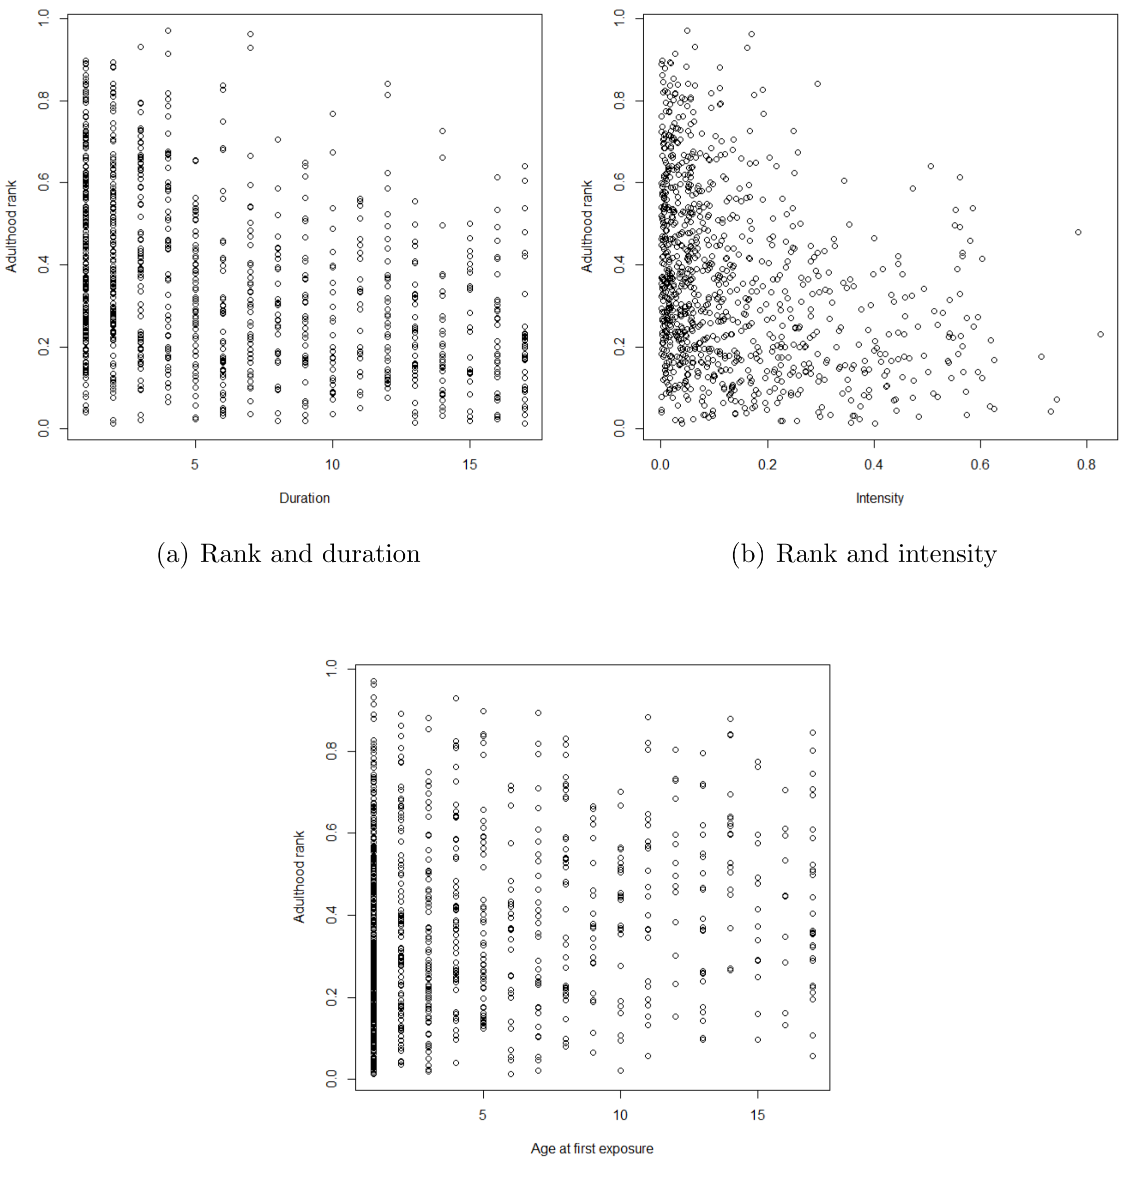

(a) Rank and duration (b) Rank and intensity

(c) Rank and age

Figure 4.3.: Bivariate scatterplots of rank by treatment variables

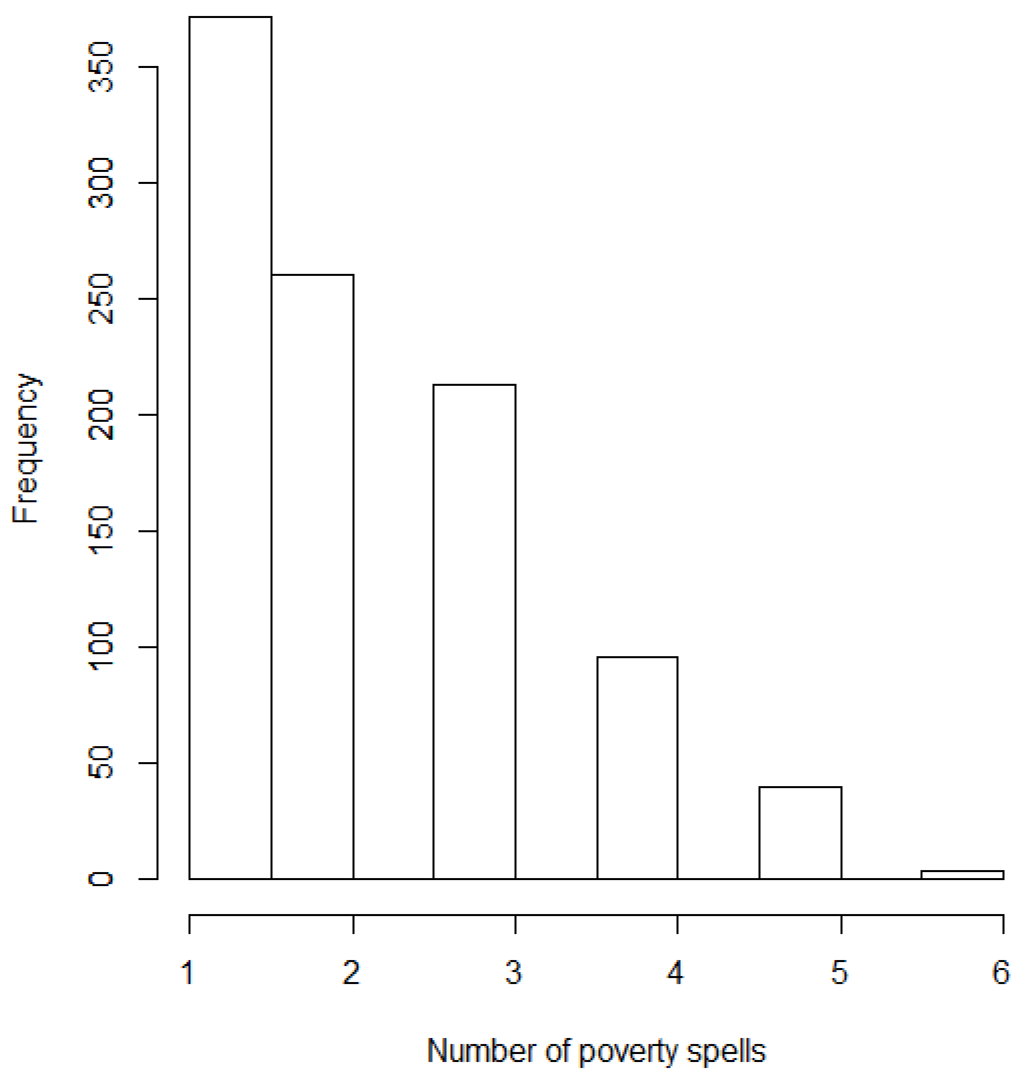

Figure 4.4.: Frequency of numbers of poverty spells during childhood

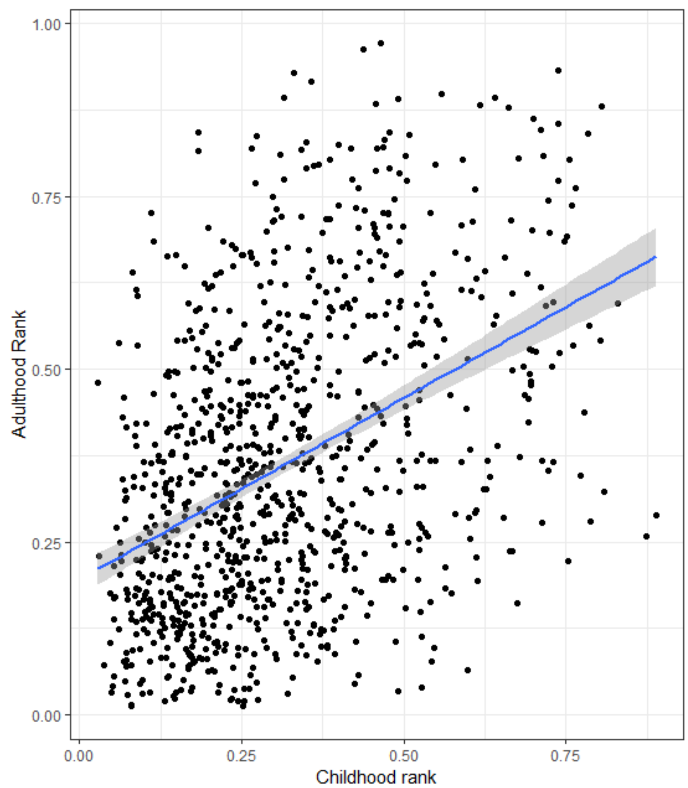

Figure 4.5.: Regression of rank on parental rank with 95 percent CI and scatterplot

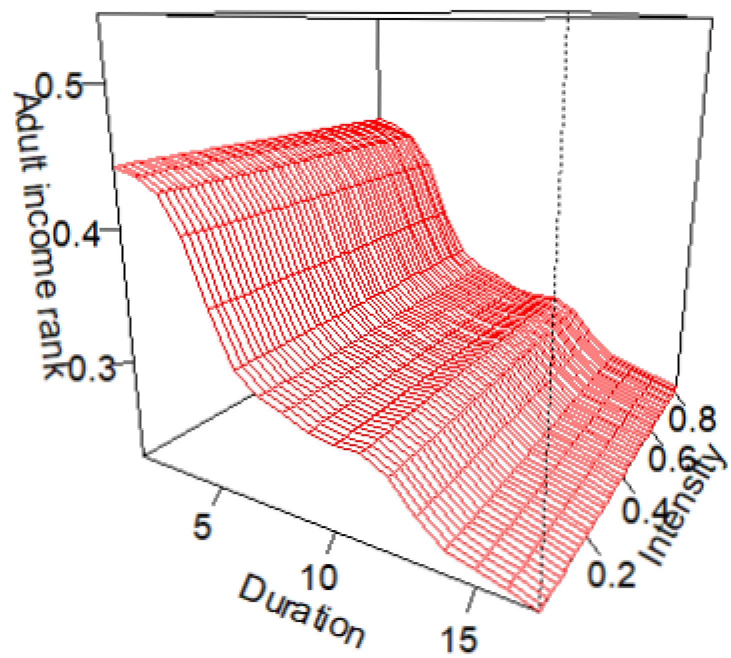

Figure 4.6.: Fitted values by intensity and duration, from nonparametric regression of rank on treatments

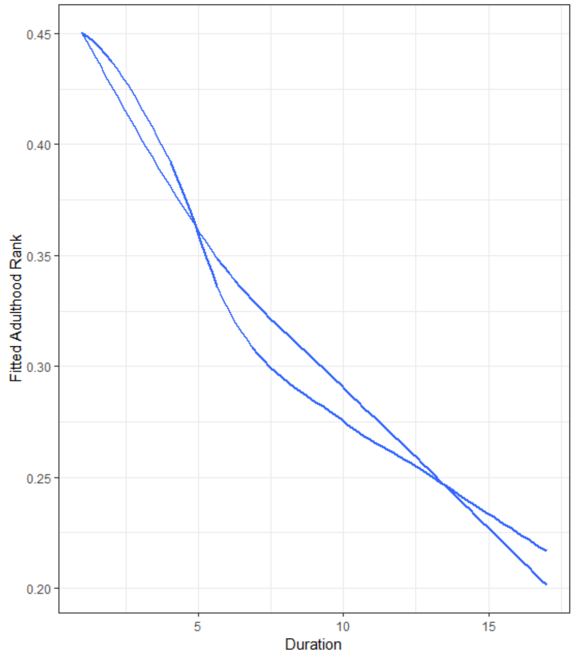

(a) Linear parametric

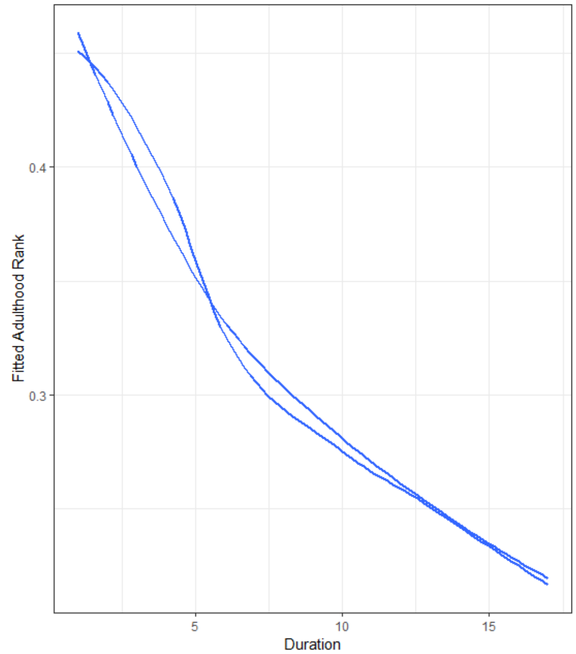

(b) Quadratic parametric

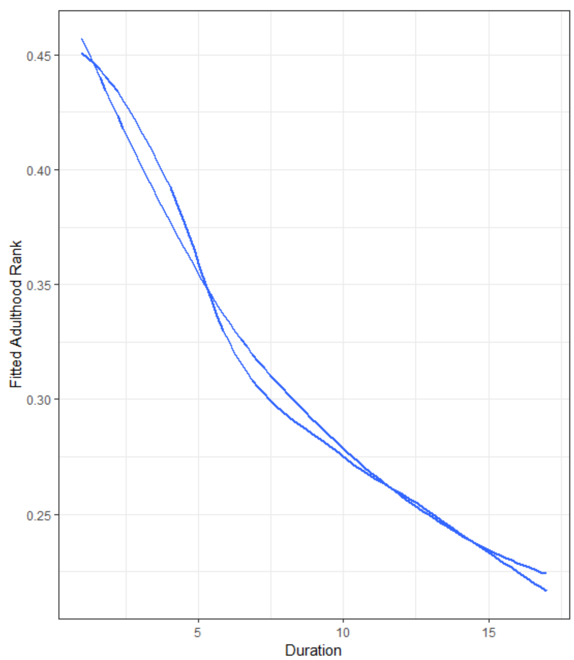

(c) Cubic parametric

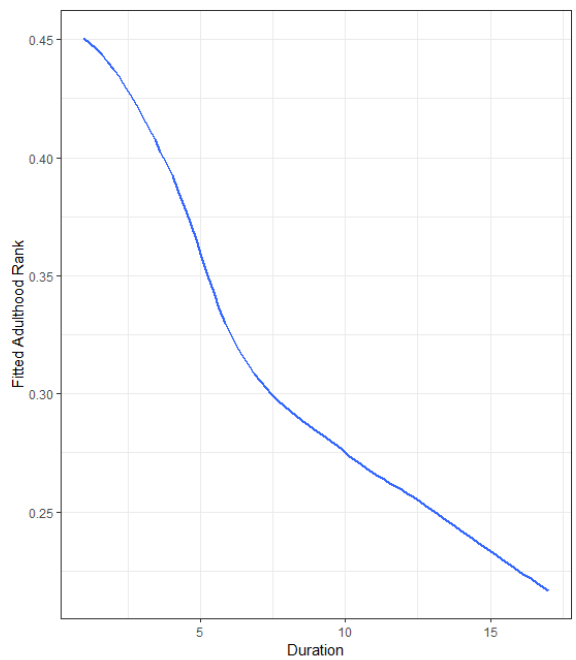

(d) Partially linear model only

Figure 4.7.: Conditional duration models by polynomial degree, overlayed with the partially linear model

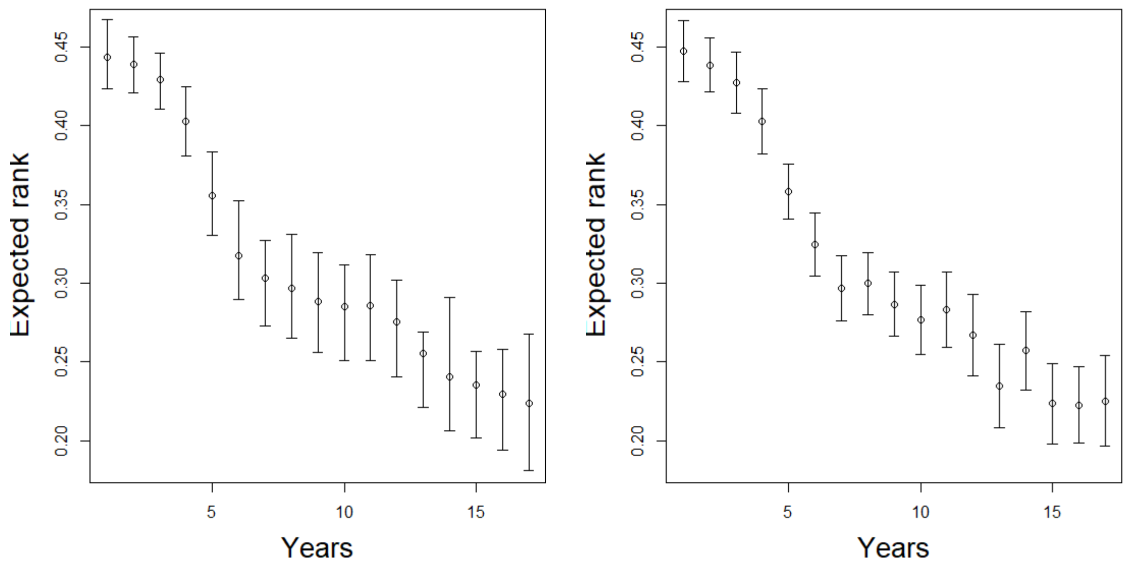

(a) Bivariate nonparametric duration model (b) Partially linear model

Figure 4.8.: Mean fitted values by years of duration with bootstrapped confidence intervals

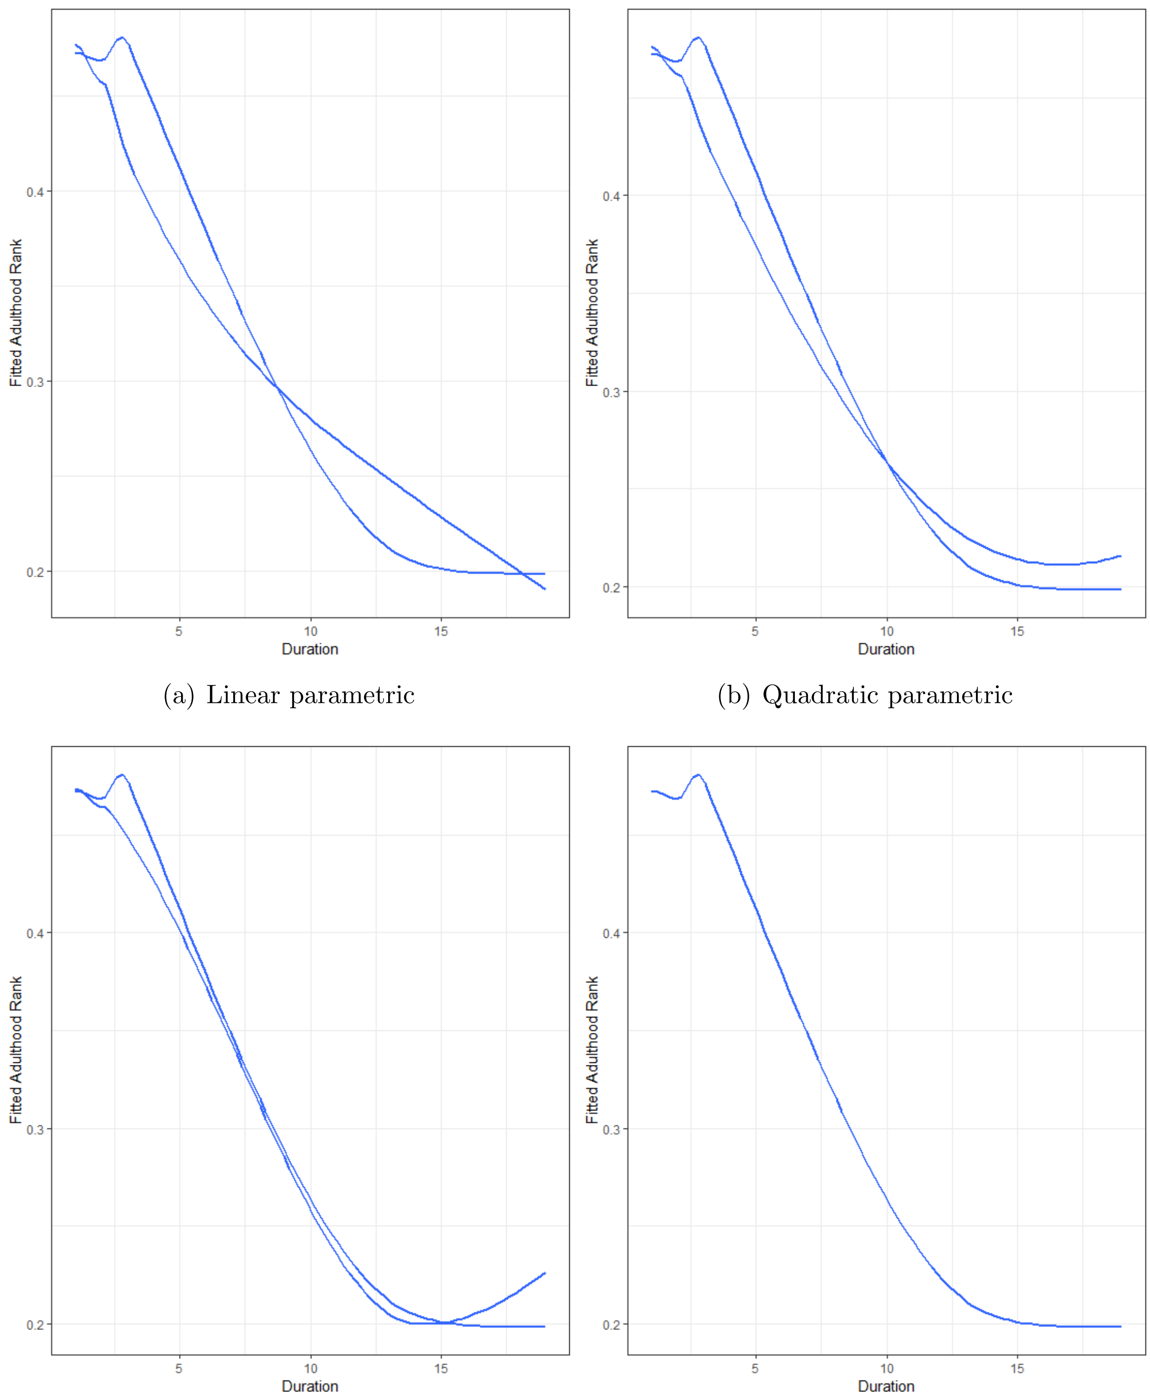

(a) Linear parametric

(b) Quadratic parametric

(c) Cubic parametric

(d) Partially linear model

Figure 4.9.: Conditional duration models by polynomial degree, overlayed with partially linear model (contiguous sample)

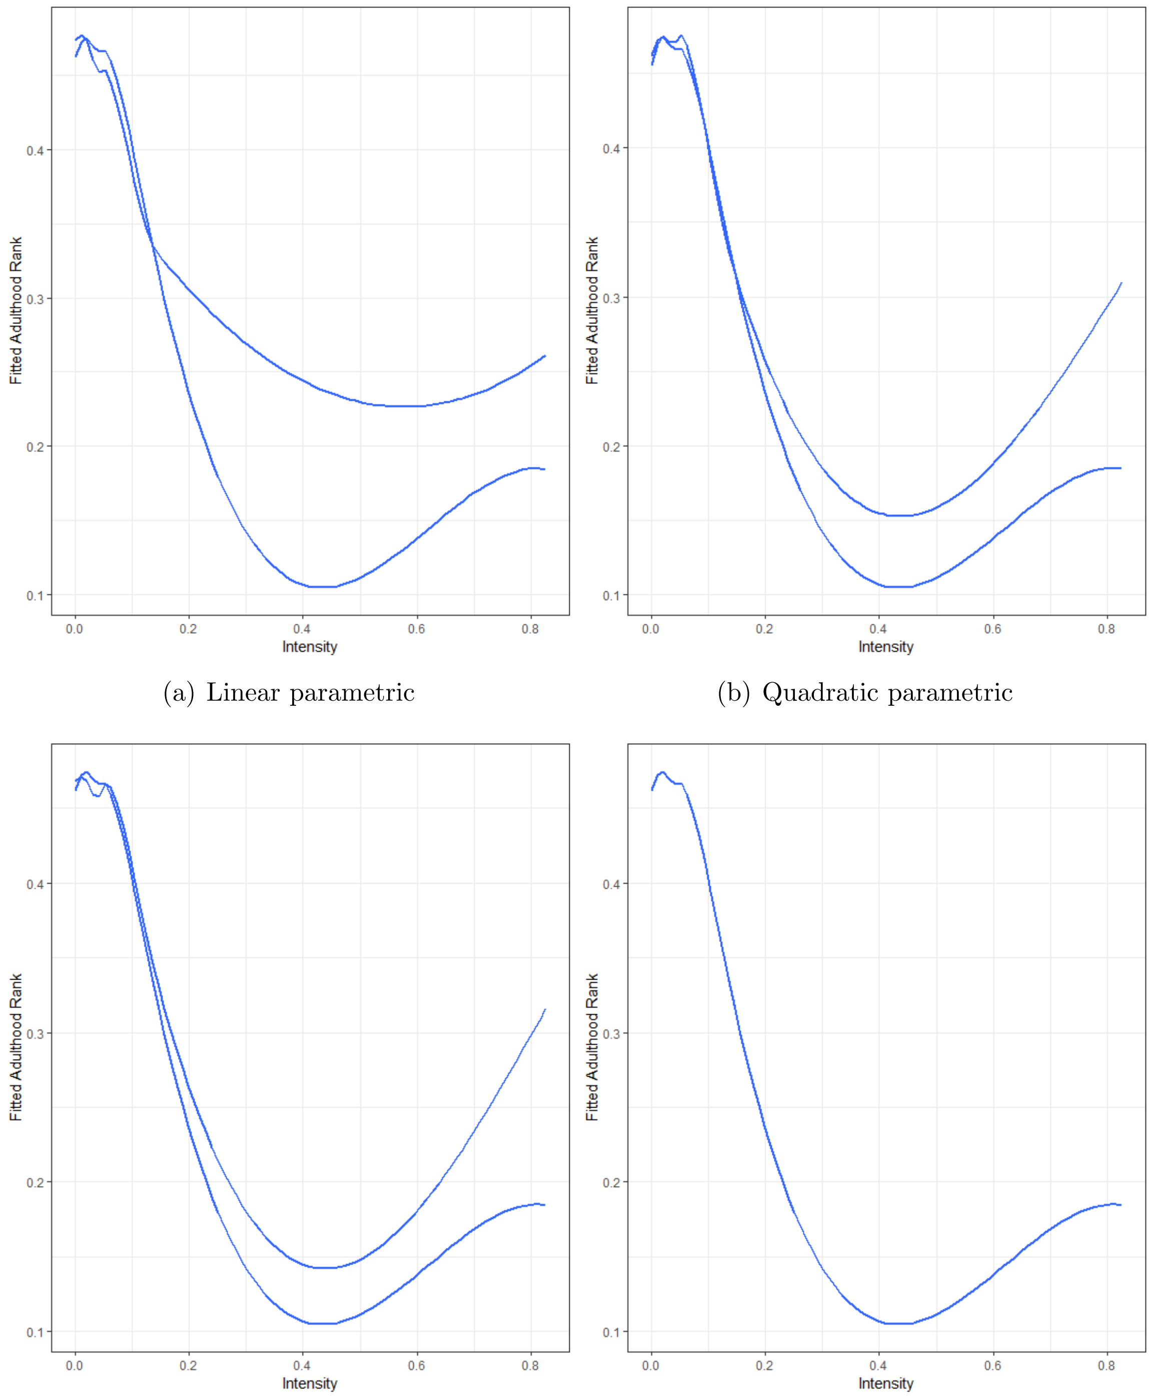

(a) Linear parametric

(b) Quadratic parametric

(c) Cubic parametric

(d) Partially linear model

Figure 4.10.: Conditional intensity models by polynomial degree, overlayed with partially linear model (contiguous sample)

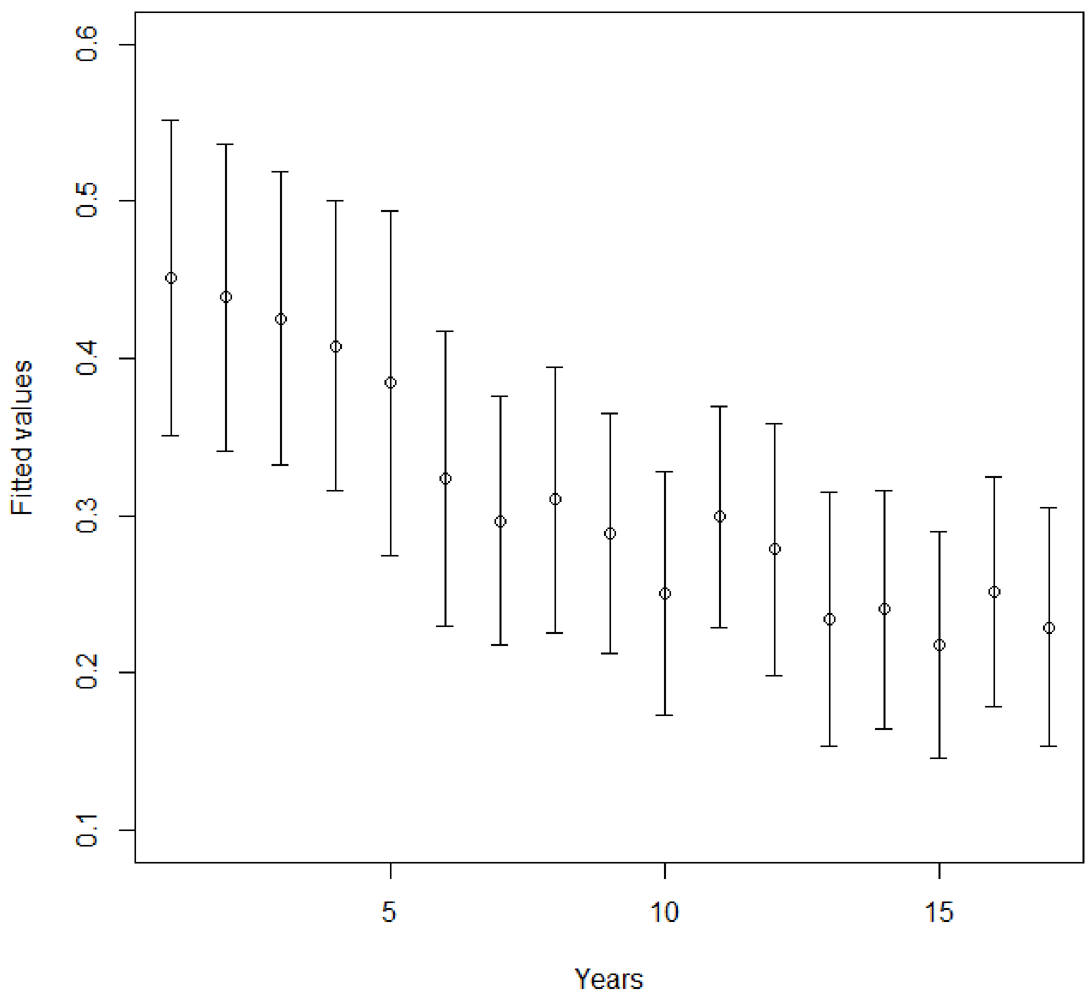

Figure 4.11.: Partially linear model including spells, averaged by year

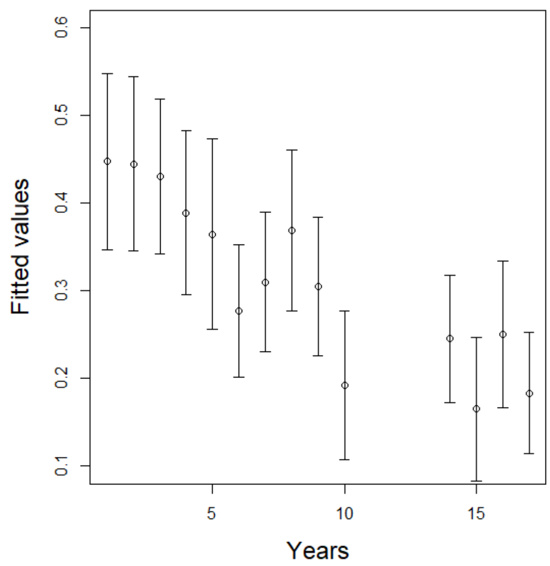

(a) One spell

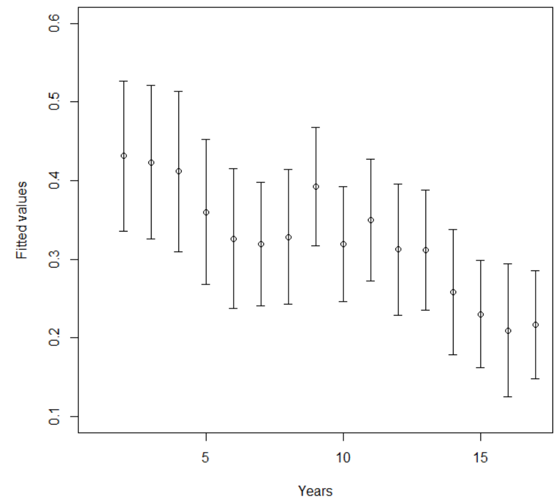

(b) Two spells

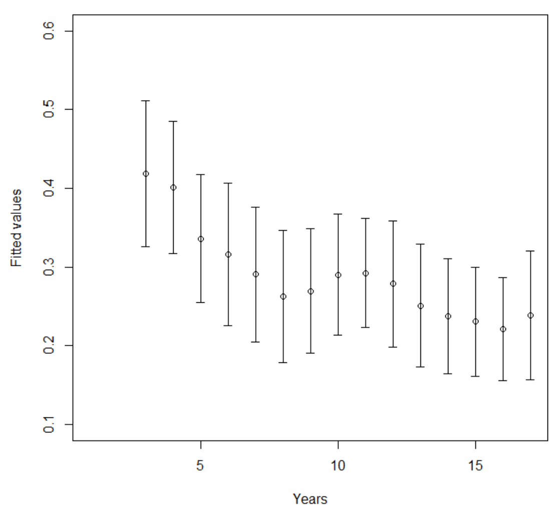

(c) Three spells

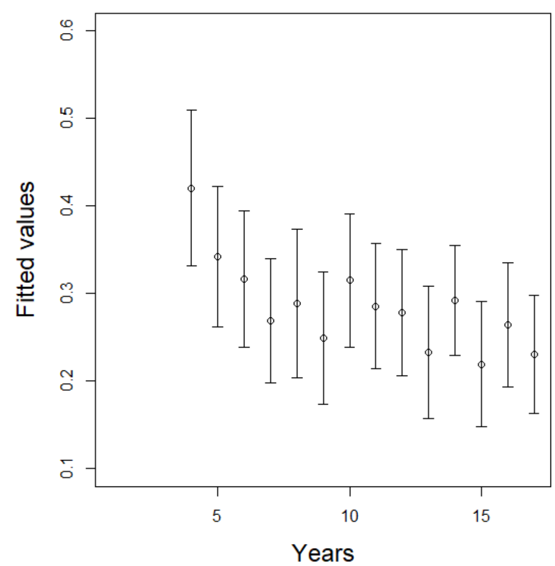

(d) Four spellsl

Figure 4.12.: Partially linear model means, by spell and duration

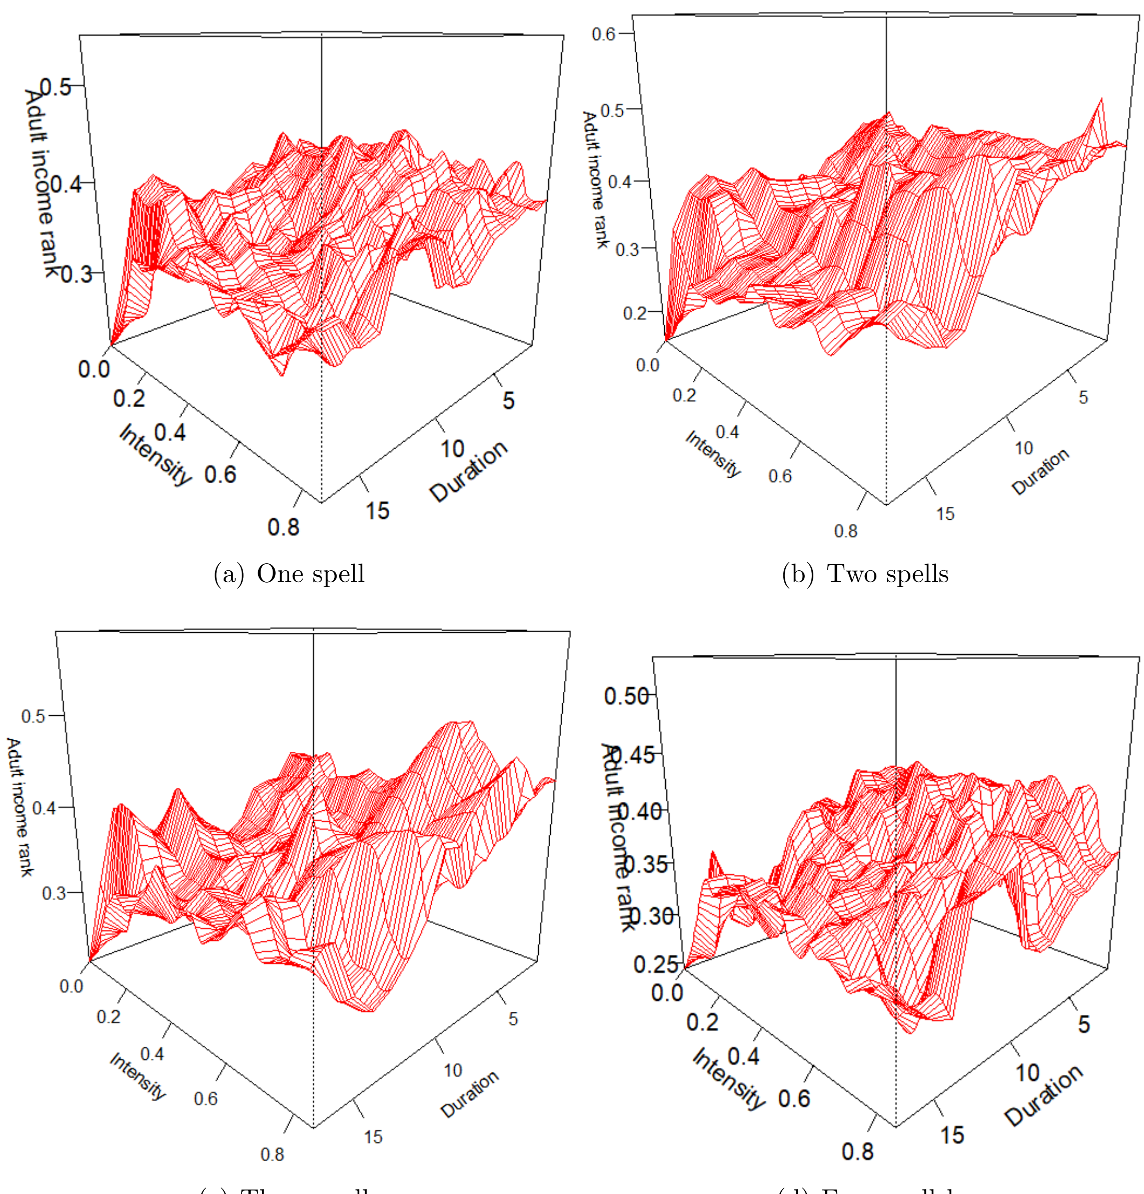

Figure 4.13.: Partially linear model intensity-duration surfaces, by spells

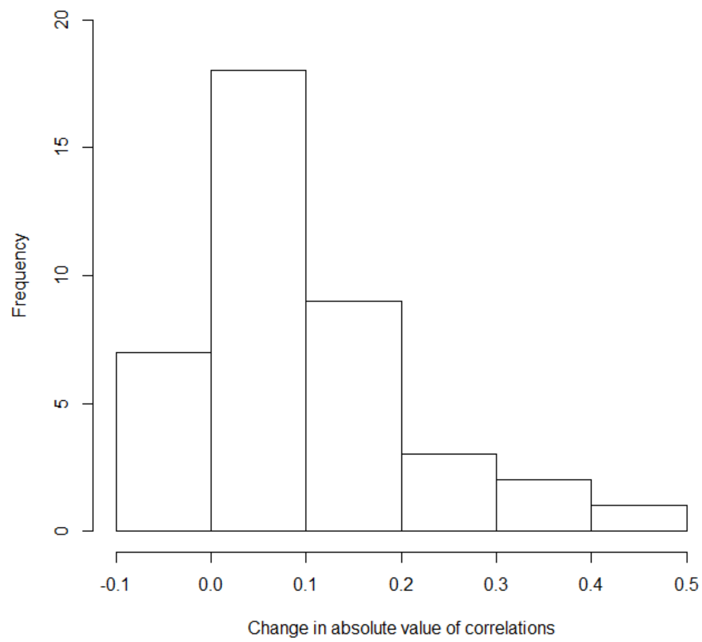

Figure 4.14.: Change before and after CBGPS balancing, in the correlation of rank and the control variables

Table 4.1.: Key variable descriptive statistics

	Minimum	Mean	Maximum	St. Dev.
Age	1.00	4.27	17.00	4.30
Duration	1.00	6.07	17.00	5.12
Intensity	0.00	0.14	0.78	0.15
Family income (adulthood)	131.90	50120	568200	37845.18
Family income (childhoood)	4054	41300	235100	23528.14
Rank (adulthood)	0.01	0.22	0.99	0.16
Rank (childhood)	0.03	0.31	0.86	0.17

Table 4.2.: Descriptive statistics of the control variables

	Minimum	Mean	Maximum	SD
Demographic variables:				
Black Head of Household (HOH)	0.00	0.54	1.00	0.49
Other Race HOH	0.00	0.02	1.00	0.12
Hispanic HOH	0.00	0.03	1.00	0.14
Female	0.00	0.56	1.00	0.50
Married	0.00	0.63	1.00	0.38
Number of children in family	1.00	2.65	14.00	2.04
Individual's birth order	1.00	1.34	7.00	0.68
Education variables:				
High School Attainment	0.00	0.15	1.00	0.36
High School plus some additional Attainment	0.00	0.23	1.00	0.42
College Attainment	0.00	0.06	1.00	0.17
Occupation variables:				
Professional Occupation	0.00	0.08	1.00	0.18
Manager Occupation	0.00	0.03	1.00	0.10
Sales Occupation	0.00	0.09	1.00	0.19
Clerical Occupation	0.00	0.15	1.00	0.26
Craftsmen Occupation	0.00	0.16	1.00	0.26
Operatives Occupation	0.00	0.06	1.00	0.17
Transportation Occupation	0.00	0.08	1.00	0.20
Laborer Occupation	0.00	0.00	0.38	0.01
Farm Manager Occupation	0.00	0.02	1.00	0.12
Farm Labor Occupation	0.00	0.17	1.00	0.27
Service Occupation	0.00	0.03	1.00	0.13
Household Work Occupation	0.00	0.01	0.86	0.05
Industry Variables:				
Agricultural Industry	0.00	0.01	0.82	0.05
Mining Industry	0.00	0.06	1.00	0.17
Construction Industry	0.00	0.22	1.00	0.32
Manufacturing Industry	0.00	0.07	1.00	0.18
Transportation Industry	0.00	0.15	1.00	0.25
Finance Industry	0.00	0.02	1.00	0.09
Business Service Industry	0.00	0.04	1.00	0.14
Personal Service Industry	0.00	0.05	1.00	0.16
Entertainment Service Industry	0.00	0.01	0.73	0.05
Professional Service Industry	0.00	0.15	1.00	0.26
Employment and welfare variables:				
Employment	0.00	0.72	1.00	0.28
Sum of AFDC transfers during childhood ($)	0.00	309701	3316162	55879
Location Variables:				
Northeast Residence	0.00	0.11	1.00	0.29
North Central Residence	0.00	0.24	1.00	0.41
South Residence	0.00	0.53	1.00	0.47
West Residence	0.00	0.12	1.00	0.30

Table 4.3.: Baseline IGE and rank-rank regressions

	Dependent variable:	
	Log Family Income (adulthood)	Adulthood Rank
	(1)	(2)
Log Family Income (childhood)	0.553***	
	(0.040)	
Family Rank (childhood)		0.526***
		(0.036)
Constant	4.766***	0.196***
	(0.422)	(0.013)
Observations	984	984
R^2	0.162	0.182
Adjusted R^2	0.161	0.181
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 4.4.: Duration and age baseline regression

	Dependent variable:			
	Family Income ($)			
	(1)	(2)	(3)	(4)
Duration	−2,351.426***	−2,127.944***	−1,099.100***	−902.837**
	(229.586)	(288.553)	(328.760)	(374.375)
Age at first exposure	−132.492	130.431	−343.796	−118.859
	(258.287)	(330.159)	(256.885)	(328.841)
Duration×Age		−131.486		−111.802
		(102.897)		(102.056)
Constant	64,168.240***	63,673.270***	239,288.900***	243,587.900***
	(2,411.867)	(2,442.007)	(44,713.500)	(44,880.650)
Marginal effect of duration	−2,351.426***	−2390.915***	−1,099.100***	−1126.44**
	(229.586)	(299.37)	(328.760)	(381.118)
Observations	984	984	984	984
Controls	NO	NO	YES	YES
R^2	0.118	0.120	0.243	0.244
Adjusted R^2	0.116	0.117	0.209	0.210

Note: *p<0.1; **p<0.05; ***p<0.01.
Marginal effects calculated at median age; SE's calculated analytically.

Table 4.5.: Parametric linear in parameters income rank models

	Dependent variable:		
	Adulthood income rank		
	(1)	(2)	(3)
Age	0.0001	0.002	0.0009
	(0.002)	(0.003)	(0.005)
Duration	-0.021^{***}	-0.02^{***}	-0.035^{***}
	(0.003)	(0.004)	(0.007)
Intensity	0.204^{*}	0.151	0.321
	(0.092)	(0.282)	(0.223)
Age×Duration		-0.0005	
		(0.001)	
Age×Intensity		-0.03	
		(0.055)	
Duration×Intensity		0.002	
		(0.016)	
Age×Duration×Intensity		0.003	
		(0.005)	
Age^2			-0.00008
			(0.0003)
Duration^2			0.0009^{**}
			(0.0004)
Intensity^2			-0.31
			(0.3)
Constant	0.458^{***}	0.459^{***}	0.360^{***}
	(0.015)	(0.019)	(0.006)
Marginal effect of Duration	-	-0.0207^{***}	-0.0115^{***}
		(0.003)	(0.005)
Marginal effect of Intensity	-	0.159	-0.0495
		(0.178)	(0.19)
Marginal effect of Age	-	-0.0009	-0.0034
		(0.0025)	(0.004)
Observations	984	984	984
R^2	0.149	0.151	0.154
Adjusted R^2	0.147	0.144	0.148

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Marginal effects calculated at median age; SE's calculated analytically, except age.

Table 4.6.: Parametric linear in parameters income rank models, with controls

	Dependent variable:		
	Adulthood income rank		
	(1)	(2)	(3)
Duration	-0.01^{***}	-0.0094^{*}	-0.015^{**}
	(0.003)	(0.0.004)	(0.0074)
Age	-0.0012	0.0004	-0.0041
	(0.0016)	(0.003)	(0.0052)
Intensity	0.082	-0.099	-0.0067
	(0.096)	(0.29)	(0.23)
Duration^2			0.0005
			(0.0004)
Age^2			0.00017
			(0.0003)
Intensity^2			0.117
			(0.305)
Duration$\times$Age		-0.0005	
		(0.001)	
Duration$\times$Intensity		0.009	
		(0.016)	
Age$\times$Intensity		-0.036	
		(0.055)	
Duration$\times$Age$\times$Intensity		0.0045	
		(0.0051)	
Constant	0.911^{***}	0.94^{***}	0.968^{***}
	(0.277)	(0.279)	(0.279)
Marginal effect of Duration	-	-0.009^{***}	-0.028^{***}
		(0.0034)	(0.005)
Marginal effect of Intensity	-	-0.062	0.277
		(0.18)	(0.19)
Marginal effect of Age	-	-0.003	0.0006
		(0.0034)	(0.004)
Observations	984	984	984
Controls	YES	YES	YES
R^2	0.257	0.259	0.259
Adjusted R^2	0.223	0.222	0.223

Note: Controls include demographic, occupation, industry, employment, welfare, and locational characteristics, from Table 4.2.
Marginal effects calculated at median age; SE's calculated analytically, except age.
*p<0.1; **p<0.05; ***p<0.01

Table 4.7.: Parametric models with concentration

	Dependent variable:		
	Adulthood Rank		
	(1)	(2)	(3)
Years	-0.021^{***}	-0.019^{**}	-0.010^{***}
	(0.003)	(0.008)	(0.003)
Age	0.00002	0.002	-0.002
	(0.002)	(0.004)	(0.002)
Intensity	0.200^{**}	0.353	0.084
	(0.092)	(0.439)	(0.095)
Concentration	-0.012	-0.027	-0.009
	(0.009)	(0.045)	(0.009)
Years×Age		0.001	
		(0.003)	
Years×Intensity		0.013	
		(0.046)	
Age×intensity.conditional		-0.146^{*}	
		(0.079)	
Years×Concentration		0.003	
		(0.021)	
Age×Concentration		0.003	
		(0.016)	
Intensity×Concentration		0.353	
		(0.691)	
Years×Age×Intensity		0.019	
		(0.023)	
Years×Age×Concentration		-0.003	
		(0.008)	
Years×Intensity×Concentration		-0.110	
		(0.157)	
Age×Intensity×Concentration		0.269	
		(0.199)	
Years×Intensity×Age×Concentration		-0.046	
		(0.075)	
Controls	NO	NO	YES
F test p. val., vs. no concentration	-	.999	.999
Observations	984	984	984
R^2	0.151	0.164	0.256
Adjusted R^2	0.147	0.151	0.224
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 4.8.: Parametric models with controls, weighted by CBGPS

	Dependent variable:		
	Adulthood income rank		
	(1)	(2)	(3)
Duration	-0.022^{***}	-0.014^{***}	-0.026^{***}
	(0.003)	(0.005)	(0.008)
Age	-0.004^{***}	0.001	-0.020^{***}
	(0.002)	(0.003)	(0.005)
Intensity	0.300^{**}	0.460	0.577^{**}
	(0.121)	(0.314)	(0.260)
Duration^2			0.0002
			(0.0005)
Age^2			0.001^{***}
			(0.0003)
Intensity^2			-0.548
			(0.433)
Duration$\times$Age		-0.003^{**}	
		(0.001)	
Duration$\times$Intensity		-0.025	
		(0.019)	
Age$\times$Intensity		0.019	
		(0.059)	
Duration$\times$Age$\times$Intensity		0.003	
		(0.006)	
Constant	1.006^{***}	1.023^{***}	1.095^{***}
	(0.311)	(0.311)	(0.311)
Marginal effect of Duration	-	-0.021^{***}	-0.025^{***}
		(0.0036)	(0.005)
Marginal effect of Intensity	-	0.354^{*}	0.498^{**}
		(0.2)	(0.21)
Marginal effect of Age	-	-0.008^{***}	-0.016^{***}
		(0.0025)	(0.004)
Observations	984	984	984
Controls	YES	YES	YES
R^2	0.318	0.324	0.328
Adjusted R^2	0.287	0.290	0.295

Note: *p<0.1; **p<0.05; ***p<0.01
Marg. effects evaluated at medians; SE's computed analytically, except age in Model 2, which is bootstrapped

CHAPTER 5. SOCIOECONOMIC HETEROGENEITY OF RACIAL INCOME AND MOBILITY GAPS IN THE UNITED STATES

5.1 Introduction

Different racial and ethnic groups in the United States rank differently in the overall income distribution and these differences have persisted over time. In 1967, the median real household income of whites was 44,000 U.S. dollars, compared to only 24,000 for blacks and 34,000 for Hispanics. Half a century later, in 2014, white median household income had increased to 71,000 dollars whereas that of blacks and Hispanics had barely reached the 1967 median income level of whites. While the white-Hispanic gap widened during the past half-century as the black–white gap shrunk, racial and ethnic income gaps have persisted for decades. The picture looks equally bleak in terms of poverty: blacks and Hispanics are more than twice as likely to be poor as whites. This number has fallen from a high of four times as high for blacks in the early seventies, and although it fluctuated for Hispanics, their poverty prevalence is now about the same as for blacks, unchanged from its 1970's level.[1]

Black and Hispanic children born two generations ago found themselves in a disadvantaged position, because their parents were comparatively poor. For racial gaps to persist today, these groups must have experienced a level of intergenerational economic mobility that was not sufficient to close the original gaps in intervening decades (Hertz, 2005, 2008; Bhattacharya and Mazumder, 2011). Scholars have proposed distinct, but not necessarily exclusive, economic explanations for these observed

[1] See Pew Research Center, June 27, 2016, "On Views of Race and Inequality, Blacks and Whites Are Worlds Apart". Household income is standardized to a household size of three and is reported in constant 2014 prices. Race and ethnicity are determined by the race and ethnicity of the head of the household. Whites and blacks include only those who reported a single race. Data from 1970 to 2014 include only non-Hispanic whites and blacks; data prior to 1970 include Hispanics. Data source: Integrated Public Use Microdata Series (IPUMS),1968–2015 Current Population Survey Annual Social and Economic Supplement.

differences in the extent to which individuals are able to improve on their parents' economic status. These explanations emphasize externalities arising from social interactions within and between different groups, structural group disparities in socio-economic position and access to economic opportunities, and frictions arising from place-based factors tied to specific geographical locations where particular groups are clustered.[2] Along with an individual's innate ability and talents, these externalities provide a detailed account, rooted in economic theory, of how economic mobility may differ among individuals and groups as well as across space.

Recent work using field experiments and high-quality observational data provides valuable insights into the importance of childhood context and the degree of intergenerational mobility. In a series of recent papers, Raj Chetty and collaborators show that there is evidently spatial heterogeneity in economic mobility patterns, and these spatially differentiated mobility differences depend on a family's starting point in terms of their position on the income ladder during childhood.[3] The extent to which the interplay of individual/family and social, economic and place-based factors contributes to these mobility differences remains unclear, however, both in general and with respect to the observed race gaps. For instance, do strong social ties and segregation play out differently in terms of racial mobility gaps in economically advantageous as compared to economically deprived areas? Even more recently, Chetty et al. (2018) have found that blacks are moving downward in the income distribution across generations in contrast to whites, in part because black children who grow up in high-income families are often downwardly mobile.

In this paper, we focus on the effects of social, economic and place-based externalities for mobility variations across individuals and families, and we characterize the extent to which these individual/family and externalities are substitutes or complements. Instead of providing a general assessment, we tailor our investi-

[2] The foundational papers are Borjas (1992) for social interactions, Loury (1977) and Wilson (1987) for economic structure, and Kain (1968) for frictions related to the geographical location of people and jobs.

[3] See, for instance, Chetty et al. (2014a), Chetty and Hendren (2015) and Chetty et al. (2016).

gation around the observed black-white mobility gap. The econometric models we suggest allow for individual, household and place-based factors to influence economic mobility, and accounts for social, economic and place-based externalities. This approach allows us to test well-known hypotheses, such as the claim that exposure to highly educated members of one's own racial group during childhood enhances mobility (Borjas, 1992), but it also allows us to consider the possibility that the strength of this relationship varies over parental education or the occupational structure of the local economy. Furthermore, this framework lets us consider a variety of person–place counterfactuals, such as a comparison of how black children in black neighborhoods fare compared to white children in black neighborhoods, or whether rurality affects the mobility of black children in the same way that it affects the mobility of white children.

5.2 Communities, Interactions, and Economic Opportunity

As we discuss in detail in Chapter 2, the economic opportunities an individual faces throughout his or her life depend in part on the individual's family background and socioeconomic environment. An individual who was raised in a family with limited resources, received an education in a community that lacked access to education-related welfare programs such as free school meals, lived in a region that was subject to economic deprivation, or grew up in a racially segregated neighborhood, typically enjoys less economic success as an adult than an individual who was raised under relatively better conditions. Childhood exposure to these types of disadvantage negatively impacts the individual's social and educational development, thereby leading (on average) to relatively less economic success in adulthood. These facts lead to the longstanding empirical questions: to what extent do these family and socioeconomic factors that characterize an individual's childhood influence the economic opportunity experienced by the individual later in life (in adulthood), and to what extent do these factors have interactive – that is, complementary – effects?

We expect that family and community characteristics have interactive effects on individual economic well-being, and that these interactions imply either substitutability or complementarity of the effects of these factors on the individual. If a relatively high unemployment rate is associated with poor economic opportunities, and a relatively high rate of single-mother households is associated with poverty and an absence of positive role-models, ceteris paribus, what is the effect of growing up in an environment characterized by high rates of both? Or, if ethnic capital corresponds to the availability of positive role-models, what is the combined effect of living in an environment characterized by a high unemployment rate (a negative effect) but yet with a high rate of ethnic capital (a positive effect)? My anticipation of significant, interactive effects leads me to the hypothesis that the family and socioeconomic forces we describe have substitutable and/or complementary effects on individual economic well-being. There is relatively little empirical work on how these mechanisms interact to determine the degree of mobility, Fortunately, theoretical models of intergenerational mobility, as well as some segments of the empirical literature on mobility, neighborhood effects, and segregation, provide a fairly clear framework that lets us approach these models from a conceptually coherent perspective.

As we note above, if some attribute of a household or community inhibits parents' ability to invest or reduces the returns on that investment, in terms of children's long-term earnings, a child will be less upwardly mobile.[4] Racial discrimination in labor markets or in the provision of human capital inputs, such as school quality or positive peer influences, could clearly do both.This suggests that adverse circumstances, such as childhood exposure to high unemployment areas, or geographic remoteness, would decrease mobility more for black Americans than other Americans, and that positive circumstances, such as living in a highly educated community, would increase their mobility by less. It is important, however, to be clear that, if we do find results consistent with these predictions, we cannot interpret them as evidence that the mechanism is discrimination per se. We may be able to say, for example,

[4]Note that this is distinct from the standard intergenerational elasticity measure of mobility, which reduced investment would in fact decrease, which is generally interpreted as increasing mobility.

that black Americans' mobility suffers more from local unemployment, and we may be able to put that in a broader context using other concurrent effects, but our data and empirical strategy do not let us attribute these relationships to discrimination directly, or to any other specific behavior in the labor market or in service provision.

While there is limited evidence regarding interactions between household and community factors in determining mobility itself, there is some empirical evidence regarding the interaction between residential segregation, economic structure, and individuals' locations, corresponding to the spatial versus structural explanations discussed in detail in Chapter 3. Evidence of 'spatial mismatch' - the idea that the location of minority populations relative to job opportunities in cities makes it harder for them to get and keep jobs - is perhaps the best example of this nexus, as it connects racial disparities to interlocking economic frictions. In brief, scholars working in this literature have found that black workers do not fully adjust to jobs moving by relocating themselves, so that jobs moving implies substantial disemployment (Martin, 2001, 2004; Weinberg, 2000), and that, conversely, jobs returning to inner cities, which generally have disproportionately large black populations, reduces the black unemployment differential (Weinberg, 2004). Hellerstein et al. (2008) departs slightly from these results, finding that urban disparities between black and white employment are partially determined by the prevalence of jobs which hire black workers in areas where they live, while white job density has no effect on local levels of white employment, indicating the presence of frictions which affect black workers specifically. These results are not, of course, as direct a source of evidence on the community interactions we have proposed as we would like, but it is a case in which economic structure and individuals' locations in space combine to make black Americans uniquely vulnerable to shocks, which is likely to inhibit mobility, based on our interpretation of the underlying theory. As we develop our models in subsequent sections, we do not rely heavily on the spatial mismatch literature, and we are not studying the effects of spatial mismatch on mobility, but the way it integrates person and place-based factors do provide valuable conceptual empirical support for the questions we do ask.

5.3 Data

5.3.1 PSID Sample

This analysis requires a sample of individuals who are observed as both children and as adults, so that we can measure both the individual's adult income rank as well as the income rank of his/her parents when he/she was a child. In constructing our PSID sample, we require that the individual is observed over the age range of 25 to 55, and that the individual can be linked to his/her parents for at least five years over the time at which the individual was aged between 5 to 18 years. Of this subsample of individuals, we further require the individual's parents to be aged between 25 and 55 years over the period in which the individual was a child. The reason we measure adult income rank (for the individual and his/her parents) over the 25-55 age range is to capture adulthood earnings, and to avoid measuring earnings prior to adulthood or in years of retirement as these periods of time are not representative of the individual's lifetime economic status. Any differences across individuals in age-related earnings profile that occurs within the 25-55 range is controlled for in our regressions via the inclusion of the individual's and his/her parents average age over the years that income rank is calculated. The five years of observation requirement is to minimize the likelihood that measurement error or random income fluctuations influence our income rank measurements (Mazumder, 2005; Haider and Solon, 2006). Finally, within these years of observation, income (and other variables) are averaged over time for the individual. Thus, we obtain a cross-sectional sample of individuals that are linked to their parents, for which we measure the individual's average income in both childhood (i.e., parental average income) and adulthood.

It is important to note that our income measurement is total family income. Conceptually, an individual's economic status is determined by total family income, and PSID reports both individual labor and total family income for all survey respondents.[5] These averaged family income measurements are converted into income rank measurements using the Integrated Public Use Micro-Sample (IPUMS) database, through which we construct a national income distribution for every census year since 1970. For consistency with our PSID measure, we use family total income in nominal prices, and since the IPUMS income measurements represent a sample of the national income distribution, we construct PSID percentile values from the empirical cumulative distribution function of the IPUMS income. We average the years over which the individual is included in our PSID sample to determine the closest census year for which to generate the income rank measurement; for example, if the individual is in the PSID sample as a child over the years 1968-1972, we calculate income rank based on the national income distribution from the year 1970.

Other individual and family background variables are derived similarly from PSID. All variables are averaged across the sampled years; time-varying variables are measured as the average, and time-invariant variables (e.g., the black indicator variable) remains a binary indicator. In total, we are left with a cross-sectional sample of 5,248 individuals.

5.3.2 Sample Generation and Descriptive Statistics

The PSID and harmonized census samples are merged at the census tract and county level to complete our spatially explicit, intergenerationally linked cross-section of individuals. Specifically, an individual's geographic location is defined as the location where the individual lived at the time he/she was a child; and for any individual who moved geographically over that time period, we select the geographic region where he/she lived the longest. This final sample contains 5,248 individuals

[5]PSID has different measurements income, including labor income of a family head, total family income, and total family wealth. The labor income of a family head measures a head's labor part of farm income and business income, wages, bonuses, overtime, commissions, professional practice, labor part of income from roomers and boarders or business income. The total family income is the sum of taxable income and total transfers of all the family members. Total family wealth measures the value of a family assets net of debts.

that reside in 2,000 census tracts that span 525 counties, and allows us to estimate the spatially heterogeneous versions of our empirical parameters and test our hypotheses about individual economic status.

Table 5.1 provides descriptive statistics for our sample. Of the 5,248 individuals, 2,178 individuals are black and 3,061 individuals are non-black. The average family total income rank in our entire sample is 0.44, i.e. the average individual ranks at the 44th percentile of national income. This rank is slightly lower than the average parental income rank, 0.53. Further, there is a noticeable race difference: for both individual and parental income rank income, blacks rank substantially lower than non-blacks. For individuals, blacks rank on average at the 32^{nd} percentile whereas non-blacks rank on average at the 53rd percentile.

We see racial differences with respect to the other control variables as well. On average, 46 percent of the individuals in our sample are male, the individual has just over 3 siblings, and 26 percent of the individuals grew up in mother-headed households. Yet, it is clear that the number of siblings and the incidences of growing up in mother-headed households are much higher for blacks than for non-blacks.[6] In addition, the standard deviation for these variables is larger for blacks than for non-blacks.

[6]We code mother-headed household to be one as long as an individual was ever lived in a family during childhood that was mother headed.

Table 5.1.: Descriptive statistics

	Total		Black		Non-black	
Variables	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Income measurements						
Individual income rank	0.44	0.25	0.32	0.21	0.53	0.24
Parental income rank	0.53	0.27	0.37	0.23	0.65	0.23
Individual/family characteristics						
Male	0.46	0.50	0.43	0.49	0.49	0.50
Number of siblings	3.15	1.60	3.70	1.80	2.76	1.30
Mother-headed household	0.26	0.44	0.47	0.50	0.11	0.32
Census tract and county						
Tract black population	0.29	0.35	0.64	0.27	0.05	0.12
Tract single mother	0.09	0.06	0.14	0.07	0.05	0.03
Tract ethnic capital	0.12	0.10	0.07	0.07	0.15	0.11
County unemployment rate	0.19	0.12	0.21	0.12	0.18	0.12
County poverty rate	0.13	0.07	0.17	0.07	0.11	0.05
County index of rurality	0.48	0.10	0.46	0.10	0.50	0.09
County rural-urban continuum codes	2.28	1.88	1.96	1.58	2.51	2.03
Control variables						
Age	31.45	4.36	31.46	4.52	31.44	4.24
Parental age	39.82	6.24	39.21	6.80	40.25	5.77
N	5,248		2,187		3,061	

All parental characteristics and the census tract and county variables are averaged over the years in which the individual was a child. The index of relative rurality is scaled to be between 0 (most urban) and 1 (most rural).

On average, the individuals in our sample live in census tracts in which 29 percent of the population are black, 9 percent are from single mother households, and 12 percent of the individual's racial group are college educated (ethnic capital).

The average unemployment rate is 19 percent, with a 13 percent poverty rate, and finally 48 percent of the counties are rural. It is also clear that blacks, on average, live in substantially worse socioeconomic environments: higher rates of single-mother households, lower ethnic capital, higher unemployment rates, and higher poverty rates. Blacks also tend to live in tracts with a much larger share of the population being black.

Finally, we see that the average age of the individual in our sample is 31 years, and average parental age is about 40 years. These averages do not differ by race. Additionally, we report bivariate correlations for these variables in Table 5.9 in Appendix **??**. The highest correlations occur between black, the percent of the population in the census tract that are black, and percent of households in census tract that are led by a single woman.

5.4 Methodology

5.4.1 Hypothesis and Key Parameters

Our story describes the importance of the individual's family background (including both parental income rank and race), socioeconomic environment, and geographical proximity to resources and amenities for determining the economic status of the individual later in life. These factors may directly affect individual income rank, but they may also influence the intergenerational link in income rank and they may also exacerbate racial disparities in income rank. We have also emphasized the importance of considering the interaction between these factors in order to understand realized economic outcomes across a sample of individuals. Thus, our empirical goal is to understand the nature of these interactions as they drive the baseline associations shown in Table 5.3. To formalize this goal, we posit the following empirical hypothesis.

Hypothesis *An individual's family background, socioeconomic environment, and geographical location have mutually reinforcing effects on individual economic oppor-*

tunity in general, and specifically through differences in intergenerational and racial income rank.

Embedded in our hypothesis is the notion that these three factors (the individual's family background, socioeconomic environment, and geographical location) affect individual economic status both directly and indirectly through both intergenerational and racial channels.

Our empirical hypotheses are defined in terms of three empirical parameters: the *intergenerational rank* parameter, the *race gap in mobility* parameter, and the *race gap in intergenerational rank* parameter. The *intergenerational rank* parameter defines the correlation between the income rank of an individual and the income rank of his/her parents at the time that he/she was a child (holding constant race); the *race gap in mobility* parameter defines the black/non-black income rank disparity (holding constant parental income rank); and the *race gap in intergenerational rank* parameter captures the extent to which there are racial disparities in the intergenerational link in income rank. Assuming an average parameter, from the generic model these parameters are:

$$\begin{aligned} \text{IG Rank Parameter} &= \mathbb{E}\left[\frac{\partial rank}{\partial rank^{\pi}}\right] \\ \text{Race Gap in Mobility} &= \mathbb{E}[rank|black=1] - \mathbb{E}[rank|black=0] \\ \text{Race Gap in IG Rank Param.} &= \mathbb{E}\left[\frac{\partial rank}{\partial rank^{\pi}}\middle| black=1\right] - \mathbb{E}\left[\frac{\partial rank}{\partial rank^{\pi}}\middle| black=0\right]. \end{aligned} \tag{5.1}$$

We expect that these factors have interactive effects on income rank, both directly and through these parameters, as the effect of one unfavorable characteristic is exacerbated through the presence of another. At the same time, we recognize that many of these factors are correlated with one another in observational data, and it is not clear which of these different factors bears the strongest association with individual income rank. Consequently, in testing this hypothesis we test three sub-hypothesis: whether and to what extent each of these factors affect individual income rank (i) directly, holding the other factors constant; (ii) indirectly, through intergenerational and racial differences in income; and (iii) whether these factors have mutually reinforcing effects on individual income rank, either directly or indirectly.

5.4.2 Hierarchical Linear Model

To test our empirical hypothesis, we develop a hierarchical linear model in which individuals are linked intergenerationally with their parents and are nested spatially within census tracts and counties. The outcome variable is the individual's income rank spanning adulthood, and all spatial factors correspond to the time that the individual was a child (not the characteristics of the individual's location at the time of adulthood). This model allows us to capture the varied sources of heterogeneity discussed previously by letting us analyze an individual's family background and spatial environment concurrently. We measure social interactions at the census tract level because a census tract is a sufficiently small geographic area (of about 4,000 inhabitants) to capture the individual's neighborhood, and therefore is more likely to reflect the social influences that he or she receives during childhood. At the county level, we shift our emphasis to the variables that reflect economic structure and geographic remoteness, because counties implement local public policies that lead to differences in local industry structure and labor markets, and because the degree of geographic remoteness associated with infrastructure, amenities, and public services is likely to differ more meaningfully among counties than among tracts.

The Baseline Specification In fitting these models, we start from the baseline model that is a linear version of the generic model, Equation **??**:

$$\begin{aligned} rank &= \alpha_0 + \alpha_1 rank^{\pi} + \alpha_2 black + \alpha_3 (rank^{\pi} \times black) \\ &\quad + \gamma_1 age + \gamma_2 \text{age}^2 + \lambda_1 age^{\pi} + \lambda_2 (age^{\pi})^2 + X + \varepsilon. \end{aligned} \tag{5.2}$$

We augment the model to include the interaction between $rank^{\pi}$ and $black$, and control for both the individual's average age and the average age of the individual's parents over the years that income rank is measured to account for life cycle patterns in earning capacity.[7] We are primarily interested in estimating $\alpha = (\alpha_0, \alpha_1, \alpha_2, \alpha_3)'$, and treat $\gamma = (\gamma_1, \gamma_2)'$ and $\lambda = (\lambda_1, \lambda_2)'$ as nuisance.

Adding Individual and Family Background We introduce heterogeneity in terms of individual and family background characteristics by allowing α to be functions of these family and individual factors. That is:

$$\alpha_s = \beta_s + \sum_m \beta_{sm} x_m, \quad s = 0, 1, 2, 3 \tag{5.3}$$

where x_m is a vector of individual and family characteristics.

Adding Equation 5.3 to Equation 5.2 generates a set of interactions between each of the main variables in the model ($rank^{\pi}$, $black$, and $rank^{\pi} \times black$) and the variables in x_m. These interactions imply that the marginal effects of $rank^{\pi}$ and $black$ on individual income rank now depend on the characteristics of the individual and his/her family in addition to each other. Consider, for instance, the effect of a unit change in parental income rank: β_1 captures the average direct effect of parental income rank on individual income rank, while β_{1m} adjusts that average effect according to the individual's characteristics. This marginal effect further varies according to race via $rank^{\pi} \times black$, which itself varies according to x_m via α_3. These different

[7] Traditionally, age is accounted for via a quartic polynomial specification (Lee and Solon, 2009). In our models, the third and fourth order polynomial terms are not significant (via both t-tests and F-tests), and so our final specification is quadratic in age.

interactions allow us to measure the extent to which the baseline parameter estimates vary across to individual and family background.

Adding Socioeconomic Factors and Geographic Location We add the socioeconomic factors and geographical location in much the same way as the individual and family background variables were added, by allowing the intercept and slope parameters to vary with respect to these factors. The individual and family characteristics were added first because there is less cross-sectional variation in the socioeconomic and geographic factors; the result is that the parameters of the individual and family factors are now specified to vary according to the socioeconomic and geographic factors. We specify the coefficients in Equation (5.3) as:

$$\begin{aligned}\beta_s &= \bar{\bar{\beta}}_s + \sum_k \bar{\beta}_{sk}\bar{o}_k + \sum_j \bar{\bar{\beta}}_{sj}\bar{\bar{o}}_j, \quad s = 0,1,2,3 \\ \beta_{sm} &= \bar{\bar{\beta}}_{sm} + \sum_k \bar{\beta}_{smk}\bar{o}_k + \sum_j \bar{\bar{\beta}}_{smj}\bar{\bar{o}}_j, \quad s = 0,1,2,3\end{aligned} \tag{5.4}$$

where $\bar{o}_k$ represents the k^{th} census tract factor, and $\bar{\bar{o}}_j$ denotes j^{th} county factor (that includes geographical remoteness); recall that both are defined according to the individual's childhood location (not adulthood). The coefficients of these spatial variables now capture the extent to which individual income rank and the associated channels we have already described depend on the socioeconomic environment that the individual was raised in as well as the geographic location. More specifically, this setup allows us to test whether the individual and family characteristics and socioeconomic and geographic factors have mutually reinforcing effects on individual income rank.

Full Model Specification and Estimation The complete empirical specification is defined by plugging Equations (5.3) and (5.4) into Equation (5.2):

$$
\begin{aligned}
rank \;=\;& \bar{\bar{\beta}}_0 + \bar{\bar{\beta}}_1 rank^{\pi} + \bar{\bar{\beta}}_2 black + \bar{\bar{\beta}}_3 (rank^{\pi} \times black) \\
& + \sum_m \bar{\bar{\beta}}_{0m} x_m + \sum_k \bar{\bar{\beta}}_{0k} \bar{o}_k + \sum_j \bar{\bar{\beta}}_{0j} \bar{\bar{o}}_j + \sum_m \sum_k \bar{\bar{\beta}}_{0mk} x_m \bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{0mj} x_m \bar{\bar{o}}_j \\
& + \left(\sum_m \bar{\bar{\beta}}_{1m} x_m + \sum_k \bar{\bar{\beta}}_{1k} \bar{o}_k + \sum_j \bar{\bar{\beta}}_{1j} \bar{\bar{o}}_j + \sum_m \sum_k \bar{\bar{\beta}}_{1mk} x_m \bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{1mj} x_m \bar{\bar{o}}_j \right) rank^{\pi} \\
& + \left(\sum_m \bar{\bar{\beta}}_{2m} x_m + \sum_k \bar{\bar{\beta}}_{2k} \bar{o}_k + \sum_j \bar{\bar{\beta}}_{2j} \bar{\bar{o}}_j + \sum_m \sum_k \bar{\bar{\beta}}_{2mk} x_m \bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{2mj} x_m \bar{\bar{o}}_j \right) black \\
& + \left(\sum_m \bar{\bar{\beta}}_{3m} x_m + \sum_k \bar{\bar{\beta}}_{3k} \bar{o}_k + \sum_j \bar{\bar{\beta}}_{3j} \bar{\bar{o}}_j + \sum_m \sum_k \bar{\bar{\beta}}_{3mk} x_m \bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{3mj} x_m \bar{\bar{o}}_j \right) \\
& \times (rank^{\pi} \times black) + \gamma_1 age + \gamma_2 \mathrm{age}^2 + \lambda_1 age^{\pi} + \lambda_2 (age^{\pi})^2 + \varepsilon \qquad (5.5)
\end{aligned}
$$

in which the α parameters in the baseline equation are now functions of the individual and family background factors, the spatial (socioeconomic and geographic) factors, and the interaction between the two sets of parameters. The model remains conditional on the age of the individual and his/her parents. The model includes 115 parameters in total. I fit the models using ordinary least squares with heteroskedasticity robust standard errors to account for clustering in the errors across households or space.[8] Equation 5.5 provides the parameters necessary to fully define our three parameters of interest, each of which is, in turn, a function of an observation's data:

[8] We also estimate a maximum likelihood generalized least squares version of the model with heterogeneous error components across the individual, tract, and county levels, but find that the higher levels of error components are not statistically significant. Furthermore, these auxiliary estimates are nearly identical to our robust ordinary least squares estimates, and so we focus only on the latter.

$$
\begin{aligned}
\text{IG Rank Parameter} &= \frac{\partial rank}{\partial rank^{\pi}} \\
&= \bar{\bar{\beta}}_1 + \bar{\bar{\beta}}_3 black \\
&\quad + \left(\sum_m \bar{\bar{\beta}}_{1m} x_m + \sum_k \bar{\bar{\beta}}_{1k} \bar{o}_k + \sum_j \bar{\bar{\beta}}_{1j} \bar{\bar{o}}_j + \sum_m \sum_k \bar{\bar{\beta}}_{1mk} x_m \bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{1mj} x_m \bar{\bar{o}}_j \right) \\
&\quad + \left(\sum_m \bar{\bar{\beta}}_{3m} x_m + \sum_k \bar{\bar{\beta}}_{3k} \bar{o}_k + \sum_j \bar{\bar{\beta}}_{3j} \bar{\bar{o}}_j + \sum_m \sum_k \bar{\bar{\beta}}_{3mk} x_m \bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{3mj} x_m \bar{\bar{o}}_j \right) black \\
\\
\text{Race Gap in Mobility} &= [rank | black = 1, x_m, \bar{o}_k, \bar{\bar{o}}_j] - [rank | black = 0, x_m, \bar{o}_k, \bar{\bar{o}}_j] \\
&= \bar{\bar{\beta}}_2 + \bar{\bar{\beta}}_3 rank^{\pi} \\
&\quad + \left(\sum_m \bar{\bar{\beta}}_{2m} x_m + \sum_k \bar{\bar{\beta}}_{2k} \bar{o}_k + \sum_j \bar{\bar{\beta}}_{2j} \bar{\bar{o}}_j + \sum_m \sum_k \bar{\bar{\beta}}_{2mk} x_m \bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{2mj} x_m \bar{\bar{o}}_j \right) \\
&\quad + \left(\sum_m \bar{\bar{\beta}}_{3m} x_m + \sum_k \bar{\bar{\beta}}_{3k} \bar{o}_k + \sum_j \bar{\bar{\beta}}_{3j} \bar{\bar{o}}_j + \sum_m \sum_k \bar{\bar{\beta}}_{3mk} x_m \bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{3mj} x_m \bar{\bar{o}}_j \right) rank^{\pi} \\
\\
\text{Race Gap in IG Rank Parameter} &= \left[\frac{\partial rank}{\partial rank^{\pi}} \middle| black = 1, x_m, \bar{o}_k, \bar{\bar{o}}_j \right] - \left[\frac{\partial rank}{\partial rank^{\pi}} \middle| black = 0, x_m, \bar{o}_k, \bar{\bar{o}}_j \right] \\
&= \bar{\bar{\beta}}_3 + \sum_m \bar{\bar{\beta}}_{3m} x_m + \sum_k \bar{\bar{\beta}}_{3k} \bar{o}_k + \sum_j \bar{\bar{\beta}}_{3j} \bar{\bar{o}}_j + \sum_m \sum_k \bar{\bar{\beta}}_{3mk} x_m \bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{3mj} x_m \bar{\bar{o}}_j . \qquad (5.6)
\end{aligned}
$$

Variables and Definitions at Each Tier We choose specific variables to capture the individual, family, and spatial characteristics according to the hypotheses described in our literature review on intergenerational mobility and neighborhood effects. Table 5.2 lists the variables we include in each category, and and briefly summarizes the mechanism that each measures.

The individual/family characteristics include the individual's own gender, the number of siblings the individual has, and whether the individual grew up in a single-mother household.[9] The impacts of social influence are captured by the following three factors measured at the census tract level: the fraction of the population in the tract that is black, the fraction of households in the tract that are single-mother households, and the fraction of individuals in the tract of the same race as the individual that has a college degree or higher. In particular, this last variable measures the individual's ethnic capital within the tract, and is defined analogously to Borjas (1992, 1994) in that 'ethnic capital' refers to a positive social influence. To measure economic structural factors we select the unemployment rate and the poverty rate, and geographic remoteness is measured by a continuous index of relative rurality (Waldorf and Kim, 2015).[10] Both the economic structural factors and geographic remoteness are defined at the county level. All empirical measurements are defined in the following section.

[9]The PSID database always identifies the head of the household to be the male in a married-couple family. Only when a female is unmarried can she be listed as the head of the household. Thus, any mother-headed household can be taken to indicate that a single-mother household.

[10]The county rurality is based on four most frequently used dimensions of rurality: population size, population density, built-up area, and geographical remoteness.

Table 5.2.: List of variables

Definition	Label	Mechanism
Key Variables		
Individual family income rank	rank	Economic status in adulthood
Parental family income rank	rank^{π}	Monetary resources
Individual race dummy	Black	Racial disparities
Individual and Family Characteristics (x_m)		
Individual's gender	Male	Individual control
Number of siblings	Number of siblings	Household resource allocation
Mother-headed household	Mother headed hh	Family structure
Age	Age	Life-stage earning profile
Parental age	Age^{π}	Parental earning profile
Census Tract Factors ($\bar{o}$)		
Fraction black	Tract black	Segregation
Fraction of single-mother household	Tract single mother hh	Role model and resource allocation
Fraction college degree by race	Tract ethnic capital	Role model
County Factors ($\bar{\bar{o}}$)		
Unemployment rate	County employment	Job opportunities
Poverty rate	County poverty	Income levels
Geographic remoteness	County rurality	Access to public resources and amenities

5.5 Results

Our hierarchical specification allows us to compute individual-specific versions of our three empirical parameters – the *intergenerational rank* parameter, the *race gap in mobility* parameter, and the *race gap in intergenerational rank* parameter. In the following subsections, we present those results.

Table 5.3.: Baseline estimates of the link between parental income/race and individual income

	Dependent Variable: Individual Income Rank			
	Model 1	Model 2	Model 3	Model 4
Parental Rank	0.44***	-	0.32***	0.37***
	(0.01)	-	(0.01)	(0.02)
Black	-	−0.21***	-0.12***	−0.07***
	-	(0.01)	(0.01)	(0.01)
Parental Rank × Black	-	-	-	−0.11***
	-	-	-	(0.03)
Constant	0.21***	0.53***	0.32***	0.29***
	(0.01)	(0.004)	(0.01)	(0.01)
Adjusted R^2	0.22	0.18	0.27	0.27
N	5,248	5,248	5,248	5,248

Table 5.4.: Summary of the estimates of the IG rank parameter, the race gap in mobility parameter, and the race gap in IG rank parameter

	Full sample			Significant subsample			
	Mean	SD	Range	Mean	SD	Range	Number
IG Rank	0.26***	0.02	[0.02, 0.46]	0.31***	0.02	[0.17, 0.48]	4,066
Race Gap in Mobility	−0.09***	0.02	[−0.25, 0.08]	−0.15***	0.02	[−0.31, −0.06]	2,071
Race Gap in IG Rank	-0.05	0.09	[−0.64, 0.35]	-0.06	0.08	[−0.79, 0.58]	732

Significance: $^{***}p < .01,^{**} p < .05,^{*} p < .1$; significant subsample is group with statistically significant parameter estimates at the 10 percent leve.

5.5.1 Summary of Key Parameters

In Table 5.4, we report the average parameter estimate across the entire sample, and the $[0.05, 0.95]$ percentile range of values for each parameter.[11] On average, the intergenerational rank parameter is 0.26, the race gap in mobility parameter is -0.09, and the race gap in the intergenerational rank parameter is -0.05. The intergenerational rank parameter and the race gap in mobility parameter are statistically significant, and the race gap in the intergenerational rank parameter is not significant. Together, these estimates imply that a one percentile increase in parental income rank correlates with about a 0.26 percentile increase in individual income rank, and that blacks average approximately 9 percent lower in income rank than non-blacks. These estimates are qualitatively similar to the baseline estimates in Table 5.3, except we no longer find evidence that there is a significant racial difference in the intergenerational rank parameter, consistent with the Chetty et al. (2018) finding that all racial groups in the U.S. have similar rates of relative mobility. We expect that this change is driven by bias in the baseline models, which could easily misattribute socioeconomic heterogeneity to racial differences. As in Table 5.3, the expected rank for black children is lower than for white children; in Model 3, it is 12 percentage points lower based on the 'Black' intercept shift, and in Model 4 it is 13.5 percentage points lower, based on an evaluation of the linear combination of the 'Black' and '$ParentalRank \times Black$' coefficients, evaluated at a median rank of 0.5. Both of these values are comparable to the race gap in mobility values of 9 and 15 percentage points presented in Table 5.4, for the full and significant-only samples, respectively.

Table 5.4 also shows that there is substantial heterogeneity in these parameters across individuals. We focus on the subsample of observations for which the heterogeneous parameter estimates are significant, and find that within the significant subsample, $100 \times (4,066/5,248) = 77$ percent of observations have a significant

[11] We report the $[0.05, 0.95]$ percentile range of parameters to eliminate several extreme outlier values that appear for different parameters.

intergenerational rank parameter, $100 \times (2,071/5,248) = 39$ percent show a statistically significant race gap in mobility parameter, and only $100 \times (732/5,248) = 14$ percent show a significant racial gap in the intergenerational rank parameter.

Figures 5.1 provides more insight into the nature of the parameter heterogeneity. Figure 5.1 shows that nearly all of the negative intergenerational rank parameter estimates are insignificant, while most of the positive estimates are significant (the left panel of kernel densities). Furthermore, the significant estimates correspond to individuals with relatively higher parental income rank, while the insignificant intergenerational rank estimates correspond to a much greater extent to individuals with relatively lower parental income rank (the right hand panel of overlaid histograms). At the 20th percentile of parental income rank, the number of observations in the insignificant subsample is nearly identical to the number of observations in the significant subsample, while at much higher ranks of parental income the majority of observations have a significant parameter estimate. It is clear that the links in intergenerational income rank come primarily from higher income households (at the time of childhood), but much less to from low ranked families. We suspect that this pattern emerges because socioeconomic forces are strong enough to substantially weaken the link between the individual's income rank and the income rank of his/her parents for poor children, but not for rich ones.

The second panel of Figure 5.1 reveals that nearly all of the positive race gap in mobility parameter estimates are insignificant (the left hand panel of kernel densities) and that most of the statistical significance comes directly from black individuals. That is, $100 \times (1,441/2,071) = 70$ percent of the individuals with a significant race gap in mobility parameter are black (the right hand panel of bar charts). And, interestingly, the third panel of Figure 5.1 shows that the distribution of significant race gap in intergenerational rank parameter estimates is bimodal and centered at zero (the left hand panel of kernel densities), without much indication that the significance or insignificance is driven by blacks or non-blacks in the sample (the right hand panel of bar charts).

Throughout the rest of this chapter, we investigate whether there are specific individual, household, and place characteristics that drive the three parameters. We summarize the individual/family, socioeconomic, and geographic factors for the subpopulation for which the the three parameters are significant, and another subpopulation with insignificant parameters. As shown in Table 5.5, the two subpopulation differ significantly in most of the factors.

Table 5.5.: Descriptive statistics of variables by subsamples with significant and insignificant IG rank parameter, the race gap in mobility parameter, and the race gap in IG rank parameter

	IG Rank			Race Gap in Mobility			Race Gap in IG Rank		
Mean	[1]	[2]	t-test	[1]	[2]	t-test	[1]	[2]	t-test
Male	0.46	0.47	−0.33	0.36	0.53	−11.89***	0.24	0.50	−14.50***
No. of siblings	2.99	3.68	−11.19***	2.75	3.41	−16.02***	3.48	3.09	5.23***
Mother-headed household	0.19	0.51	−20.51***	0.17	0.31	−12.00***	0.13	0.28	−10.70***
Individual age	31.56	31.07	3.41***	31.81	31.22	4.80***	31.91	31.37	3.11***
Parent age	39.68	40.32	−3.01***	39.92	39.76	0.89	40.87	39.66	5.04***
Tract black population	0.23	0.53	−25.95***	0.20	0.36	−17.49***	0.24	0.30	−5.24***
Tract single mother	0.08	0.13	−20.73***	0.07	0.10	−15.29***	0.08	0.09	−4.64***
Tract ethnic capital	0.11	0.13	−2.64***	0.12	0.11	3.89***	0.12	0.12	1.80*
County unemployment rate	0.18	0.21	−4.83***	0.20	0.18	4.43***	0.19	0.19	0.36
County poverty rate	0.12	0.17	−17.64***	0.12	0.14	−12.05***	0.13	0.13	0.69
County rurality	0.49	0.46	6.43***	0.48	0.48	1.51	0.49	0.48	2.16**
N	4,066	1,182		2,071	3,177		732	4,516	

In each panel, column [1] denotes means of the variables in the sample with significant effects, and column [2] denotes those in the sample with insignificant effects.

5.5.2 The Importance of Individual and Socioeconomic Heterogeneity

In the preceding section, we describe the three estimated parameters and the nature of the heterogeneity within those parameters that is induced by the individual/family, socioeconomic, and geographic factors. In this section, we explore the relationshi between these factors and heterogeneity in income rank and the gap parameters. Specifically, we evaluate the partial derivative of the regression function and each of the three estimated parameters, with respect to each of the individual/family,

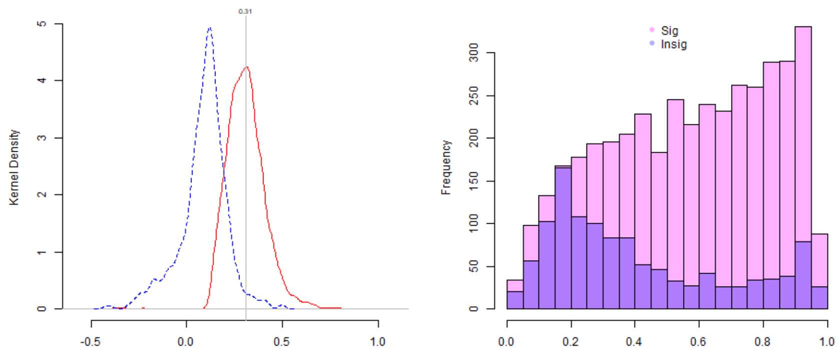

(a) L: IG rank parameter; R: Distribution of parental income rank (Blue:insignificant; Red: significant)

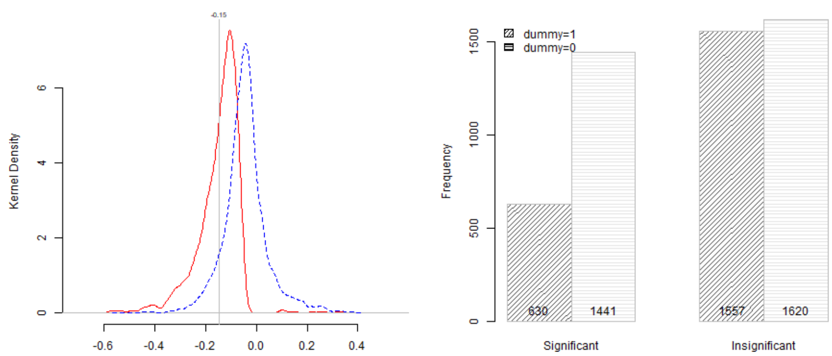

(b) L: Race gap in mobility parameter; R: Distribution of the black indicator (Blue:insignificant; Red: significant)

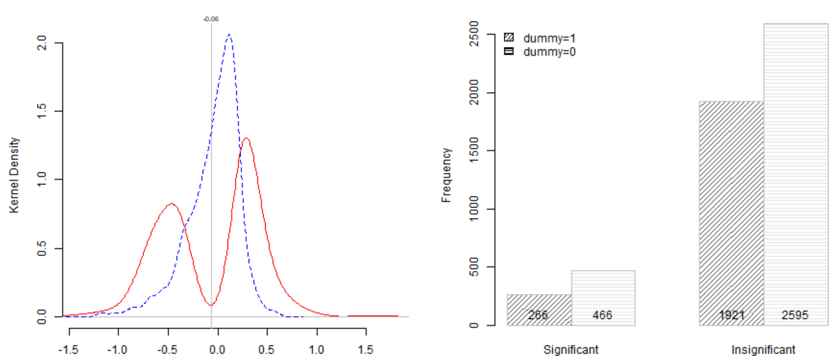

(c) L: Race gap in IG rank parameter; R: Distribution of the black indicator (Blue:insignificant; Red: significant)

Figure 5.1.: Summaries of the estimated intergenerational rank parameter, the race gap in mobility parameter, and the race gap in the intergenerational rank parameter

socioeconomic, and geographic variables in the model. We report the results in Table 5.6, and we report additional distributional summaries of these marginal effects in the appendix. As before, we summarize both the average marginal effect and the distribution of marginal effects both for the entire sample and specifically for the subset of observations for which th marginal effects are significant. Most of the distributions of the marginal effects across individuals span both negative and positive quadrants, and so the average marginal effect is often zero as a result, so we believe that the subsample of individuals for which the marginal effects are significant merits particular consideration.

Panel I in Table 5.6 summarizes the marginal effect of each variable on individual income rank, irrespective of any particular channel. It evaluates the hypothesis *(i) whether and to what extent each of these factors affects individual income rank directly, holding the other factors constant*. We find males have higher income rank than females in both the entire sample and the subsample of observations with a significant effect. We also find that individual income rank is negatively associated with a higher number of siblings, the share of households in the census tract that are single mother headed, and the county unemployment and poverty rates, while there is a positive correlation between individual income rank and census tract ethnic capital. It is noteworthy that the magnitude of negative effect of the percent of households in the census tract that are single mother households indicates that a 1 percent increase leads to 0.29 percentage point decrease in income rank, and that this effect increases to -0.85 when considering only the subsample of observations with a significant effect. Furthermore, while each variable has a subset of the sample with a significant marginal effect, we find very little evidence that heterogeneity in individual income rank is driven by whether or not the individual is raised in a single-mother household, the percent of the tract that is black, or rurality. The significance of tract ethnic capital is consistent with Borjas (1992): a one percent increase in tract ethnic capital is associated with a 0.11 percent increase in the individual's income rank. Overall, these estimates point towards social role modeling via single-motherhood (a negative effect) and ethnic capital (a positive effect) as well as economic forces (the number of siblings within the family, and the unemployment and poverty rates in the county) as driving factors.

Panels II, III, and IV in Table 5.6 summarize the heterogeneity in the intergenerational rank parameter, the race gap in mobility parameter, and the race gap in the intergenerational rank parameter. They evaluate *(ii) whether and to what extent each of theses factors affect individual income rank indirectly through intergenerational and racial differences in income.* Across all three parameters, the number of siblings, the percent of households in the census tract with a single-mother head, and the county unemployment rate are significant for the largest number of individuals in the sample. We therefore take this as evidence that the most fundamental factors that drive the intergenerational link in income gap, the race gap in mobility, and the race gap in the intergenerational rank parameter come from interhousehold resource allocation (i.e., the number of children in the household), the importance of role modeling (i.e., the share of the population that has a single-mother headed household), and structural economic opportunities (i.e., the county unemployment rate). It is worth emphasizing that these channels are those that appear statistically significant while controlling for other possible channels of effect that are often discussed in both theoretical and empirical literatures.

5.5.3 Interaction Effects between Individual and Socioeconomic Factors

It is important to investigate whether and to what extent the three estimated parameters would be different if an individual was exposed to different socioeconomic/geographic environment. This evaluates the hypothesis *(iii) whether and to what extent individual/family factors and socioeconomic/geographic factors have mutually reinforcing effects on individual income rank, either directly or indirectly.* We test this hypothesis by assessing the marginal effect of each individual/family variable on the three estimated parameters. As shown in equations (5.7) below, the

Table 5.6.: Summary of the effects of individual, socioeconomic, and geographic factors on individual income rank and the three empirical parameters

	Entire sample			Subsample of significant effect			
	Mean	Std. Err.	Range	Mean	Std. Err.	Range	No.
	Panel I: Marginal Effects on Individual Income Rank						
Male	0.02**	0.01	[−0.04, 0.07]	0.05***	0.01	[−0.06, 0.10]	1470
Number of siblings	−0.01**	0.00	[−0.03, 0.01]	−0.02***	0.00	[−0.04, −0.01]	2389
Mother-headed hh	0.00	0.01	[−0.09, 0.09]	−0.01	0.01	[−0.15, 0.16]	720
Tract black population	0.03	0.03	[−0.16, 0.22]	0.06	0.05	[−0.30, 0.32]	858
Tract single mother hh	−0.29**	0.12	[−1.52, 0.50]	−0.85***	0.16	[−2.06, 0.60]	1705
Tract ethnic capital	0.07	0.05	[−0.26, 0.30]	0.11*	0.06	[−0.48, 0.34]	1227
County unemployment	−0.06	0.04	[−0.42, 0.18]	−0.26***	0.05	[−0.73, 0.26]	1318
County poverty rate	−0.11	0.07	[−0.46, 0.32]	−0.36***	0.11	[−0.77, 0.38]	1275
County rurality	−0.03	0.04	[−0.31, 0.28]	−0.03	0.07	[−0.37, 0.48]	729
	Panel II: Marginal Effects on Intergenerational Rank Parameter						
Male	0.00	0.03	[−0.16, 0.14]	−0.03	0.04	[−0.28, 0.35]	443
Number of siblings	0.00	0.01	[−0.08, 0.07]	−0.01	0.01	[−0.10, 0.10]	1134
Mother-headed hh	0.02	0.05	[−0.21, 0.24]	−0.07	0.06	[−0.41, 0.39]	236
Tract black population	−0.11	0.14	[−0.63, 0.45]	−0.44**	0.21	[−0.86, 1.12]	601
Tract single mother hh	0.28	0.51	[−2.51, 3.32]	0.95	0.64	[−2.54, 3.59]	2784
Tract ethnic capital	−0.30	0.23	[−1.47, 0.86]	−1.47**	0.58	[−1.93, −1.01]	542
County unemployment	−0.08	0.17	[−2.01, 1.00]	−0.25	0.22	[−2.27, 1.13]	2506
County poverty rate	−0.27	0.31	[−1.96, 1.65]	−0.88*	0.48	[−2.28, 1.72]	936
County rurality	−0.36	0.22	[−1.78, 0.21]	−1.85**	0.83	[−2.28, −1.60]	435
	Panel III: Marginal Effects on Race Gap in Mobility						
Male	0.01	0.04	[−0.13, 0.14]	0.12***	0.03	[0.06, 0.22]	675
Number of siblings	0.02**	0.01	[−0.03, 0.09]	0.05**	0.02	[−0.05, 0.13]	1149
Mother-headed hh	−0.07	0.05	[−0.32, 0.11]	−0.23***	0.06	[−0.53, 0.13]	1033
Tract black population	−0.04	0.06	[−0.32, 0.49]	0.04	0.07	[−0.39, 0.82]	1405
Tract single mother hh	0.01	0.26	[−2.11, 1.87]	0.25	0.29	[−3.05, 2.43]	1944
Tract ethnic capital	−0.17	0.12	[−0.67, 0.40]	−0.56***	0.20	[−0.86, −0.39]	1185
County unemployment	−0.23**	0.09	[−0.85, 0.33]	−0.53***	0.12	[−1.09, 0.42]	2048
County poverty rate	0.00	0.16	[−0.80, 0.68]	0.01	0.21	[−0.82, 0.68]	1911
County rurality	0.12	0.11	[−0.39, 0.60]	0.50*	0.20	[0.27, 0.92]	646
	Panel IV: Marginal Effects on Race Gap in Intergenerational Rank Parameter						
Male	0.04	0.12	[−0.15, 0.28]	0.43**	0.21	[0.30, 0.62]	40
Number of siblings	−0.01	0.05	[−0.21, 0.18]	0.02	0.04	[−0.26, 0.24]	1418
Mother-headed hh	−0.14	0.17	[−0.56, 0.42]	−0.53**	0.22	[−0.79, −0.35]	587
Tract black population	−0.04	0.28	[−1.37, 0.74]	−0.93*	0.51	[−2.31, 0.90]	501
Tract single mother hh	−1.99*	1.13	[−4.25, 1.49]	−3.43***	1.15	[−4.55, −2.45]	2639
Tract ethnic capital	0.08	0.49	[−1.74, 1.90]	−1.83**	0.93	[−2.35, −1.49]	425
County unemployment	−1.07***	0.36	[−2.83, 1.14]	−1.51**	0.38	[−2.96, 1.47]	3549
County poverty rate	−0.58	0.65	[−1.62, 0.53]	–	–	–	–
County rurality	−0.44	0.46	[−1.69, 0.47]	–	–	–	–

Statistical significance is denoted via $^{***}p < .01,^{**} p < .05,^{*} p < .1$.

marginal effects of individual/family factors on the three parameters are functions of socioeconomic and geographic factors.

$$\begin{aligned}
\frac{\partial(\text{IG rank para})}{\partial x_m} &= (\sum_m \bar{\bar{\beta}}_{1m} + \sum_m \sum_k \bar{\bar{\beta}}_{1mk}\bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{1mj}\bar{\bar{o}}_j) \\
&+ (\sum_m \bar{\bar{\beta}}_{3m} + \sum_m \sum_k \bar{\bar{\beta}}_{3mk}\bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{3mj}\bar{\bar{o}}_j)\text{black} \\
\frac{\partial(\text{Race gap in mobility})}{\partial x_m} &= (\sum_m \bar{\bar{\beta}}_{2m} + \sum_m \sum_k \bar{\bar{\beta}}_{2mk}\bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{2mj}\bar{\bar{o}}_j) \\
&+ (\sum_m \bar{\bar{\beta}}_{3m} + \sum_m \sum_k \bar{\bar{\beta}}_{3mk}\bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{3mj}\bar{\bar{o}}_j)\text{rank}^{\pi} \\
\frac{\partial(\text{Race gap in IG rank para})}{\partial x_m} &= (\sum_m \bar{\bar{\beta}}_{3m} + \sum_m \sum_k \bar{\bar{\beta}}_{3mk}\bar{o}_k + \sum_m \sum_j \bar{\bar{\beta}}_{3mj}\bar{\bar{o}}_j), \quad m = 1,2,3
\end{aligned} \tag{5.7}$$

The coefficients in equations (5.7) are listed in Table 5.7, where Panels I – III represent the three estimated parameters respectively, and columns [1] – [3] represent the expected difference between male and female, and the marginal effects of number of siblings and mother headed household on the three parameters, respectively. The relationship could be *complementary* in a case where the effect of an individual/family factor is enlarged along with one of the socioeconomic/geographic factors, or *substituted* if the effect shrinks along with the socioeconomic/geographic factor. The interacted effect of an individual/family variable and a socioeconomic/geographic variable has two equivalent interpretations. The interacted effect represents effect of the individual/family variable depends on the degree of the socioeconomic/geographic variable, but this effect could also arise if the effect of the socioeconomic/geographic variable depends on the individual/family variable. We present the results with an emphasis on the first interpretation, however.

Table 5.7.: Coefficients in the marginal effects of individual/family factors x_m (full sample)

	[1]		[2]		[3]	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
	Panel I: Intergenerational Rank Parameter					
	Male		*Number of siblings*		*Single mother hh*	
Intercept	−0.06	0.29	0.14	0.12	0.02	0.53
Black	0.34	0.47	−0.10	0.16	0.62	0.75
Tract black population	−0.25	0.39	0.31**	0.16	0.20	0.66
Tract single mother hh	−0.38	1.35	−0.98**	0.47	−3.77	2.97
Tract ethnic capital	−0.06	0.38	−0.13	0.17	0.36	0.69
County unemployment	0.41	0.35	−0.27**	0.14	0.00	0.60
County poverty rate	0.30	0.72	−0.23	0.27	1.86	1.44
County rurality	0.01	0.42	0.01	0.17	−0.06	0.72
Black×tract black population	−0.13	0.44	−0.39**	0.17	−0.06	0.72
Black×tract single mother hh	0.70	1.62	0.90*	0.54	2.44	3.20
Black×tract ethnic capital	−0.37	0.78	0.48	0.32	−1.61	1.01
Black×county unemployment	−0.13	0.55	0.64***	0.19	−1.29	0.84
Black×county poverty rate	1.58	1.01	0.12	0.34	−0.29	1.74
Black×county rurality	−0.97	0.72	−0.14	0.24	−1.00	1.13
	Panel II: Race Gap in Mobility					
	Male		*Number of siblings*		*Single mother hh*	
Intercept	0.01	0.25	0.02	0.10	0.02	0.30
Rank$^{\pi}$	0.34	0.47	−0.10	0.16	0.62	0.75
Tract black population	0.00	0.25	0.26**	0.09	0.35	0.26
Tract single mother hh	−1.05	0.96	−0.80**	0.35	−3.45***	1.24
Tract ethnic capital	−0.11	0.45	−0.15	0.16	0.86*	0.52
County unemployment	0.24	0.32	−0.29**	0.12	−0.11	0.37
County poverty rate	0.00	0.57	−0.15	0.20	−0.14	0.72
County rurality	0.13	0.38	0.17	0.14	0.32	0.45
Rank$^{\pi}$×tract black population	−0.13	0.44	−0.39**	0.17	−0.06	0.72
Rank$^{\pi}$×tract single mother hh	0.70	1.62	0.90*	0.54	2.44	3.20
Rank$^{\pi}$×tract ethnic capital	−0.37	0.78	0.48	0.32	−1.61	1.01
Rank$^{\pi}$×county unemployment	−0.13	0.55	0.64**	0.19	−1.29	0.84
Rank$^{\pi}$×county poverty rate	1.58	1.01	0.12	0.34	−0.29	1.74
Rank$^{\pi}$×county rurality	−0.97	0.72	−0.14	0.24	−1.00	1.13
	Panel III: Race Gap in intergenerational Rank Parameter					
	Male		*Number of siblings*		*Single mother hh*	
Intercept	0.34	0.47	−0.10	0.16	0.62	0.75
Tract black population	−0.13	0.44	−0.39**	0.17	−0.06	0.72
Tract single mother hh	0.70	1.62	0.90*	0.54	2.44	3.20
Tract ethnic capital	−0.37	0.78	0.48	0.32	−1.61	1.01
County unemployment	−0.13	0.55	0.64**	0.19	−1.29	0.84
County poverty rate	1.58	1.01	0.12	0.34	−0.29	1.74
County rurality	−0.97	0.72	−0.14	0.24	−1.00	1.13

Statistical significance is denoted via: $^{***}p < .01, ^{**}p < .05, ^{*}p < .1$.

Decomposition of Intergenerational Rank Parameter

As we shown in Panel II in Table 5.6, the marginal effects of gender, number of siblings, and mother-headed households on the intergenerational rank parameter are negligible for majority of the sample. This null effect may arise from the socioeconomic/geographic factors and/or their interplay with race, as described in equation (5.7). This is especially relevant in the role of the number of siblings, as column [2] in Panel I in Table 5.7 suggests that the effect of the number of siblings on the intergenerational rank parameter depends on census tract shares of black population and single mother households and county unemployment rate, which vary by racial group as well. The effects of gender and whether living in a mother-headed household seem not relate to race and the socioeconomic/geographic factors, however. The details regarding to what extent the effect of number of siblings is affected by the three socioeconomic/geographic factors and race status are discussed as below:

The effect of siblings number on the intergenerational rank parameter varies in the share of census tract black population, at a magnitude of ($0.31^{*} - 0.39^{**} \times$ black dummy), conditional on the other socioeconomic factors. This implies that if a group of blacks and a group of non-blacks both moved to census tracts that have 1 more percentage point black population, the marginal effect of number of siblings on the intergenerational rank parameter would change by 0.31 for the non-blacks, but by -0.08 for the black group ceteris paribus. This suggests that the effect of number of siblings on the intergenerational rank parameter acts as a complement with census tract black population for non-blacks, but as a substitute blacks.

Second, the effect of siblings number on the intergenerational rank parameter is affected by the share of single mother household within census tracts, specifically ($-0.98^{*} + 0.90^{**} \times$ black dummy). This means that to the strength of the relationship between the number of siblings matters and the intergenerational rank parameter depends on the share of census tract single mother households, and on the individual's own race identity. In concrete terms, if a black individual moved to a census tract that has 1 percentage point more single mother households, these results suggest that adding one more sibling to this individual would change the intergenerational rank parameter by -0.08 percentage point, while if this individual was non-black, this effect is -0.98. This heterogeneity suggests that limited intra-house resources (number of siblings) on the intergenerational rank parameter is weaker when the negative role modeling is the neighborhood (census tract single mother households) is high, additionally, even more weakly for non-blacks. Finally, county unemployment rate also appears to matter, as it affects the effect of sibling numbers by $(-0.27^{**} + 0.64^{**} \times \text{black dummy})$, implying that the impact of the number of siblings on the intergenerational rank parameter is changed by 0.37 for blacks, and by -0.27 for non-blacks, when one more percentage point of county unemployment rate occurs.

All of above findings tell us racial identity matters for the impact of the number of siblings and county unemployment is complemented or substituted in explaining the role of parent income rank on individual's. This, in turn, implies that the intergenerational rank parameter may arise from these interactions between race identity and the socioeconomic factors.

Decomposition of Race Gaps in Mobility and Intergenerational Rank Parameter

Panel II in the Table 5.7 describes how the effects of the individual factors on the race gap in mobility, as shown in panel III in Table 5.6, depend on socioeconomic/geographic factors and parental income rank. Here we note the similar findings as in the intergenerational rank parameter regarding the effect of number of siblings: the large range of the effect of siblings number on the race gap in mobility (shown in Panel III in Table 5.6) comes from the differences in census tract shares of black population and single mother households, and county unemployment rate, and additionally these relationships depend on how wealthy of the parents as well. For instance, county unemployment rate affects the role of sibling numbers in the race gap in mobility by: $(-0.29^{***} + 0.64^{***} \times \text{rank}^{\pi})$. Overall, this effect is positively associated with parental income rank. Interestingly, by solving rank$^{\pi}$, this hints that at greater than 45th percentile of parental income rank, this effect will be positive, otherwise, this effect is negative. This implies that for a wealthy family, the effects of siblings number and county unemployment on the race gap in mobility is complementary, while it is substituted for a low ranked family. Similar implications can be drawn regarding the effect of census tract shares of single mother households on the role of number of siblings on the race gap in mobility, as the estimated effect is measured at $(-0.80^{***} + 0.90^{***} \times \text{rank}^{\pi})$. The cutoff of parental income rank is 0.89: when parental income rank is above 89th percentile, this effect is positive, otherwise this effect is negative. This relatively higher cutoff suggests that only for the bottom ten percentile ranked families the effects of siblings number and census tract share of single mother household is substituted.

However, we find different implications regarding the effect of census tract black population. It affects the effect of sibling numbers on the racial mobility gap by: $(0.26^{**} - 0.39^{**} \times \text{rank}^{\pi})$. This suggests that, conditional on parental income rank, the effect of the sibling numbers on the intergenerational rank parameter would increase by 0.26 percentage point along with one percentage point increase in census tract black population, but eventually attenuates by 0.39 percentage point along with one percentage point increase in parental income rank. By solving rank$^{\pi}$, we note that the cutoff is 67th percentile: when parental income rank is lower than the cutoff, this effect is positive, otherwise the effect is negative.

Finally, Panel III in Table 5.7 describes how the effects of the individual factors on the race gap in intergenerational rank parameter depend on socioeconomic/geographic factors and parental income rank. Since the race gap in intergenerational rank parameter can be computed as the marginal effect of parental income rank on race gap in mobility or the expected difference of intergenerational rank parameter between black and non-black, the implications of these effects of the the socioeconomic/geographic factors are the same as in Panel I and II, as we discussed above. Generally, these relationships underscore the importance of intra-house resource allocation, social interactions, and local economic conditions in influencing individual's income rank, intergenerational rank parameter, race gap in mobility and intergenerational rank parameter.

5.6 Conclusion

The implications that arise from these results remain somewhat unclear. We have decomposed racial differences in both relative and absolute mobility into several components, many of which include interactions occurring at several levels, and while we find substantial racial variation in the size of these gaps across the data, no one clear story emerges. We do, however, find several interesting stylized facts, arising from our results regarding the racial mobility gap, the most easily interpretable of our parameters. First, the mobility gap varies meaningfully over community characteristics: black children are substantially less mobile compared to other children in tracts with high rates of single motherhood, and in counties with high levels of unemployment, for example. These gaps also vary across households, suggesting that black and white households differ in their ability to mitigate certain challenges, such as household-level single motherhood. These household variables also vary as community characteristics change, but not by as much as I had hypothesized. Overall, these results underscore the evidence for large and persistent racial mobility gaps, and for significant geographic variation in those gaps, presented in Hertz (2005) and Chetty et al. (2018), with the added insight that these differences appear to vary across different kinds of households and neighborhoods. More work remains to be done to apply the approach developed here to answering more specific questions about cross-level interactions in models of the geography of mobility gaps, but these results represent a valuable step in that project.

5.7 Appendix

Table 5.8.: Significant coefficients in the HLM regression

	Coef.	Std. Err.
Intercept	-0.74^{**}	0.32
Tract black population	0.82^{**}	0.35
Tract single mother hh	-4.86^{***}	1.14
County unemployment	-0.68^{*}	0.37
Number of siblings × Tract black population	-0.24^{***}	0.09
Number of siblings × Tract single mother hh	0.88^{***}	0.34
Mother-head hh × Tract single mother hh	3.12^{***}	1.15
Number of siblings × County unemployment	0.20^{*}	0.11
Parent income rank × Tract black population	4.92^{***}	1.62
Parent income rank × County unemployment	0.99^{**}	0.48
Parent income rank × Number of siblings × Tract black population	0.31^{*}	0.16
Parent income rank × Number of siblings × Tract single mother hh	-0.98^{**}	0.47
Parent income rank × Number of siblings × County unemployment	-0.27^{*}	0.14
Black × Tract black population	-0.96^{***}	0.37
Black × Number of siblings	5.18^{***}	1.27
Black × County unemployment	1.13^{**}	0.45
Black × Number of siblings × Tract black population	0.26^{***}	0.09
Black × Number of siblings × Tract single mother hh	-0.80^{**}	0.35
Black × Mother-head hh × Tract single mother hh	-3.45^{***}	1.24
Black × Number of siblings × County unemployment	-0.29^{**}	0.12
Parent income rank × Black × Tract black population	1.26^{**}	0.62
Parent income rank × Black × Number of siblings	-5.77^{***}	2.01
Parent income rank × Black × County unemployment	-2.69^{***}	0.75
Parent income rank × Black × Number of siblings × Tract black population	-0.39^{**}	0.17
Parent income rank × Black × Number of siblings × Number of siblings	0.90^{*}	0.54
Parent income rank × Black × Number of siblings × County unemployment	0.64^{***}	0.19
Own adult age	0.07^{***}	0.01
(Own adult age)^2	-0.00^{***}	0.00
N	5,248	
Adjusted R^2	0.34	

Includes only statistically significant coefficients; $^{***}p < .01, ^{**}p < .05, ^{*}p < .1$.

Table 5.9.: Bivariate correlations between independent variables

	1	2	3	4	5	6	7	8	9	10	11	12	13
1. Parental income rank	1												
2. Black	−0.51	1											
3. Male	0.06	−0.06	1										
4. Number of siblings	−0.17	0.29	−0.03	1									
5. Mother-headed hh	−0.56	0.40	−0.03	0.04	1								
6. Tract black population	−0.48	0.82	−0.05	0.25	0.42	1							
7. Tract single mother hh	−0.46	0.63	−0.06	0.11	0.43	0.76	1						
8. Tract ethnic capital	0.44	−0.41	0.02	−0.26	−0.20	−0.40	−0.31	1					
9. County unemployment	0.04	0.12	0.02	0.39	0.01	0.15	−0.14	−0.30	1				
10. County poverty rate	−0.42	0.43	−0.02	0.12	0.14	0.40	0.35	−0.27	−0.06	1			
11. County index of rurality	−0.05	−0.19	0.03	−0.05	−0.19	−0.33	−0.34	−0.08	−0.20	0.26	1		
12. Age	0.11	0.00	−0.12	0.14	−0.07	−0.01	−0.18	−0.15	0.50	−0.08	0.00	1	
13. Parental age	0.24	−0.08	0.02	0.11	−0.21	−0.08	−0.21	0.03	0.37	−0.08	−0.01	0.23	1

5.7.1 Additional Summary Statistics – IG Rank Parameter, Race Gap in Mobility, and Race Gap in IG Rank Parameter, by Significance

Table 5.10.: Descriptive statistics of the three parameters by significance [a]

	IG Rank Parameter								
	sample with sig values (N=4066)				sample with insig values (N=1182)				
	Mean	St. Dev.	Min	Max	Mean	St. Dev.	Min	Max	t-test
male	0.461	0.499	0	1	0.466	0.499	0	1	-0.334
No. of siblings	2.991	1.430	0.600	10.000	3.683	1.983	0.800	9.714	-11.192***
mother-headed household	0.185	0.388	0	1	0.508	0.500	0	1	-20.514***
Individual age	31.560	4.339	25.000	52.000	31.066	4.404	25.000	50.100	3.405***
Parent age	39.681	6.145	26.167	54.000	40.322	6.522	27.200	53.400	-3.012***
Tract black population	0.226	0.313	0.000	0.999	0.527	0.362	0.000	1.000	-25.953***
Tract single mother	0.076	0.052	0.005	0.430	0.129	0.084	0.009	0.558	-20.731***
Tract ethnic capital	0.114	0.085	0.000	0.614	0.126	0.153	0.000	0.805	-2.638***
County unemployment rate	0.184	0.116	0.019	0.488	0.205	0.135	0.022	0.495	-4.834***
County poverty rate	0.122	0.057	0.021	0.445	0.166	0.081	0.025	0.527	-17.644***
County IRR	0.485	0.095	0.267	0.772	0.464	0.101	0.267	0.777	6.427***
	Race Gap in Mobility								
	sample with sig values (N=2071)				sample with insig values (N=3177)				
	Mean	St. Dev.	Min	Max	Mean	St. Dev.	Min	Max	t-test
male	0.363	0.481	0	1	0.527	0.499	0	1	-11.891***
No. of siblings	2.748	1.221	0.600	9.714	3.407	1.753	0.800	10.000	-16.024***
mother-headed household	0.173	0.378	0	1	0.313	0.464	0	1	-11.997***
Individual age	31.807	4.375	25.000	51.833	31.216	4.331	25.000	52.000	4.802***
Parent age	39.920	6.277	26.167	53.600	39.763	6.211	27.000	54.000	0.887
Tract black population	0.198	0.282	0.000	0.999	0.356	0.372	0.000	1.000	-17.486***
Tract single mother	0.072	0.053	0.005	0.505	0.098	0.070	0.009	0.558	-15.291***
Tract ethnic capital	0.124	0.096	0.000	0.805	0.113	0.109	0.000	0.764	3.888***
County unemployment rate	0.198	0.118	0.022	0.488	0.183	0.123	0.019	0.495	4.426***
County poverty rate	0.119	0.054	0.021	0.377	0.140	0.071	0.022	0.527	-12.050***
County IRR	0.483	0.094	0.267	0.777	0.479	0.098	0.270	0.769	1.511
	Race Gap in IG Rank Parameter								
	sample with sig values (N=732)				sample with insig values (N=4,516)				
	Mean	St. Dev.	Min	Max	Mean	St. Dev.	Min	Max	t-test
male	0.243	0.429	0	1	0.497	0.500	0	1	-14.503***
No. of siblings	3.475	1.874	0.750	9.714	3.094	1.542	0.600	10.000	5.230***
mother-headed household	0.128	0.335	0	1	0.279	0.448	0	1	-10.696***
Individual age	31.911	4.333	25.000	52.000	31.374	4.358	25.000	51.833	3.105***
Parent age	40.868	5.999	27.429	53.000	39.656	6.259	26.167	54.000	5.041***
Tract black population	0.238	0.304	0.000	0.994	0.303	0.354	0.000	1.000	-5.238***
Tract single mother	0.078	0.060	0.015	0.430	0.089	0.065	0.005	0.558	-4.638***
Tract ethnic capital	0.123	0.095	0.001	0.571	0.116	0.106	0.000	0.805	1.799*
County unemployment rate	0.191	0.133	0.024	0.495	0.189	0.119	0.019	0.478	0.355
County poverty rate	0.133	0.074	0.022	0.445	0.131	0.065	0.021	0.527	0.689
County IRR	0.488	0.098	0.267	0.750	0.479	0.096	0.267	0.777	2.163***

[a] In each panel, [1] denotes subsample with significant parameter, and [2] denotes subsample with insignificant parameter.

5.7.2 Compare with Chetty et al. (2014): Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States

To further validate our estimation, we compare with Chetty et al. (2014a). We use slightly different calculation to replicate Figure VI in Chetty et al. (2014a). Given data availability, our calculation differs from Chetty et al. (2014a). In terms of the magnitudes, our estimation is about a quarter of the value from Chetty et al. (2014a).

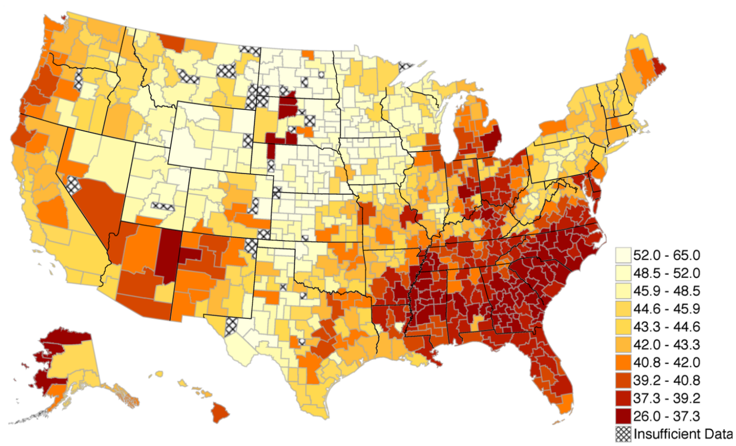

Figure 5.1.: Mean Child Percentile Rank for Parents at 25th Percentile (Y_{25}) (Chetty et al., 2014a)

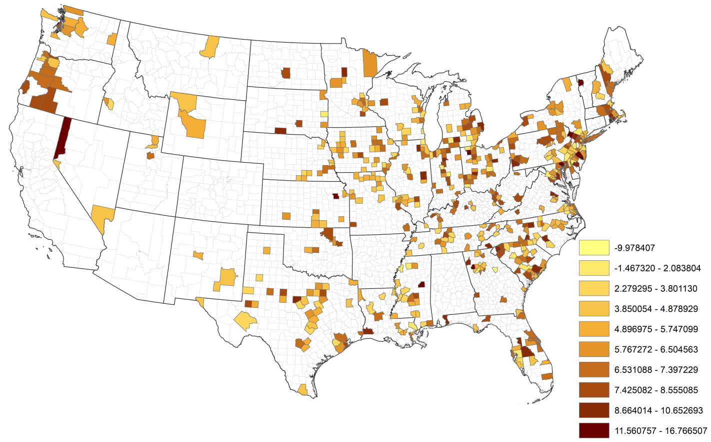

Figure 5.2.: Mean Child Percentile Rank for Parents at 25th Percentile (Y_{25}) using Our Sample

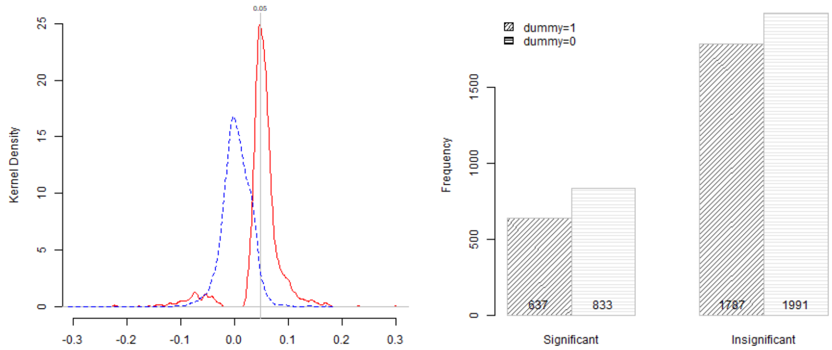

(a) L: Male-female difference; R: Distribution of male (Blue:insignificant; Red: significant)

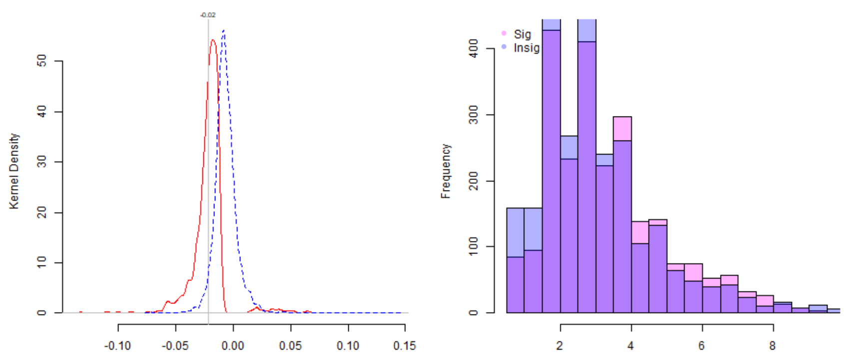

(b) L: Marginal effect of number of siblings; R: Distribution of Number of siblings (Blue:insignificant; Red: significant)

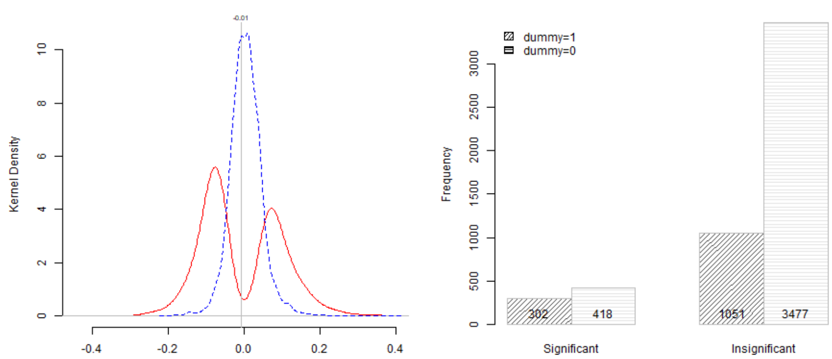

(c) L: Marginal effect of growing up in mother-head hh; R: Distribution of Mother-headed hh (Blue:insignificant; Red: significant)

Figure 5.3.: Marginal effects of individual factors on <u>income rank</u>

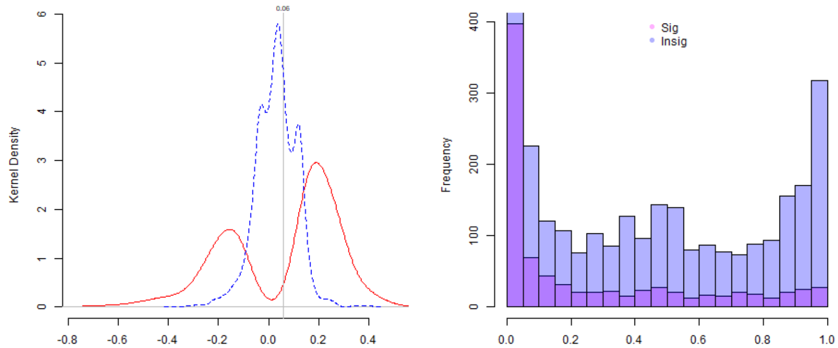

(a) L: Marginal effect of tract black population; R: Distribution of tract black population (Blue:insignificant; Red: significant)

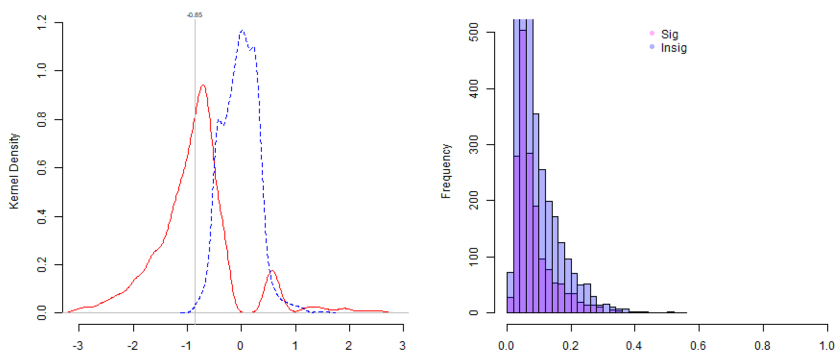

(b) L: Marginal effect of tract single mother hh; R: Distribution of tract single mother hh (Blue:insignificant; Red: significant)

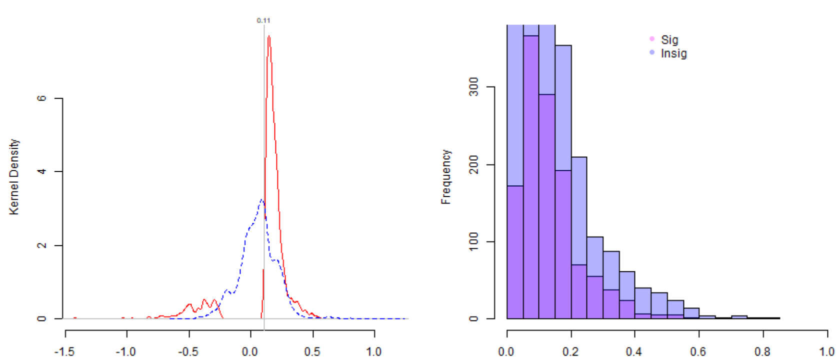

(c) L: Marginal effect of tract ethnic capital; R: Distribution of tract ethnic capital (Blue:insignificant; Red: significant)

Figure 5.4.: Marginal effects of tract factors on <u>income rank</u>

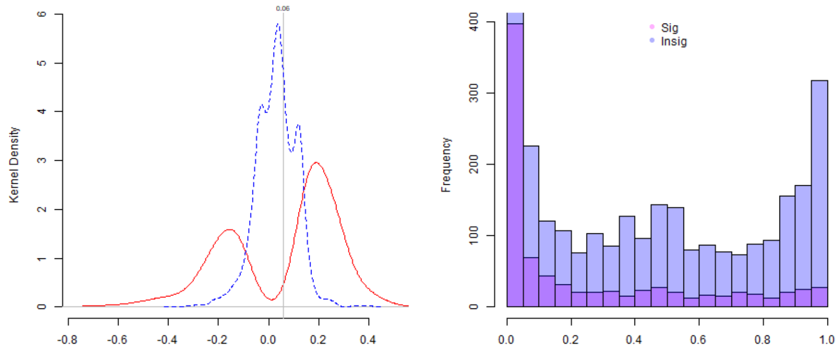

(a) L: Marginal effect of county unemployment; R: Distribution of county unemployment (Blue:insignificant; Red: significant)

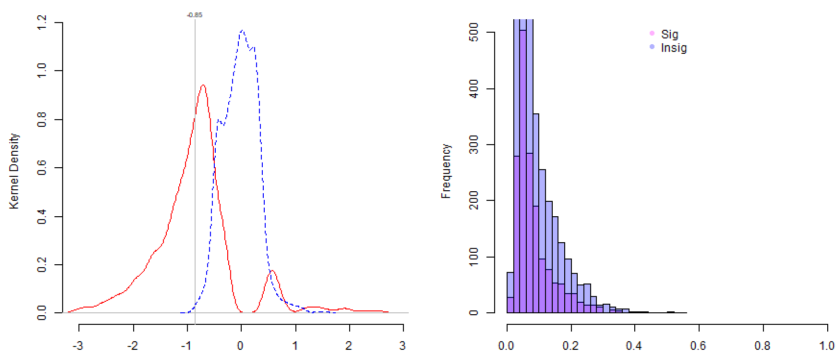

(b) L: Marginal effect of county poverty rate; R: Distribution of county poverty rate (Blue:insignificant; Red: significant)

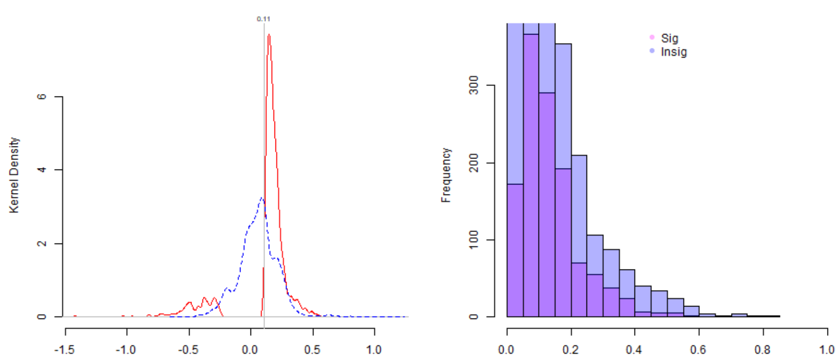

(c) L: Marginal effect of county rurality; R: Distribution of county rurality (Blue:insignificant; Red: significant)

Figure 5.5.: Marginal effects of county factors on income rank

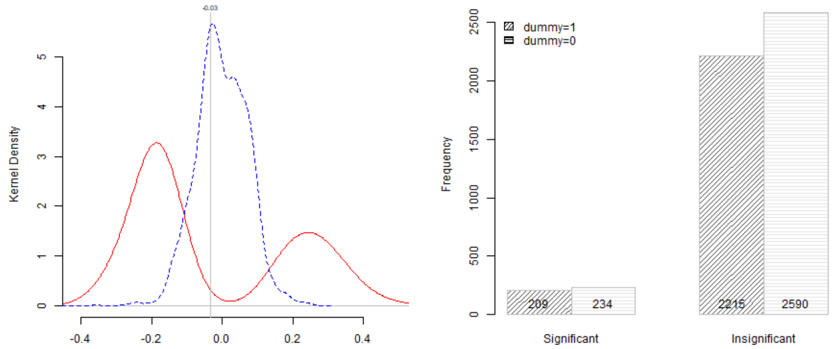

(a) L: Male-female difference; R: Distribution of Male (Blue:insignificant; Red: significant)

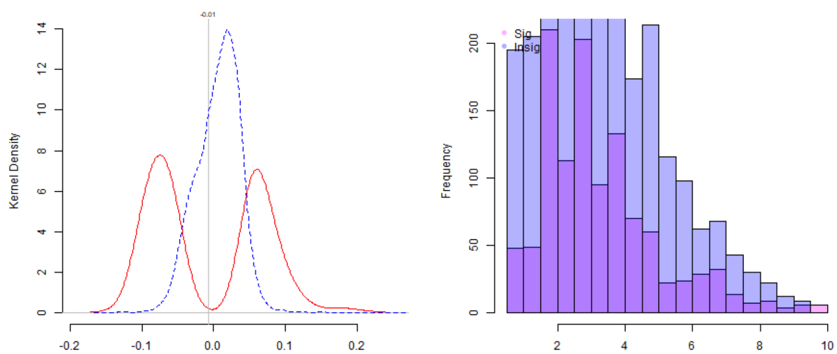

(b) L: Marginal effect of number of siblings; R: Distribution of Number of siblings (Blue:insignificant; Red: significant)

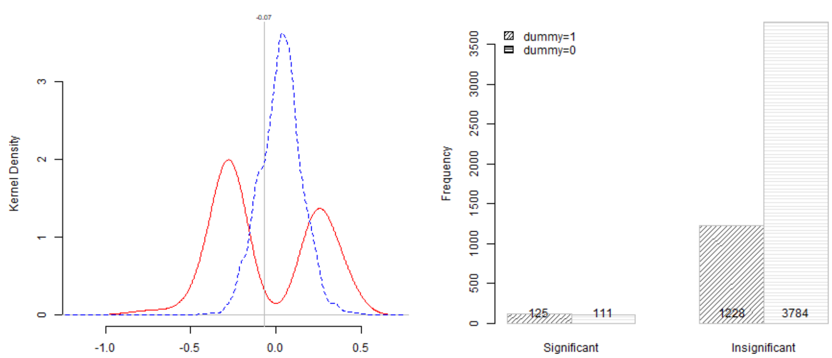

(c) L: Marginal effect of growing up in mother-head hh; R: Distribution of Mother-headed hh (Blue:insignificant; Red: significant)

Figure 5.6.: Marginal effects of individual factors on intergenerational rank parameter

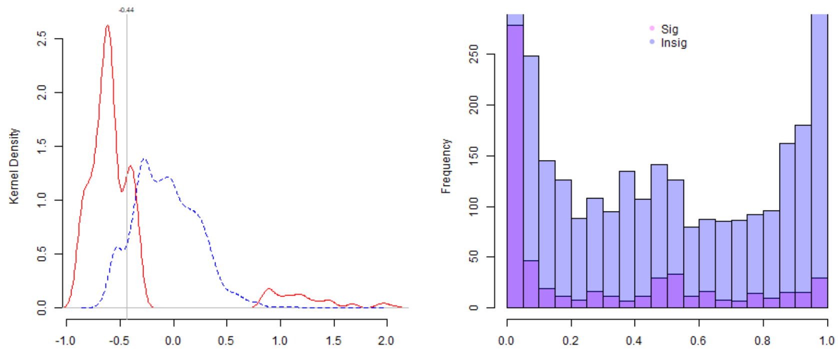

(a) L: Marginal effect of tract black population; R: Distribution of tract black population

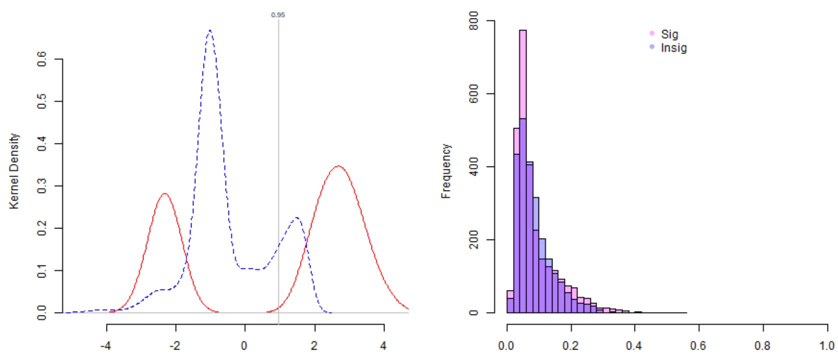

(b) L: Marginal effect of tract single mother hh; R: Distribution of tract single mother hh

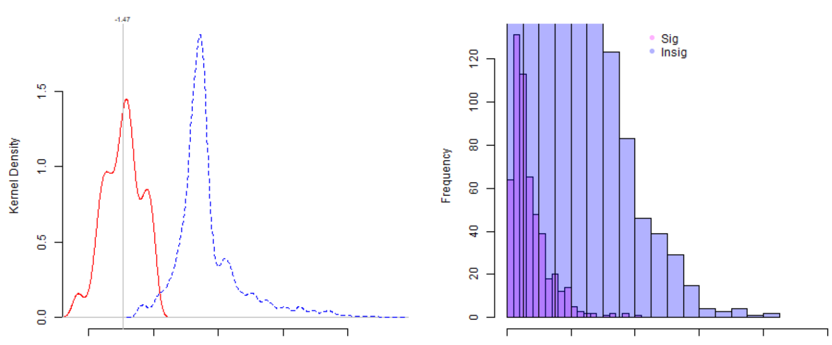

(c) L: Marginal effect of tract ethnic capital; R: Distribution of tract ethnic capital

Figure 5.7.: Marginal effects of tract factors on <u>intergenerational rank parameter</u>

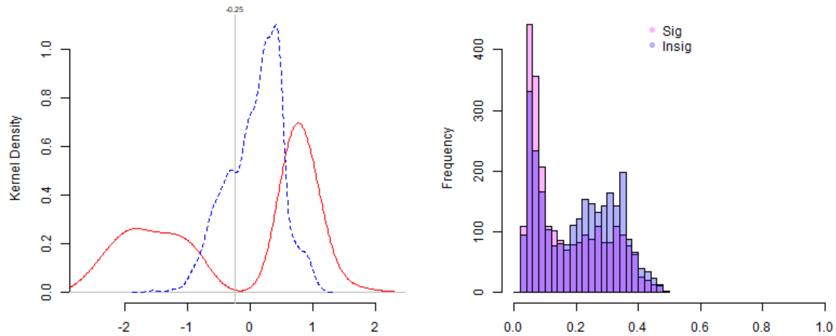

(a) L: Marginal effect of county unemployment; R: Distribution of county unemployment (Blue:insignificant; Red: significant)

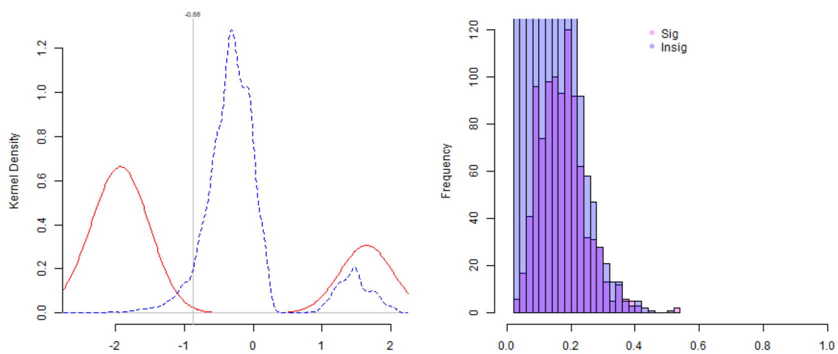

(b) L: Marginal effect of county poverty rate; R: Distribution of county poverty rate (Blue:insignificant; Red: significant)

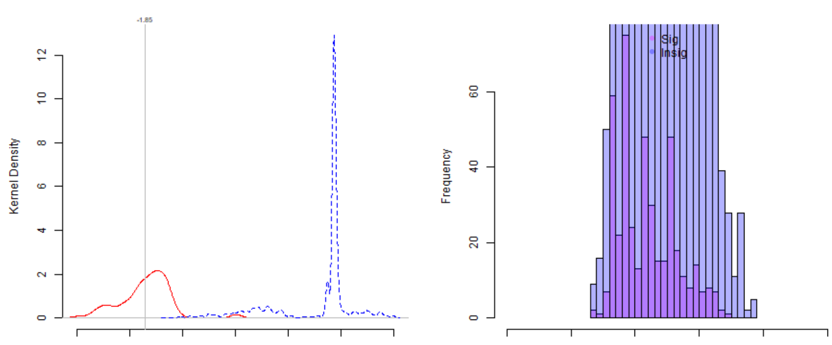

(c) L: Marginal effect of county rurality; R: Distribution of county rurality (Blue:insignificant; Red: significant)

Figure 5.8.: Marginal effects of county factors on intergenerational rank parameter

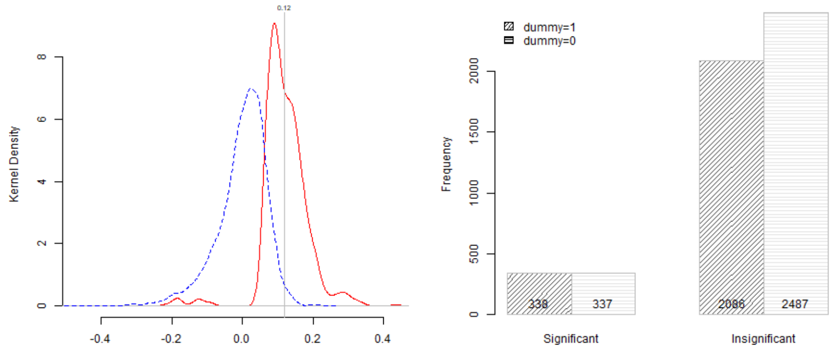

(a) L: Male-female difference; R: Distribution of Male (Blue:insignificant; Red: significant)

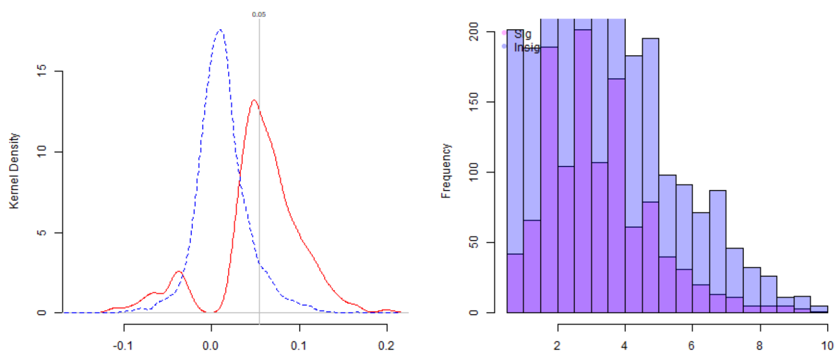

(b) L: Marginal effect of number of siblings; R: Distribution of Number of siblings (Blue:insignificant; Red: significant)

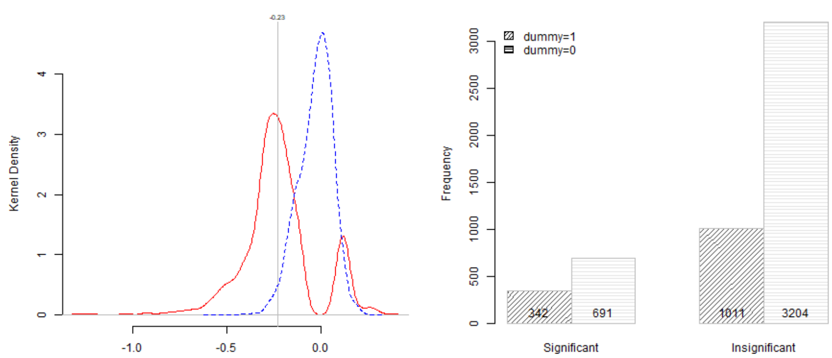

(c) L: Marginal effect of growing up in mother-head hh; R: Distribution of Mother-headed hh (Blue:insignificant; Red: significant)

Figure 5.9.: Marginal effects of individual factors on race gap in mobility

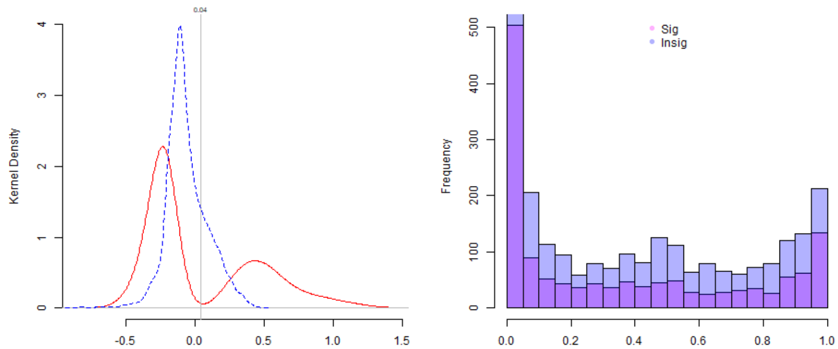

(a) L: Marginal effect of tract black population; R: Distribution of tract black population

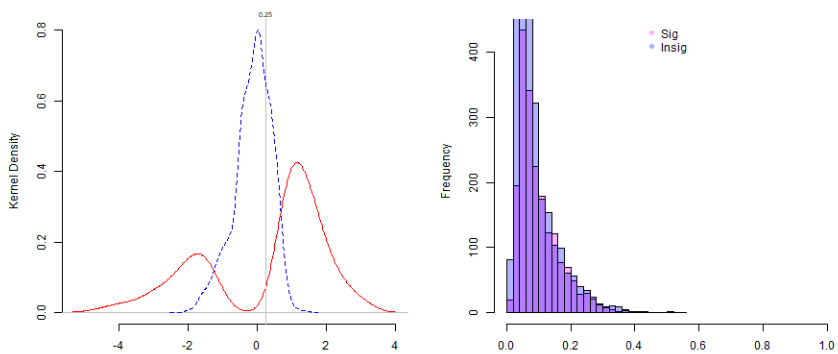

(b) L: Marginal effect of tract single mother hh; R: Distribution of tract single mother hh

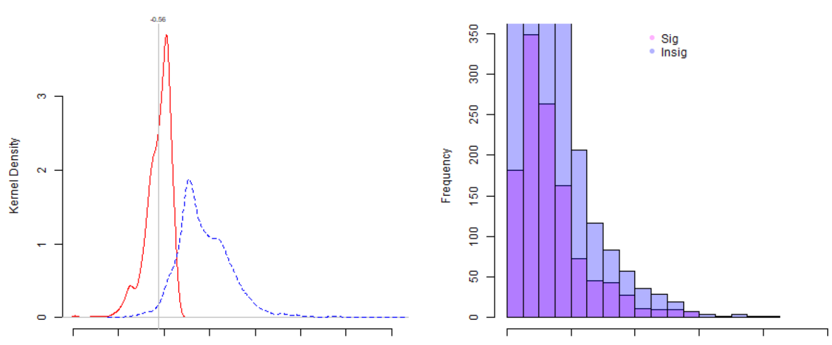

(c) L: Marginal effect of tract ethnic capital; R: Distribution of tract ethnic capital

Figure 5.10.: Marginal effects of tract factors on race gap in mobility

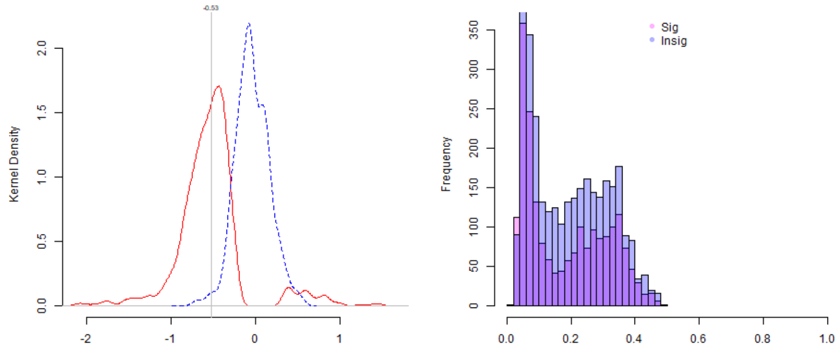

(a) L: Marginal effect of county unemployment; R: Distribution of county unemployment (Blue:insignificant; Red: significant)

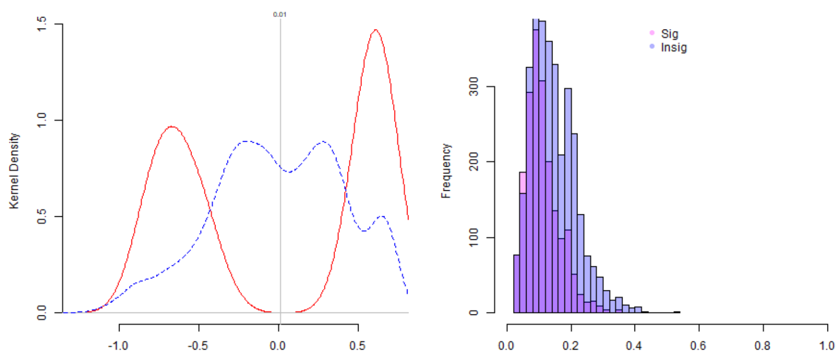

(b) L: Marginal effect of county poverty rate; R: Distribution of county poverty rate (Blue:insignificant; Red: significant)

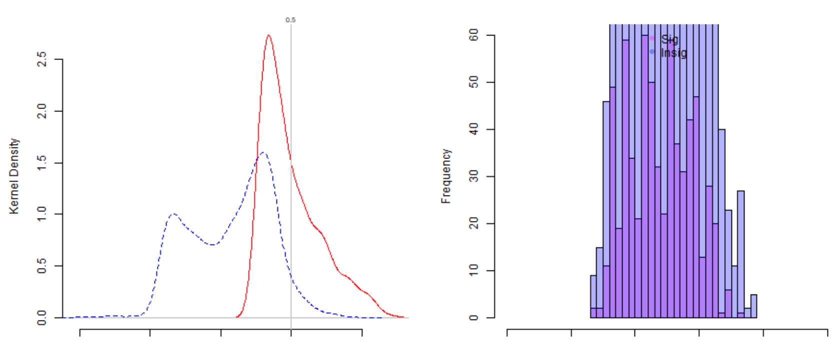

(c) L: Marginal effect of county rurality; R: Distribution of county rurality (Blue:insignificant; Red: significant)

Figure 5.11.: Marginal effects of county factors on race gap in mobility

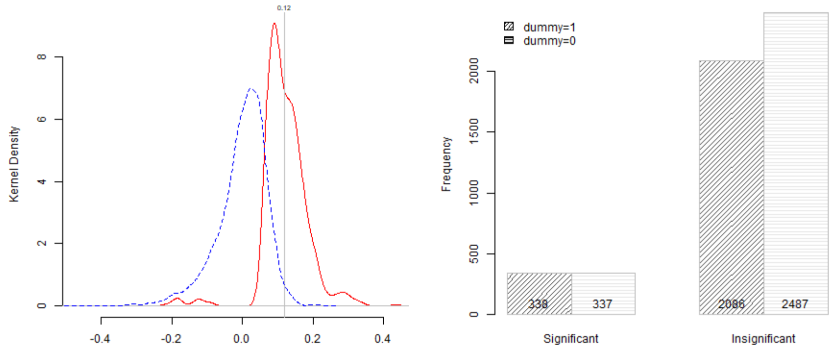

(a) L: Male-female difference; R: Distribution of Male (Blue:insignificant; Red: significant)

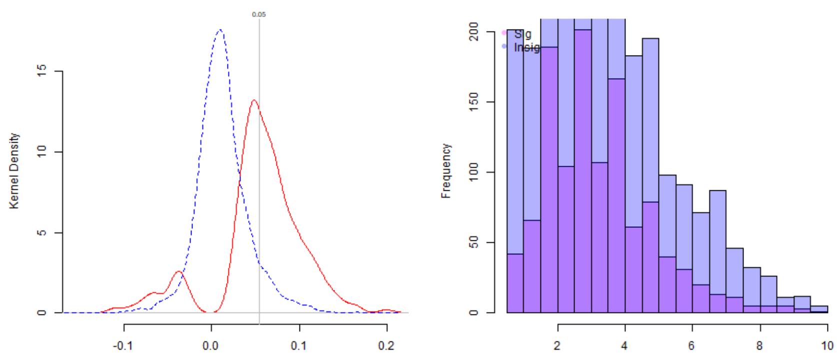

(b) L: Marginal effect of number of siblings; R: Distribution of Number of siblings (Blue:insignificant; Red: significant)

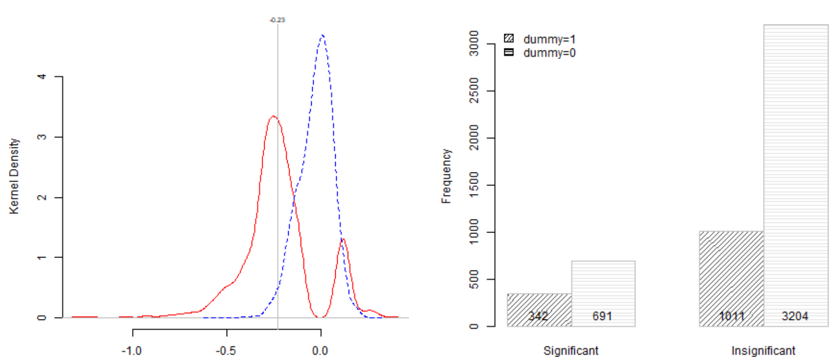

(c) L: Marginal effect of growing up in mother-head hh; R: Distribution of Mother-headed hh (Blue:insignificant; Red: significant)

Figure 5.12.: Marginal effects of individual factors on race gap in IG rank parameter

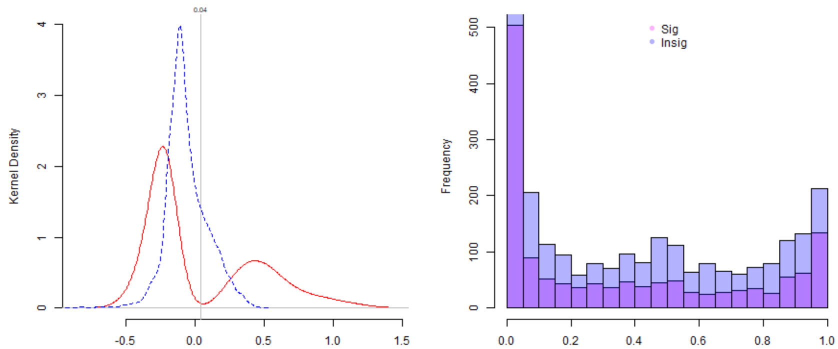

(a) L: Marginal effect of tract black population; R: Distribution of tract black population

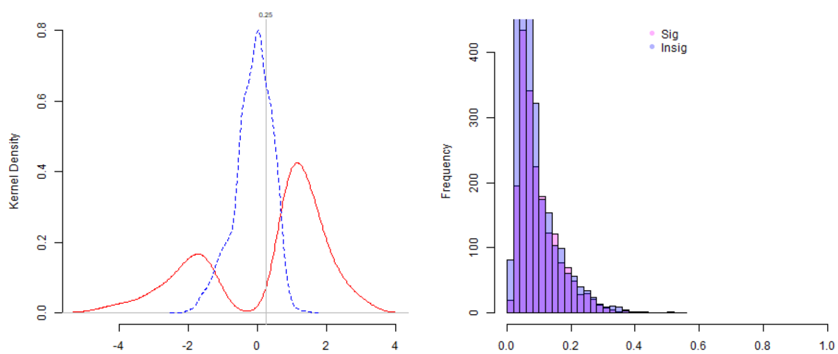

(b) L: Marginal effect of tract single mother hh; R: Distribution of tract single mother hh

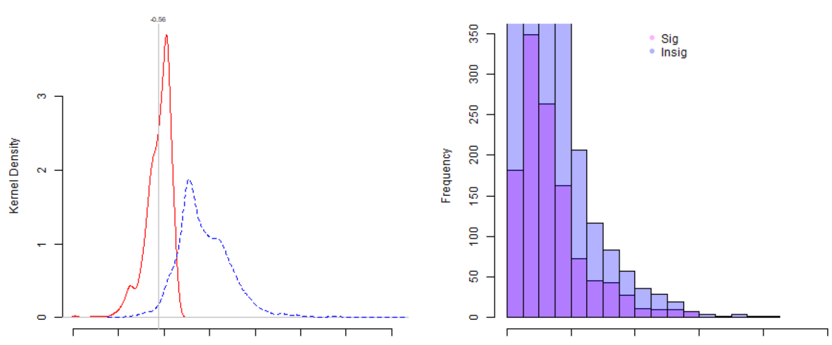

(c) L: Marginal effect of tract ethnic capital; R: Distribution of tract ethnic capital

Figure 5.13.: Marginal effects of tract factors on race gap in IG rank parameter

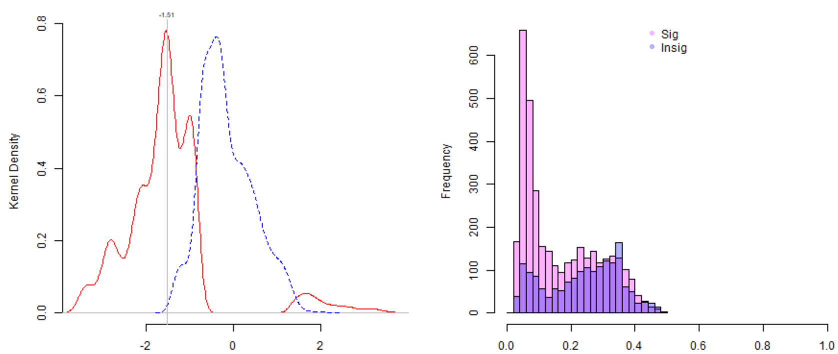

(a) L: Marginal effect of county unemployment; R: Distribution of county unemployment

Figure 5.14.: Marginal effects of county factors on race gap in IG rank parameter

CHAPTER 6. AGRICULTURAL TRANSFORMATION AND INTERGENERATIONAL MOBILITY IN THE UNITED STATES

6.1 Introduction

The United States' economy changed dramatically throughout the 20^{th} century, with new technologies emerging, the service sector growing dramatically, and human capital becoming increasingly specialized, well-compensated, and common. The speed of these changes meant that for much of this century, each successive generation faced a set of economic conditions and opportunities that diverged significantly from those their parents had faced. Many scholars have wondered how these changes affected intergenerational mobility, that is, the strength of the relationship between parents and children in economic outcomes. Much of the early contributions to the literature as it stands today are motivated by attempts to validate the standard narrative of increasing 20^{th} century opportunity (e.g. Solon (1992), Zimmerman (1992)), and more recent work has found that the U.S. population was more mobile early in the 20^{th} century than it is today (Feigenbaum, 2018; Hilger, 2017a), although literature studying mobility trends in the later decades of the century suggests that mobility remained relatively steady (Hertz, 2007; Lee and Solon, 2009; Chetty et al., 2014a).

One important trend, well-documented in the economic history and agricultural economics literatures, is the so called "structural transformation of agriculture", which saw a steady substitution of capital for labor over time in the agriculture sector, leading to large gains in productivity, profitability, and the world food supply. In the context of labor markets and mobility, this was a particularly important shift in the South, which lagged behind the North in mobility throughout this period and today, but which saw dramatic gains in productivity alongside a massive exodus from smallholder agriculture in the middle of the 20^{th} century. Hilger (2017a) sets out to explain trends in mobility between 1940 and 2000 with an emphasis on changes in the South, focusing on the expansion of educational and economic opportunities for African Americans, and finds that only mean state income and inequality are consistently associated with increased mobility. This result is interesting because it is counterintuitive: the mobility literature dating back to Becker and Tomes (1979) suggests that limiting racial disparities should increase mobility, but they make ambiguous predictions about average incomes, so explaining this result can lead to new insights about long-term changes in mobility. This standard model of mobility suggests that when the parameters that affect parental investment decisions change, mobility levels will also change, but this only implies that income and mobility will move together when incomes rise as a result of shifting returns to skills, easier access to credit, broader access to education, or similar structural changes, not that they should be robustly related over time and space. There are certainly conditions under which theory predicts mobility and incomes to rise together, most obviously the case in which the entire income distribution shifts rightward, so that all children to relatively better compared to the previous generation. Given the complexity of the changes during this period, however, we have no real reason to assume this is the case. There is no reason to expect a negative relationship between average incomes and mobility, so these results do not contradict established theoretical or empirical results, and because the mobility measurements include children who moved away from their state of birth, they do not imply regional divergence in income. This pattern of facts raises questions about the relationship between agricultural development and economic mobility, both because the release of labor from the agriculture sector as industry and service sectors grew likely affected children's prospects, and because agricultural changes could explain Hilger's result.

In this paper, I study these relationships, investigating how the transformation of agriculture since 1940 affected intergenerational mobility, and whether changes in agricultural production and labor markets can explain the income-mobility result

Hilger finds. I proceed by presenting a series of empirical exercises which connect mobility to different aspects of the structural transformation of agriculture. In brief, I study the relationship between mobility, income, agricultural employment, and agricultural productivity by applying a modified version of Hilger's empirical strategy to a series of exercises in modelling variation in state-level mobility parameters, and in modelling intergroup differences in those parameters. Taken together, these exercises provide valuable evidence regarding the relationship between agricultural development and intergenerational mobility in the long-term, between 1940 and 2000.[1]

This work is important because understanding the sources of mobility in the middle of the 20^{th} century provides crucial context for explaining the levels of mobility scholars have observed in recent decades and today. This is valuable, in part, because it extends the trend that research relying on recent longitudinal surveys studies, which may provide a totally different impression of how long-term trends developed. This is important on its own, because it enhances our understanding of this important period in U.S. history, but it also has policy salience today, however, because a proper understanding of historical trends changes our the policy prescription for responses to research on mobility patterns today. This work suggests that mobility in many parts of the country, and for particular groups, is very low (Chetty et al., 2014b, 2018). For example, if, as Davis and Mazumder (2016) argue, mobility dropped starting in the 1980's due to policies that favored high earners, policy interventions that tax these high earners to invest in targeted child development programs would be an ideal instrument for improving mobility (if, indeed, this is what policymakers and the public want). If, instead, this apparent drop in mobility comes, in part, from exhausting the low-hanging fruit of modernization, or of the dynamic effects of previous changes, as Nybom and Stuhler (2016) suggest, these policies will have very different results, even when judged on their intended effects.

[1]Some scholars, notably Long and Ferrie (2015) and Collins and Wanamaker (2017) use linked Census microdata to look back to 1850 and 1880 respectively, but I prefer the Hilger method despite the cost in years, because it allows much more representative links between mobility and agricultural trends.

6.2 Background

6.2.1 Agricultural Development, Labor Markets, and Human Capital

The claim that massive changes in technology would lead to a shift in the labor market is uncontroversial, but the connection of these changes to changes in intergenerational mobility, and to the apparent link between income levels and mobility, requires more justification. The economic history literature presents strong evidence that these changes did affect labor markets in ways that appear to be related to mobility, however, and it provides evidence pointing towards specific mechanisms that provide this connection. Mundlak (2005) provides a valuable overview of this literature, integrating the broad literature on technological change with the specifics of agricultural innovations and the movement of labor from rural to urban markets. He summarizes the affects of this structural transformation on labor markets concisely: "...[O]utput growth was triggered largely by new technology, which in large part was labor-saving. This, together with the development of nonagriculture, resulted in off-farm occupational migration of labor. The decline in food prices improved consumers' welfare, and labor mobility to nonagriculture contributed to overall economic development. The off-farm migration was a major factor in the alleviation of rural poverty." This narrative is broadly consistent with my hypothesis that the transformation of agriculture enhanced mobility by releasing labor from the agricultural sector to new opportunities, but a number of additional papers add to this point by considering the effects of technological and institutional change on markets for agricultural labor.

One important strand of this literature emphasizes the labor market institutions that grew up around agriculture, particularly in the South. Much of this literature relies, implicitly or explicitly, on the framework developed in Whatley (1985) and Whatley (1987), which show that foregoing mechanization of cotton in the south, and instead relying on annual labor contracts which helped maintain high profits for landowners, alongside low wages and human capital for laborers, was optimal for many landowners. This system was, in part, supported by the institutional legacy of slavery, which had been particularly important in cotton production, due to the prevalence of sharecropping and limitations on legal rights and labor market opportunities for African Americans, prevalent in the South during the early 20^{th} century. Jung (2018) extends this thinking into the middle of the 20^{th} century, arguing that the institutions that grew up around cotton should have long-term effects, and he finds a series of long-term labor market effects that could affect mobility. Most notably for my purposes, he finds that areas with more cotton production in 1879 had less productive agricultural laborers in the middle of the 20^{th} century, due largely to the positive selection of immigrants during the Great Migration. In my framework, this would imply that the transition out of labor-intensive agriculture would be associated with increased mobility through releasing labor to other places and other local sectors as Mundlak suggests, but it provides a very precise, and clearly testable, mechanism.

Other scholars have built on this work, focusing on ecological shocks that forced landowners to modify their strategies, leading to long-term shifts in local markets for agricultural labor. Hornbeck and Naidu (2014) find that the loss of labor due to the displacement of black workers by the Great Missssissippi Flood of 1927 contributed to the mechanization of agriculture in treated counties, as the upward trend in capital improvements was greater in treated counties. These results reinforce the claim that mechanization was labor saving, as well as Whatley's argument that labor intensive southern agriculture delayed these investments, as it took a natural shock to induce them in this case. Land values did not rise in response to these investments, however, which suggests that white landowners were able to mitigate the induced immigration of black agricultural workers, as it implies that they did not expect capital improvements to enhance the value of land, instead betting on the status quo, which worked in their favor. These results focus primarily on the early 20^{th} century, rather than the period I study in this paper, but while they do not provide direct evidence regarding the relationships I study, the broad point – that adoption of new agricultural technology massively displaced labor, and that the laborers they displaced were disproportionately low-skilled, and facing very limited economic prospects – is relevant nonetheless.

6.2.2 A Model of Mobility and Structural Transformation

The model of intergenerational mobility proposed in Nybom and Stuhler (2016) provides clear and concise intuition for the mechanisms that I expect to connect changes in the agricultural sector to intergenerational mobility. In contrast to the original Becker and Tomes (1979) model of mobility, which focuses on developing a tractable model that relates generations in a stable equilibrium, the Nybom and Stuhler (2016) model is designed to account for the immediate and dynamic effects of structural economic change on intergenerational mobility. The core of their model is a pair of equations,

$$y_t = \gamma_t y_{t-1} + \rho' e_t + \sigma_t u_t \tag{6.1}$$

$$e_t = \Lambda_t e_{t-1} + \Phi_t v_t \tag{6.2}$$

, in which y_t and e_t represent individual income and human capital, their lagged analogs represent parental values of the same variables, γ captures the strength of income relationships across generations, roughly equivalent to the standard IGE, ρ represents the returns to human capital, Λ represents the heritability of human capital, Φ and σ represent returns to random human capital and income shocks, and v_t and u_t represent those shocks. We should think of e_t broadly as an index containing a wide variety of skills. This model is a simplification of a three-equation model, in which parental incomes affect children directly and through investments and in human capital, and in which children's human capital is a function of both realized skills and endowments. Simplifying this model by treating γ as an aggregation of direct and investment effects, and ρ as an aggregation of parameters governing returns to both types of human capital, is much more convenient, and does not detract from its utility.

The authors denote the set of parameters facing a child of generation t with ξ_t, which, for our purposes, captures the relevant economic environment, as it allows us to relate structural changes in the economy to intergenerational mobility. They define a 'structural change' as "a permanent change in any of the features in generation $t = T$ such that $\xi_t = \xi_1 \neq \xi_{T+s} = \xi_2 \forall_S \geq 0$". In other words, the kinds of change that occur in one period and then persist indefinitely are the type of changes that I reference when I talk about structural changes, although it is certainly true that other types of changes could exist, and could matter. Nybom and Stuhler (2016) discuss their models implication for changes past the standard parent-child transmission mechanism, and these are important for our analysis, but using this simple case of the model helps to build some microeconomic intuition about the mechanism I propose to connect agricultural change and mobility. I expect that the transition from labor intensive agriculture to capital and technology intensive agriculture changed both ρ and γ for many participants in the agricultural economy. For farm workers who had, or could easily acquire, the technology-biased skills that grew in value in response to the introduction of new technology, ρ rose, and for children whose skills outside the agricultural sector grew in value in part because of the erosion of institutions that would otherwise have kept them on the farm, it rose as well. In contrast, for children who will inherit farmland and agricultural equipment, γ should rise, as innovations increase the land's value even if one's family lacks the liquidity or expertise to take advantage of it by increasing current productivity. A far larger number of children work in the off-farm economy than have parents who own farmland, however, so it is reasonable assume that the effect through ρ, which should increase mobility as I measure it, will dominate, leading to a rise in mobility as agriculture transforms, because this change produces unambiguously higher incomes for the children I have described, compared to what they would have expected before the structural change.

The primary contribution of Nybom and Stuhler (2016) is in predicting effects of structural changes outside this standard parent-child paradigm, and because my interest is in long term trends, this is germane to my analysis as well. The authors show that the IGE may shift multiple times over subsequent generations as a response to a structural change, as it converges to a steady state, that the IGE's response may be non-monotonic over time, and that a change in the strength of one transmission mechanism relative to another (i.e., human capital vs. income) tends to temporarily increase mobility. They apply these results to several scenarios, including the case of changes in returns to certain skills and endowments, the type of structural change most relevant in this context. In this case, they find that the model predicts an increase in mobility in the generation in which the change occurs, but a drop in mobility from that heightened value, eventually converging to a new equilibrium mobility level, lower than the initial condition but higher than the value in the first shocked generation. They argue that this makes sense intuitively, because the structural change allows poorer people with skills to move up in the income distribution, which leaves richer people with relatively less productive skills relatively worse off, which increases mobility by definition, while in later generations, people who moved up in the distribution pass skills onto their children attenuating the original effect.

This nonmonotonicity is interesting in part because it may help to explain the rise and fall in mobility that Hilger (2017a) finds, and which Feigenbaum (2018), Zimmerman (1992), and Solon (1992) suggest, as it provides a compelling reason that this stylized fact could be the result of one coherent process. I do not, of course, expect that the effects of agricultural transformation explain this pattern on their own, as the agricultural sector had already undergone substantial changes by the start of my observations, and because it accounted for a relatively small share of national production. Despite this, however, providing a new partial explanation for these facts impacts our understanding of their policy implications.

6.3 Data

I describe the procedure I use to derive state-level estimates of intergenerational mobility parameters, originally proposed in Hilger (2017a) in Section 3.3. While these data are at the core of this study, and require far more explanation than the other data sources I rely on here, the data I use to measure agricultural change merits discussion as well, as do the descriptive statistics that warrant some of the stylized facts that have helped to motivate the questions I ask here.

6.3.1 Agricultural Data

My agricultural data comes from the United States Census of Agriculture, the International Science and Technology Practice and Policy (INSTEPP) project, and from the United States Department of Agriculture's Economic Research Service (ERS). We use a state-level series on multifactor productivity in agricultural production from INSTEPP as a measure of of productivity, which can be interpreted as being comparable to total-factor productivity of agriculture. This series is measured as an index, with values in 1949 acting as the base year, and it displays substantial cross-sectional and temporal variation. Price-weighted input and output quantities are represented by a chained Fisher Index for each year, state, and quantity bin, using data from a variety of sources, and these quantities are then used to compute the multifactor productivity measure that I use as an indicator of overall agricultural productivity. I am particularly interested in the aggregate labor and capital usage series, and in the multifactor productivity series, which aggregate many types of price-adjusted labor and capital quantities to form their state-year index values.

I also use TFP and input level data compiled by Eldon Ball and coauthors and distributed by the ERS. These data, described in detail in Wang et al. (2015), are measured somewhat differently than the INSTEPP productivity series, as it uses a Törnquvist index rather than a Fisher Index, but they have essentially the same interpretation; for my purposes, the main difference is that the INSTEPP data go back to 1949 while the ERS data series begins in 1960.[2] The input and TFP measures in each sources are positively correlated, but they are not correlated so closely as to be redundant. We include both series, rather than relying on the longer INSTEPP series, primarily for robustness,

Finally, I draw on the U.S. Census of Agriculture for data on land value and crop production data. I compute agricultural land values at the state-year level by taking the mean of value of county-level mean land value, in year 2000 dollar terms, weighted by county agricultural land area, so that average land values in counties with little agricultural land are not overcounted. My state-level average value for an individual state can thus be thought of as the total state-level land value divided by the total state-level land, equivalent to the average value of an acre of land, which is not distorted by the fact that some counties may have very high land values despite having very little land.[3] I also use these data to compute the prevalence of cotton production, necessary to test the Jung (2018) narrative of cotton's long-run influence on labor markets, by dividing the state-level acres of land in cotton production by the state-level number of acres harvested in each year, to compute the percentage of harvested in acres in cotton. I do the same thing to produce the 1879 values, which Jung uses as his measure of historical cotton intensity.

6.3.2 Control Variables

Finally, I rely heavily on Census data, accessed from Ruggles et al. (2017) and the National Historical Geographic Information System (NHGIS) database. I use IPUMS data to compute state level mobility, as discussed in Section 3.3, and I also compute state-level wage income, state black population percentage, agricultural employment percentage, interquartile range of logged wage income, and location

[2] Diewert (1992) compares the two indices in detail, and while he finds that the Fisher index passes more of the tests he proposes than alternative indices, including the Törnquvist index, he also notes that the two indices lead to the same conclusions in most circumstances.

[3] I also compute this measure for farms rather than acres, dividing the total value of farms by the number of farms.

quotients for industry and services from these data. Income and interquartile range serve as key measures of wages and inequality in Hilger (2017a), and I compute them following his formulas, using only individuals between ages 16 and 64, and measure income, which is used to compute the IQR, in constant year 2000 dollars. I measure agricultural employment percentage by dividing state-level number of employees working in the agricultural industry by the total number of individuals aged 16 to 64 with non-missing industry values.

My data are summarized in Table 6.1. The most important takeaway is the fact that I have substantial variation in my variables of primary interest, the IGE estimate, agricultural percentage, wage, and productivity measures. This includes both cross-sectional and intertemporal variation, although in each of these variables except the IGE, the intertemporal variation dominates, despite deflating all money variables to constant year 2000 dollars.

6.3.3 Trends and Descriptive Statistics

Much of this analysis relies on state and region level differences, both within single year cross-sections and in trends across my study period. I summarize several important trends – in mobility slopes, agricultural employment, and agricultural productivity (measured using the INSTEPP data) – in Figures ?? through ??. Figure ?? shows the change in state-level IGE from 1940 to 2000 as a percentage, that is, the values that determine the shading are $1 - \frac{IGE_{2000}}{IGE_{1940}}$, so a large value – corresponding to a darker shade – reflects a large increase in mobility, while a lighter value reflects a smaller increase. Figure ?? maps percentage point changes in agricultural employment share between 1940 and 2000, with darker shades corresponding to smaller absolute changes, and Figure ?? maps the ratio of INSTEPP multifactor productivity in 2000 to its value in 1960. Taken together, these figures show that while mobility grew across the country, it grew the least in the South, where agricultural employment was most prevalent in 1940, but dropped the most by 2000, and it grew the most in the west, where both agricultural employment and productivity changes were less uniform. Apart from the cluster of high productivity gains in the south, INSTEPP trends are fairly evenly distributed.

6.4 Results

6.4.1 Agricultural Change and the Income-Mobility Link

Modeling Strategy

I build my empirical analysis around some more pointed hypotheses, each of which relate to one or both of these hypotheses fairly directly. First, I test the basic hypothesis regarding the relationship between Hilger's state-level mobility measure and the transition away from agriculture. To do this, I replicate Hilger's mobility measure, and follow his approach to studying the effect of historical shifts on mobility, regressing these values[4] on measures of states' speed in transitioning away from agricuture. Because I are following Hilger, I focus on the period from 1930 to 1970. Specifically, this hypothesis states that:

Hypothesis 1: The speed of the drop in agricultural employment accounts for a large portion of the relationship between economic growth and aggregated intergenerational mobility.

The logic behind this claim is, as I suggest above, that a rapid transition away from the agricultural sector should increase mobility, because selecting to leave the farm will generally happen in the presence of better opportunities. Furthermore staying on a parent's farm was likely, absent a very rapid increase in returns to farming, to produce a similar position in the income distribution to that parent during this time period. I test this hypothesis with the following regression model,

[4]That is, mobility statistics at the state-census year level, built from microdata within each of these state-year bins.

$$\rho_{it} = \beta_0 + \overline{income}\beta_1 + agpct\beta_2 + agpct * \overline{income}\beta_3 + X\beta_4 + \epsilon \tag{6.3}$$

where ρ_{it} denotes the mobility parameter in state i and year t, $\overline{income}$ refers to state $i's$ income level, *agpct* denotes the levelof employment in agriculture (as a percentage), and X is a vector of controls. If β_3 is positive and significant, a higher *agpct* value is associated with a stronger marginal effect on income, consistent with our hypothesis that the speed of movement away from agriculture explains the income-mobility relationship. If β_1 is approximately zero, the result reflects our hypothesis more strongly, because the marginal effect of state income level depends on the prevalence of agricultural employment completely, rather than having a main effect and an offset depending on the speed of change.

It is worth noting that this test of our hypothesis should not be interpreted as commentary on whether agriculture causes an individual or a group of individuals to reach a certain point on the income distribution. Rather, it is a test of how a set of population-level descriptive statistics have changed together over a lengthy period, which saw a great deal of change in the US economy. Looking back to a period when I had very little high resolution household survey data – the Panel Study of Income Dynamics (PSID) did not begin until 1968, and the National Longitudinal Survey of Youth did not begin until much later, in 1979 – means using techniques amenable to the kind of data I have, and in this case, that means studying trends in the IGE, itself a descriptive statistic. Hilger (2017a) demonstrates that this approach can be informative, however, by answering questions about linkages between sharp changes and mobility levels in ways that help us understand where mobility comes from, and our investigation of the role of changes in the agricultural sector can generate similar knowledge.

These regressions, which focused on agricultural employment share offer a valuable baseline, but omitting measures of productivity is clearly inappropriate both theoretically and econometrically. The theory discussed above suggests that additional productivity should be associated with increased mobility directly, but that more productive states should see a weaker relationship between income and mobility gains, as well as a weaker inverse relationship between the income effect and agricultural percentage. In other words:

Hypothesis 2: States with more productive agriculture experienced a slower rise in mobility as agricultural employment declined, but more mobility unconditionally.

I test this hypothesis by fitting a modified version of Equation 1, taking the following form:

$$\begin{aligned} \rho_{it} = \beta_0 + \overline{income}\beta_1 + agpct\beta_2 + productivity\beta_3 + agpct * \overline{income}\beta_4 + \\ agpct * productivity\beta_5 + \overline{income} * productivity * \beta_6 + \\ \overline{income} * agpct * productivity * \beta_7 + X\beta_8 + \epsilon, \end{aligned} \tag{6.4}$$

where notation is the same as in Equation 1, except that *productivity* measures agricultural TFP in state i and time t. In this regression, the strength of the relationship between income and mobility depends on the speed of the reduction in agricultural employment, as before, but now that the effect of agricultural employment is moderated by the the *productivity* indicator, through both the two and three way interaction terms. Only β_6 tests the hypothesis, however, because our hypothesis makes a prediction about the marginal effect of mean income - $\beta_1 + agpct * \beta_4 + productivity\beta_6 + agpct * productivity * \beta_7$ - depends on *agpct* only through β_4 and $productivity * \beta_7$. This means that Hypothesis 2 depends on β_7: if it is positive and significant, states with less productive agriculture experienced a stronger relationships between the reduction in agricultural employment and the strength of the income-mobility association. The marginal effect of *productivity* is also interesting, however, because if β_5 is positive, the direct relationship between movement away from agricultural employment is stronger in less productive, also consistent with the claim that differences in agricultural dynamism contribute to the income-mobility relationship.

As in the discussion of my empirical strategy for testing Hypothesis 1, I do not claim to have identified a causal effect. Instead, I set out to study trends in intergenerational mobility, and exploit cross-sectional differences between states and counties, as well as in the obvious temporal dimension, to estimate parameters that reflect these trends. This does not mean, of course, that I can ignore concerns over bias in my parameters of interest caused by endogeneity, but establishing identification with techniques like instrumental variables, difference in differences, matching, etc. is difficult due to a variety of data constraints. My overall objective is to come to a better understanding of the puzzle in Hilger (2017a) to improve inform our perception of mobility trends, and ideally to enhance our understanding of what kinds of social change create mobility, and how they do so. This requires using data with inherent limitations, and being clear about what this study can and cannot accomplish.

Results

These regressions are summarized in Table 6.2, where I present three models, starting with the model proposed in Equation 1, and then adding additional variables and interactions to strengthen the model and consider additional heterogeneity. In Model 1, more income is associated with more mobility, albeit relatively weakly, while more agricultural employment is associated with a large drop in mobility, although this is clearly an overestimate, driven in part by a lack of very high values of agricultural employment in the dataset. The interaction is negative and significant, suggesting that a 10 percentage point change in agricultural employment would lead to an absolute increase in the marginal effect of income of approximately 0.02, a fairly small, but potentially meaningful, amount. Adding basic controls for location, other industries, and percentage black leaves only income significant, but accounting for the possibility that the interaction varies regionally, taking a different value in the south than elsewhere, leads to substantially different results. I include this interaction primarily because the literature on southern agriculture suggests that institutions around agriculture were different in the South than in the rest of the country, independent of contemporaneous crop patterns or other readily-available observables. The maps in Section 6.3 also help to motivate this exercise, because across all three figures, the South seems to follow a particular pattern. It is not entirely clear how the south ought to differ from the rest of the country, but it seems reasonable to allow this particular regional difference in my models based on this trend, rather than allowing state-specific or region specific trends for all possible groups. In these models, the income-employment interaction retains the negative sign and magnitude from Model 1, but the interactions with the south indicator change the interpretation substantially. They suggest that, while the qualitative interpretation of Model 1 remains in Model 3 for the rest of the country, incomes are associated with increased mobility much less than in the rest of the country, and that places in the South have weaker links between increased income and increased mobility, rather than the stronger links the negative sign on the base interaction implies.

In Table 6.3, I add measures of agricultural productivity and input use, which let me test the hypothesis proposed in Equation 2. I begin, in Model 1, by allowing them to enter the model additively, alongside the core income and agricultural variables, and then add controls and interactions, as in Table 6.2. Note that the number of observations has changed between these two tables. This occurs because, these data series only begin in 1949, in the case of the INSTEPP data, and 1960, in the case of the ERS data, so I lose the 1940 cross-section included in the Table 6.2 regressions. The three-way interaction term is insignificant, so I do not find support for my core productivity hypothesis. I do, however, find some support for the hypothesis that states with more productive agriculture benefitted more from income independent of agricultural percentage, as the coefficient on the income-TFP interaction is negative and significant. Because the TFP measure is standardized, this figure implies that a 1 SD increase in productivity is associated with an increase of about 1/3 in the absolute value of the income coefficient, a fairly large increase.[5] At the same time, more productivity is strongly associated with less mobility.

These coefficients are informative insofar as they reflect variation in the marginal effect of income, but since the marginal effect varies across values of two other variables in the three-way interaction models, the extent to which this is the case is not immediately clear from the coefficients I report. In Figures ?? through 6.6, I summarize the marginal effects of income by evaluating them at the minimum through the maximum values of interaction variables, using parameters from the second and third models in Tables 6.2 and 6.3. These results do not substantially contradict those presented in the corresponding tables, but they do provide context that is not obvious from simply interpreting the coefficients. Figure ?? and the first panel of Figure 6.5 show that in a pooled model, and in the non-south observations of the south indicator interaction model (the third column of Table 6.2), the marginal effect of income remains negative and significant across values of agricultural employment, although this variation is not significant across values of agricultural employment share in either case. The fact that this income relationship is sufficiently weaker in the South that it loses significance at all points is interesting, however, because while it does not provide a conclusive point estimate, it does establish a difference between the marginal effect between the two groups, suggesting that the south benefitted substantially less from income growth than the rest of the country.[6] The second panel of Figure 6.6 shows the marginal effect of income across values of agricultural employment and agricultural TFP, which constitutes weak evidence that agricultural employment strengthens the income-mobility link in relatively unproductive states, while it appears to have no effect in highly productive states. As in the other marginal

[5]I standardize primarily to ensure that I can compute the standard errors of marginal effects; in the unstandardized case, covariances between parameters are often negative, and in computing the weighted variances that determine the standard errors of marginal effects, this can lead to negative variance estimates.

[6]This interpretation would also arise from evaluating the marginal effects implied by the coefficients themselves, but it need not be the case that composite marginal effects maintain the sign, significance, or magnitude of their constituent coefficients.

effect figures, however, this result is tempered by the variance of the marginal effects: at the minimum TFP value, all marginal effects are insignificant (although in the remainder of the evaluation points in this figure, the marginal effect is significant in over 90%), but the vast majority of confidence intervals are indistinguishable, so it is unclear whether the variation in marginal effects that the interaction term coefficients imply reflects true variation or simply random noise.

So far, the primary source of variation in the marginal effect is between spatial regions, which the literature suggests display meaningful institutional differences that may not be captured in my measures of agricultural development. I examine this difference further in Tables 6.7 and 6.8, based on three-way interaction models equivalent to the models in the third columns of Table 6.3 in all ways except for the structure of the interaction terms.[7] In Figure 6.7, we see that more productive areas display a stronger mobility-income association in the South and elsewhere, and that the trend is more or less the same across regions. This eliminates one explanation for the previous result, i.e. that particular southern institutions hindered gains from TFP growth.

I continue to evaluate this possibility by interacting income and the south indicator with the logged value of agricultural land acres, measured in constant year 2000 dollars. I use these measures as proxies for the degree of capital intensity and inequality in agricultural capital (between county averages) respectively, following Hornbeck and Naidu (2014) and Barkley (1990). This is helpful because the direct measure of TFP, and of labor and capital use, adjust for labor and capital productivity, so it is possible that, in the presence of the kind of skill-biased technical change that could enhance mobility, a change in productivity would not cause a change in these labor or capital measures (although it would shift 'real' TFP). In these models, fitted using the same controls included in the third column of Table 6.3 and summa-

[7] I compute the standard errors for these models, and for models I present in throughout the rest of this paper, by using a residual bootstrap to compute standard errors for the marginal effect of income at each evaluation point, rather than computing them analytically as in previous figures. This is necessary because the covariance terms dominate in the analytical formula, resulting in negative variances, which imply complex standard errors.

rized in Figure 6.8, the marginal effect of income is consistently negative outside the South, but insignificant for the upper half of the land value range in the south, with a steeper increase in the South than elsewhere. Insofar as this land value measure captures capital intensity in the agricultural sector, this result implies that Southern capital is associated with larger decreases in the mobility-income relationship, to the point that the relationship disappears entirely in states with substantial land investments. This provides some support for my initial hypothesis, but given the lack of support for alternative versions, and the fact that the mechanism in this case is unclear, more investigation is needed. One possibility is that places with higher values are also more unequal in the distribution of these values, for example with many productive plantations in one part of the state, and many smallholders elsewhere, in which case the marginal effect of income would rise in the Gini coefficient of county-level land values, but in the third and fourth panel of Figure 6.8, the opposite occurs. This very likely has to do with the fact that Gini coefficients and mean values are in fact positively correlated, so we are left without a staightforward explanation for this pattern of evidence.

6.4.2 Cotton Production and Mobility

The results in the previous section are more or less inconclusive; the income slope remains negative and significant across a wide array of specifications, reinforcing the robustness of the result in Hilger (2017a), and it varies very little across attributes of state-level agricultural economies. The factor that seems to matter most is region, i.e. whether a state is in the south or not, and much of the literature suggests that this could have to do with the institutions surrounding southern agriculture, particularly cotton production. To test this hypothesis – that more cotton production is associated with a weaker income-mobility relationship and with less mobility overall – I begin by following Jung (2018) in estimating models including the 1879 prevalence of cotton production, measured as the percentage of harvested acres planted in cotton. This approach is reasonable because Jung's argument, and my justification for adapting it here, hinges on long-lasting institutional factors – such as the rigid labor contracts Whatley, Hornbeck and Naidu, and others have studied – and so looking back many years into the past is helpful conceptually, in addition to reducing the likelihood that outcome, controls, and this variable of interest are codetermined by later unobserved economic changes. Figure 6.4 provides some additional motivation for examining this relationship; the first panel shows county-level mobility slopes for children born in the late 1970's and early 1980's from Chetty et al. (2014b) and the second shows cotton shares from Jung (2018), and the overlap – particularly in the Mississippi Delta region – is substantial.

Jung has good reasons to choose a specific year for his outcome, but this choice is less clear in my application, and so I convert my outcome – the mobility slope – and controls to 1940 through 2000 trends, so that the coefficient on the 1879 cotton percentage variable is interpreted as the change in the 1940 to 2000 mobility trend (negative for all states) for a change in cotton percent, measured decimally. Furthermore, Jung's variable of interest is observed at the county level within states that grew some amount of cotton, but in my case, omitting non-cotton states would yield too few observations. This threatens my analysis primarily because my cotton intensity measured is left censored in most cases: only 13 states grew cotton in 1879, and this number is fairly cosistent over time. To mitigate this concern, I apply the Heckman correction procedure (Heckman, 1979), using United Nations' Food and Agricultural Organization (FAO) data on the agro-ecological suitability of land for cotton production as the first-stage instrument, following Jung (2018), measured as an index. This is helpful because this variable is uncorrelated with institutional and demographic factors that could affect long-term development, so the component of cotton selection correlated with those factors is not captured in the first stage. The suitability measure is excludable from the second stage regression because it is not correlated with the outcome and because there is no reason to suspect that it is necessary to adjust for it in theory or the literature. I present two models in Table

6.4: a baseline model in which I regress the slope trend on cotton percentage and controls, and a model in which I control for the Inverse Mills Ratio from the first-stage regression of a cotton production indicator on the FAO Suitability Index value. These models suggest that more cotton production is associated with slower mobility growth in the 20^{th} century, and that the association is strong; at average values of mobility and cotton percentage, these coefficients imply that cotton production reduces the mobility trend by about 30%.

While following Jung directly in estimating the long-term impacts of 19^{th} century cotton production yields interesting results, I am only able to marshall a small fraction of my total data in answering this question. I have limited cross-sectional variation in any given year due to the lack of county-level mobility data, so unless I select a specific year, I need to aggregate already limited data. To partially mitigate this limitation, I fit a series of models that are similar to the models in the previous exercise, using 10-year lags of cotton percentage as my variable of interest. I expect that higher cotton intensity will be associated with less mobility, and in models with interaction effects, will suppress the the relationship between income and mobility. As in the 1879 case, I correct for left-censoring of cotton share by applying the Heckman correction, but using the suitability index interacted with year indicators instead of the suitability index alone, in order to both adjust for time trends and ensure that linear predictors differ across years, as the index does not.

I report the coefficients in Table 6.5, and the marginal effects of both income and cotton share in Figures 6.10 and 6.11. In the first two specifications, cotton is not associated with mobility levels, and marginal effects, but in the third, in which I allow three-way linear interactions between income, agricultural share, and cotton share, several relevant coefficients and marginal effects become significant and large. This specification is necessary because of the relationship between this exercise and the south interaction in the previous section: evaluating the marginal effect of average income is necessary if I want to comment on the original Hilger finding, and allowing both income and agriculture effects to vary in cotton is necessary if I want to evaluate the claim that institutions surrounding cotton drive the regional pattern apparent in Figure 6.5. Figure 6.10 shows that marginal effects of income do not vary across agricultural shares, but that they vary across large shifts in cotton shares, as the marginal effects evaluated at high values fall outside the confidence intervals of marginal effects evaluated at lower levels. Similarly, marginal effects of cotton are insignificant across most of the agricultural share distribution, as Figure 6.11 demonstrates, but for high values of agricultural percentage and income, larger cotton shares are associated with higher slopes, and thus with less mobility. Taken together, these results are consistent with the institutional explanation: higher cotton shares mitigate the mobility gains from income, albeit with fairly wide confidence bounds, and in places with particularly large agricultural share, where the influence of a particular crop is more likely to have a substantive impact on the broader economy. In the context of theory and empirical evidence regarding cotton-related institutions, these results suggest that cotton production helped to limit gains in mobility from rising incomes primarily through its institutional legacy, as the strength of the general agriculture-mobility relationship is fairly limited across a variety of models.

6.4.3 Evidence from Inter-Group differences

Black-White Differences

The literature's emphasis on low-skilled agricultural labor suggests that black workers would have fared particularly badly in states with a stronger institutional legacy from cotton production, that is, their mobility should be lower in places with larger agricultural and cotton shares. I test this hypothesis by computing separate mobility statistics for black and white populations, and fitting a series of regressions, similar to those discussed so far in this section, to the difference between the two. I compute the difference as white mobility less black mobility, so positive values of marginal effects imply divergence of mobility levels, while negative values imply convergences, rather than the more straightforward interpretation of more or less mobility thus far. I expect that larger cotton shares will be associated with larger gaps, and will mitigate positive signs on income, and enhance negative signs on agricultural share. Note that these signs have inverted implications for income convergence: if slopes are converging by white slopes becoming larger than black slopes by less, black mobility is decreasing relative to white mobility, so the speed of income convergence between the two groups is decreasing.

I examine these differences by repeating much of my analysis, using this difference as the dependent variable. These results are presented, in somewhat abridged form, in Tables 6.6 and 6.7 and in Figures ?? through 6.14. Across the specifications, the coefficient on income is positive, implying that as incomes rose, black mobility rose more quickly than white mobility, consistent with Hilger's findings, but while the marginal effects of income generally take the same signs as this base coefficient, they very relatively little over the variables I expected to mediate them.

Capturing trends in convergence between black and white mobility levels is informative, but it is not conducive to estimating cross-population differences in marginal effects; I can easily estimate the marginal effect of income on the gap, but not the gap in the marginal effect of income on mobility. The possibility that black and white mobility slopes could respond differently to change is, however, consistent with the history literature, with my hypotheses regarding the opportunities afforded by smallholder farming, and with the mobility literature relying on more recent data, e.g. Hertz (2005). To estimate cross-group differences in these marginal effects, I repeat my analysis on each measure at the state-year level, and I present the results in Figures 6.15 through 6.19.[8] Black mobility is, on average, lower across the sample, but while regressions on the size of the black-white gap generate tenuous results, evaluating marginal effects of income separately shows that black and white populations exhibit different mobility responses to income levels. Figure 6.15 suggests that the marginal effect of income falls, in absolute terms, as agricultural share rises for blacks, while it is consistently negative and steady for whites, although these black

[8] I do not present tables of regression coefficients because this exercise is solely concerned with differences in marginal effects across the two supgroups.

estimates are very noisy, and the 95% confidence intervals do not rule out the white pattern. The differences between the South and the rest of the country add a bit more context: in the South, the marginal effects are insignificant across all values of agricultural share, but elsewhere, black Americans in more agricultural states see smaller mobility gains from income, while white mobility gains from income either rise or remain steady as agricultural share increases. This is not quite consistent with my hypotheses, as the story about southern institutions would suggest that black mobility would benefit less than the rest of the country, and I instead find that the relationship is insignificant, but these results do suggest a substantive cross-group difference. The cross-group returns to productivity growth, summarized in Figures 6.17 and 6.18, display larger differences: as the INSTEPP productivity metric rises from its 20^{th} to its 80^{th} percentile, the marginal effect at each level of agricultural share becomes more positive, although this leads to mostly insignificant results in the fourth panel of Figure 6.17, and the slope across levels of agricultural share becomes steeper. In contrast, the marginal effects for the white population become more negative and less steep in agricultural share as productivity rising, suggesting that overall, increased productivity is associated with more mobility for white workers across levels of agricultural employment compared to black workers, whose mobility returns to income seem to decline as productivity rises. The marginal effects of income for the black population becomes insignificant at high values, so it is important to not rely heavily on the point estimates in Figure 6.17, but there is clearly a substantial difference between black and white marginal effect levels and patterns, consistent with the narrative of black smallholders having trouble accessing the improvements that drove increased productivity. Finally, I summarize the marginal effects of income by cotton share in Figure 6.19, which shows very limited differences between the black and white populations.[9]

[9] I omit the 40^{th} and 60^{th} percentiles presented in previous four-panel figures because they show virtually no variation, so presenting only the endpoints makes conveys the same information in less space.

Were farmers less mobile?

In the previous exercises, I have explored the relationship between state-year mobility levels as a function of aggregate data on state economies. I do this because I can only observe mobility at the state level, because the one variable Hilger finds to consistently explain mobility changes, income level, falls in this category, and because I am concerned with changes in the agricultural economy, which I can capture in aggregates. The Hilger mobility measures also allow me to capture some individual variation by breaking down state-year mobility statistics by individual characteristics, e.g. racial, gender, or immigration status groups, however. Because I am interested in agriculture, I exploit this option to test the hypotheses that farmers are less mobile, and that farmers were mobile in more productive states.

I do this by computing mobility statistics by state-year combinations for farmers and non-farmers separately, where I define farmers as people working in the agricultural industry in adulthood[10], and computing regressions similar to those discussed in the previous section. As in that section, this procedure also provides a gap between the two mobility levels, which helps to motivate and contextualize this discussion. In this case, however, the differences in mobility slopes are small: as Figure ?? shows, the majority of differences are between -0.1 and 0.1, and the mean of the differences is 0.012. A two-sample t-test fails to reject the null hypothesis of no difference in means between the two distributions, so analyzing differences, as I do in the previous section, is unlikely to yield interesting or informative results. Consequently, I move directly into between-group comparisons of marginal effect patterns, with the same interpretation as the black-white marginal effects in the previous section. As above, I present only figures of marginal effects by group, rather than beginning with a table presenting the coefficients that determine these marginal effects, as those are of limited interest.

[10]IPUMS also captures occupation, but using occupation rather than industry makes very little difference – the lowest percentage of individuals 16 to 64 with agricultural occupations to work in the agricultural industry is approximately 98%, and in most years, it is 100%.

Figure 6.20 presents the marginal effects of income on mobility slopes for agricultural and nonagricultural workers. While the confidence bounds vary across agricultural shares by only small amounts, the trend does suggest that mobility gains from income for agricultural workers are nonincreasing in agricultural share, while the gains for non-agricultural workers are nondecreasing in agricultural share. At each evaluation point, however, the confidence intervals overlap, so while I detect a difference in the pattern of marginal effects in terms of agricultural share, these results are inconclusive with regard to differences in the income-mobility relationship given agricultural share, and to overall levels of mobility. Differences between the South and the rest of the country – summarized in Figure 6.21 – present a similarly opaque picture, but lead to similar provisional conclusions: in the South, the income-mobility relationship is insignificant for both agricultural and nonagricultural workers, but elsewhere, the non-agricultural marginal effects are consistently more negative and downward sloping in agricultural share, although many confidence intervals overlap between the groups.

Interpreting regressions on these slopes and gaps is complicated by the fact that the groups are endogenous to outcome. In many cases, very likely including agriculture, occupation is codetermined with educational attainment and income, so between-occupation mobility differences are more likely to be the result of these human capital factors, or of preferences, than mobility differences between racial, ethnic, or sex groups. I lack the microdata necessary to evaluate the size of this bias, or to develop a quasi-experiment to mitigate it, but the IPUMS data provide several data series that allow me to execute a series of placebo tests. I rely on IPUMS data on occupational income level and occupational prestige indices, repeating the agricultural vs. non-agricultural analysis I discuss here with each variable, computing mobility statistics for the group encompassing agriculture with all other groups. If results for the comparison group display the same pattern as the results for the agricultural group, it seems likely that the human capital and compensation components that these classifications have in common are driving my agricultural results, but if they display significant differences, my results are credible.

In this test, I rely heavily on an IPUMS series of "Occupational Income Scores", which captures the median income of workers in each occupation, by year. This is not an ideal measure of income in general, but it does provide a reasonable metric for classifying farmers' peers in the income distribution. I define a comparison group based on this metric, classifying all workers outside of agriculture whose occupational income scores fall within 10 percentage points of the occupational income score of farmers in a given year as the comparison group.[11] I also define two additional groups for comparison with the non-agricultural group in my original analysis: the 10% of the distribution immediately above the highest occupational income score that falls in the comparison group, and the entirety of the occupational income score distribution that falls above the maximum of the comparison group distribution. I compute the mobility statistics for each group, by state-year bins, and then repeat my previous cross-group analyses to compare the comparison and placebo groups. If the mobility statistics and marginal effects of income, and of other key variables, on the mobility statistics display the same patterns in the comparison group, which excludes farmers, as in the agricultural group in the main sample, it would provide a good reason to believe that my results for the agricultural group are driven by the fact that that group has lower incomes, rather than by its participation in agriculutural labor per se. Similarly, if the next 10% of the distribution is substantially different from the comparison group in mobility statistics or marginal effect patterns, my original analysis is threatened by the possibility that apparent agriculture vs. non-agriculture differences are driven by idiosyncracies in this region of the income distribution, as it would suggest that moving just slightly outside of the comparison region changes mobility substantially. This is particularly true if this placebo region

[11]The 10% band is somewhat arbitrary, as I am not aware of a more systematic method for segmenting the income distribution for occupational groups. It is attractive because it is a relatively small region of the distribution that nevertheless provides reasonably precise estimates of my mobility statistics, as it encompasses large number of individuals in each year.

displays patterns similar to the nonagricultural mobility patterns I estimate above, as it would suggest that outside the comparison region, mobility patterns converge to one 'normal' state, and that the income region farmers happen to fall into happens to be an idiosyncratic exception, rather than something substantively different, as I have argued above.

I begin with a t-test comparing the agricultural slopes to the comparison group, (i.e., the 10% of the IPUMS sample in each year with incomes within five percentage points of the agricultural occupational income) which fails to reject the null of no difference in means between the groups, but as in other exercises, I am concerned with groups' responses to changes, rather than mean mobility levels alone. I then repeat the analyses I present in Figures 6.20 through 6.23 in Figures 6.25 through 6.28. The imprecision of my marginal effect estimates for the agricultural group complicates comparison to a degree, but these results suggest that agricultural mobility is distinct from the mobility of the occupation's peers in income terms. I reach this conclusion primarily because the comparison group marginal effects are larger and in many cases significant at the 95% confidence level, suggesting that non-agricultural workers in occupations with average incomes similar to agricultural workers enjoyed a large increase in mobility from rising incomes at each level of agricultural employment share. The 'placebo' group also consistently displays a modest downward shift in the slope and instercept of the line connecting marginal effect estimates, relative to the comparison group, suggesting that the next 10% of the distribution is more mobile than the comparison decile, but that it does not fully converge to the overall non-agricultural pattern as it would if the comparison decile results were totally idiosyncratic. Interpreting the differences broken down by geographic location and productivity levels is more difficult due to the lack of precision, as the marginal effects for the comparison decile are consistently insignificant, but the pattern of the placebo decile differing from the comparison group only modestly, and in the direction of benefitting more from rising incomes, is consistent with the conclusion that my agricultural comparison results are not driven by farmers falling on the left tail of the income distribution. The endogeneity of selection into agricultural employment remains concerning, but these results provide some evidence that the differences, as limited as they often are, that I find between agricultural and non-agricultural mobility responses to rising incomes and other changes are not driven primarily by the position in the income distribution farmers happen to occupy. This does not prove that these results are driven by the evolution of the agricultural sectors, as I and others have contended, but it does show that accounting for one obvious threat to their credibility does not invalidate them.

6.5 Discussion

In the previous section, I outline the results of several exercises devised to determine the role of agricultural changes in determining patterns of economic mobility across time and space in the 20^{th} century United States. Because I am motivated by the Hilger (2017a) finding that mean household income is the most robust correlate of rising mobility, I focus on variation in the marginal effect of income across a variety of agricultural variables, controlling for demographic and economic factors that also changed substantially over this time period. In many cases, the results from these exercises are imprecise and inconclusive, but taken together, these results are consistent with several stylized facts about the relationship between agricultural development and trends in intergenerational mobility.

First, accounting for changes in the agricultural system – measured through agricultural labor share, the INSTEPP productivity index, and both historical and contemporaneous cotton shares – does not explain away the mobility-income association Hilger finds, reinforcing the robustness of that somewhat puzzling result. Across specifications, the relationship between mobility slopes and income is negative, significant, and comparable in magnitude to the results Hilger reports, and while the marginal effects vary substantially in size and significance, they are generally fairly similar to the Hilger estimates. This pattern of results refutes a particularly strong version of my hypothesis, i.e. that changes in agriculture – which I expect to be correlated with both rising incomes and increased mobility – can explain away the relationship Hilger finds as the result of an unobserved 'lurking variable'. This is important because it places an upper bound on the explanatory power of my analysis with regard to explaining mobility trends themselves, but it leaves room for explanations of differences in the strength of the marginal effects of income.

The original marginal effects are remarkably robust to the inclusion of agricultural variables, however. The income marginal effects maintain their negative sign and significance across a variety of specifications, and more agricultural employment, along with less productivity, is associated with stronger marginal effects, directly contradicting my hypothesis. This effect varies between the south and the rest of the country, with the south effect indistinguishable from zero and the rest of the country displaying a pattern of large (in absolute value) marginal effects of income at higher levels of agricultural share, but this difference does not negate the contradiction of the hypothesis. This pattern of results would be consistent with more agricultural states

I expand on this ambiguous set of results by trying to explain the south versus non-south difference in the marginal effect of income across agricultural share, considering the possibility that cotton production confounds the relationship, as it was rare to nonexistent outside the south in this period. I also consider the possibility that black and white workers could experience differential mobility gains from rising mean incomes and from agricultural changes. After adjusting for selection into cotton production, I find that states with more cotton production saw a much slower increase in mobility between 1940 and 2000, and that states with relatively high incomes and agricultural labor shares had large and positive, if imprecise, marginal effects of cotton share, implying that an increase in cotton share is associated with a decrease in mobility. At average values of income and cotton share, this predicted decrease would counteract mobility gains from income in cotton states. These differences in marginal effects across levels of cotton share raise the issue of black-white differences in responses to agricultural development, as there is evidence that labor intensive cotton agriculture was more persistent in predominantly black labor markets, and was disproportionately bad for these workers. The marginal effects for both blacks and whites in the south are measured with very low precision, and are thus always insignificant, but their contrast with the significant and negative effects outside the south is instructive. As Figure 6.16 shows, blacks outside the South benefitted from income more than whites at low levels of agricultural employment, but much less so at higher levels, suggesting that the transition away from agriculture increased black mobility. Similarly, my results suggest that individuals working in agriculture saw less of a mobility increase from rising mean incomes, and that these workers gained less in more productive states, while non-agricultural workers gained more mobility from rising incomes in more productive states at all levels of agricultural employment. Taken together, these results suggest that trends away from labor intensive agriculture, along with the proliferation of off-farm opportunities, helped to strengthen the relationship between income levels and intergenerational mobility, but that this effect was weaker for those who remained in agriculture, in the south, and for black workers. The cotton results, combined with the economic history literature on the topic, suggest that the legacy of cotton's long-term effects on labor market institution contributed to this dampening, but without finer mobility data that can be matched to finer variation in cotton intensity, as in Hornbeck and Naidu (2014) and Jung (2018), I lack the data to evaluate this line of argument with confidence.

6.6 Conclusion

In this paper, I have examined the connection between the structural transformation of agriculture and an apparent link between rising incomes and rising income mobility in the 20^{th} century United States. I find that while these agricultural changes cannot explain away this relationship, even in the southeast, where agriculture remained labor intensive through the 1940's and mobility and incomes remain

remain relatively low to this day, variation in agricultural labor shares and productivity affect the strength of that relationship, which in turn varies between black and white workers, agricultural and nonagricultural workers, and across values of cotton production share. These results suggest that the structural transformation of agriculture did contribute to rising mobility in the 20^{th} century, but that this had more to do with changing labor markets in the south than with gains from increased productivity being widely distributed. In the context of the broader mobility literature, this suggests that discussions of declining mobility in recent decades should account for the possiblity that mobility was high in earlier decades, as Nybom and Stuhler suggest, due to rapid economic change, and has declined in part because of a slowdown in those changes.

Table 6.1.: Descriptive statistics of key variables and controls

	Minimum	Mean	Maximum	SD
IGE	0.06	0.29	0.69	0.14
Ag. percent	0.00	0.12	0.66	0.13
Industrial LQ	0.13	0.86	2.13	0.40
Service LQ	0.57	0.98	1.64	0.14
Black percent	0.00	0.10	0.71	0.13
IQR of logged wage	0.00	2.51	8.26	2.20
Wage (constant 2000 $)	4051.49	30234.39	63604.22	12526.22
INSTEPP TFP	104.33	192.40	445.32	65.61
ERS TFP	0.33	0.78	1.69	0.26
ERS capital index	0.02	1.88	9.41	1.68
ERS labor index	0.03	2.84	12.59	2.56
INSTEPP labor index	17.55	51.08	97.18	19.16
INSTEPP capital index	34.63	98.36	222.33	28.12
IQR of county land values (state level)	35.92	799.82	3936.29	685.66
Average land value (constant 2000 $)	59.13	1373.05	6362.92	1144.13
Cotton percentage	0.00	0.05	0.61	0.11
Percent population in nonmetro areas	0.00	0.45	1.00	0.29

Table 6.2.: Income and agricultural employment models

	Dependent variable:		
	slope		
	(1)	(2)	(3)
Mean income (logged)	-0.215^{***}	-0.202^{***}	-0.195^{***}
	(0.017)	(0.018)	(0.019)
Ag. employment %	0.016^{**}	0.007	0.021^{**}
	(0.008)	(0.008)	(0.010)
Industrial LQ		0.017	0.018
		(0.018)	(0.018)
Service LQ		0.080	0.083^{*}
		(0.050)	(0.050)
Black %		0.150^{***}	0.150^{***}
		(0.051)	(0.051)
South indicator		0.037^{**}	-1.999^{**}
		(0.017)	(0.954)
Mean income (logged)×Ag. employment %	-0.002^{**}	-0.001	-0.002^{**}
	(0.001)	(0.001)	(0.001)
Mean income (logged)×South indicator			0.190^{**}
			(0.090)
Ag. employment×South indicator			-0.049^{***}
			(0.019)
Mean income (logged)×Ag. employment ×South indicator			0.007^{***}
			(0.002)
Constant	2.514^{***}	2.247^{***}	2.173^{***}
	(0.179)	(0.230)	(0.234)
Observations	302	302	302
R^2	0.592	0.638	0.653
Adjusted R^2	0.588	0.629	0.641

Note: *p<0.1; **p<0.05; ***p<0.01

Table 6.3.: Income, productivity, and agricultural employment models

	Dependent variable:		
	slope		
	(1)	(2)	(3)
Mean income (logged)	−0.249***	−0.167***	−0.180***
	(0.033)	(0.031)	(0.031)
Ag. employment %	0.144***	0.073***	0.051
	(0.029)	(0.026)	(0.054)
Industrial LQ		0.031	0.032
		(0.021)	(0.020)
Service LQ		0.026	0.048
		(0.058)	(0.058)
Black %		0.374***	0.357***
		(0.065)	(0.065)
South indicator		0.041**	0.035*
		(0.018)	(0.018)
INSTEPP TFP (standardized)	0.032***	0.038***	0.611*
	(0.008)	(0.007)	(0.334)
INSTEPP Labor Index	−0.0004	0.001**	0.0005
	(0.0005)	(0.0005)	(0.0005)
INSTEPP Capital Index	0.0003	−0.0004	−0.0002
	(0.0003)	(0.0003)	(0.0003)
Mean income (logged)×Ag. employment %	−0.015***	−0.008***	−0.006
	(0.003)	(0.003)	(0.005)
Mean income (logged)×INSTEPP TFP			−0.056*
			(0.032)
Ag. employment×INSTEPP TFP			−0.047
			(0.062)
Mean income (logged)×Ag. employment %×			
INSTEPP TFP			0.005
			(0.006)
Constant	2.918***	1.927***	2.045***
	(0.352)	(0.332)	(0.338)
Observations	240	240	240
R^2	0.499	0.640	0.663
Adjusted R^2	0.487	0.625	0.644
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 6.4.: Mobility trends and historical cotton intensity

	Dependent variable:	
	slope	
	(1)	(2)
Cotton percent 1879	0.343**	0.319**
	(0.129)	(0.126)
Service LQ	0.066	0.098
	(0.178)	(0.174)
Agricultural %	0.006	0.005
	(0.004)	(0.004)
Industry LQ	−0.051	−0.108
	(0.091)	(0.093)
Black %	−0.061	−0.098
	(0.299)	(0.291)
Mean income (logged)	0.202	0.173
	(0.122)	(0.119)
Inverse Mills Ratio		−0.072*
		(0.038)
Constant	−0.528***	−0.401**
	(0.156)	(0.165)
Observations	49	49
R^2	0.330	0.386
Adjusted R^2	0.235	0.281
Note:		*p<0.1; **p<0.05; ***p<0.01
		All controls measured in 2000 less 1940 trends.

Table 6.5.: Mobility and 20^{th} century cotton production

	Dependent variable:		
	slope		
	(1)	(2)	(3)
Mean income (logged)	-0.201^{***}	-0.202^{***}	-0.201^{***}
	(0.018)	(0.019)	(0.019)
Agricultural employment %	0.008	-0.001	0.020^{**}
	(0.008)	(0.002)	(0.010)
Cotton share	-0.024	-0.618	-5.266^{*}
	(0.057)	(0.614)	(2.817)
Service LQ	0.081	0.065	0.051
	(0.050)	(0.054)	(0.055)
Industrial LQ	0.016	0.016	-0.0003
	(0.018)	(0.019)	(0.020)
Black %	0.152^{***}	0.177^{***}	0.141^{***}
	(0.051)	(0.052)	(0.053)
South indicator	0.040^{**}	0.032^{*}	0.041^{**}
	(0.018)	(0.019)	(0.019)
Mean income (logged)×Agricultural employment %	-0.001		-0.002^{**}
	(0.001)		(0.001)
Mean income (logged)×Cotton share		0.063	0.489^{*}
		(0.063)	(0.265)
Agricultural employment %×Cotton share			-0.094^{**}
			(0.047)
Mean income (logged)×Agricultural employment % ×Cotton share			0.013^{**}
			(0.006)
Inverse Mills Ratio		0.001	0.001
		(0.012)	(0.012)
Constant	2.243^{***}	2.258^{***}	2.299^{***}
	(0.230)	(0.242)	(0.243)
Observations	302	294	294
R^2	0.638	0.623	0.634
Adjusted R^2	0.628	0.611	0.619
Note:			$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Table 6.6.: White-Black slope difference models

	Dependent variable:		
	slope.diff		
	(1)	(2)	(3)
Mean income (logged)	0.139***	0.156**	0.187***
	(0.041)	(0.077)	(0.049)
Agricultural employment %	0.082***	0.390**	0.198***
	(0.021)	(0.170)	(0.052)
INSTEPP TFP		−1.305	
		(0.879)	
Service LQ	−0.279**	−0.163	−0.317***
	(0.112)	(0.131)	(0.110)
Industrial LQ	−0.077*	−0.029	−0.091**
	(0.043)	(0.046)	(0.042)
Black population %	0.224**	0.119	0.161
	(0.103)	(0.151)	(0.102)
South indicator	0.010	0.035	−2.009
	(0.031)	(0.036)	(1.866)
Mean income (logged)×Agricultural employment %	−0.009***	−0.040**	−0.021***
	(0.002)	(0.016)	(0.005)
Mean income (logged)×INSTEPP TFP		0.118	
		(0.083)	
Agricultural employment %×INSTEPP TFP		−0.044	
		(0.200)	
Mean income (logged)×Agricultural employment %:INSTEPP TFP		0.006	
		(0.019)	
Mean income (logged)×South indicator			0.178
			(0.176)
Agricultural employment %×South indicator			−0.221***
			(0.061)
Mean income (logged)×Agricultural employment %×South indicator			0.026***
			(0.007)
Constant	−1.128**	−1.428*	−1.539***
	(0.513)	(0.815)	(0.578)
Observations	210	187	210
R^2	0.213	0.245	0.276
Adjusted R^2	0.185	0.198	0.240
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table 6.7.: White-Black slope difference cotton share models

	Dependent variable:	
	Black/white slope difference	
	(1)	(2)
Mean income (logged)	0.110**	0.139***
	(0.043)	(0.046)
Agricultural employment %	−0.011**	0.104**
	(0.005)	(0.041)
Cotton share	6.449***	−6.664
	(1.746)	(6.563)
Service LQ	0.081	−0.322***
	(0.115)	(0.114)
Industrial LQ	−0.067	−0.103**
	(0.045)	(0.045)
Black population %	0.225**	0.161
	(0.104)	(0.105)
South indicator	−0.007	0.010
	(0.035)	(0.035)
Mean income (logged)×Agricultural employment %		−0.012***
		(0.004)
Mean income (logged)×Cotton share	−0.623***	0.580
	(0.173)	(0.614)
Agricultural employment %×Cotton share		−0.400***
		(0.127)
Mean income (logged)×Agricultural employment %×Cotton share		0.052***
		(0.017)
IMR	−0.005	−0.005
	(0.006)	(0.006)
Constant	−0.834	−1.015*
	(0.539)	(0.547)
Observations	208	208
R^2	0.217	0.264
Adjusted R^2	0.182	0.219
Note:	*p<0.1; **p<0.05; ***p<0.01	

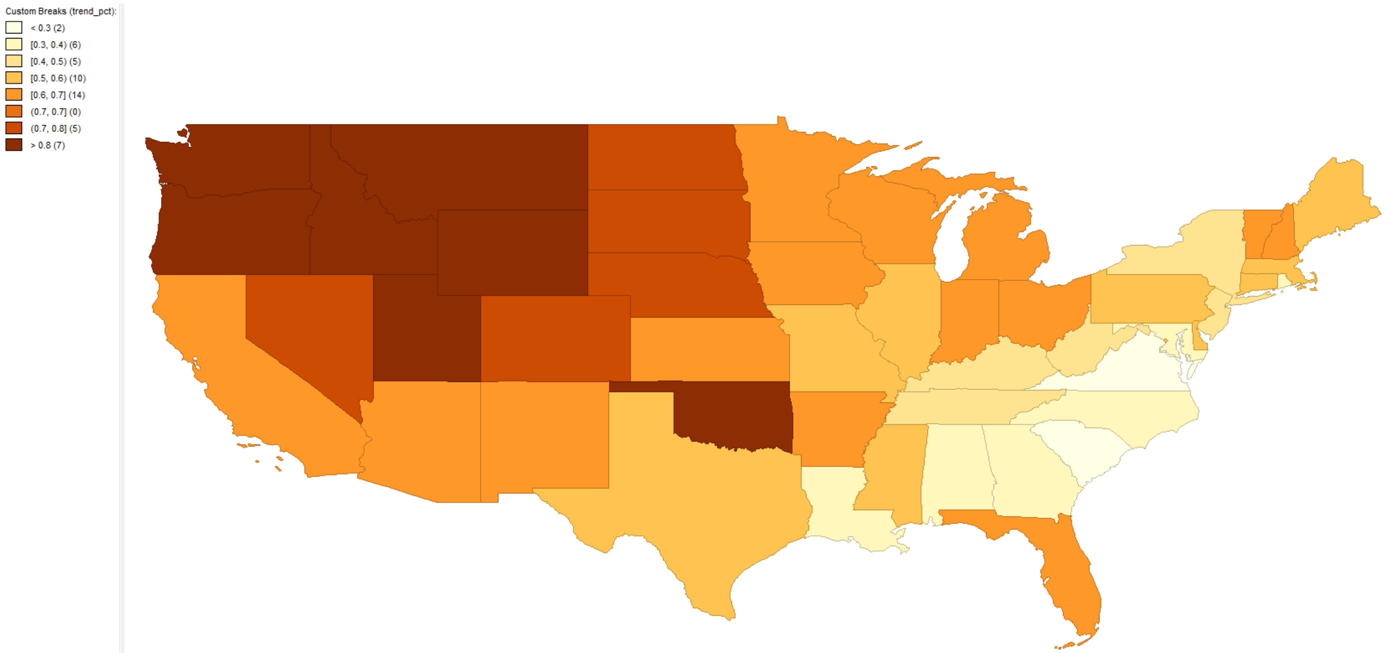

Figure 6.1.: Mobility trends by state, drop in IGE as percentage of 1940 value

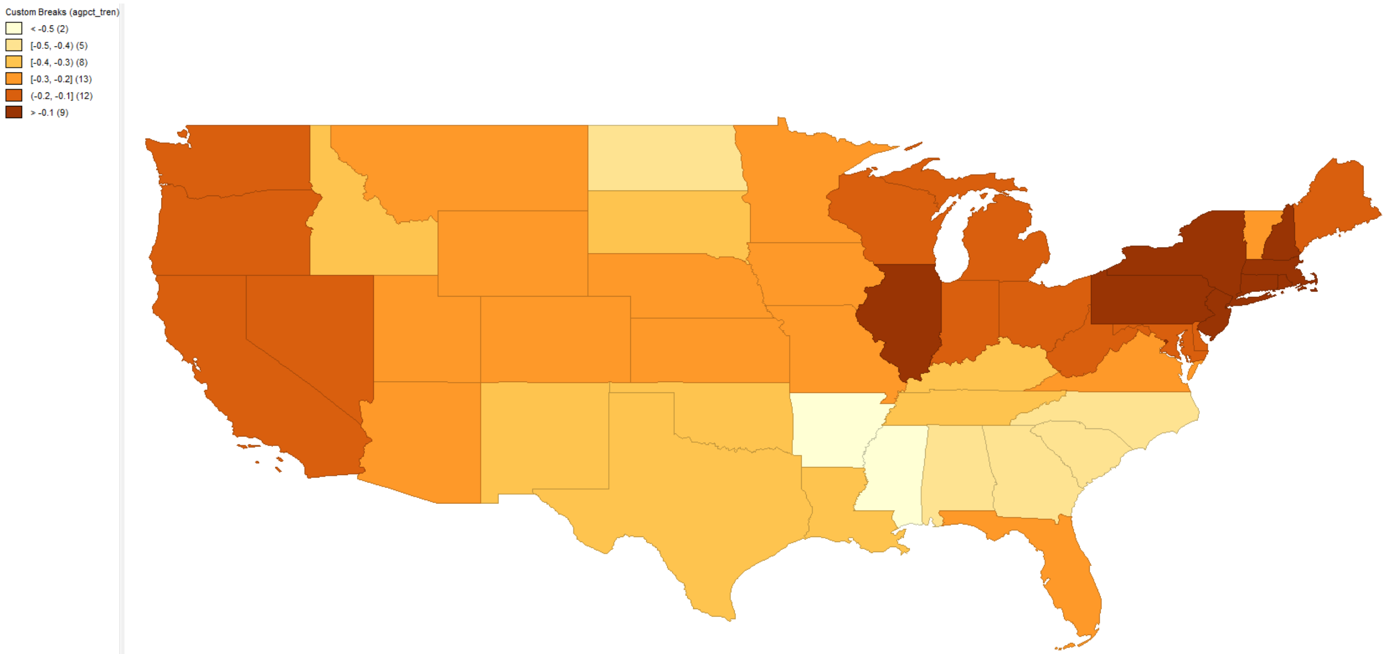

Figure 6.2.: Agricultural share trend by state

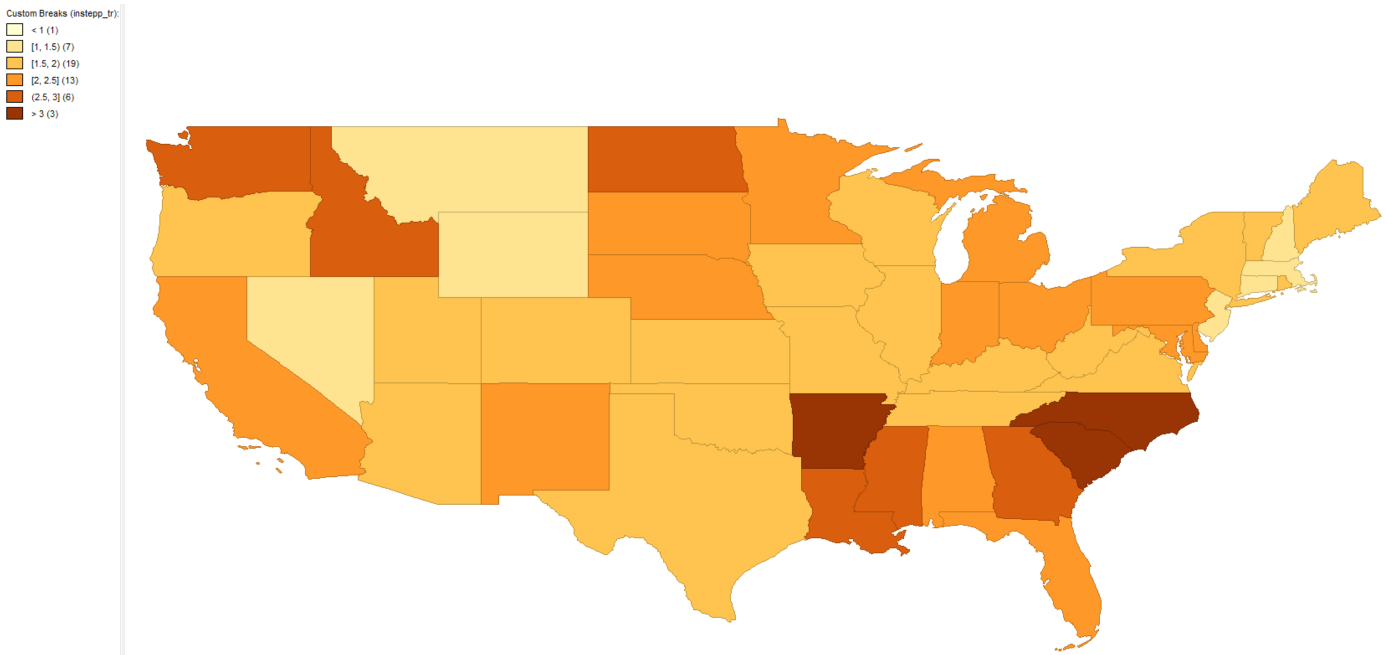

Figure 6.3.: Productivity trend by state, as percentage of 1960 value

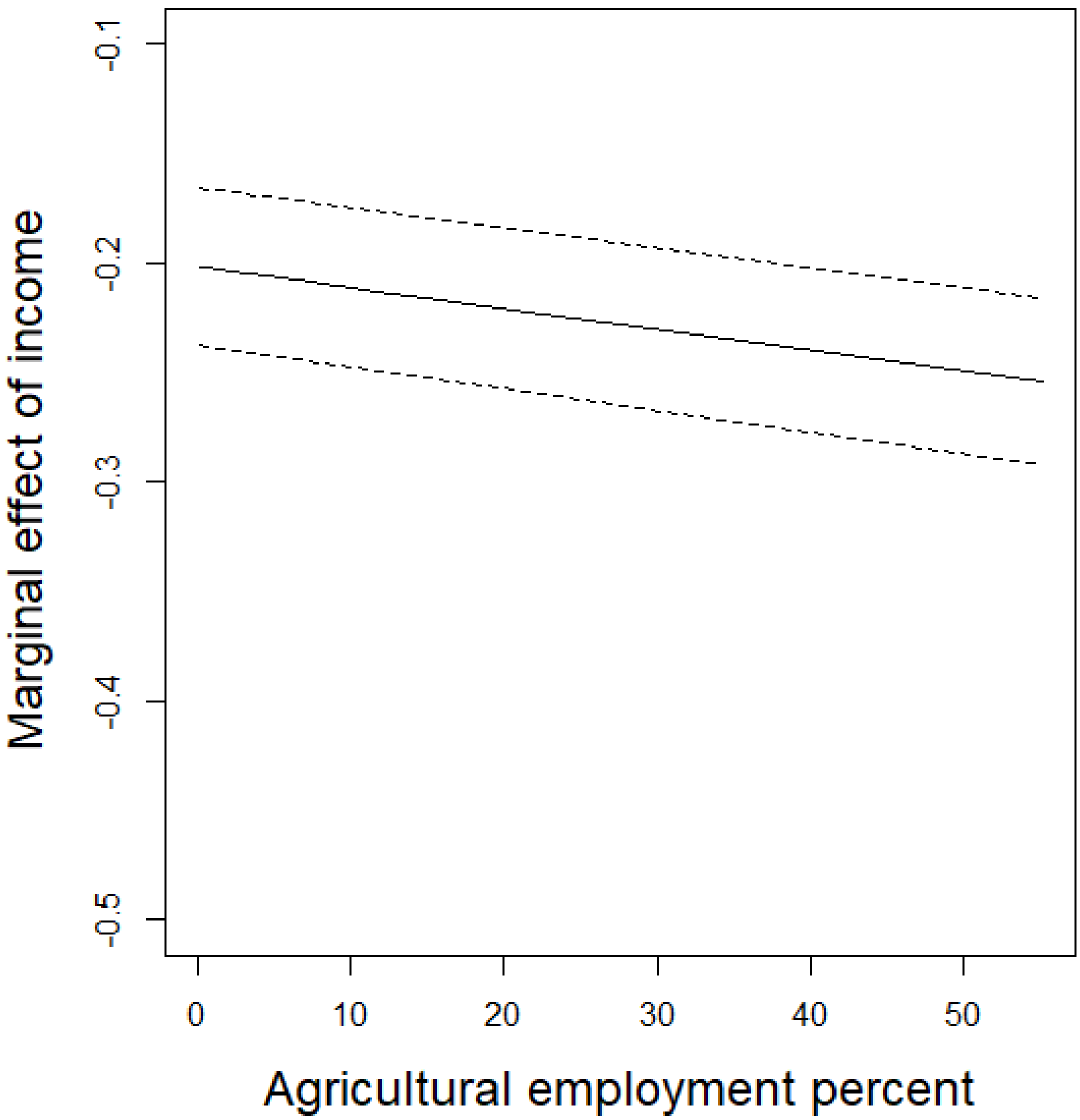

Figure 6.4.: Marginal effects of income by agricultural percent, two way interaction model

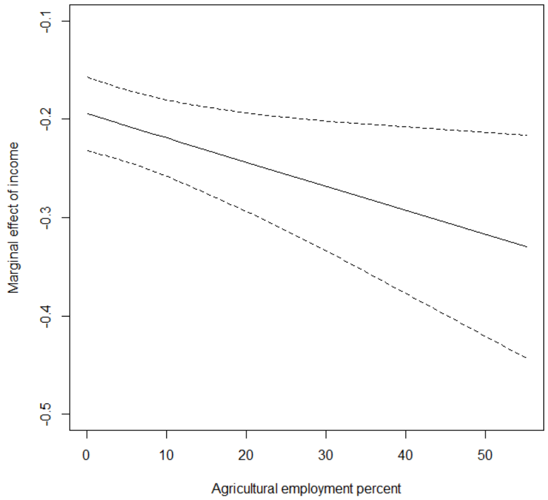

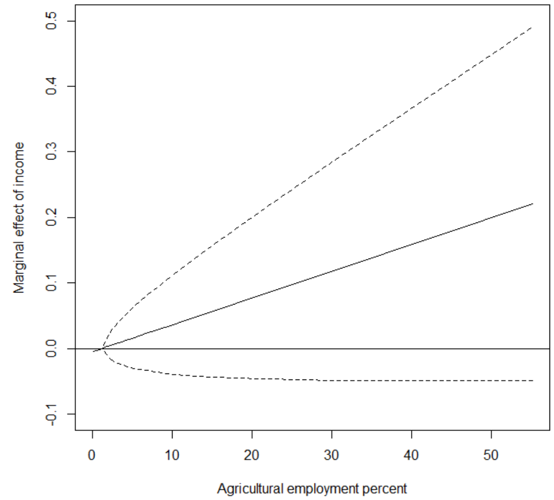

(a) Marginal effect by agricultural percent, non-South

(b) Marginal effect by agricultural percent, South

Figure 6.5.: Marginal effects of income by agricultural percent, three way interaction model

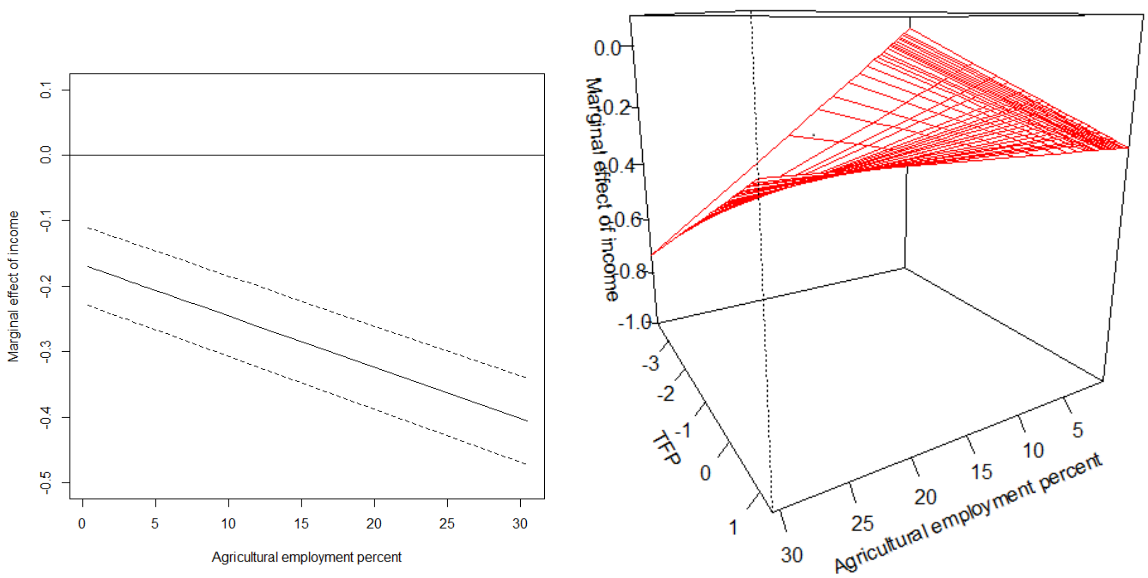

(a) Marginal effect by agricultural percent given TFP

(b) Marginal effect by agricultural percent and TFP

Figure 6.6.: Marginal effect by agricultural percent and TFP

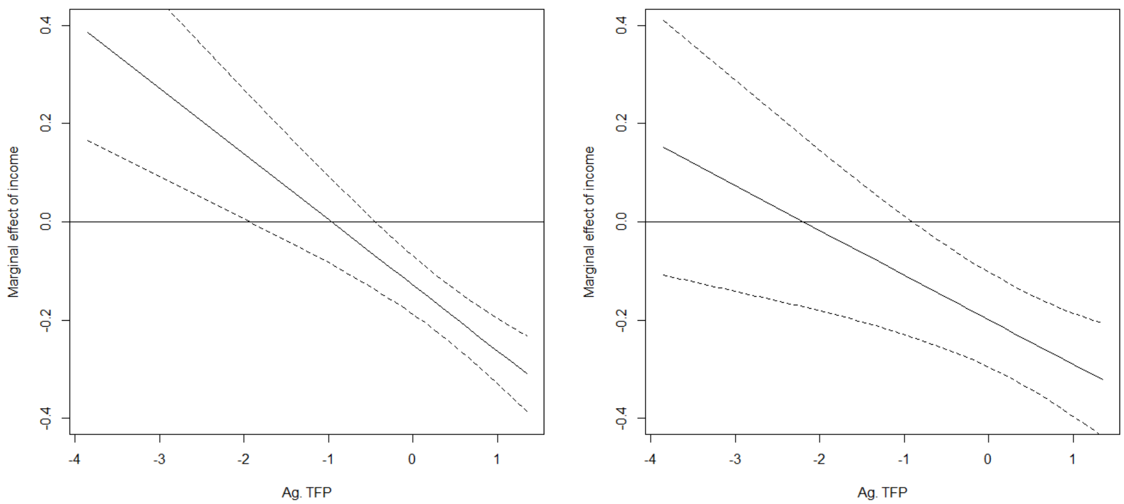

(a) Marginal effect by TFP, non-south

(b) Marginal effect by TFP, South

Figure 6.7.: Marginal effects of income by TFP and south indicator

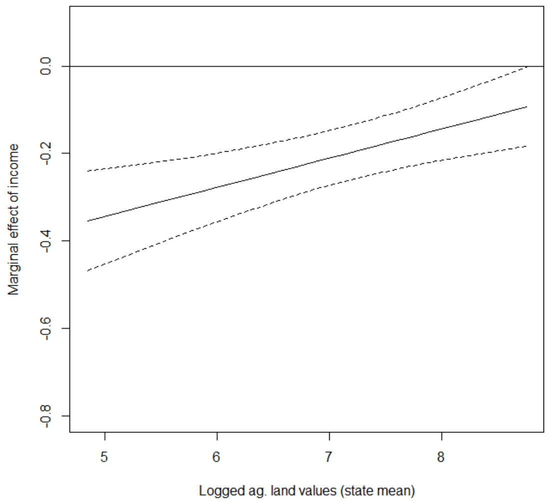

(a) Marginal effect by land value, non-south

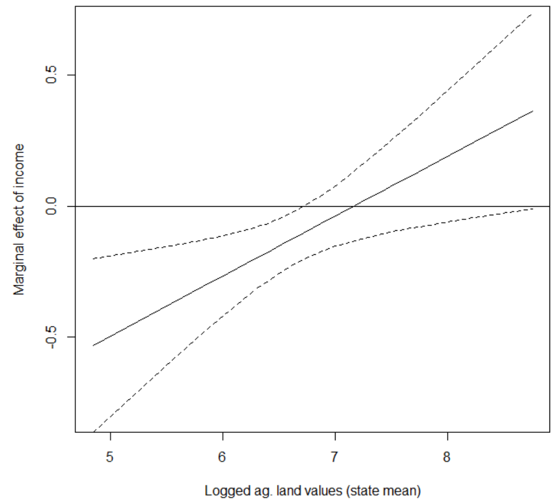

(b) Marginal effect by land value, South

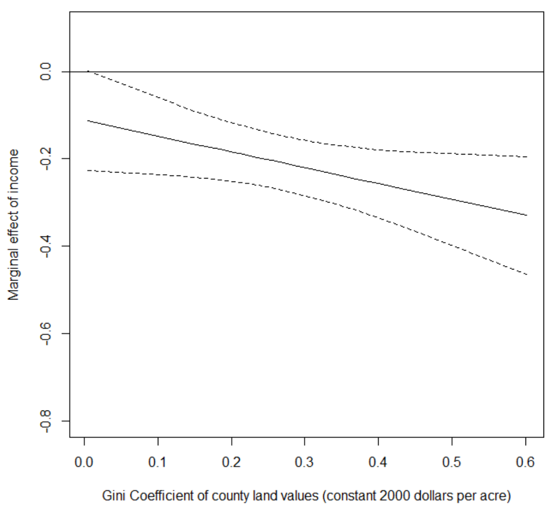

(c) Marginal effect by land value Gini, non-south

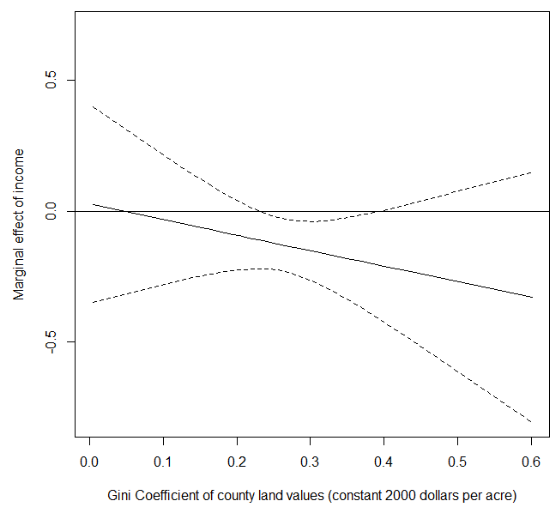

(d) Marginal effect by land value Gini, South

Figure 6.8.: Marginal effects of income by land value, land value Gini coefficient, and south indicator

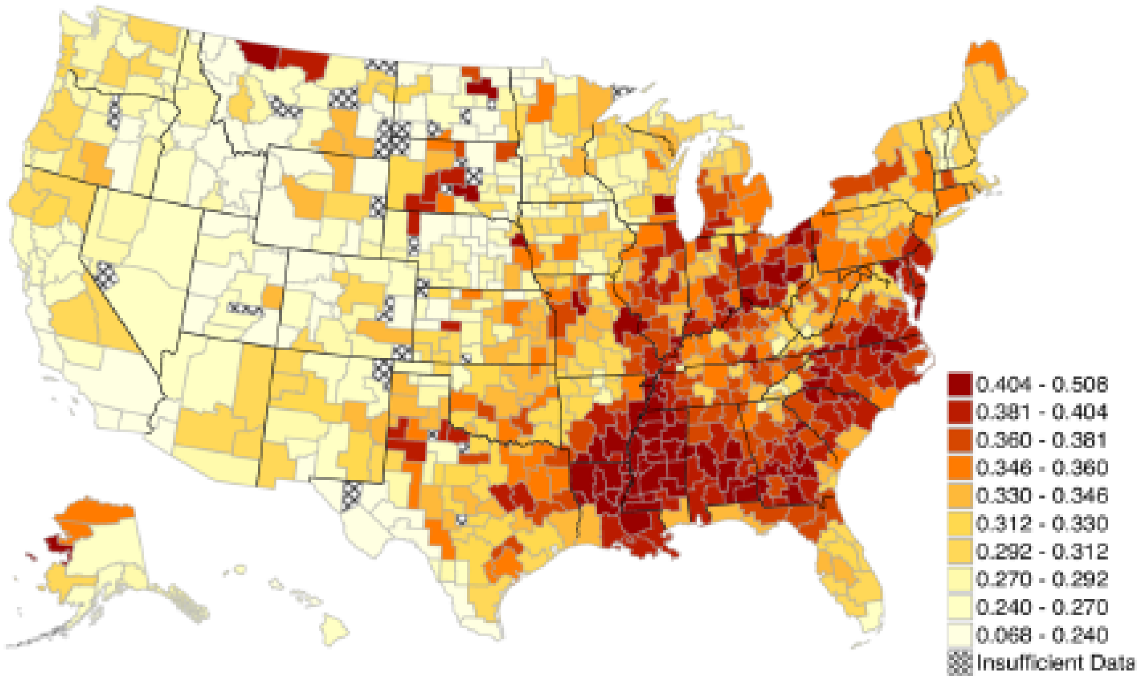

(a) Mobility slope by county (Chetty et al., 2014b)

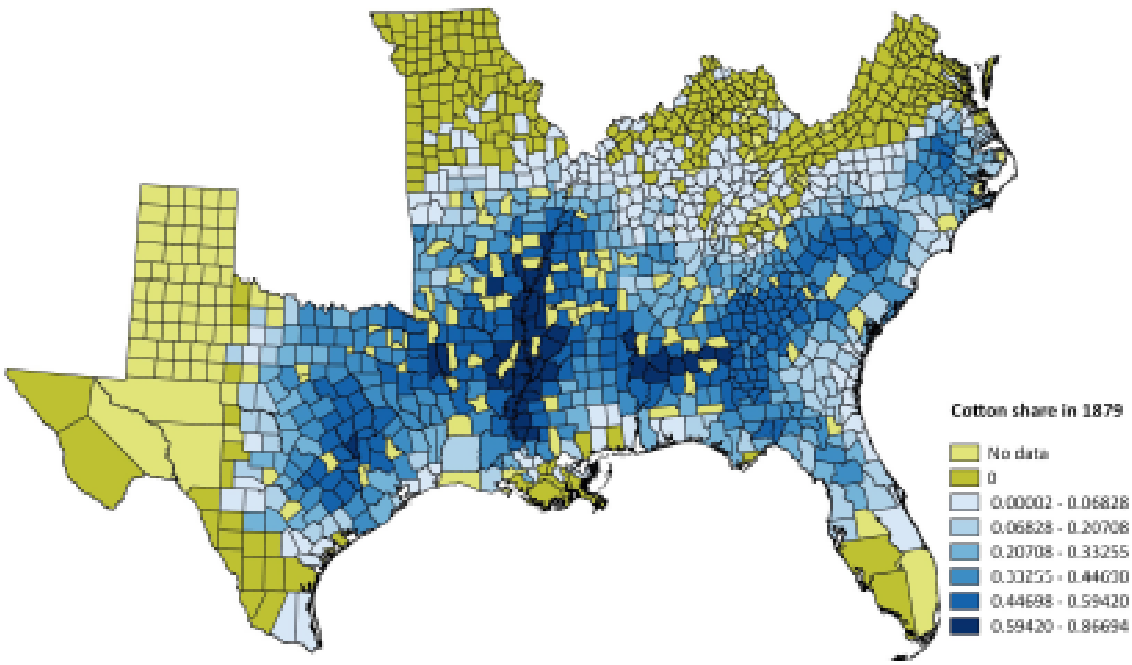

(b) Cotton cultivation % by county (Jung, 2018)

Figure 6.9.: 20^{th} century intergenerational mobility and cotton production

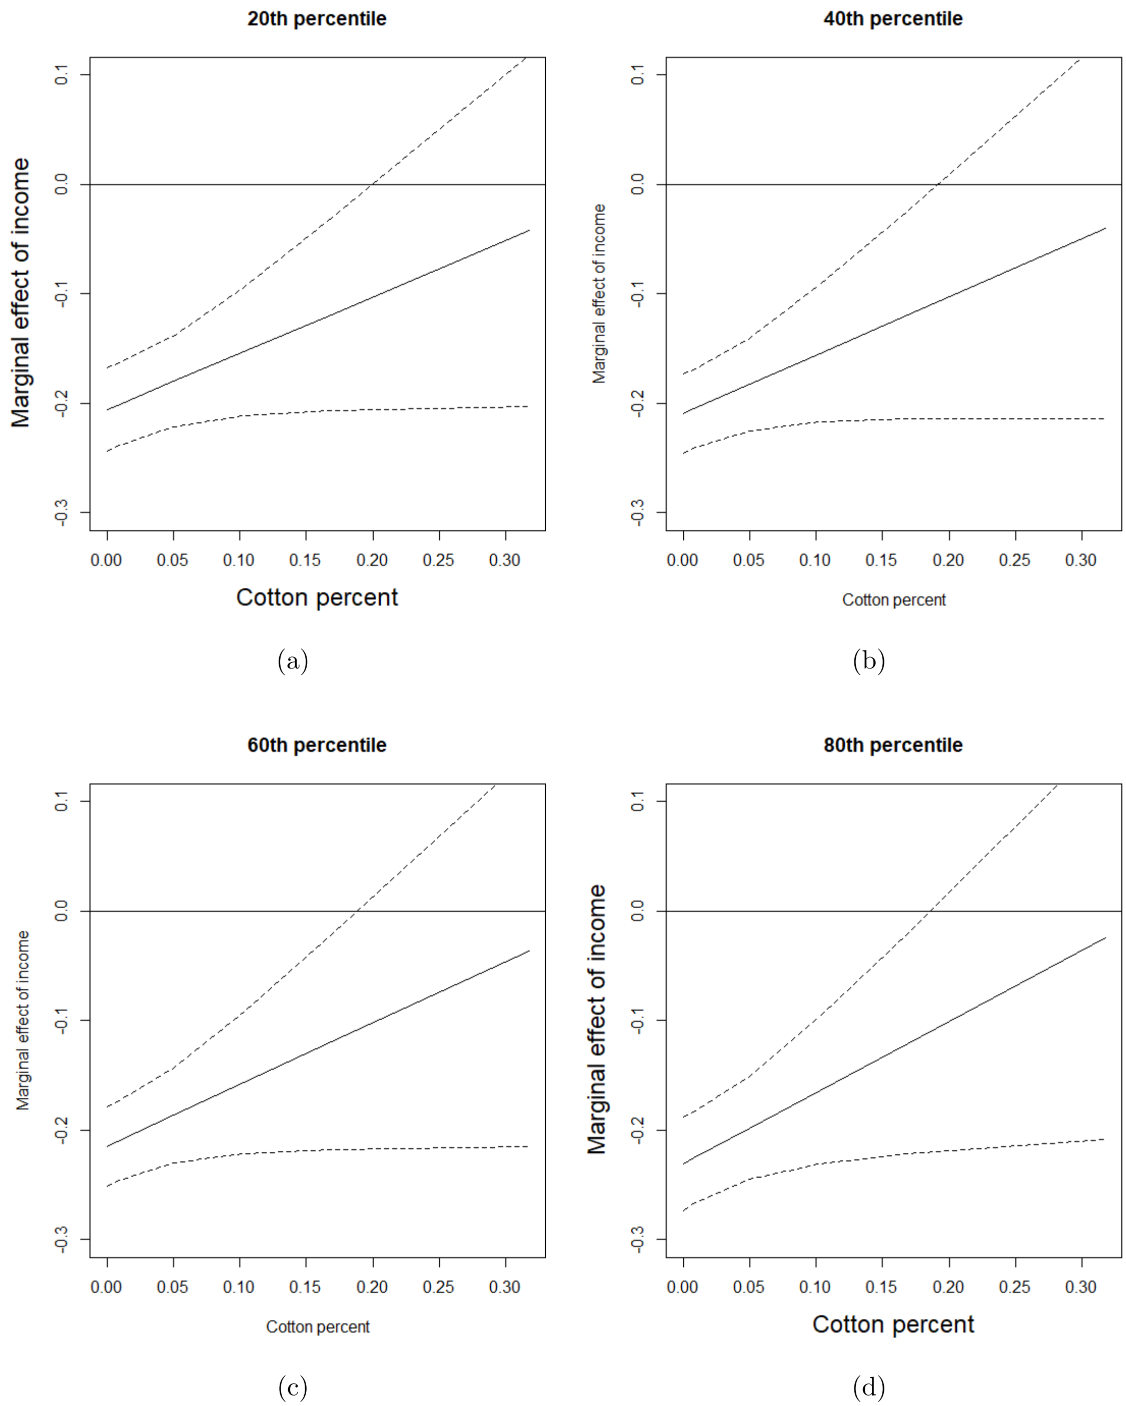

Figure 6.10.: Marginal effects of income by cotton percent and agricultural share percentile

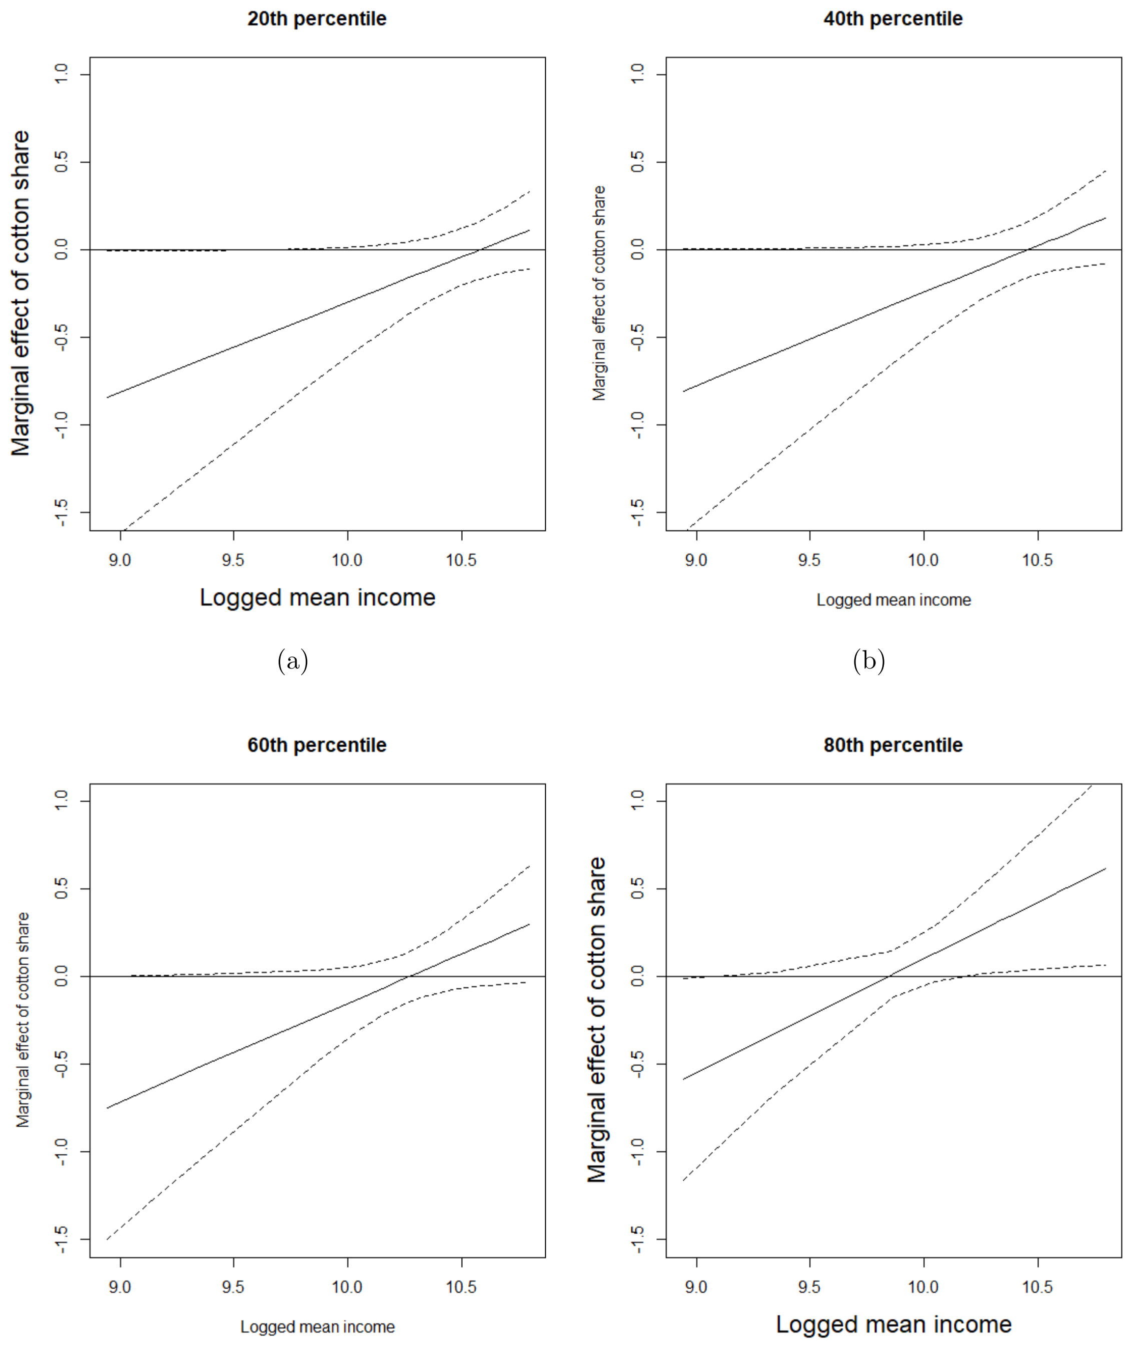

Figure 6.11.: Marginal effects of cotton share by logged income and agricultural share percentile

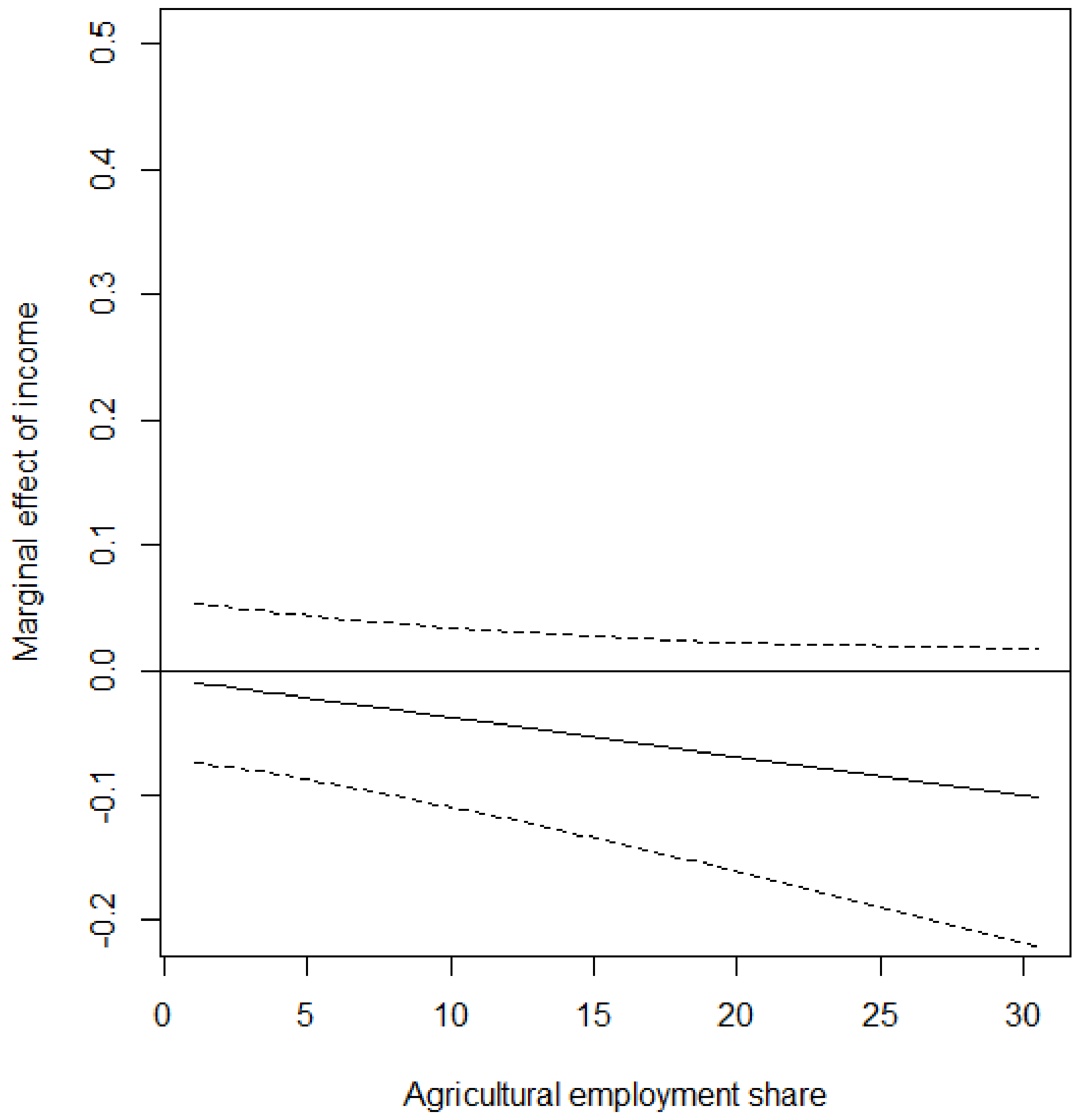

Figure 6.12.: Marginal effects of income on white-black slope difference, by agricultural percent, two way interaction model

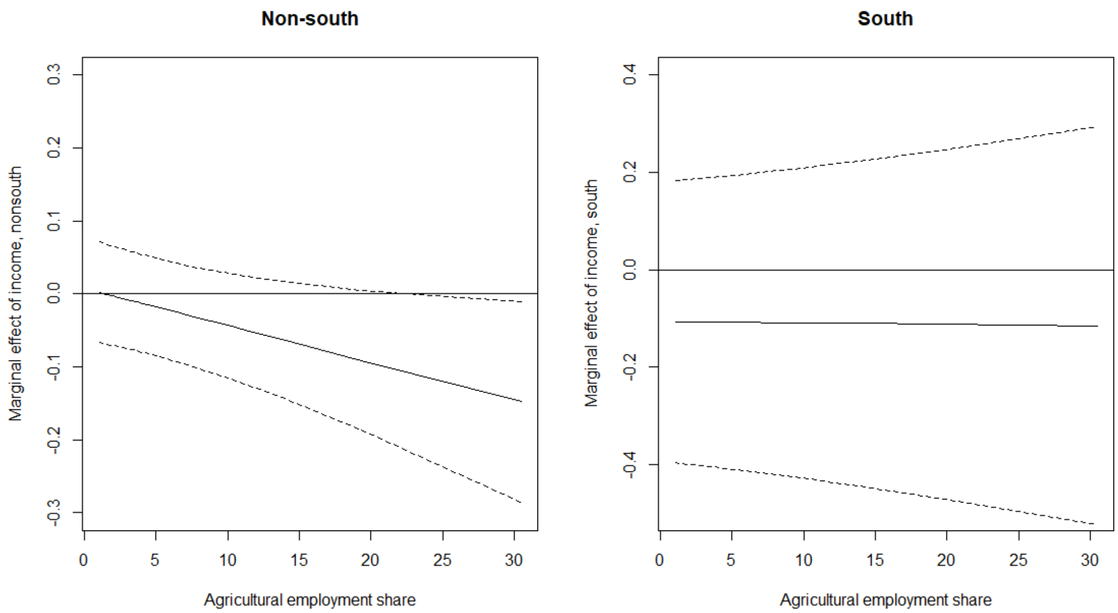

(a) Marginal effect by agricultural percent, Non-south

(b) Marginal effect by agricultural percent, South

Figure 6.13.: Marginal effects of income on white-black slope difference, by agricultural share and south indicator

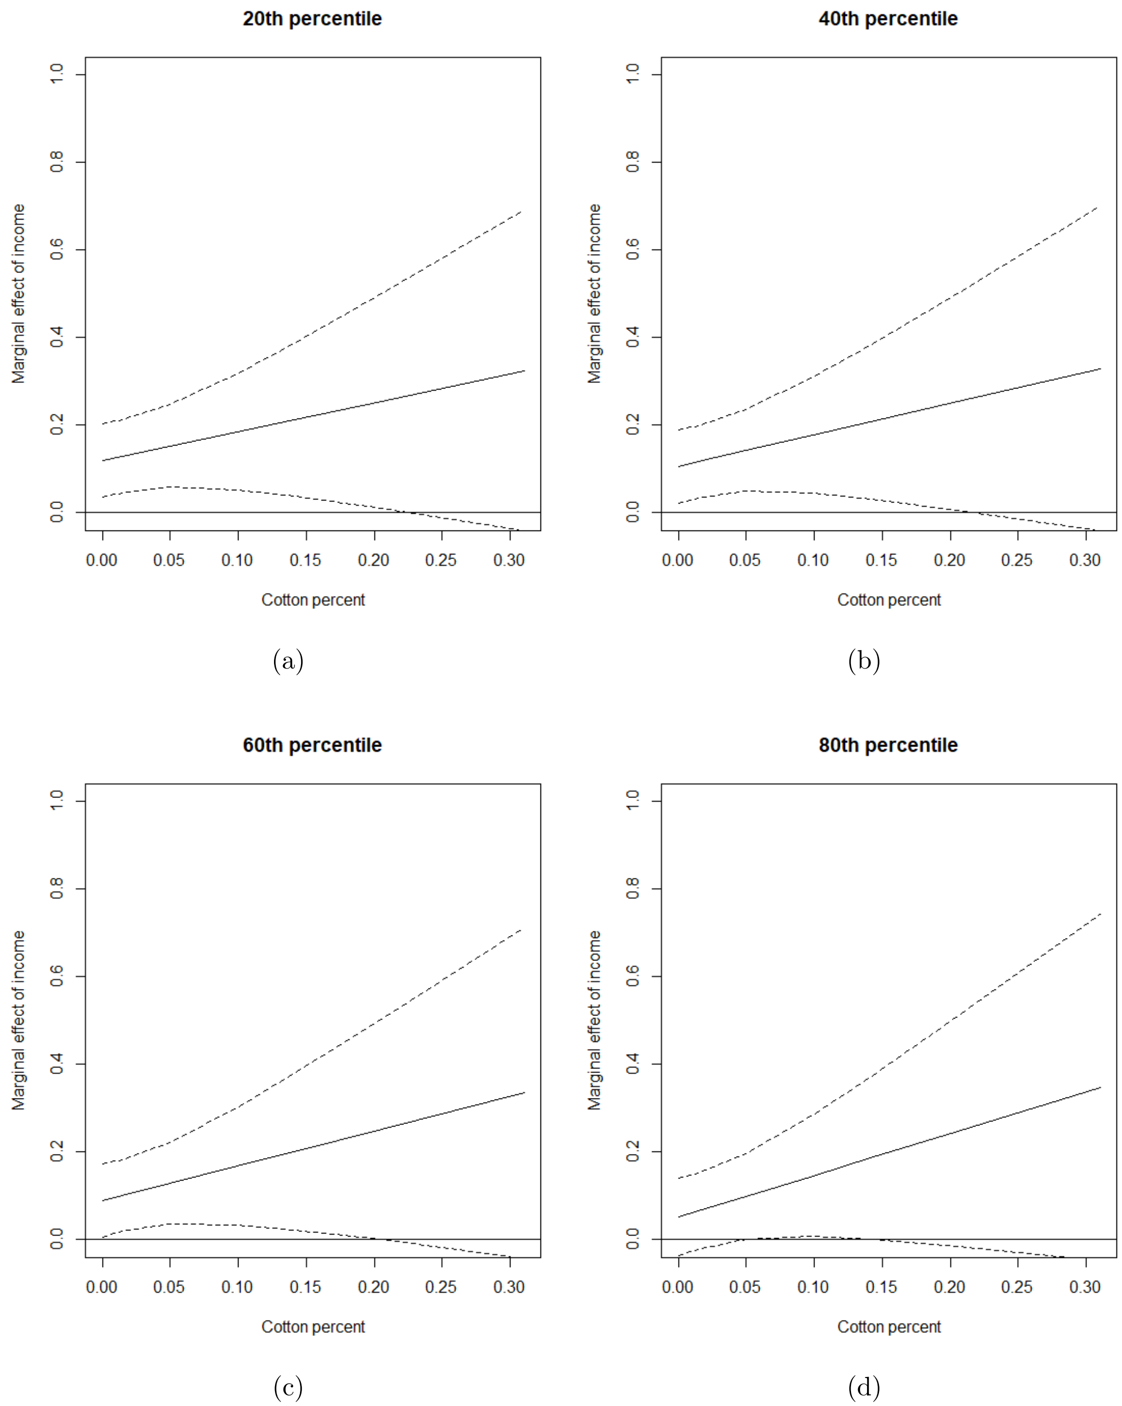

Figure 6.14.: Marginal effects of income on white-black difference by cotton percent and agricultural share percentile

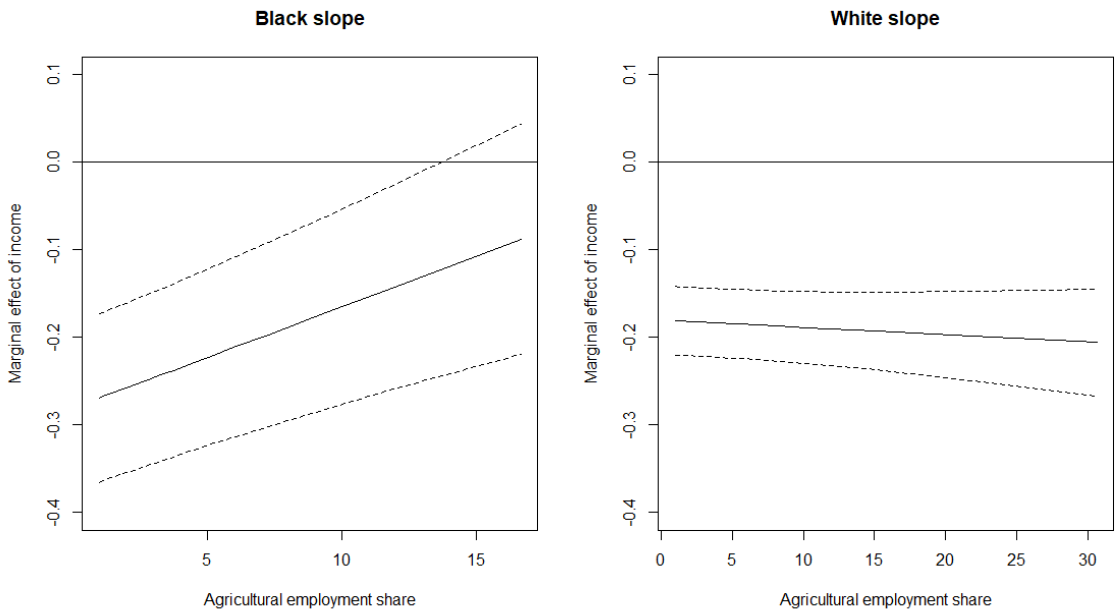

(a) Marginal effect by agricultural percent (b) Marginal effect by agricultural percent

Figure 6.15.: Marginal effects of income on black and white slopes, by agricultural share

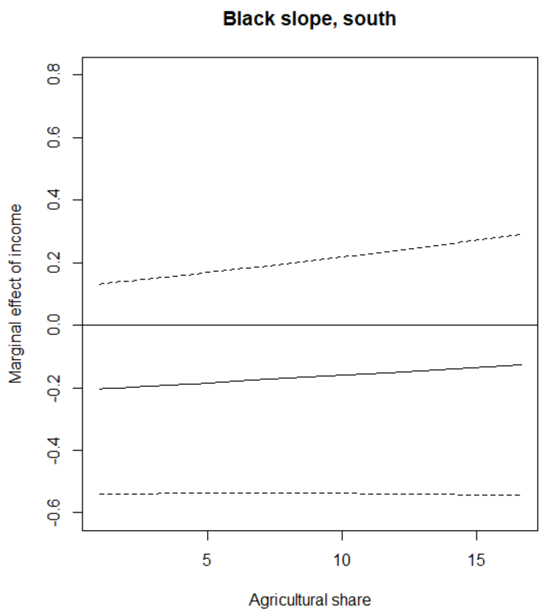

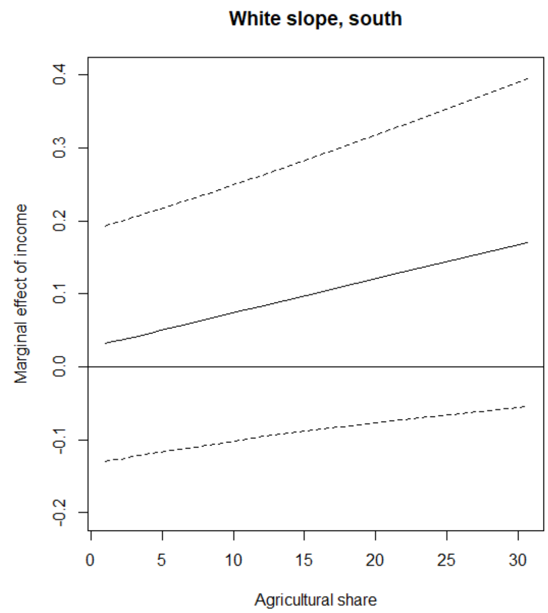

(a) Marginal effect by agricultural percent, South (b) Marginal effect by agricultural percent, Non-South

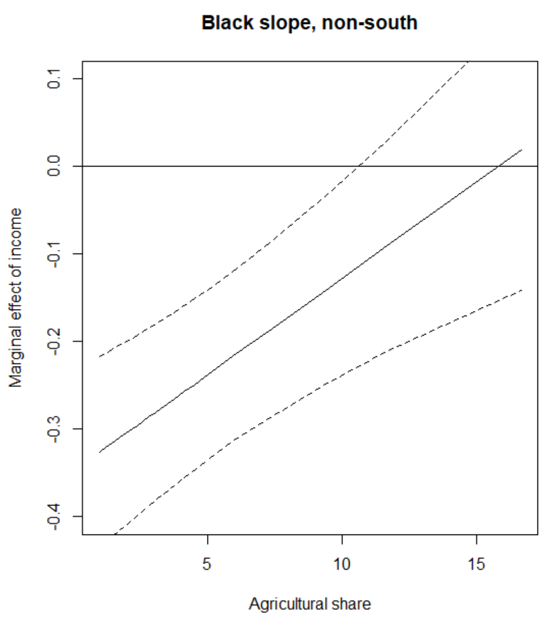

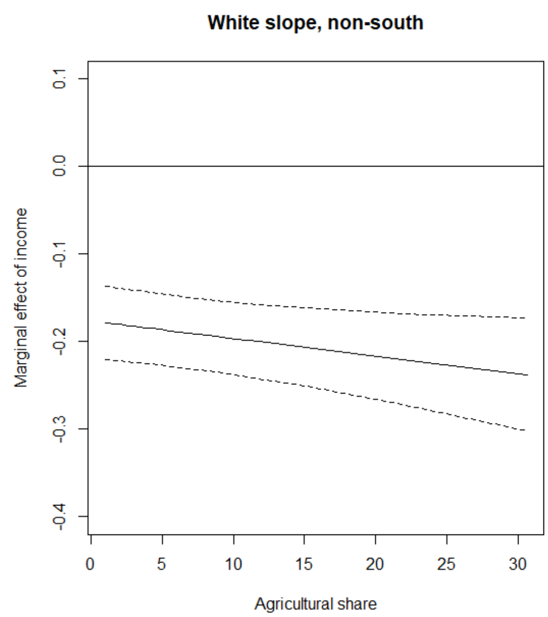

(c) Marginal effect by agricultural percent, South (d) Marginal effect by agricultural percent, Non-South

Figure 6.16.: Marginal effects of income on black and white income slopes, by agricultural share and south indicator

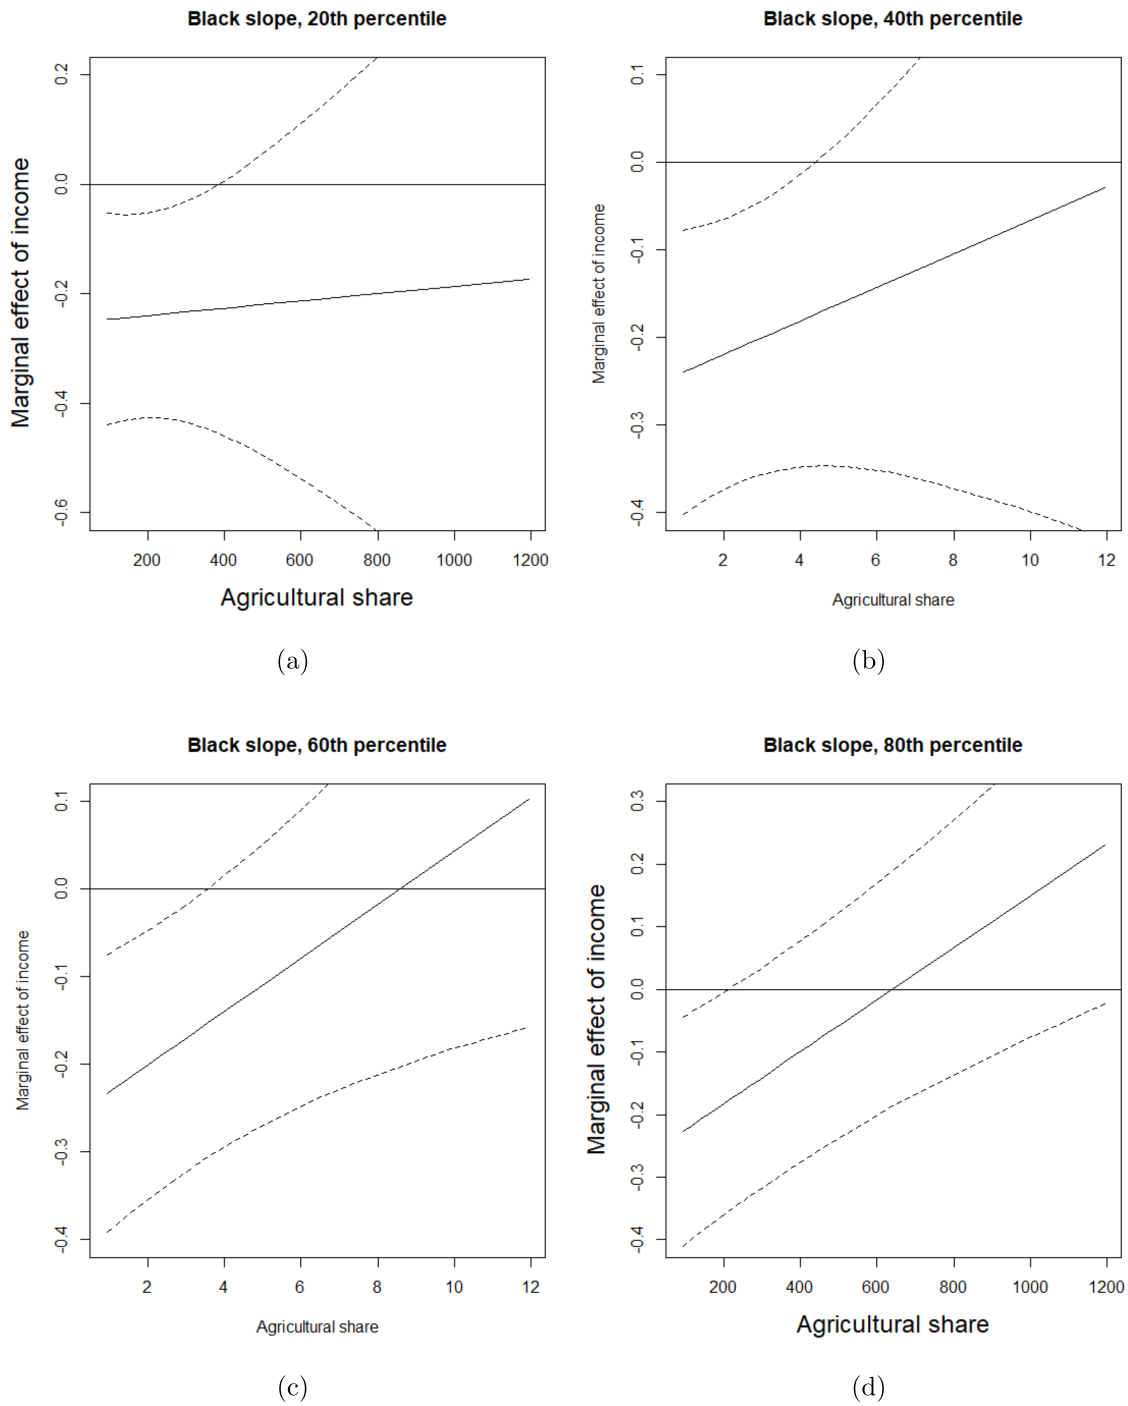

Figure 6.17.: Marginal effects of income on black slopes by agricultural share and INSTEPP percentile

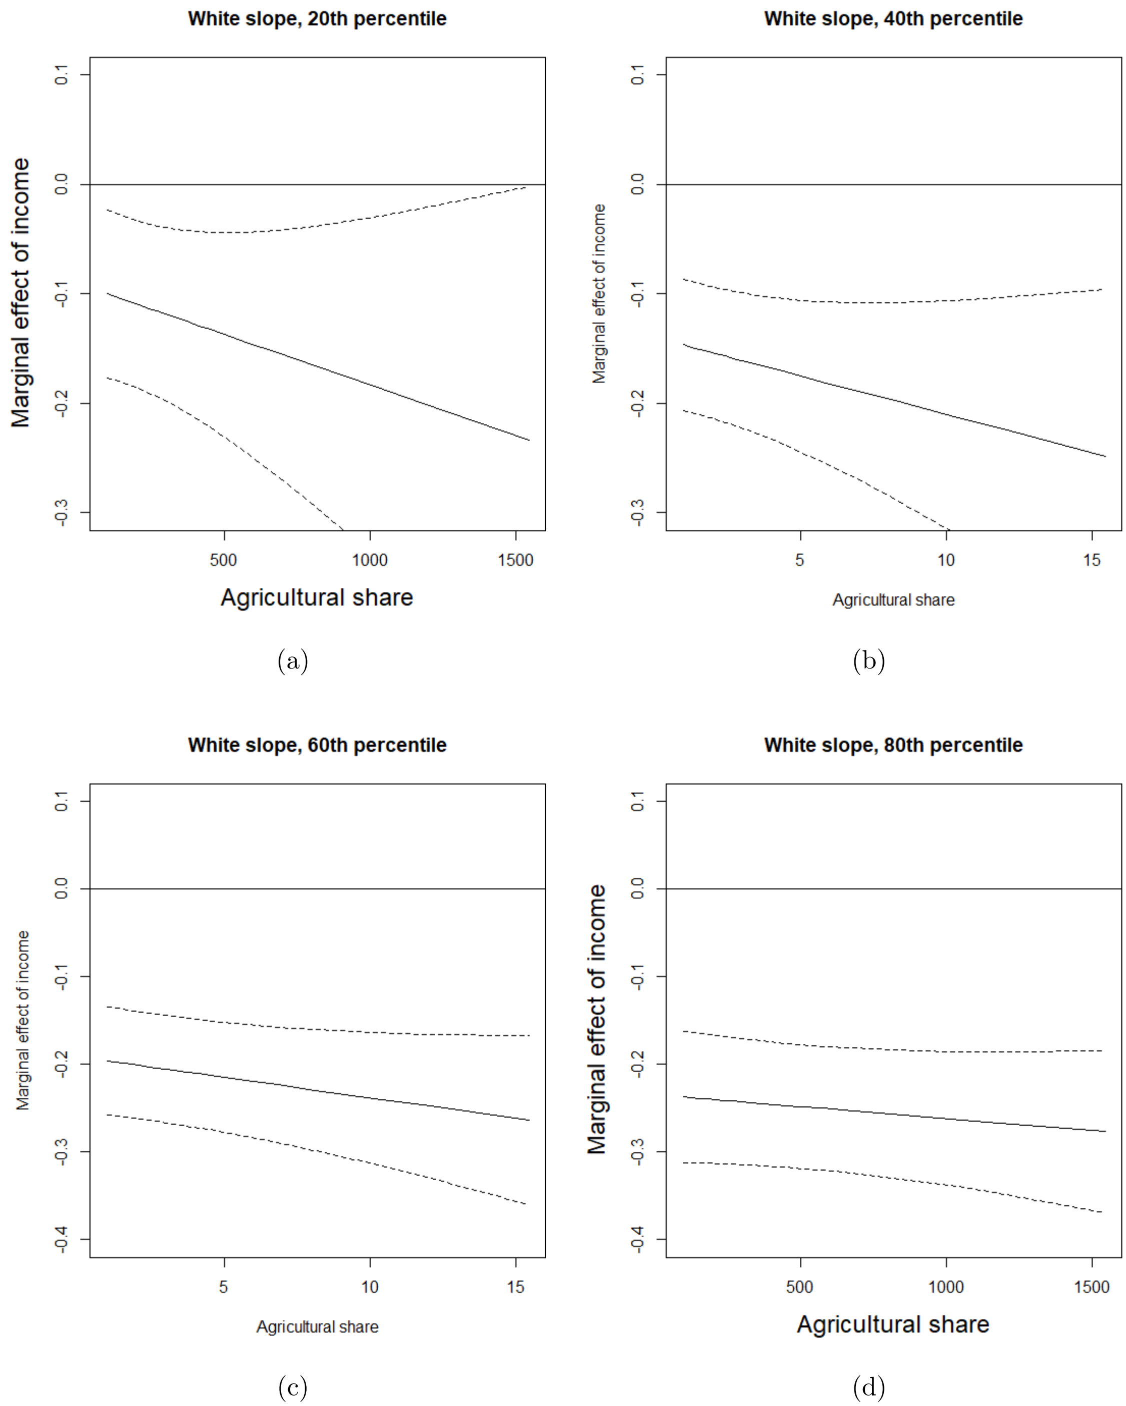

Figure 6.18.: Marginal effects of income on white slopes by agricultural share and INSTEPP percentile

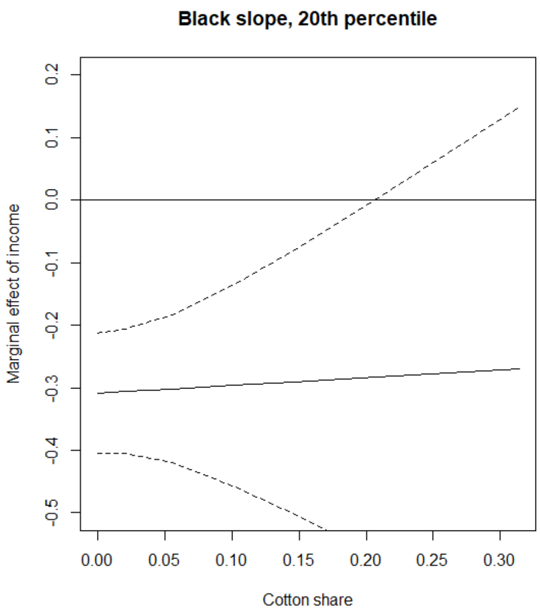

(a) Black, 20^{th} percentile

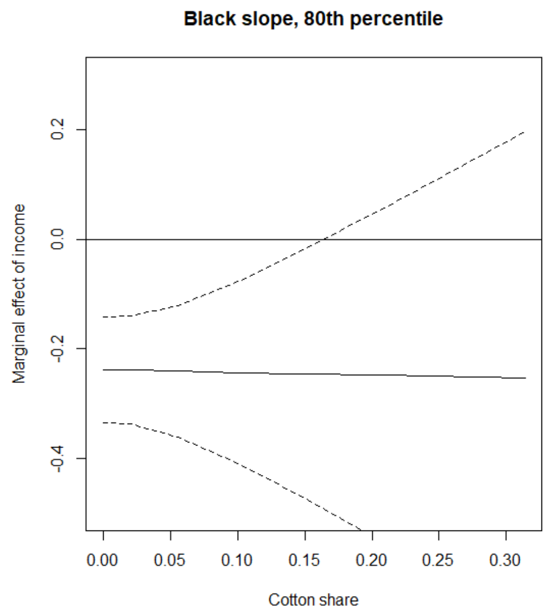

(b) Black, 80^{th} percentile

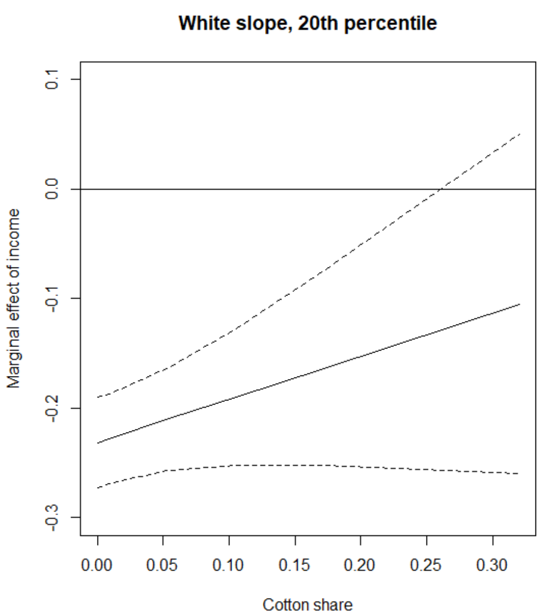

(c) White, 20^{th} percentile

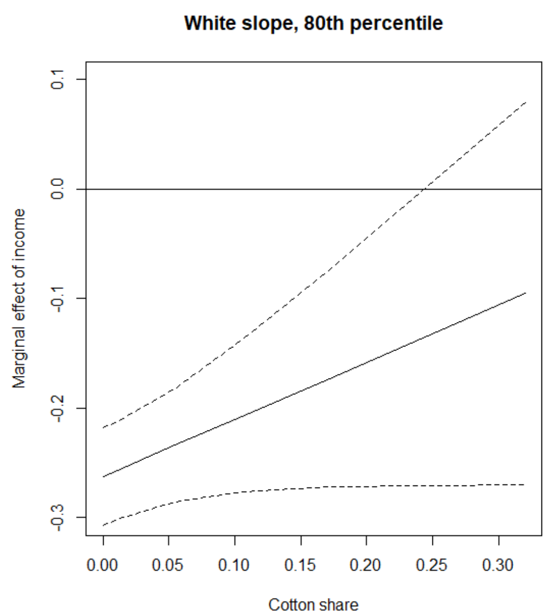

(d) White, 80^{th} percentile

Figure 6.19.: Marginal effects of income on black and white slopes by cotton share and agricultural share percentile

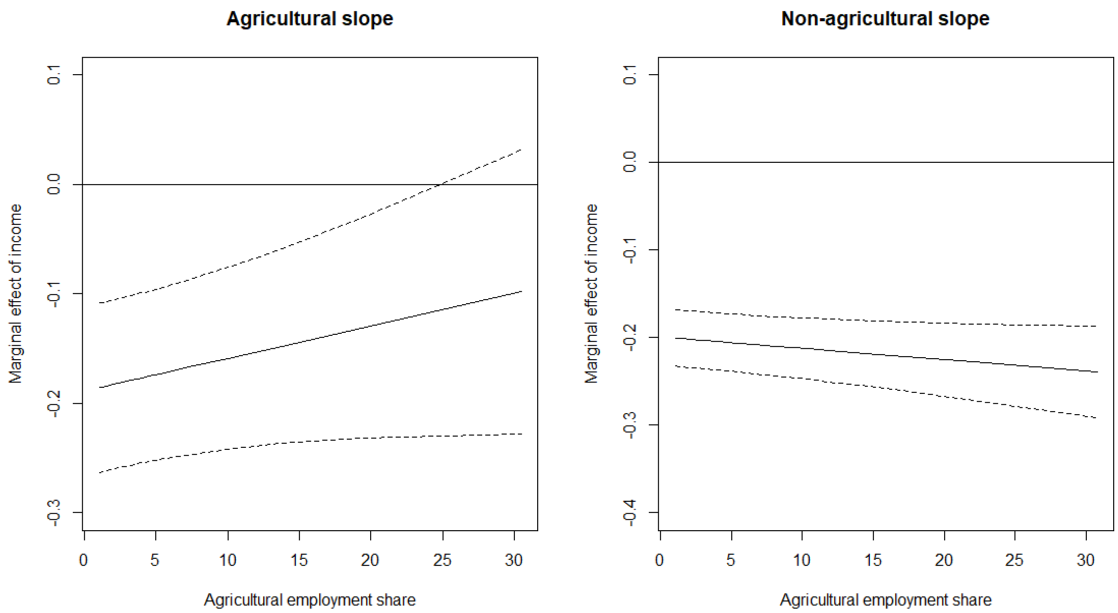

(a) Marginal effect by agricultural percent

(b) Marginal effect by agricultural percent

Figure 6.20.: Marginal effects of income on agricultural and non-agricultural slopes, by agricultural share

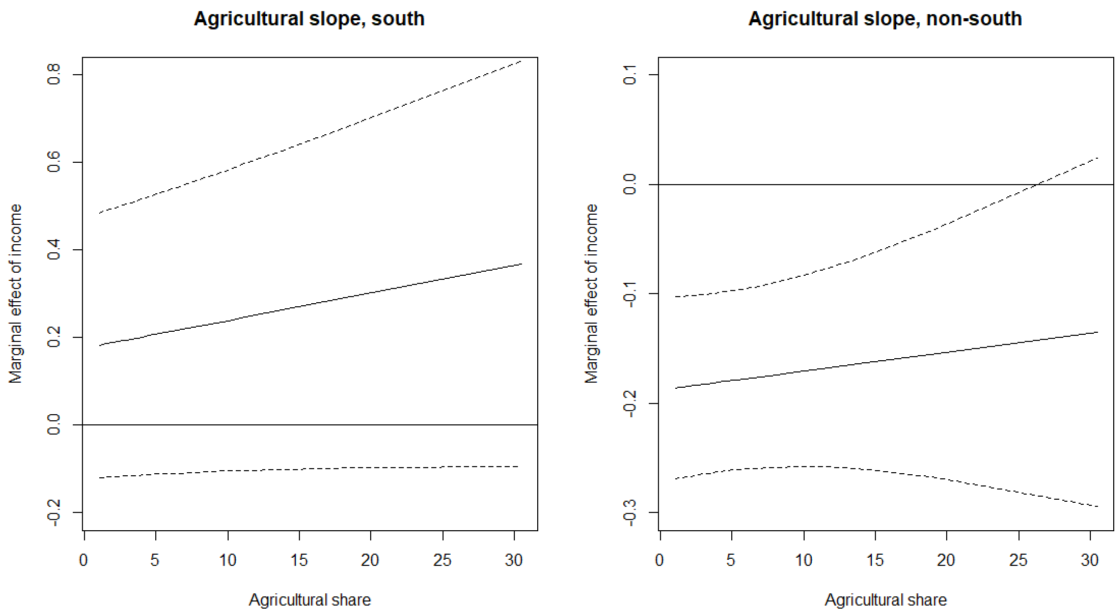

(a) Marginal effect by agricultural percent, South (b) Marginal effect by agricultural percent, Non-south

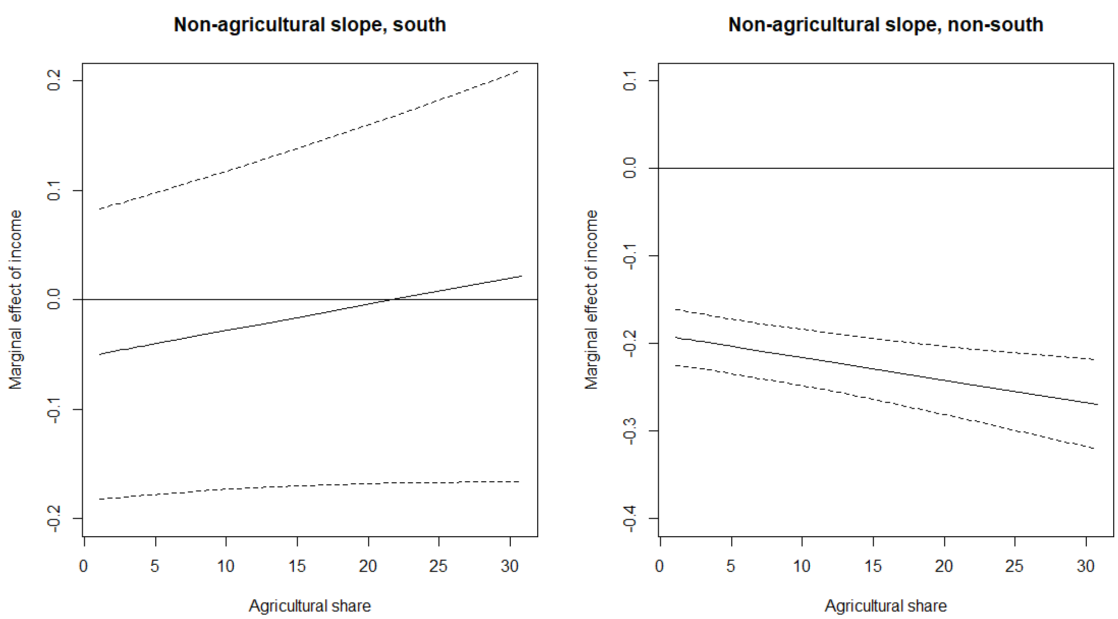

(c) Marginal effect by agricultural percent, South (d) Marginal effect by agricultural percent, Non-south

Figure 6.21.: Marginal effects of income on agricultural and non-agricultural slopes, by agricultural share and south indicator

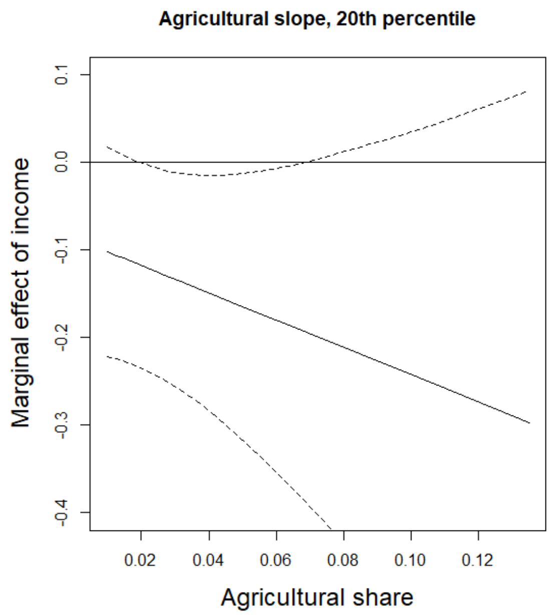

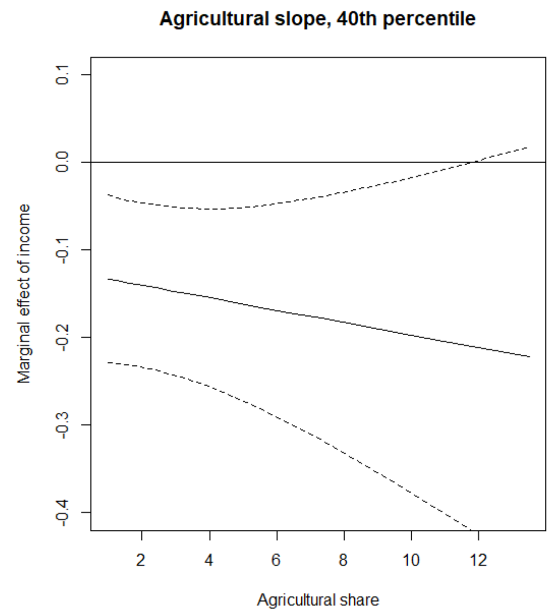

(a) Marginal effect by agricultural percent, South

(b) Marginal effect by agricultural percent, Non-south

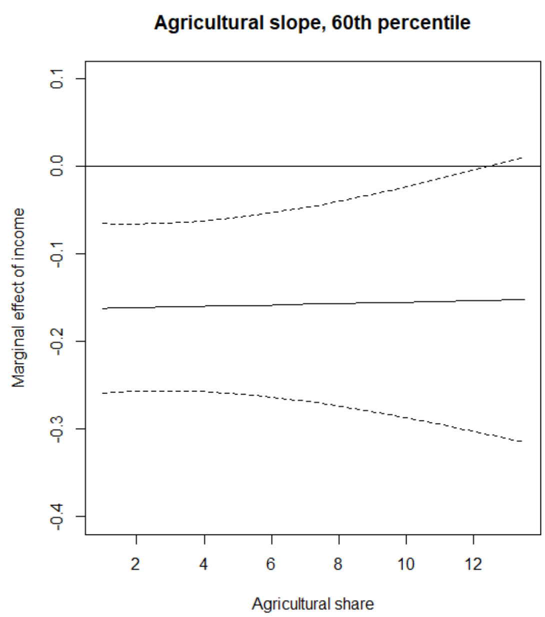

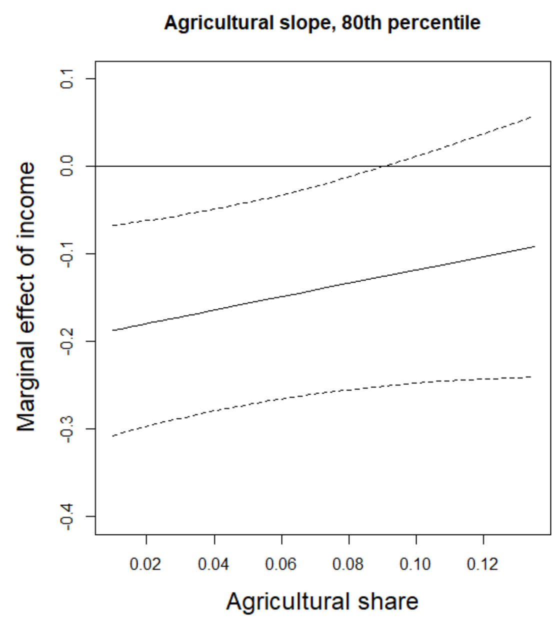

(c) Marginal effect by agricultural percent, South

(d) Marginal effect by agricultural percent, Non-south

Figure 6.22.: Marginal effects of income on agricultural slopes by agricultural share and INSTEPP percentile

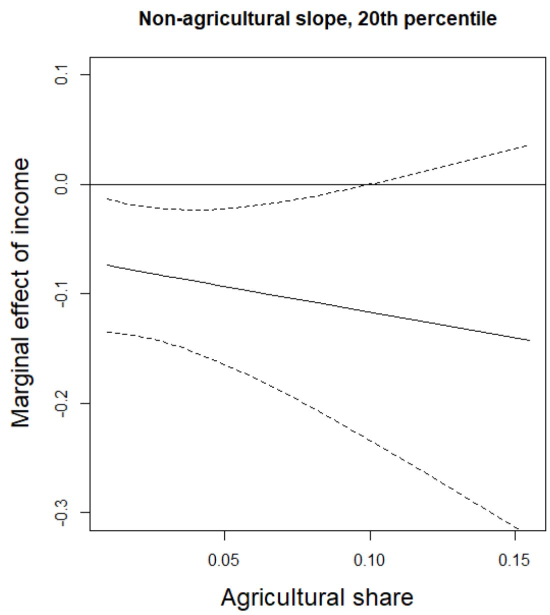

(a) Marginal effect by agricultural percent, South

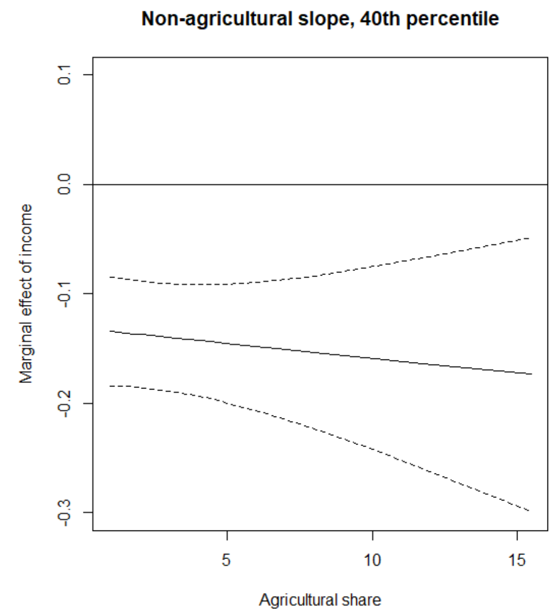

(b) Marginal effect by agricultural percent, Non-south

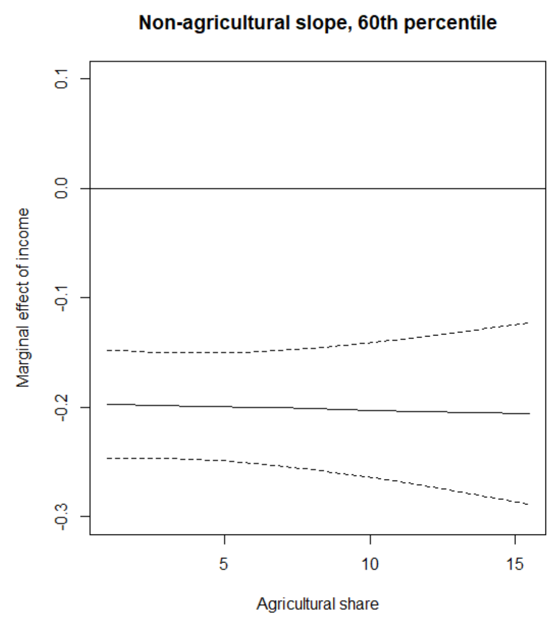

(c) Marginal effect by agricultural percent, South

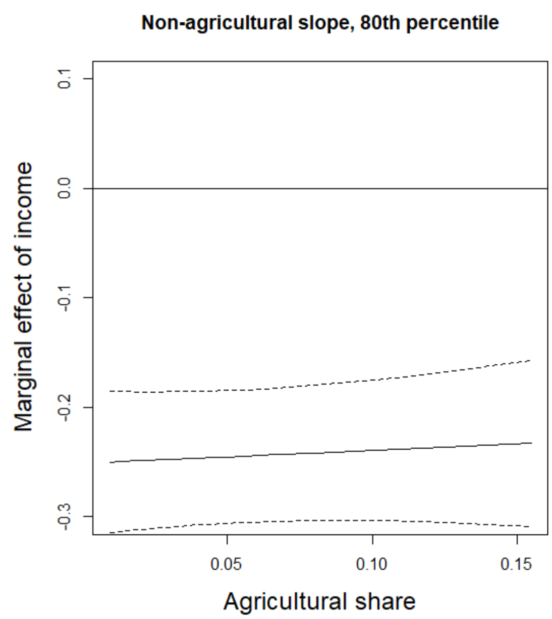

(d) Marginal effect by agricultural percent, Non-south

Figure 6.23.: Marginal effects of income on non-agricultural slopes by agricultural share and INSTEPP percentile

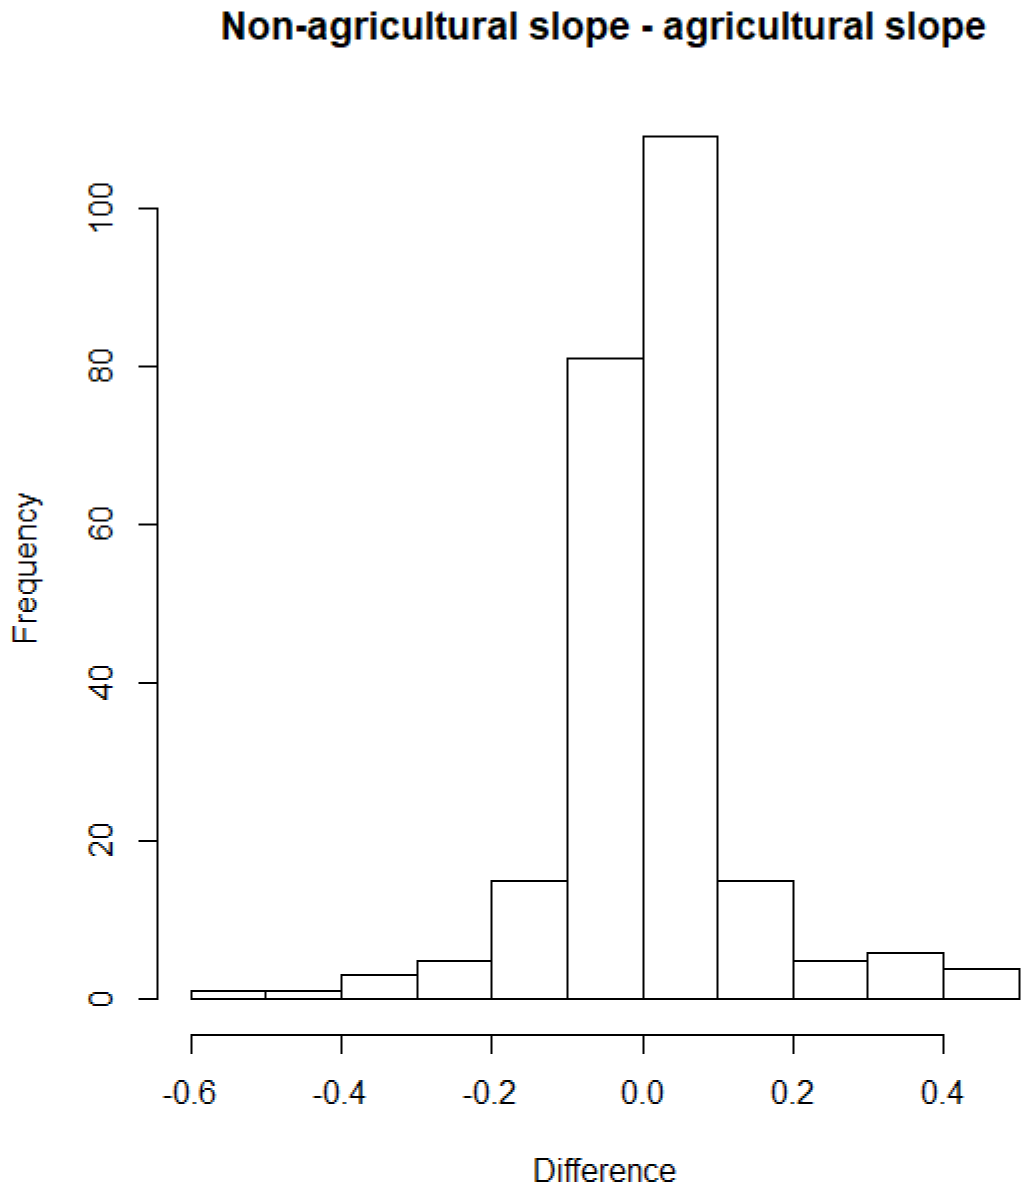

Figure 6.24.: Differences between non-agricultural and agricultural slopes

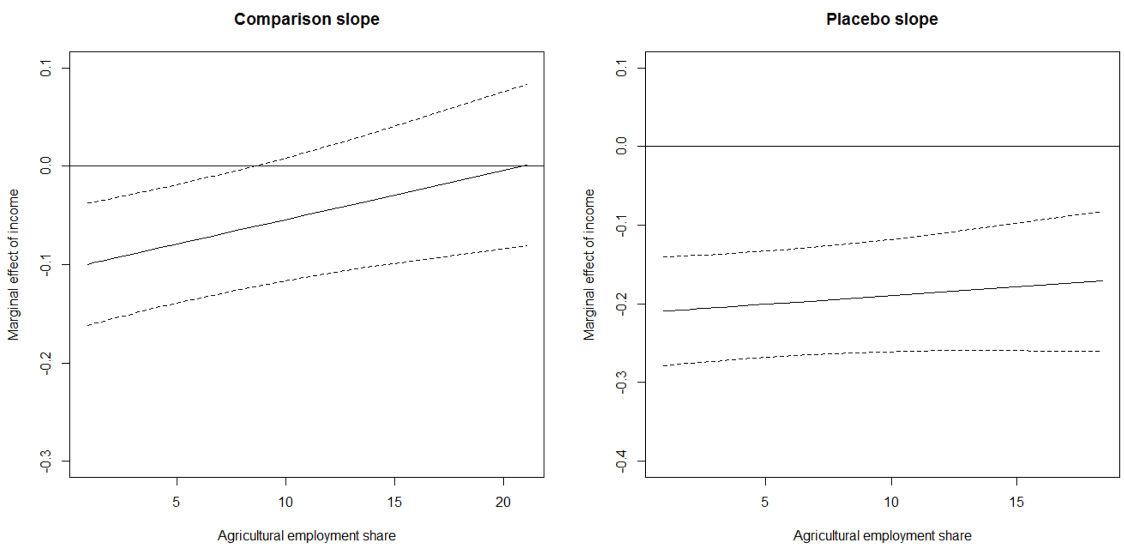

(a) Marginal effect by agricultural percent (b) Marginal effect by agricultural percent

Figure 6.25.: Marginal effects of income for comparison and placebo groups, by agricultural share

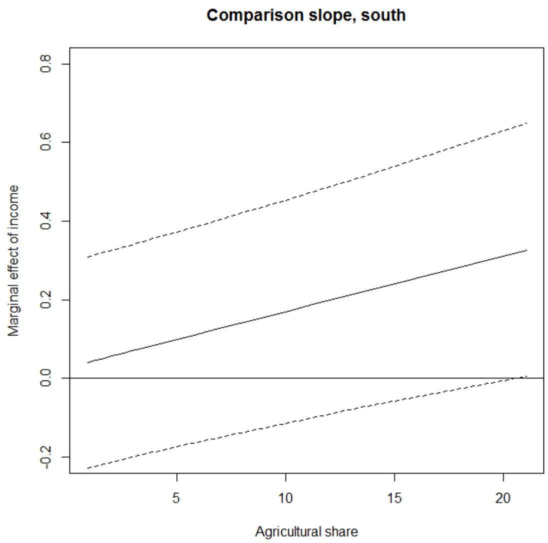

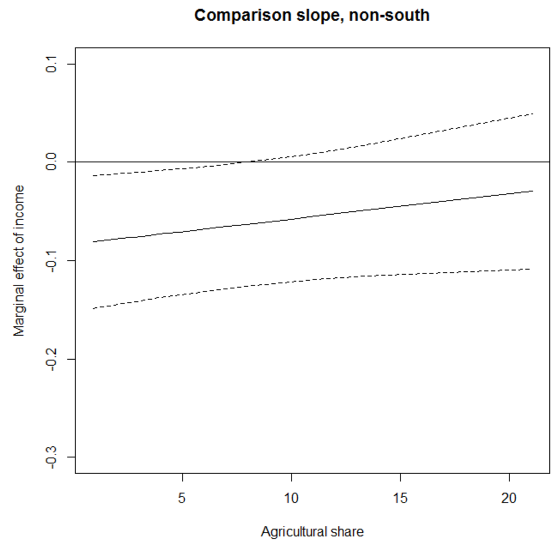

(a) Marginal effect by agricultural percent, South (b) Marginal effect by agricultural percent, Non-south

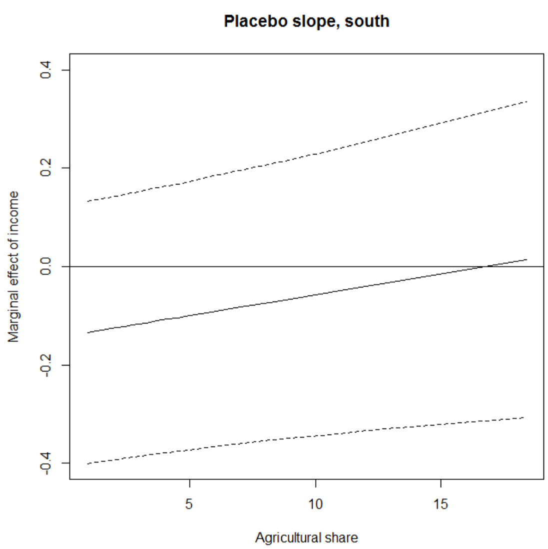

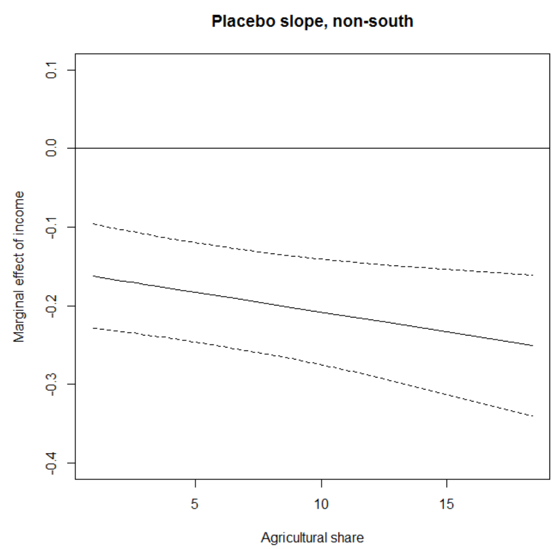

(c) Marginal effect by agricultural percent, South (d) Marginal effect by agricultural percent, Non-south

Figure 6.26.: Marginal effects of income for comparison and placebo groups, by agricultural share and south indicator

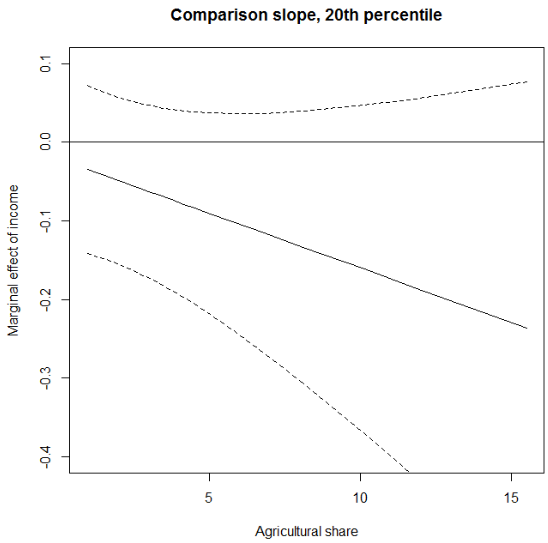

(a)

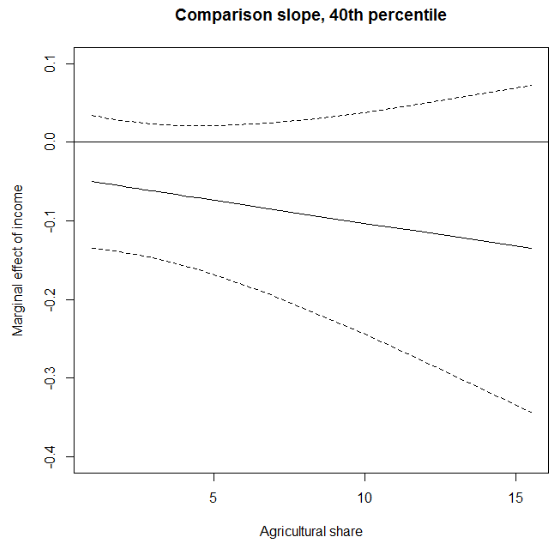

(b)

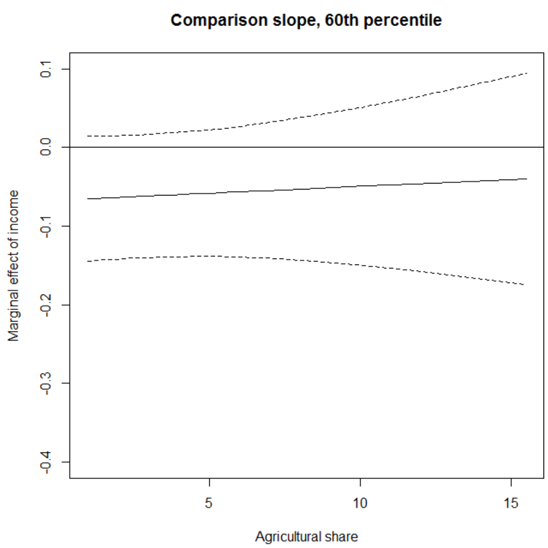

(c)

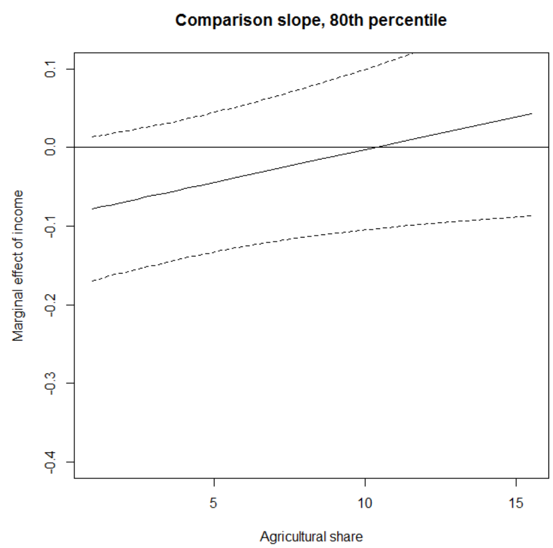

(d)

Figure 6.27.: Marginal effects of income on comparison group slopes by agricultural share and INSTEPP percentile

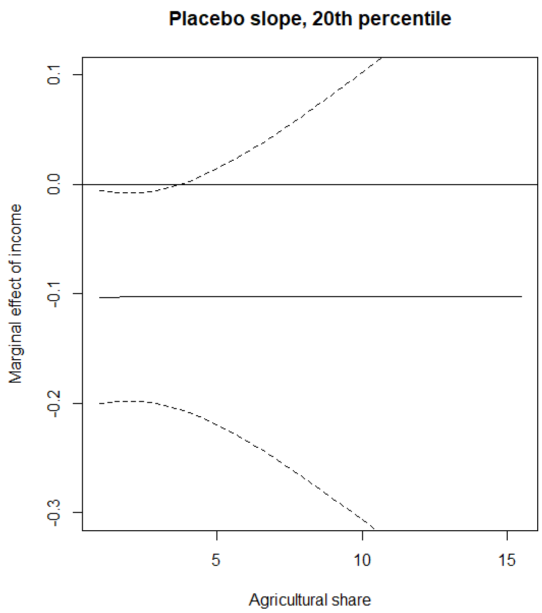

(a)

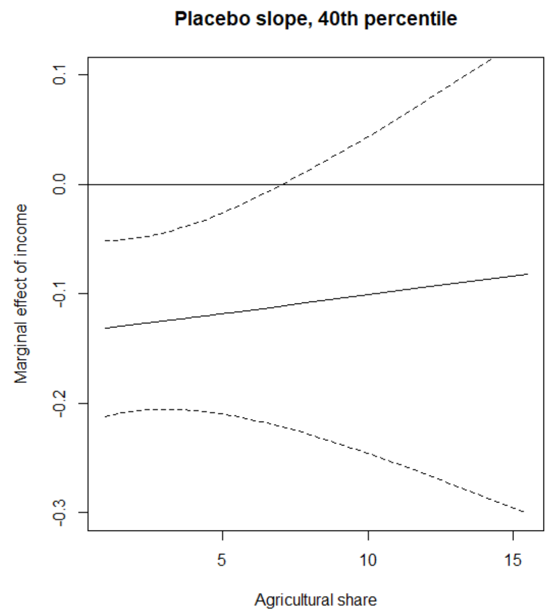

(b)

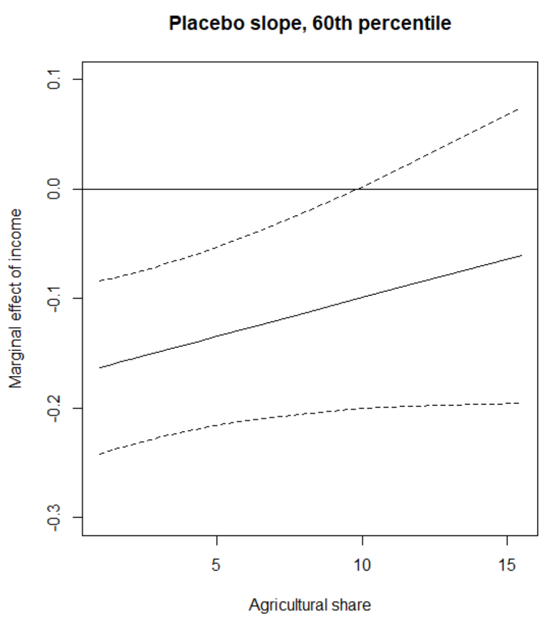

(c)

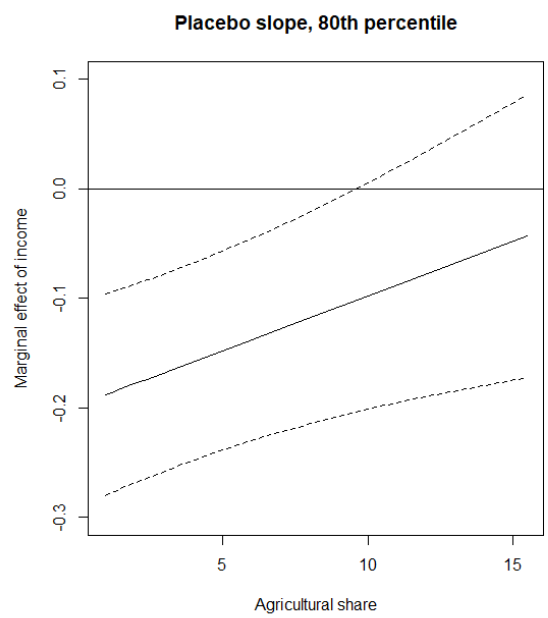

(d)

Figure 6.28.: Marginal effects of income on placebo group slopes by agricultural share and INSTEPP percentile

CHAPTER 7. CONCLUSION

In this dissertation, I have reviewed the evolution of the literature on intergenerational mobility in the economics literature, focusing on its development and key findings, as well as its overlaps with other important literatures, and developed three empirical projects which fill gaps I have identified in this literature. Overall, this dissertation contributes to this literature both by answering a set of empirical questions that inform our understanding of mobility, and by employing a conceptual and empirical approach that emphasizes explaining mobility over measuring it, and which integrates data occurring at several levels of aggregation, relying on insights from several related strands of literature. Neither of these contributions is definitive, as I am hardly the first, and will certainly not be the last, to make these points, but I am confident this perspective has value nonetheless.

My first empirical chapter asks how childhood poverty affects adulthood rank in the income distribution, motivated by concerns about the uniformity of mobility statistics the literature often assumes and by tensions in the child development literature regarding the shape of these effects. I complicate the standard definition of poverty by treating it as a composite of four distinct and continuous features, which allows me to detect features of the production function for adulthood rank that previous studies could not. At the same time, my emphasis on poverty keeps my work fairly grounded in metrics that policymakers can observe easily, which makes translating my results into policy implications, and into future studies engaged even more directly with policy interventions, fairly straightforward. I find a strong and highly nonlinear relationship between poverty duration and adulthood rank, with children experiencing one year of poverty reaching approximately the 44^{th} percentile of the income distribution, while children with fifteen years of poverty reach approximately the 23^{rd}. Other factors, which I had expected to help explain tensions in the literature, appear to matter very little, however, either on their own or through interactions with duration, which suggests that targeting programs based on children's cumulative disadvantage could be efficient.

In my second empirical chapter, I examine spatial heterogeneity in black-white racial gaps in levels of relative and absolute mobility. This approach to studying mobility is important because it combines two widely acknowledged facts about mobility – the large gap between blacks and whites and the degree of geographic variation in mobility – that we can still learn a great deal about, which can in turn inform our understanding of mobility more generally. While my results remain fairly preliminary, I find substantial variation in the relationship between household variables and gaps in expected mobility for blacks and whites, and in the degree to which community variables exacerbate or reduce the strength of these relationships. The association between mobility gaps and the number of siblings in the household, for example, is positive for blacks but negative for whites, and that negative local economic conditions – measured with the unemployment rate – increase the absolute value for both groups. This result, like many of my results in this chapter, is somewhat puzzling, and more work must be done, but it does underscore the size of the differences in mobility responses to local stimuli between groups, which highlights how much remains to be learned about these racial gaps in the community and regional setting.

Finally, I study the role of agricultural development in defining 20^{th} century trends in U.S. mobility. The data that allow me to answer questions about these processes do not allow me to study individual mobility in the same ways that I do in the previous chapters, but they let me compare trends in a very powerful way, as I can estimate state-level mobility for several years, and then explain variation in those trends. This lets me provide direct estimates of relationships between economic changes, in this case the structural transformation of agriculture, and mobility, which is a very direct way to learn about the historical sources of mobility. This application of this trend-oriented approach is particularly salient because the mobility literature is largely motivated by a reconsideration of assumptions about high mobility in the U.S., and studying trends in the middle of the 20^{th} century – the period that created this perception – makes evaluating certain sources of that mobility possible, in addition to allowing me to compare and contrast the middle and late decades of the century. I find that agricultural employment dynamics do relatively little to explain these trends, but that cotton cultivation does, an effect I attribute to the well-documented relationship between cotton cultivation and institutions that hinder human capital accumulation and labor mobility, which I expect to also hinder intergenerational mobility.

As I have noted throughout the dissertation and in this section, I believe each of these papers makes a contribution to the literature, but I believe my emphasis on learning about the sources of mobility has potential value as well. I am far from the first scholar to be interested in thinking about mobility in this way, as I believe my discussion of numerous papers on mobility and closely related topics, dating back to the 1990's, makes clear. This dissertation does, however, contribute to the movement towards using data from different sources and at different levels, most notably historical data, to explain mobility rather than to measure it. Work focusing on standard mobility topics, such as educational attainment and occupational choice, remains extremely important, and in fact my first chapter does more or less fall into this tradition, but that strand of the literature can only go so far in explaining the role of the major economic and institutional changes many scholars in this literature acknowledge as relevant. Understanding how we got to where we are, wherever that is, will require an understanding of mechanisms at a variety of scales, but it may also provide a sense of which policy levers might allow stakeholders to shift the equilibrium, as well as avenue for making important discoveries about broader aspects of human capital and family dynamics, all of which are very exciting possibilities, to which I sincerely hope my work here can contribute.

REFERENCES

Aaronson, D. and Mazumder, B. (2008). Intergenerational economic mobility in the United States, 1940 to 2000. *Journal of Human Resources*, 43(1):139–172.

Agostinelli, F. (2018). Investing in children's skills: An equilibrium analysis of social interactions and parental investments. *Working Paper*.

Agostinelli, F. and Wiswall, M. (2016). Estimating the technology of children's skill formation. *NBER Working Paper no. 22442*.

Almond, D. and Currie, J. (2010). Human capital development before age five. *NBER Working Paper No. 15827*.

Ananat, E. (2011). The wrong side(s) of the tracks: the causal effects of racial segregation on urban poverty and inequality. *American Economic Journal: Applied Economics*, 3(2):34–66.

Atkinson, A. (1980). On intergenerational income mobility in Britain. *Journal of Post Keynesian Economics*, 3(2):194–218.

Azariadis, C. (2006). *The Theory of Poverty Traps: What Have We Learned?*, chapter 1, pages 17–40. Poverty Traps. Russell Sage.

Azariadis, C. and Stachurski, J. (2005). *Poverty Traps*, volume 1A of *Handbook of Economic Growth*, chapter 5, pages 296–383. Elsevier.

Barkley, A. P. (1990). The determinants of the migration of labor out of agriculture in the United States, 1940-95. *American Journal of Agricultural Economics*, 72(3):567–573.

Becker, G. and Tomes, N. (1979). An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of Political Economy*, 87(6):1153–1189.

Becker, G. and Tomes, N. (1986). Human capital and the rise and fall of families. *Journal of Labor Economics*, 4(3):1–39.

Behrman, J. and Taubman, P. (1985). Intergenerational earnings mobility the United States: some estimates and a test of Becker's intergenerational endowments model. *Review of Economics and Statistics*, 67(1):144–151.

Bhattacharya, D. and Mazumder, B. (2011). A nonparametric analysis of black-white differences in intergenerational income mobility in the United States. *Federal Reserver Bank of Chicago Working paper WP 2007-12*.

Björklund, A. and Jäntti, M. (1997). Intergenerational income mobility in Sweden compared to the United States. *American Economic Review*, 87(5):1009–1018.

Björklund, A., Jäntti, M., and Lindquist, M. (2009). Family background and income during the rise of the welfare state: brother correlations in income for Swedish men born 1932-1968. *Journal of Public Economics*, 93(5):671–680.

Björklund, A., Lindahl, M., and Plug, E. (2006). The origins of intergenerational association: lessons from Swedish adoption data. *Quarterly Journal of Economics*, 121(3):999–1028.

Black, S., Devereux, P., Loken, K., and Salvanes, K. (2014). Care or cash? the effect of child care subsidies on student performance. *Review of Economics and Statistics*, 96(5):824–837.

Black, S. and Devereux, P. J. (2011). *Recent Developments in Intergenerational Mobility*, volume 4B of *Handbook of Labor Economics*, chapter 16, pages 1487–1541. Elsevier.

Blanden, J. (2013). Cross-country rankings in intergenerational mobility: a comparison of approaches from economics an sociology. *Journal of Economic Surveys*, 27:38–73.

Borjas, G. (1992). Ethnic capital and intergenerational mobility. *Quarterly Journal of Economics*, 107(1):123–150.

Borjas, G. (1994). Ethnicity, neighborhoods, and human capital externalities. *American Economic Review*, 85.

Borjas, G. (1995). Ethnicity, neighborhoods, and human capital externalities. *American Economic Review*, 85(3):365–390.

Bosquet, C. and Overman, H. (2019). Why does birthplace matter so much? *Journal of Urban Economics*, 110:26–34.

Brittain, J. (1977). *The Inheritance of Economic Status*. Brookings Institution.

Brooks-Gunn, J., Duncan, G., Klebanov, P. K., and Sealand, N. (1993). Do neighborhoods influence child and adolescent development? *American Journal of Sociology*, 99(2):353–395.

Brooks-Gunn, J. and Duncan, J. (1997). The effects of poverty on children. *The Future of Children*, 7(2):55–71.

Case, A. and Katz, L. (1991). The company you keep: the effects of family and neighborhood on disadvantaged youths. *NBER Working Paper No. 3705*.

Chamberlain, G. and Griliches, Z. (1975). Unobersvables with a variance-components structure: Ability, schooling, and the economic success of brothers. *International Economic Review*, 16(2).

Chetty, R., Friedman, J., Hendren, N., Jones, M., and Porter, R. (2018). The opportunity atlas: Mapping the childhood roots of social mobility. *Census Bureau Center for Economic Studies Working Paper 18-42*.

Chetty, R. and Hendren, N. (2015). The impacts of neighborhoods on intergenerational mobility: childhood exposure effects and county-level estimates. *Working paper*.

Chetty, R., Hendren, N., and Katz, L. (2016). The effects of exposure to better neighorhoods on children: new evidence from the Moving to Opportunity experiment. *American Economic Review*, 106(4):855–902.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014a). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Quarterly Journal of Economics*, 129(4):1553–1623.

Chetty, R., Hendren, N., Kline, P., Saez, E., and Turner, N. (2014b). Is the United States still a land of opportunity? Recent trends in intergenerational mobility. *No. w19844. National Bureau of Economic Research.*

Chyn, E. (2016). Moved to opportunity: The long-run effect of public housing demolition on labor market outcomes of children. *Working Paper.*

Clark, G. (2014). *The son also rises: Surnames and the history of social mobility.* Princeton, NJ: Princeton University Press.

Clark, G. and Cummins, N. (2015). Intergenerational wealth mobility in England, 1858-2012: Surnames and social mobility. *Economic Journal*, 125(582):61–85.

Collins, W. and Wanamaker, M. (2017). Up from slavery? African American intergenerational mobility since 1880. *NBER Working Paper 23395.*

Collins, W. J. and Margo, R. A. (2000). Racial segregation and socioeconomic outcomes: when did ghettos go bad? *Economics Letters*, 69:239–243.

Corak, M. (2006). Do poor children become poor adults? Lessons from a cross country comparison of generational earnings mobility. *Institute for the Study of Labor Working Paper.*

Crandall, M. and Weber, B. (2004). Local social and economic conditions, spatial concentrations of poverty, and poverty dynamics. *American Journal of Agricultural Economics*, 86(5):1276–1281.

Cunha, F. and Heckman, J. (2010). Investing in our young people. *IZA Discussion Paper Series, No. 5050.*

Cunha, F., Heckman, J., and Schennach, S. (2010). Estimating the technology of cognitive and noncognitive skill formation. *Econometrica*, 78(3):883–931.

Cutler, D. and Glaeser, E. (1997). Are ghettos good or bad? *Quarterly Journal of Economics*, 112(3):827–872.

Dahl, G. B. and Lochner, L. (2012). The impact of family income on child achievement: evidence from the earned income tax credit. *American Economic Review*, 102(5):1927–1956.

Davis, J. and Mazumder, B. (2016). The decline in intergenerational mobility after 1980. *Working Paper.*

Dearden, L., Machin, S., and Reed, H. (1997). Intergenerational mobility in Britain. *Economic Journal*, 107(440):47–66.

Diewert, W. E. (1992). Fisher ideal output, input, and productivity indexes revisited. *Journal of Productivity Analysis*, 3:211–248.

Duncan, G., Yeung, W. J., Brooks-Gunn, J., and Smith, J. (1998). How much does childhood poverty affect the life chances of children? *American Sociological Review*, 63(3):406–423.

Duncan, G., Ziol-Guest, K., and Kalil, A. (2010). Early-childhood poverty and adult attainment, behavior, and health. *Child Development*, 81(1):306–325.

Durlauf, S. (1996). A theory of persistent income inequality. *Journal of Economic Growth*, 1:75–93.

Durlauf, S. N. (2004). *Neighborhood Effects*, volume 4 of *Handbook of Regional and Urban Economics*, chapter 50, pages 2174–2230. Elsevier.

Durlauf, S. N. (2006). *Groups, Social Influences, and Inequality*, chapter 6, pages 141–175. Poverty Traps. Russell Sage.

Farrigan, T. and Parker, T. (2012). The concentration of poverty is a growing rural problem. *Amber Waves*.

Feigenbaum, J. (2015). Intergenerational mobility during the great depression. *Working Paper*.

Feigenbaum, J. (2018). Multiple measures of historical intergenerational mobility: Iowa 1915 to 1940. *Economic Journal*, 128(612):1–36.

Ferrie, J., Massey, C., and Rothbaum, J. (2016). Do grandparents and great-grandparents matter? Multigenerational mobility in the US, 1910-2013. *NBER Working Paper no. 22635*.

Fong, C., Hazlett, C., and Imai, K. (2018). Covariate balancing propensity score for a continuous treatment: application to the efficacy of political advertisements. *Annals of Applied Statistics*, 12(1):156–177.

Foundation, A. E. C. (2018). Children in poverty. https://datacenter.kidscount.org/data/tables/43-children-in-poverty-100-percent-poverty?loc=1&loct=1#detailed/1/any/false/870,573,869,36,868/any/321,322.

Gebremariam, G., Gebremedhin, T., and Schaeffer, P. (2011). Employment, income and migration in Appalachia: a spatial simultaneous equations approach. *Journal of Regional Science*, 51(1):102–120.

Ghatak, M. (2015). Theories of poverty traps and anti-poverty policies. *World Bank Economic Review*, 29:1–29.

Gorseline, D. (1932). *The Effect of Schooling upon Income*. PhD thesis, Indiana University.

Grawe, N. (2004). *Intergenerational Mobility for Whom? The Experience of High and Low Earnings Sons in International Perspective*, chapter 4, pages 58–89. Generational Income Mobility in North America and Europe. Cambridge: Cambridge University Press.

Haider, S. and Solon, G. (2006). Life-cycle variation in the association between current and lifetime earnings. *American Economic Review*, 96(4):1308–1320.

Hall, P., Li, Q., and Racine, J. (2007). Nonparametric estimation of regression functions in the presence of irrelevant regressors. *Review of Economics and Statistics*, 89(40):784–789.

Hall, P., Racine, J., and Li, Q. (2004). Cross-validation and the estimation of conditional probability densities. *Journal of the American Statistical Association*, 99(468):1015–1026.

Hanson, A. and Rohlin, S. (2013). Do spatially targeted redevelopment programs spillover? *Regional Science and Urban Economics*, 43(1):86–100.

Heckman, J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1):153–161.

Heckman, J. (2008). Schools, skills, and synapses. *Economic Inquiry*, 46(3):289–352.

Heckman, J. and Mosso, S. (2014). The economics of human development and social mobility. *Annual Review of Economics*, (6):689–733.

Hellerstein, J., Neumark, D., and McInerney, M. (2008). Spatial mismatch or racial mismatch? *Journal of Urban Economics*, 64(2):464–479.

Henderson, D. and Parmeter, C. (2015). *Applied Nonparametric Econometrics*. Cambridge University Press.

Hertz, T. (2005). *Unequal chances: Family background and economic success*, chapter Five: Rags, Riches, and Race., pages 165 – 191. Princeton University Press.

Hertz, T. (2007). Trends in the intergenerational elasticity of family income in the United States. *Industrial Relations: A Journal of Economy and Society*, 46(1):22–50.

Hertz, T. (2008). A group-specific measure of intergenerational persistence. *Economics Letters*, 100(3):415–417.

Hilger, N. (2017a). The great escape: Intergenerational mobility in the United States since 1930. *Working Paper*.

Hilger, N. (2017b). Upward mobility and discrimination: The case of Asian Americans. *Working Paper*.

Hornbeck, R. and Naidu, S. (2014). When the levee breaks: Black migration and economic development in the american south. *American Economic Review*, 104(3):963–990.

Jacob, B., Kapustin, M., and Ludwig, J. (2015). The impact of housing assistance on child outcomes: evidence from a randomized housing lottery. *Quarterly Journal of Economics*, 130(1):465–506.

Jencks, C. and Mayer, S. (1990). *The Social Consequences of Growing Up in a Poor Neighborhood*, chapter 4, page 186. Inner-City Poverty in the United States. National Academy Press.

Jung, Y. (2018). The long reach of cotton in the u.s. south: Tenant farming, mechanization, and low–skill manufacturing. *Working Paper*.

Kain, J. (1968). Housing segregation, Negro employment, and metropolitan decentralization. *Quarterly Journal of Economics*, 82(2):175–197.

Katz, L., Kling, J., and Liebman, J. (2001). Moving to opportunity in Boston: early results of a randomized mobility experiment. *Quarterly Journal of Economics*, 116(2):607–654.

Kling, J., Liebman, J., and Katz, L. (2007). Experimental analysis of neighborhood effects. *Econometrica*, 75(1):83–119.

Knapp, T. and White, N. (2016). The effect of youth poverty rates and migration on adult wage. *Journal of Regional Science*, 56(2):239–256.

Koball, H. and Jiang, Y. (2018). Basic facts about low income children. http://www.nccp.org/publications/pdf/text_1194.pdf.

Kumar, I., Kim, Y. J., Beaulieu, L., Delgado, M., Florax, R. J. G. M., Hertz, T., Li, H., Smith, T., Waldorf, B., Wilcox, M., and Zhalnin, W. (2019). Shed: A spatially harmonized county and tract socioeceonomic database for the U.S.A. *Working Paper*.

Lavergne, P. and Vuong, Q. (2000). Nonparametric significance testing. *Econometric Theory*, 16(4):576–601.

Lee, C. and Solon, G. (2009). Trends in intergenerational income mobility. *Review of Economics and Statistics*, 91(4):766–772.

Lefgren, L., Lindquist, M., and Sims, D. (2012). Rich dad, smart dad: decomposing the intergenerational transmission of income. *Journal of Political Economy*, 120(2):268–303.

Lefranc, A., Ojima, F., and Yoshida, T. (2014). Intergenerational earnings mobility in Japan among sons and daughters: levels and trends. *Journal of Population Economics*, 27(1):91–134.

Leventhal, T. and Brooks-Gunn, J. (2000). The neighborhoods they live in: the effects of neighborhood residence on child and adolescent outcomes. *Psychological Bulletin*, 126(2):309–337.

Levine, D. and Mazumder, B. (2002). Choosing the right parents: changes in the intergenerational transmission of inequality - between 1980 and the early 1990's. *Federal Reserve Bank of Chicago Working Paper Series, 2002-08*.

Levine, D. and Mazumder, B. (2007). The growing importance of family: evidence from brothers' earnings. *Industrial Relations*, 46(1):7–21.

Li, Q. and Racine, J. S. (2007). *Nonparametric Econometrics: Theory and Practice*. Princeton University Press.

Liu, H. and Zeng, J. (2009). Genetic ability and intergenerational earnings mobility. *Journal of Population Economics*, 22:75–95.

Lockwood-Reynolds, C. and Rohlin, S. (2015). The effects of location-based tax policies on the distribution of household income: evidence from the federal Empowerment Zone program. *Journal of Urban Economics*, 88:1–15.

Long, J. and Ferrie, J. (2015). Intergenerational occupational mobility in Great Britain and the United States since 1850. *American Economic Review*, 103(4):1109–1137.

Loury, G. (1977). *A Dynamic Theory of Racial Income Differences*, chapter 8, pages 153–188. Women, Minorities and Employment Discrimination. Lexington Books.

Loury, G. (1981). Intergenerational transfers and the distribution of earnings. *Econometrica*, 49(4):843–867.

Lucas, R. and Kerr, S. (2013). Intergenerational income immobility in Finland: contrasting roles for parental earnings and family income. *Journal of Population Economics*, 26(3):1057–1094.

Ludwig, J., Duncan, G., Gennetian, L., Katz, L., Kessler, R., Kling, J., and Sanbonmatsu., L. (2013). Long-term neighborhood effects on low-income families: evidence from Moving to Opportunity. *American Economic Review*, 103(3):226–231.

Martin, R. (2001). The adjustment of black residents to metropolitan employment shifts: how persistent is spatial mismatch? *Journal of Urban Economics*, 50(1):52–76.

Martin, R. (2004). Can black workers escape spatial mismatch? Employment shifts, population shifts, and black unemployment in American cities. *Journal of Urban Economics*, 5(1):179–194.

Massey, D. (1990). American apartheid: segregation and the making of the underclass. *American Journal of Sociology*, 96(2):329–357.

Massey, D. and Denton, N. (1993). *American Apartheid: Segregation and the Making of the Underclass*. Harvard University Press.

Massey, D., Gross, A., and Eggers, M. (1991). Segregation, the concentration of poverty, and the life chances of individuals. *Social Science Research*, 20(4):397–420.

Massey, D., Gross, A., and Shibuya, K. (1994). Migration, segregation, and the geographic concentration of poverty. *American Sociological Review*, 59(3):425–445.

Mayer, S. E. (1997). *What Money Can't Buy: Family Income and Children's Life Chances*. Harvard University Press.

Mazumder, B. (2005). Fortunate sons: new estimates of intergenerational mobility in the United States using social security earnings data. *Review of Economics and Statistics*, 87(2):235–255.

Mazumder, B. (2008). Status similarities and economic inequality in the U.S. *Journal of Population Economics*, 21(3):685–701.

Mazumder, B. (2015). Estimating the intergenerational elasticity and rank association in the U.S.: Overcoming the current limitations of tax data. *Federal Reserve Bank of Chicago Working Paper 2015-04*.

Milligan, K. and Stabile, M. (2011). Do child tax benefits affect the well-being of children? evidence from Canadian child benefit expansions. *American Economic Journal: Economic Policy*, 3(3):175–205.

Mitnik, P., Bryant, V., Grusky, D., and Weber, M. (2015). New estimates of intergenerational mobility using administrative data. *working paper*.

Modalsli, J. (2016). Multigenerational persistence: evidence from 146 years of administrative data. *Statistics Norway Research Department Working Paper 850*.

Modalsli, J. (2017). Intergenerational mobility in Norway, 1865-2011. *Scandinavian Journal of Economics*, 119(1):34–71.

Mundlak, Y. (2005). Economic growth: Lessons from two centuries of american agriculture. *Journal of Economic Literature*, 63:989–1024.

Nybom, M. and Stuhler, J. (2016). Interpreting trends in intergenerational mobility. *Working Paper*.

OECD (2018). OECD family database: child well-being module. http://www.oecd.org/els/soc/oecdfamilydatabasechildwell-beingmodule.htm.

Olivetti, C. and Paserman, M. D. (2015). In the name of the son (and the daughter):Intergenerational mobility in the United States, 1850-1940. *American Economic Review*, 105(8):2695–2724.

Olivetti, C., Paserman, M. D., and Salisbury, L. (2018). Three-generation mobility in the United States, 1850-1940: the role of maternal and paternal grandparents. *Explorations in Economic History*, 70:73–90.

Oreopoulos, P. (2003). The long-run consequences of living in a poor neighborhood. *Quarterly Journal of Economics*, 118(4):1533–1575.

Page, M. and Solon, G. (2003). Correlations between brothers and neighboring boys in their adult earnings: the importance of being urban. *Journal of Labour Economics*, 21(4):831–855.

Partridge, M. and Rickman, D. (2005). High-poverty nonmetropolitan counties in America: can economic development help? *International Regional Science Review*, 28(4):415–440.

Partridge, M. and Rickman, D. (2007a). Persistent pockets of extreme American poverty and job growth: is there a place-based policy role? *Journal of Agricultural and Resource Economics*, 32(1):201–224.

Partridge, M. and Rickman, D. (2007b). Persistent rural poverty: is it simply remoteness and scale? *Review of Agricultural Economics*, 29(3):430–436.

Partridge, M. and Rickman, D. (2008a). Does a rising tide lift all metropolitan boats? Assessing poverty dynamics by metropolitan size and county type. *Growth and Change*, 39(2):283–312.

Partridge, M. and Rickman, D. (2008b). Place-based policy and rural poverty: insights from the urban spatial mismatch literature. *Cambridge Journal of Regions, Economy and Society*, 1(1):131–156.

Popkin, E., Rosenbaum, J., and Meaden, P. (1993). Labor market experiences of low-income black women in middle-class suburbs: evidence from a survey of Gautreaux program participants. *Journal of Policy Analysis and Management*, 12(3):556–573.

Ruggles, S., Flood, S., Goeken, R., Grover, J., Meyer, E., Pacas, J., and Sobek, M. (2018). IPUMS USA: Version 8.0 (dataset).

Ruggles, S., Genadek, K., Goeken, R., Grover, J., and Sobek, M. (2017). Integrated public use microdata series: Version 7.0.

Sacerdote, B. (2007). How large are the effects from changes in family environment? A study of Korean American adoptees. *Quarterly Journal of Economics*, 122(1):119–157.

Sampson, R. F. and Morenoff, J. D. (2006). *Durable Inequality: Spatial Dynamics, Social Processes, and the Persistence of Poverty in Chicago Neighborhoods*, chapter 7, pages 176–203. Poverty Traps. Russell Sage.

Sewell, W. and Hauser, R. (1975). *Education, Occupation, and Earnings. Achievement in the Early Career*. Academic Press Inc.

Shea, J. (2000). Does parents' money matter? *Journal of Public Economics*, 77(2):155–184.

Solon, G. (1992). Intergenerational income mobility in the United States. *American Economic Review*, 82(3):393–408.

Solon, G. (1999). Intergenerational mobility in the labor market. *Handbook of labor economics*, III:1761–1800.

Solon, G. (2002). Cross-country differences in intergenerational earnings mobility. *Journal of Economic Perspectives*, 16(3):59–66.

Solon, G. (2014). Theoretical models of inequality transmission across multiple generations. *Research in Social Stratification and Mobility*, 35:13–18.

Solon, G. (2015). What do we know so far about multigenerational mobility? *NBER Working Paper 21053*.

Solon, G., Corcoran, M., Gordon, R., and Laren, D. (1991). A longitudinal analysis of sibling correlations in economic status. *Journal of Human Resources*, 26(3):509–534.

Solon, G., Page, E., and Duncan, G. (2000). Correlations between neighboring children in their subsequent educational attainment. *Review of Economics and Statistics*, 82(3):383–392.

Torche, F. and Corvalan, A. (2016). Estimating intergenerational mobility with grouped data: a critique of Clark's *The Son Also Rises*. *Sociological Methods and Research*, pages 1–26.

Vigdor, J. (2002). Locations, outcomes, and selective migration. *Review of Economics and Statistics*, 84(4):751–755.

Wagmiller, R., Lennon, M. C., Kuang, L., Alberti, P., and Aber, J. L. (2006). The dynamics of economic disadvantage and children's life chances. *American Sociological Review*, 71:847–866.

Waldorf, B. and Kim, A. (2015). Defining and measuring rurality in the US: from typologies to continuous indices.

Wang, S. L., Heisey, P., Schimmelpfennig, D., and Ball, E. (2015). Agricultural productivity growth in the United States: Measurements, trends, and drivers. *ERS Economic Research Report 189*.

Weinberg, B. (2000). Black residential centralization and the spatial mismatch hypothesis. *Journal of Urban Economics*, 48(1):110–134.

Weinberg, B. (2004). Testing the spatial mismatch hypothesis using inter-city variations in industrial composition. *Regional Science and Urban Economics*, 34:505–532.

Whatley, W. (1985). A history of mechanization in the cotton south: The institutional hypothesis. *Quarterly Journal of Economics*, 100(4):1191–1215.

Whatley, W. (1987). Southern agrargian labor contracts as impediments to cotton mechanization. *Journal of Economic History*, 47(1):45–70.

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

Wodtke, G. T., Harding, D. J., and Elwert, F. (2016). Neighborhood effect heterogeneity by family income and developmental period. *American Journal of Sociology*, 121(4):1168–1222.

Wu, J. and Gopinath, M. (2008). What causes spatial variations in economic development in the United States? *American Journal of Agricultural Economics*, 90(2):392–408.

Zimmerman, D. (1992). Regression toward mediocrity in economic stature. *American Economic Review*, 82(3):409–429.

VITA

Tim Smith is an applied economist interested in labor economics, regional economics, and applied econometrics, with a particular interest in using the conceptual and methodological tools of all three literatures to understand intergenerational mobility and economic opportunity. He received a Bachelor's Degree in Public Policy and Global Studies in May 2011 from the University of North Carolina at Chapel Hill, and a Master's Degree in Agricultural Economics from Purdue University in May 2014. In August 2014, he began his Ph.D. in Agricultural Economics at Purdue. His current research agenda combines expanding his dissertation work on the sources of intergenerational mobility with more applied work on scale decisions in public policy, economic opportunity in the rural U.S., and the role of nutrition in human capital accumulation in the developing world.