Intergenerational Mobility and Housing Wealth in the United States*

Ariel Binder† Max Risch‡ John Voorheis§

March 2025

Preliminary Draft
(Please do not cite or circulate)

Abstract

Housing represents the largest asset owned by all but the very wealthiest households, yet we still face large gaps in our knowledge of the distribution of this housing wealth and its transmission over time. Using a novel linked dataset combining Decennial Census data, property tax records and administrative tax data from the Internal Revenue Service, we provide new estimates of the intergenerational transmission of housing wealth. We document how patterns in intergenerational housing wealth mobility vary over time and across demographic groups, and provide a theoretical framework decomposing housing wealth mobility into direct channels working through parental housing assets themselves, and indirect channels working through parental housing asset's relationship with children's adult labor market outcomes. We present evidence that between metro area differences in housing supply elasticity are an important determinant of the patterns in intergenerational housing wealth mobility we observe.

*We thank conference participants at the Opportunity Insights Annual Conference, the AEA Annual Meetings, and seminar participants at the Norwegian School of Economics and ETH Zurich for valuable feedback. Any opinions and conclusions expressed herein are those of the author(s) and do not reflect the views of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product (Data Management System (DMS) number: 7511151, Disclosure Review Board (DRB) approval numbers: CBDRB-FY24-CES014-005, CBDRB-FY24-CES023-001, CBDRB-FY25-0171 and CBDRB-FY25-0186).

†Center for Economic Studies, U.S. Census Bureau

‡Tepper School of Business, Carnegie Mellon University

§Center for Economic Studies, U.S. Census Bureau

1. Introduction

Homeownership is widely considered to be a key part of the "American Dream". For families that are homeowners, the home they live in is in many cases the most important financial asset they hold. Housing is *important*, and yet our understanding of even relatively basic facts about the distribution of the stock of wealth held in housing, or how this stock of wealth is transmitted over time – remains limited. Using a novel dataset linking administrative tax records, household surveys, and residential property ownership and valuation records, we provide new evidence on the intergenerational relationship between parent and child housing wealth in the contemporary United States.

A long literature has estimated the intergenerational mobility (IGM) of income in the United States and mechanisms driving mobility.[1]. Yet, housing is not reflected in the income concepts measured in administrative records and is the most important capital asset for most U.S. households. In the U.S. there is no comprehensive administrative data on individual wealth, so it has been difficult to make precise claims about intergenerational persistence of resources beyond income, particularly with sufficient statistical power to do heterogeneity analysis across race or geography.[2] Our unique linked data approach provides comprehensive information on residential property values for a large representative sample of adults, giving a new picture of IGM of housing in the U.S. with sufficient power to investigate differential mobility across racial groups and geography and analyze drivers of IG relationships. A link between detailed income tax data and housing value information also allows us to analyze the joint and independent relationships between income and housing wealth mobility.

Building a novel dataset linking administrative tax records, property ownership and valuation records, and the 2000 Decennial Census Long Form, we provide new evidence on the intergenerational relationship between parent and child housing assets, as well as the

[1] See recent surveys by Black and Devereux (2011) and Mogstad and Torsvik (2023)

[2] A recent literature has used administrative data to estimate differences in IGM across racial groups and geography (Collins and Wanamaker (2022), Chetty et al. (2019),Chetty and Hendren (2018), and Chetty et al. (2014)), providing new evidence drivers of intergenerational inequality in the U.S..

interplay between asset and income transmission. Our data consist of 3.4 million families and cover children born between 1978 and 1986. We develop a reweighting procedure to correct for imperfect assignment of owner identifiers to the property records that leverages additional ownership information from the American Community Survey. We show that our procedure delivers unbiased estimates under a weak set of assumptions. We then use rank-rank regressions to estimate the intergenerational mobility of parent and child resources in the U.S., on average and across subgroups. In the first part of the paper we explore the relationship between parent and child housing assets in detail, developing statistics on aggregate mobility in the U.S., investigating variation across race and geography. In the second part, we investigate how this new information on intergenerational mobility of housing wealth informs our broader understanding of intergenerational mobility of resources in the U.S. by investigating the joint relationship between parent and child income and assets.

In the first part of the empirical investigation, we establish new patterns in the intergenerational persistence of housing assets. Understanding the intergenerational correlations of housing is somewhat distinct from income relationships because of the mass of parents and children that do not own a home (renters). Therefore, we non-parametrically estimate relationships between parent renters and child outcomes and parent home owners, and we separately investigate the relationship between parent resources and child home ownership and home values conditional on ownership.

First, we present the first evidence rank-rank relationship between parent and child housing assets in the U.S. The non-parametric rank-rank relationship shows a linear and relatively steep slope of parent and child housing assets conditional on parent home ownership, but also that children of parent renters have similar average housing ranks as children of parents in the bottom rank of home values conditional on home ownership. Conditional on parent home ownership, there is a linear slope of parent and child housing assets of 0.424, which reflects substantially more persistence than recent estimates of income persistence in Chetty et al. (2020) and in our estimates. However, we also see that children's average housing

outcomes are not better from parents owning a low value home relative to renting.[3] Further, we decompose the rank-rank slope into the portion attributable to increasing child home ownership rates and increasing child home values increasing. In the bottom portion of the parent housing distribution, most of child average gains with increased parental rank come from higher probabilities of home ownership. The share attributable home values increase throughout the parental distribution and in the top third of the parental distribution, most of the persistence comes from increases in average child home value.

Next, we compare the intergenerational housing mobility of Black and White families, and show substantial differences in these relationships. There is high IG persistence for White families (rank-rank slope of 0.425), while there is high mobility for Black families (rank-rank slope of 0.26) which is driven by extreme downward mobility, even at the top of the parent asset distribution. For example, the children of Black parents in the top 5% of the housing asset distribution are on average in the bottom half of the housing distribution of their generation. There are also substantial differences in average housing assets for children of parents at the bottom of the housing distribution (i.e. differences in absolute mobility. This combination implies that there are large differences in child outcomes across racial groups conditional on parent housing rank, and that this gap grows considerally across the parental distribution; the difference in White and Black average ranks is about 5 ranks in the 25^{th} percentile, about 15 ranks at the 75^{th} percentile and about 20 ranks at the 98^{th} percentile.

We show that large disparities in home ownership rates among Black and White children across the parent distribution are important drivers of these racial gaps. Children from Black households in the top 10% of the parental distribution are less likely to own a home than children of white parents in the 30^{th} percentile of the distribution. Given this, we explore the roles of extensive margin ownership relative to the intensive margin of asset value (the value of the home conditional on ownership) in explaining the Black-White IG gap. To do so, we

[3]This result is robust to various controls, including detailed local area and parental income controls.

develop a counterfactual exercise where we scale-up the ownership rates of Black children to equal those of White children within each parent asset bin. This exercise shows that the extensive margin of homeownership is a primary driver of the observed gap. Yet, the economic implications of this exercise depend on who the marginal black homeowner would be, i.e. how much their home would be worth if they were to buy a home or if a policy brought them into homeownership. We use various methods to impute a counterfactual housing value for the "new" Black homeowners, including capitalizing observed rent payments and estimating the value of first homes conditional on observables using observed first-time purchases in our data. Together, these exercises suggest that if homeownership rates were equalized, average ranks of Black and White children would be quite similar in the bottom half of the parental distribution, but that for Black and White children of parents at the top of housing distribution there would remain substantial conditional gaps in housing values.

Next, we examine differences in the IGM of housing assets across counties within the U.S. In doing so, we develop a public-use dataset and data tool recording these various wealth mobility statistics at the county level, covering at least 3,100 counties for the full population and the White population, and 1,250 counties for the Black and Hispanic populations. These data show substantial heterogeneity in the IG persistence of housing assets across the U.S., and in the absolute mobility of housing. We separately show IG relationships as related to home ownership rates and home values.

In the next part of the empirical analysis we use our unique link between asset and income information from the tax records to investigate the relationship between intergenerational mobility of income and wealth. While parental income and housing are highly correlated on average, in the middle of the distribution, there is substantial variation in parent housing conditional on their income. When controlling for parent housing assets, the intergenerational persistence of income significantly decreases, showing a substantial independent role of parental assets in explaining child outcomes. We complement this with a simple statistical model showing that IG income-income relationships alone will understate IG persistence in

total resources when reported income flows represent a mix of returns of assets with various realized rates of return. For example, housing produces no regular realized income flow on tax returns so will be effectively missed as a portion of parent resources.

Building on this, we use a static model of intergenerational capital accumulation to motivate a simultaneous equation model estimation approach. In the model parent resources can impact child resources in various ways. First, an "indirect effect" reflects that more parental resources may allow children to earn more in the labor market (perhaps through human capital accumulation or access to networks) and children with more labor market earnings are able to accumulate more capital. Second, through a "propensity to save;" children of parents different levels of wealth may inherit different propensities to save out of each additional dollar earned. Third, parental resources may have a "direct effect" on child wealth; children of wealthier parents may have more wealth on average even conditional on their earnings.

Using a simultaneous equation model, we find empirically that on average parent resources have important indirect and direct effects, and that the "propensity to save" channel is less important for explaining intergenerational persistence. We estimate that the direct effect is of a similar magnitude as the indirect effect and the direct effect is completely driven by non-labor capital wealth (as opposed to wealth generated by parental earnings). That is, even conditional on having similar earnings, children of wealthier parents have substantially more assets on average and this drives about half of intergenerational persistence in housing. We decompose these channels into their influence on child home ownership rates and home values conditional on ownership. The correlation between parent assets and child home ownership is almost fully associated with the indirect effect, that children of wealthier parents have higher earnings. On the other hand, the large direct effect is associated with children having more valuable homes conditional on ownership. That is, conditional on child income, children of wealthier parents are not substantially more likely to own a home, but they do own significantly more valuable homes.

We also use this framework to investigate the Black-White IGM gaps in housing. This framework shows that the gap in absolute mobility is driven by indirect effects, that is, for children of parents at the bottom of the distribution, Black-White gaps in housing are fully associated with the Black-White differences in labor market earnings. But, we also find that almost the entire difference in persistence is driven by direct effects. That is, if child labor market outcomes were equalized conditional on parent assets, there would still be a growing gap between the assets of Black and White children across the parental distribution. At the top of the parental distribution, only about 35% of the Black-White asset gap is explained by differences in child labor market earnings.

This underscores the independent role of capital wealth for understanding intergenerational persistence of resources in the U.S., with implications for how policy and macroeconomic trends can affect intergenerational inequality. While much prior work rightly discusses the role of equalizing access to human capital accumulation and resulting labor market returns in reducing intergenerational inequality, our results show that there is also a substantial independent role of equalizing access to capital markets and capital accumulation. The final section of the paper will investigate the economic and policy implications of the empirical results.

In the final section, we examine drivers of local variation in housing wealth mobility across the United States. While the previous discussion shows that housing assets have an important independent role of explaining child outcomes, this allows us to examine an important potential driver in variation on asset allocation. We leverage new estimates of housing supply elasticities in metro areas (from Baum-Snow and Han (2024)) combined with observed cohort variation in our dataset, to construct a shift-share difference-in-difference design at the county level. In an epoch of rapidly expanding aggregate demand for housing, our approach aims to identify differential price shocks across counties by interacting the national quantity change with the county-specific slope of the supply curve. For example, the oldest children in our sample experienced ages 30-36 during 2008-2014, while the youngest

experienced these ages during 2016-2022. The Case-Shiller national home price index rose by 48.1% between these two periods, while median personal income rose by 30.6%, according to the Federal Reserve Economic Database. Our difference-in-difference regression framework allows us to control for nonparametric cohort trends in mobility according to initial county characteristics that may be correlated with local demand shocks not already netted out in the mobility calculus (i.e. controls for parental wealth), such as geography, city size, and county-wide parental ownership rates and average valuations.

Relative to the between cohort trend in housing wealth mobility experienced by children growing up in more elastic markets, we find a trend toward more immobility—i.e. lower absolute mobility and a higher dependence of child wealth on parent wealth—among children growing up in more inelastic markets. We are in the process of scaling these effects into responses of mobility to asset price, as well as understanding how these responses are generated by movements in direct versus indirect effects. Our reduced-form estimates so far suggest that housing supply expansions can help stabilize housing price growth and promote intergenerational wealth mobility in the population at large, without much affecting Black-White mobility gaps.

2. Data

We leverage the Census Bureau's internal data linkage infrastructure to develop a sample of several million parent-child pairs where we can observe both generations at prime home-ownership and wealth accumulation ages. Such a sample requires three key sources of information: i) parent-child links, either through household survey rosters or tax filing records; ii) information on the parent(s)' housing and other assets; and iii) information on the children's adulthood levels of housing and other assets.

2.1. Overview of intergenerational data

We use the 2000 Census Long Form (2KLF) as a primary source of the first two pieces of information. The Long Form sampled 1 in 6 U.S. households during the year 2000, and was the precursor to the American Community Survey (ACS). It provides information on household and family structure, sociodemographic variables, and total family income along with its breakdown into various components. It also provides housing asset information, i.e. whether the householder owns the given housing unit, and if so, whether the householder possesses a mortgage on the property and the householder's estimate of the value of the home. We associate these potential homeowners to their children using householder relationship identifiers. The 2KLF data, along with many other restricted-access Census datasets, also contain unique person numbers called Protected Information Keys (PIKs), which facilitate the linkage of (de-identified) personal information across different household surveys and administrative records (Layne et al., 2014).

To measure the children's housing assets accumulated in adulthood, we use the universe of deed, property tax assessment, and market valuation records for the years 2019-2021, provided to Census by the vendor Black Knight, Incorporated (BK), as described in the next subsection. We use PIKs to link information about the children's socioeconomic circumstances from the 2KLF to their housing records in BK. Because we require a set of individuals who can reliably be linked to their parents from 2KLF household rosters, *and* who were old enough by the year 2021 to be potential homeowners, we consider all 2KLF individuals aged 14-16 and who were children of householders. This sample yields NN parent-child pairs.

We supplement this main intergenerational sample with two additional samples containing slightly older children. This is because the the average dependent sampled from the Long Form was 36 years old in 2021. Home-ownership and housing wealth at this age, while important, may not reflect lifetime holdings in the housing market.[4] However, Census

[4]Efforts are underway to assign PIKs to the 1980 and 1990 Census Long Forms, and should be complete by 2026.

Bureau and affiliated researchers have developed the Databank, which contains a subset of federal tax return information available for research based on a joint partnership between Census and the IRS. The Databank contains sociodemographic and annual taxable-income information on everyone who filed a federal income tax return in each tax year during 1994-5 and 1998-2021,[5] as well as selected prior years. The Databank includes the PIKs of all dependents claimed by each filer beginning with the 1994 tax year, which allows researchers to very closely approximate parent-child links with claimer-dependent links.

Accordingly, we consider all dependents aged 14-16 in the 1994 Databank year whose primary or spousal claimer was *also* who a sampled householder in the 2KLF (so that we can recover information on parental housing wealth). The average dependent in this sample was 42 years old in 2021, which is close to the age at which the householder's ownership rate peaks in the 2021 ACS [back this up in public data]. This sample contains NN parent-child pairs. Finally, we construct an analogous Databank sample based on the 1998 year, which contains an additional NN parent-child pairs not found in either of the previous samples. Although not all families file taxes in every tax year, we verified that including the 1995 and 1999 years would barely raise our sample size. This is due to i) the narrow age range of the dependents we need to consider, ii) the child tax credit, which provides a strong incentive for claimers to file taxes, and iii) the requirement that one of the claimers appears as a householder in the 2KLF.

To consider relationships between income and wealth, we use the 2KLF and Databank to construct a composite rank measure of childhood family income (CFI). We replicate the "tax-unit" concept of family income introduced by Chetty et al. (2014) and used in following studies of intergenerational income mobility in tax records. Based on how adjusted gross income (AGI) is summarized on federal tax returns, this is the sum of the primary filer's personal income and that of their legal co-residing spouse (or zero if none exists).[6] To

[5] Census did not receive federal tax return information for tax years 1996 and 1997.

[6] Footnote about married filing jointly versus separately, cohabitations, etc. A small share of intact marriages with children do not file jointly.

smooth out transitory income variations as well as differences in income measurement across sources, we average two years of AGI information in the Databank together with the sum of householder and spousal (if householder is married to a coresident spouse) total personal incomes as reported in the 2KLF.[7] For the main 2KLF sample, we consider the householder's AGI from the 1999 and 2001 tax years in the Databank along with 2000 "tax-unt" income as reported in the 2KLF itself. For the 1994 (1998) Databank sample, we consider 1994 and 1995 (1998 and 1999) AGI information along with 2KLF tax-unit income. *Within* each sample (i.e. to adjust for inflation and other cohort effects, so that our rank measure is comparable across samples), we average these three years of childhood family income (CFI) together and then construct a national CFI rank measure based on the averaged value.

2.2. Overview of administrative property records

To measure homeownership and housing wealth for our sample of children, we leverage data from local property tax records. These records are collected by local county assessor's offices for the purposes of administering local property tax systems. These county property tax data are typically aggregated by commercial data providers – the Census Bureau has contracted with Black Knight to acquire these data, which has been incorporated into the data linkage infrastructure described above. The Black Knight assessment data contains property characteristics, including assessed and market value for all assessed properties, as well as estimated property valuation from an automated valuation model. Individual identifiers (PIKs) for the owners of properties have been assigned by the PVS system, which we then use for linkage to the intergenerational dataset above. We link all children in the intergenerational dataset to the Black Knight data, assigning extensive margin homeownership based on whether they can be linked to any property in the 2019-2021 Black Knight data, and intensive margin gross housing wealth based on the value of any linked properties.

[7] Following Chetty et al. (2014), we drop negative incomes before constructing averages, as these usually reflect transitory accounting choices of business owners who have high permanent incomes and asset holdings.

3. Methods

We investigate the intergenerational mobility using intergenerational rank-rank relationships.

3.1. Weighting

Our estimating sample is drawn from the 2000 Long Form Decennial Census, but requires linkage to multiple administrative and third party datasets. These linkages are not perfect – not all records in the Decennial Census can be assigned a PIK, and similarly the underlying Black Knight and IRS tax data are incompletely PIK'ed. Previous work (e.g. Bond et al. (2014)) suggest that PIK assignment rates are not random, but rather are correlated with important demographic characteristics. Thus our final linked sample, weighted only by the long form weights, will potentially suffer from selection. Additionally, because we do not observe the full lifespan of our children in adulthood, we cannot fully capture ownership histories for each birth year cohort (since we only observe ownership information from 2019-2021). To address these issues, we adjust the long form weights so that we match the cross-sectional pattern in aggregate homeownership rates for our set of birth cohorts. We use observed age variation within our dataset to discipline this adjustment according to a simple and an intuitive proportionality assumption.

3.2. Main analytical approach - rank-rank relationships

Our main approach to analyzing the intergenerational persistence of parent and child housing assets is a rank-rank analysis. Housing is particularly conducive to using rank-rank slopes, rather than intergerational elasiticites, because of the masses of parent and child renters (0 housing assets) which would not be incorporated in log transformations and ii) because of the extreme skew in the distribution of assets. Therefore, we use a rank-rank slope framework, but non-parametrically allow for differences in child outcomes for the mass of pareent renters.

We then estimate rank-rank relationships between parent and child housing using regres-

sions of the form,

$$H_i^k = \alpha_i + \beta H_i^p + \delta \mathbf{1}(parent\ owners_p) + u_i, \tag{1}$$

were β represents the rank-rank relationship between parent housing assets H_i^p and child housing assets H_i^k conditional on parent home ownership; having parents one rank higher in the distribution of H_i^p is associated with children being β ranks higher in the distribution of H_i^k on average. The coefficient δ identifies the difference in average outcomes for children of parent renters relative to the lowest rank of home owners. This allows us to define three IGM parameters of interest, i) "β," or relative mobility conditional on parent home ownership, ii) "α," the average rank of children of parents at the bottom of the home value distribution, and iii) "δ," the difference in average ranks between children of renters and children of parents at the bottom of the home value distribution.

When we investigate differential rank-rank relationships across groups, like racial groups or region in the U.S., ranks are defined at the population level, not within group. I.e. parents and kids of group j are ranked within the population distribution of H_p and H_k, not the subgroup distribution H_j^p and H_j^k.

More generally, when we estimate rank-rank relationships between other measures of parent and child resources, such as income, or examine how various factors co-determine child outcomes, we estimate a rank-rank regression of the form:

$$Y_i^k = \alpha_i + \beta Y_i^p + \Gamma X_i + u_i, \tag{2}$$

Where Y_i^p and Y_i^k are parent and child ranks in the distributions of Y^p and Y^k respectively and X_i is a vector of other covariates of interest.

4. Intergenerational Mobility of Housing Wealth

In this section, we investigate the intergenerational mobility of housing assets in the U.S.. We first discuss measurement issues before presenting aggregate estimates of intergenerational

mobility of housing wealth in the U.S.. We then investigate differential mobility by race and make comparisons.

4.1. Measuring Intergenerational Mobility of Housing Wealth

We want to understand intergenerational mobility of housing wealth using rank-rank regressions. There are two major issues for estimating the relevant aggregate parameters. First, there is an important extensive margin - many parents and kids own zero housing and there is a discontinuous jump in housing wealth upon ownership. If the relationship between parents moving from no housing to being in the bottom percentile of housing wealth conditional on ownership is different than that of moving up a general rank in the distribution, then a simple linear model will (i.e. a regression of child housing wealth rank on and intercept and parent housing asset rank) will misspecifiy the true intergenerational relationship. We account for this by estimating the relationship between parent renters and child outcomes non-parametrically, as described above.

Second, as discussed in (CITES), the estimates of IGM depend on the where parents and children are in the life-cycle of income or asset accumulation. This may be particularly salient when focusing on housing, where accumulation is not smooth and involves a large initial investment. This means there can be large re-ranking in the asset distribution when children purchase their first home.[8] We observe parents at an average of 52 years old and children at ages 35-42. If children are not at the same point of relative asset accumulation, then β may not represent the intergenerational parameter of permanent resources.

4.1.1. Life-cycle of Home ownership

We empirically investigate the importance of the age at which children are observed for measurement of the IG relationship of housing assets. We have three waves of parents and kids as described in Section 2. We directly estimate the IG housing relationships across these

[8]For example, with income it may be unlikely that a child jumps from the bottom of the income distribution to the top by measuring income just a year or two later. For housing wealth this is quite possible when a child purchases their first home.

waves; in the youngest wave children are 35-36 years old, in the middle wave children are 37-38 years old and in the oldest wave children are 40-42 years old. Panel A of Appendix Figure A.1 shows the non-parametric relationship for the youngest and oldest kids in the sample and reports the associated α, β and δ parameters. There is an evident difference in the α and β terms across these age groups. In particular intergenerational persistence increases as kids age. This is in contrast to what we see for rank-rank income relationships which are quite similar across these ages (Appendix Figure A.2).[9] The extensive margin of homeownership and the age at which wealth is measured appears important for estimating rank-rank relationships of housing assets. Ideally, we should estimate permanent IG relationships if we observe parent and children at the same point of their asset accumulation (Kopczuk cite). Since even the oldest children are younger than the average age of parents in the sample, we conduct a weighting exercise to estimate the IG relationships when the children in the sample are the age of their parents.

For this reason, we use the life-cycle-adjusted (LCA) weights discussed in Section 2 as our main specification. This weighting procedure reweights children so that relative home ownership rates conditional on race, parent income and parent housing rank, match those of the oldest cohort, but are scaled up to match aggregate ownership rates (by race) for older cohorts (targeting 52 year olds) as observed in ACS data. This rewighting technique, which targets home ownership rates, but not asset value ranks conditional on ownership, assumes that distribution of assets conditional on ownership do not change as children age, only the probability of owning a home. We check that this is a reasonable assumption by estimating rank-rank relationships across our cohorts conditional on children owning a home. Panel B of Appendix Figure A.1 shows no evidence of differential rank-rank relationships across age conditional on ownership.

Therefore, the assumption associated with the LCA weights is that relative home ownership rates old across race and parent income and housing distributions at 52 years old,

[9]That the IG income relationship is fairly stable once children reach their thirties is consistent with previous studies (CITE).

mirror those of 42 year olds, just that ownership rates are higher. The implication of this process for child ownership rates across the parental housing distribution is depicted in Panel C of Appendix Figure A.1.[10] Going forward, our main estimates apply LCA weights, but we discuss robustness where relevant.

4.2. Intergenerational Mobility of Housing Assets in the U.S.

Figure 1 shows our main estimates of the intergenerational mobility in housing wealth in the U.S. The navy blue dots show the average child rank conditional on parent rank, the light blue dashed line shows the median child rank for parents of a given rank and the red and green dashed lines shows the 25th and 75th percentile child rank conditional on parent rank, respectively. Focusing first on the average rank-rank relationship, we see that the relationship is highly linear conditional on parent ownership. We find a β parameter of 0.427. This represents a substantively higher persistence in housing asset rank than for income which has a β of 0.347, as shown in Appendix Figure B.7.[11] We also see a non-linearity at the bottom of the parental distribution. Children of parent renters have similar average ranks as those of parents at the bottom of the home value distribution. We see that on average children with parents at the bottom of the home ownership distribution are in the 36.5^{th} percentile of their distribution ($\alpha = 36.5$), and that, on average, children of parent renters are 1.74 ranks higher ($\delta = 1.74$).

The quantiles provide further context. As parent wealth increases the median diverges from the mean. In the top 5% of parental housing assets, the median child rank is substantively higher than the mean - half of children are in the top 20% of the housing wealth distribution. The 75th percentile shows that 25% of children with parents in the top 5% are in the top 5% of the distribution. Perhaps most striking is the pattern for the 25th percentile. Through most of the distribution the 25th perceintile is well below the mean

[10] We show that alternate methods of scaling-up ownership rates provide very similar estimates for all of our main results. That is, it is scaling up ownership rates that is important for deriving our IGM estimates much more so than the method used to scale-up the rates.

[11] Our rank-rank income relationship is very similar to that estimated in Chetty et al. (2018). The data used are from the same source, though for a different time period, so this is unsurprising.

rank; at the top of the distribution the 25th percentile quickly approaches the mean. Besides the intergenerational persistence, this highlights another channel of intergenerational relationship of resources - children of the wealthiest parents are much less likely to move down the distribution, which we call this the "insurance" role of parental resources, which is differential at the very top of the parent wealth distribution. This pattern also holds for income, but to a slightly lesser degree (Appendix Figure B.7).

Further, we can separately explore how the IG persistence is associated with child home ownership rates increase with parent housing and how much home values increase conditional on ownership. Appendix Figure B.4 displays these relationships. Panel A shows that the probability of ownership increases relatively with parent housing at the bottom of the parent distribution, but that the slope flattens out above the median. Conditional on parent ownership, child rank conditional on ownership increases linearly. These margins highlight an important feature for children of parent renters. The probability of child ownership is substantially lower among parent renters than for children with parents at the bottom of the home value distribution, it is just that conditional on ownership, the average rank is very high, matching those at the median of the distribution. Appendix Table XXX shows that parent renters have higher income than parents at the bottom of the ownership distribution and the standard deviation of income for parent renters is very high. Parent renters appear to be a much more heterogeneous group than conditioning on other parts of the distribution.

We combine these relationships and decompose the rank-rank slope into the portion attributable to increasing child home ownership rates and increasing child home values increasing (Figure 2). In the bottom portion of the parent housing distribution, most of child average gains with increased parental rank come from higher probabilities of home ownership. The share attributable home values increase throughout the parental distribution and in the top third of the parental distribution, most of the persistence comes from increases in average child home value.

4.3. Intergenerational Racial Gaps in Housing Wealth

Persistent cross sectional differences in wealth between Black and White households have been well documented (cite cite cite), as have institutional differences in access to home ownership (cite cite cite). Recent work has highlighted differences in intergenerational income mobility in the U.S. (cite cite). Our unique combination of data allows us to provide new evidence on the intergerational role of asset accumulation across Black and White Americans.

Figure 3 shows the rank-rank relationship in housing assets separately for Black and White families. The differences are stark. The β parameters are 0.425 and 0.260 for White and Black families, respectively. The intergenerational persistence of housing is much stronger for White than for Black households. The α parameters are also substantially different - White households experience substantially higher "absolute" mobility in housing. The differences are even larger for parent renters. We see all of the non-linearity at the bottom of the distribution is driven by White households, whereas Black families see an approximately linear relationship throughout the distribution.

This pattern is somewhat in contrast to the IG income relationship documented in Chetty et al. (2019) or using these data (Appendix Figure B.7). While the fact of an α gap is similar, income-income relationships find similar intergenerational slopes. For housing, the intergenerational gap in child outcomes, conditional on parent assets, grows across the parental distribution. For parents in the 25th percentile in the housing distribution, White children are approximately 5 ranks higher than Black children on average, at the 75th percentile the gap grows to 15 ranks, and at the very top the gap is almost 20 ranks. Even for Black families at the top of the distribution, their children are relatively likely to be in the bottom half of the distribution - this reflects large differential downward mobility between Black and White families, particularly at the top of the distribution.

Next we investigate what drives the differing intergenerational patterns. Differences in home ownership between Black and White families is well documented. We investigate differential child homeownership rates conditional on parent housing assets. Panel A of

Figure 4 shows the probability of children owning a home for parents in different parts of the housing distribution separately for Black and White families, displaying large conditional disparities on home ownership rates. Children from Black households in the top 10% of the parental distribution are less likely to own a home than children of white parents in the 30^{th} percentile of the distribution.

Given this, we explore the roles of extensive margin ownership relative to the intensive margin of asset value (the value of the home conditional on ownership) in explaining the Black-White IG gap. We conduct simulations where we equalize the home ownership rates of Black and White children conditional on parent assets. I.e., we assign Black children with parents in rank Y the same home ownership rate at White children with parents in Rank Y. To conduct the exercise, we must assume the value of homes of the newly simulated Black homeowners. In the first simulation, we assume that the simulated homeowners have housing assets equal to average housing assets for the observed Black homeowners with parents in that rank (Panel B of Figure 4).

This exercise makes it clear that differences in the extensive margin of home ownership, conditional on parent resources, is an important driver of the aggregate intergenerational differences. Panel B shows that if we assign equal ownership rates at the observed gap in asset values conditional on parent rank, the patterns in intergenerational housing assets become remarkably similar between Black and White families. Throughout the majority of the parental distribution, child outcomes would be very similar with some divergence remaining at the very top. This reflects the fact that conditional on owning a home and parental rank, the gaps in child housing assets are relatively small, compared to the unconditional gaps, as shown in Appendix Figure B.5.

It is quite likely that any policy that moved toward equalization of ownership rates would not result in asset values of new homeowners equal to those of observed homeowners, but would instead provide opportunities lower in the current distribution.[12] Therefore, we use

[12]There is likely positive selection in existing home owners and efforts to make credit more equally available or neighborhoods more equally accessible/desirable would plausibly move down the distribution of housing

various methods to impute a counterfactual housing value for the simulated "new" Black homeowners. First, we use an approach where we capitalize observed rent payments of home owners that are matched with ACS data, conditional on child an parent characteristics. We extrapolate these capitalized rent values to the marginal home value for similar children. The results of this exercise are shown in Panel C. Appendix Figure B.6 presents alternative approaches including estimating the value of first homes conditional on observables using observed first-time purchases in our data (Panel A), and assigning marginal home owners the 25th and 10th percentile of the observed distribution conditional on parent rank, Panels B and C, respectively. Together, these exercises suggest that if home ownership rates were equalized, average ranks of Black and White children would be quite similar in the bottom half of the parental distribution, but that for Black and White children of parents at the top of housing distribution there would remain substantial conditional gaps in housing values.

4.4. Regional Variation in Intergenerational Mobility

In addition to developing new facts about the intergenerational transmission of housing wealth for the country as a whole, we can also examine how this transmission varies across space and demographic groups. To examine these patterns, we estimate county-level α, δ and β statistics in a two part process. First, we estimate direct estimates of these parameters (α_D, δ_D and β_D) from the microdata described above for all counties. However, these direct estimates can be extremely noisy, particularly for very small counties. Thus, we apply a standard area-level small area estimation method (Faye-Herriot regression), which smooths the direct estimates. This smoothing has a shrinkage property such that as the number of observations underlying the direct estimates goes to zero, the Faye-Herriot estimate converges to the estimate for some larger geography (in our case, commuting zones, or groups of similarly sized commuting zones chosen to meet Census Bureau disclosure avoidance guidelines.) We produce these small area estimates for all individuals pooled, as well as for White, Black and Hispanic individuals separately. For data quality reasons, we omit counties with

values.

fewer than 10 underlying microdata observations. These county level estimates are available in a data appendix.

The visualize the geographic distribution of these measures as a series of choropleth maps. Figure 5 maps these parameters for our pooled sample across race groups. The two measures of upward mobility – α and δ, representing the predicted housing wealth rank for individuals born to the worst-off homeowner parents or to renter parents respectively – show broadly similar patterns. Broad stretches of the southeast, Appalachia, many parts of Indian country (such as the region in northwest New Mexico and northeast Arizona containing the Navajo and Hopi lands) and many "Rust Belt" areas such as Detroit, MI and Cleveland, OH show markedly lower levels of absolute housing wealth mobility. Conversely, large, high cost coastal metro areas (San Francisco, New York, Boston) and a broad swath of California, the Mountain West and Pacific Northwest show substantially higher absolute mobility. Patterns for the relative mobility measure β are less stark, but a high-to-low mobility gradient from northwest to southeast still roughly prevails.

5. Intergenerational Mobility of Resources: Income v. Wealth

5.1. Mobility of Wealth v. Income - Statistical Analysis

In this section we discuss the relationship between estimating in the intergenerational transmission of resources using income and wealth. Most studies have focused on the IGM of income because it is easier to observe and is correlated with well-being and economic paramaters of interest like skill/returns to human capital. Yet, while income is related to total resources, it may only be a noisy measure of parent or child resources. Wealth is generally believed to be a better and more stable measure of resources and well-being (cite cite cite).

First, we discuss the implications of estimating income-income (II) mobility if the measure we care about is wealth-wealth mobility, or the true model of IGM is through total resources,

not income alone. Say the reduced form IGM model is:

$$W_k = \beta W_p + u_k,$$

or child wealth is a reduced form function of parent wealth governed by parameter β. Now assume that all we observe is income of parents and kids, but we would like to recover the true relationship between parent and kid resources β.

First, recognize that wealth has various components each with different rates of return that give rise to income flows. $W = \sum_k w_k$ so that observed income $Y = \sum_k r_k w_k$. In the very special case $r_k = r$ for all k then $Y = rW$ and a regression of Y_k on Y_p would identify β. In the more realistic case of varying rates of returns, the identification depends on assumptions about the composition of assets. Let $Y_i = W_i \sum_k a_k r_k = W_i \tilde{r}_i$ for all individuals i, where a_k is the share of assets held as in category k. Then

$$\frac{Y_k}{\tilde{r}_k} = \beta \frac{Y_p}{\tilde{r}_p} + u_k$$

In the special case that everyone holds the same mix of assets, $\tilde{r}$ is constant and cancels out and β is recovered. In the more general case, the rates of return may be individual specific and may be correlated with the parents and child wealth distributions, in which case an income-income regression would include an omitted variable ($\tilde{r}_p$) and potentially non-classical measurement error in child resources through $\tilde{r}_k$, which will bias the estimate.

The understand the potential bias further, let $Y_i = rW_i + \eta_i$ where η_i represents the idiosyncratic relationship between income and assets (relative to an average rate of return r on wealth) for individual i. Concretely η_i can be thought of as the degree to which income of individual i overstates the relationship between income and wealth relative to average. In a rank-rank regression, this is the degree to which the income rank overstates the wealth rank. This gives the income-income relationship of:

$$Y_k = \beta(Y_p - \eta_p) + \eta_k + ru_k, \tag{3}$$

Bias is introduced by using income as a (noisy) proxy for wealth for parents and children. The η_p is directly related to Y_p, which leads to an underestimate of the true β. Using child income as a proxy for child wealth can lead to further bias, depending on the relationship between the η_k and parent income. [13]

A more general framework can let kid resources be a flexible reduced-form function of parent income and assets:

$$R_k = \alpha_y Y_p + \alpha_w W_p + u_k$$

and again let $Y_k = rR_k + \eta_k$ where η_k, so

$$rR_k = Y_k = r\alpha_y Y_p + r\alpha_w W_p + \eta_k + ru_k \tag{4}$$

This again highlights interpretations of parameters and potential sources of bias. The first source of bias from an income-income regression is from wealth being an omitted variable. An income-income regression will overstate the role of income for the transmission of parent resources as wealth will be an omitted variable positively correlated with income, though it will understate the role of total parental resources as it will only be a noisy proxy for parent total resources.[14] Controlling for wealth in addition to income provides more information (about total resources and the composition of assets) and the degree to which the true relationship between parent and child resources is recovered depends on the use of child income as a proxy for child resources (i.e. on η_k). In general, η_k will be negatively correlated

[13] Even if η_p is distributed randomly across the parent wealth distribution, then it will be positively correlated with Y_p. If η_k is random across the child wealth distribution, it will be negatively correlated with parent wealth, and therefore to parent income, so will also lead to an underestimate. Perhaps the more realistic case is that the η terms are negatively correlated with the wealth distributions – as wealth increases, it is less well represented by income flows (for example, housing, dividends or unrealized capital gains). In this case, mismeasurement in parent and kid wealth lead to stronger underestimates of the true β.

[14] Of course, if income is a perfect proxy for total parent resources or child resources are only a function of income, independent of wealth or the composition of income, α_w will be zero and there will be no such OVB.

with parent resources, so will lead to a downward bias in the estimation of the relationship between parent and child resources.[15] Concretely, in a rank-rank regression this is something that creates a mismatch in the ranking of kids in the income distribution relative to the wealth distribution (the degree to which child income overstates wealth). In the likely case that child income understates wealth more at the top of the parent wealth[16], this will be negatively correlated with parent resources, leading to a downward bias. Then, incorporating information on child wealth helps recover the true relationship between parent and child resources. Also note if child resources are only a function of parent wealth and not income, then generally $\alpha_y = 0$. But, if i) child outcomes are (plausibly) functions of both wealth stocks and income flows, or ii) wealth is mismeasured and income provides further information on parent total resources, or iii) the composition of wealth matters in addition to total wealth and income provides infromation on the composition of parent wealth, then $\alpha_y \neq 0$ as it provides further relevant information on parent resources.[17]

Together, this discussion shows that when total resources matter, then it is important to observe total resources for parents and kids, not just income and that the composition of income and wealth can matter for estimating the reduced form relationship. More generally, it is possible that child resources depend not only on aggregate resources but the composition of parent resources - income in addition to wealth and/or the composition of assets. Given this, we leverage our unique combination of data to flexibly estimate the association between parent and child resources.[18]

[15] See discussion above, or see $Y_k = rR_k + \eta_k$, so $R_k = (Y_k - \eta_k)/r = \alpha_y Y_p + \alpha_w W_p + u_k$.

[16] For example, wealthier parents buy cars for children, pass-down housing, art, or jewelry that is not reflected in income flows.

[17] It is difficult to distinguish between these possibilities, but this supports the idea of including both income and wealth as RHS variables if available (more here). Also, note that the same discussion is true for α_w if child resources are only a function of parent income, but wealth provides more information about permanent parent income levels or composition.

[18] Later we will look at the role of income and wealth and the composition of these in relationship to understanding differences in IGM for white and Black families.

5.2. Empirical evidence on the intergenerational mobility of total resources

We directly examine at the joint relationship between parent income and parent housing assets on child income. First, we estimate regressions following Equation (2), where Y_i^k and Y_i^p are parent and child positions in the income distribution, respectively, and we additionally include parent housing rank as a right-hand-side variable. Table V shows the regression coefficients and Figure 8 shows non-parametric rank-rank relationships. Starting with Parent income and housing wealth (column 1 and Panel A), we see that including housing wealth substantially attenuates the income-income slope, and shows a large independent role of housing assets on child income - the correlation between parent housing wealth and child income is almost as strong as that between parent and child income. The interaction term is very small implying an approximately linear relationship between each conditional on the other, as confirmed by Figure 8.

The non-parametric relationships in Figure 8 are informative. First, we see approximately linear relationships between parent and child income controlling for parent housing wealth, with relatively flat slopes, and parent housing rank shifts children up the distribution on average. Conditional on housing asset rank, increases in parent income rank more moderately increases average child income rank, and parent housing wealth shifts up the average child rank at any income level. Second, each bubble of the figure is population weighted for the parents. There most of the mass at the top and bottom of the parent income distribution is also at the top and the bottom of the housing asset distribution, so that aggregate, unconditional income-income relationship (Table V) is steeper than the conditional slope because of this correlation between income and wealth. But, while the correlation is strong at the top and the bottom of the distribution, there is substantial variation in parent housing and income in the middle of the distribution. For example, in the 40th-70th percentiles (Q5-Q8), there is a fairly even distribution of housing wealth rank among parents conditional on their income rank, and the housing assets have a strong independent association with child outcomes. There are many families in the middle of the distribution where income

substantively mis-measures household resources - conditional on parent income, kids from families with more housing wealth earn substantially more on average than those whose parents have similar income but less wealth.[19] Housing wealth is one source of resources missing or mismeasured when using parent income as a proxy for parent resources. It is this mismeasurement of income as a proxy for total resources and the strong correlation with child outcomes that i) leads income-income relationships to underestimate the IG persistence in total resources, and ii) displays how incorporating more information on income and wealth provides a more comprehensive picture of IG mobility of total resources.

Overall, slope of income-income IGM is in part a proxy for parental assets, and these results underscore the importance of accounting for assets, in addition to earned income in understanding IG transmission of resources. In the next section, we decompose this relationship even further.

5.3. Interaction between parent and child income and wealth

Building on the discussion above, and leveraging our unique combination of data, we expand the reduced form model of the relationship between parent and child resources.

Take generation g's lifetime budget constraint as:

$$c_g + b_g = s_g Y_g(1+r) + (1-s_g)Y_g + b_{g-1}$$

Where lifetime consumption c_g plus the amount of resources transferred to the next generation b_g is equal to lifetime earnings, Y_g, some share of which is saved at rate s_g and earns return r, plus any inherited wealth from the previous generation b_{g-1}. We start by assuming children imperfectly inherit productivity/human capital from parents as a function of parent resources. This could be from parent inputs into the human capital development of children,

[19] In the prior discussion, this is the variation that leads parent income to overstate the role of income flows but understate the effect of total resources, i.e. η_p in Equation (3) or wealth as an ommitted variable in Equation (4).

material or otherwise.[20] Children also have a random productivity draw independent of parent resources. The reduced from relationship can be modeled as follows:

$$Y_g = f(W_{g-1}) + u_g = \beta W_{g-1} + u_g$$

Children use their own earnings to accumulate wealth, but may also have additional wealth (above the average conditional on their earned income) as a function of their parent resources. This could be from a direct inheritance or parental resources could affect the savings rate of kids, i.e. $s_g = f(W_{g-1}) + u_g$, either through transmission of preferences for or knowledge about savings opportunities or through direct assistance like putting up collateral for savings opportunities or (ex.) down payments for mortgages that allow for savings opportunities beyond those available to counterparts earning similar incomes.[21] The reduced form relationships could be given as follows:

$$b_{g-1} = f(W_{g-1}) = \alpha_b W_{g-1}$$

$$s_g = f(W_{g-1}) + \nu_g = \gamma_s W_{g-1} + \nu_g$$

where α_b is the average share of parent resources transferred to children, γ_s is the association between parental wealth and child savings and ν_g is the idiosyncratic savings rate independent of parent wealth, with mean $\bar{\nu}$.

Then current child wealth is a function of own earnings and parent wealth given by:

$$W_g^{'} = f(Y_g, W_{g-1}) = (1+r)s_g Y_g + b_{g-1}$$

[20] For example, this could include material inputs such as private schooling or tutors, computers, college costs, or it could include non-material inputs such as time spent teaching children, entrance to the education or employment networks (ex, legacy or nepotism), or transmission of preferences over labor market earnings or careers. In this way, productivity/human capital are being defined from a realized and rewarded in the market standpoint, not a potential or ability standpoint.

[21] Again, this could be a combination of material transfers – direct transfers of assets from parents to children, parents subsidizing mortgages, buying kids their first car, trust funds, etc. – and inherited preferences for savings or capital accumulation.

$$= (1+r)(\gamma_s W_{g-1} + \nu_g)Y_g + \alpha_b W_{g-1}$$

$$= (1+r)\bar{\eta} Y_g + \alpha_b W_{g-1} + (1+r)\gamma_s (Y_g \times W_{g-1})$$

This leads to a simultaneous equation model that can be represented by the regressions:

$$Y_g = \beta W_{g-1} + u_g$$

$$W_g^{'} = \alpha_y Y_g + \alpha_b W_{g-1} + \alpha_s (Y_g \times W_{g-1}) + \epsilon_g \tag{5}$$

where α_y represents the average contribution of children's income to their assets independent of parent wealth. If children have more assets than predicted by their earnings alone because of parental wealth then α_b and/or α_s will be greater than zero - α_b represents transfers from parents that add to child wealth independent of their own earnings, and α_s represents the degree to which parent assets influences the asset accumulation behaviour or opportunities for children given their income.

If parent resources only affect child resources through child earning ability, then $\beta > 0$ and $\alpha_y > 0$ and $\alpha_b = \alpha_s = 0$. Parent wealth will be associated with child wealth, but only through it's relationship with child earnings, and conditional on observing child earnings parent resources have no further association with child wealth. But, if children have direct inheritances of wealth then $\alpha_b > 0$ and if parental wealth affects children's preferences for or ability to save conditional on their earnings then $\alpha_s > 0$ and children with wealthier parents will accumulate wealth even beyond what is predicted by their observed earnings. Because we have information on parent and child earnings and wealth, we can separately identify these reduced form parameters to learn about the IG dynamics.

5.4. Direct and Indirect Effects and IGM

Using the simultaneous equation model, we find empirically that on average parent resources have important indirect and direct effects, and that the "propensity to save" channel is less

important for explaining intergenerational persistence.

Figure 9, Panel A shows the rank-rank relationship between parent and child housing, decomposed into the indirect and direct effects, with the associated coefficients reported in column 1 of Table X. The direct affect is larger than the indirect effect, with 0.254 and 0.173 from the indirect effects. That is, about 40% of the persistence is associated with the fact that children of parents with more assets have higher earnings which allows them to purchase more housing ("indirect effect"), and about 60% comes from the fact that even conditional on children having the same earnings, children of parents with more housing assets have more assets themselves on average ("direct effect").

In Panel B, we include the "propensity to save" channel in the decomposition of the rank-rank slope. We see that it only slightly attenuates the direct effect to 0.233, and the indirect effect attenuates to 0.165. This implies that about 6.8% of the persistence is associated with a higher "propensity to save" as parent housing assets increase, with 38.6% from the pure indirect effect and 54.5% from the pure direct effect.[22] These estimates clearly underscores the independent role of parent assets in determining child assets.

Next we assess how direct and indirect effects are associated with the increases in child ownership rates and child home values across the parent distribution. Figure VVV shows that the increase in ownership rates is almost fully associated with the indirect effect while the increase in home values is almost fully associated with the direct effect. That is, child earnings explain the persistence in ownership rates, but conditional on children earning the same amount, children of wealthier parents have substantially more valuable homes.

Next we include parent income rank in addition to parent housing rank in the simultaneous equations, allowing each to independently explain child earnings and child housing conditional on child earnings. We find the parent housing and parent income have similar indirect effects — that is, they affect child earnings similarly. But, we find that almost the entire direct effect is associated with parent housing; parent income has almost no direct

[22]The propensity to save channel is conceptually a mix of the indirect effect (a function of increased child earnings) and a direct effect (higher assets conditional on earnings).

effect. Conditional on children having the same earnings, parents earning more income does not have a substantive additional impact on child assets, but parent assets do (Figure 10).

5.5. Direct and Indirect Effects and the Black-White IG Gap

We also use this framework to investigate the Black-White IGM gaps in housing. This framework shows that the gap in absolute mobility is driven by indirect effects (Figure 11). That is, for children of parents at the bottom of the distribution, Black-White gaps in housing are fully associated with the Black-White differences in labor market earnings. But, we also find that almost the entire difference in persistence is driven by direct effects. That is, if child labor market outcomes were equalized conditional on parent assets, there would still be a growing gap between the assets of Black and White children across the parental distribution. At the top of the parental distribution, only about 35% of the Black-White asset gap is explained by differences in child labor market earnings.

6. The Effects of Housing Supply on Wealth Mobility

As shown in Figure B.9, house prices grew slower than prices of other goods and median personal income from 1987-1998. Thereafter, house prices considerably outpaced the other two series–especially after the trough of the Great Recession, which is exactly when the cohorts we study reached peak homebuying ages. At the same time, popular concern has risen over housing shortages in metro areas, the restrictiveness of zoning and permitting regulations, and construction costs. We add to this debate with our novel dataset by examining the role of housing supply in shaping house prices and intergenerational mobility of housing wealth.

We also want to understand the role of asset prices on wealth accumulation and mobility. This is challenging for the standard reason of price being endogenous. If a down-payment assistance program for children from poor families increases housing demand, we will observe price and mobility increasing, but it would be wrong to conclude that price increases mobility. Alternatively, if wealthy families are advised by their financial planners to sell off stocks and

buy real estate, home prices will rise and children's housing wealth holdings will become more strongly correlated to their parents', but not because price itself has lowered mobility. We aim to identify local price changes that leverage local variation in supply constraints as opposed to demand shocks.

We have shown considerable geographic variation in housing wealth mobility. This is not surprising given similar patterns found in the income mobility literature. That literature has pointed to the importance of local exposures such as segregation, crime, school quality, and environmental factors in shaping human capital formation and eventual earnings in adulthood. While those channels surely influence wealth holdings and wealth mobility as well, we have shown that labor income is not a sufficient statistic for resource mobility, and that capital markets play a further role. Accordingly, we extend the local exposure literature by examining how local housing market fundamentals, as reflected in housing supply constraints, influence housing wealth mobility and its geographic variation.

6.1. Housing supply data

We obtain information on housing supply elasticities fromBaum-Snow and Han (2024). Building on work by Saiz (2010), these authors combine micro-geographic information on city topography (e.g. land availability and zoning restrictions) with housing transactions data to model the responsiveness of housing units, housing starts, and floor space to demand-driven price shocks. They use the model to estimate housing supply elasticities in 2001 and 2011 for every residential tract in U.S. metro areas. They replicate previous work showing considerable variation in housing supply constraints across city aggregates (e.g. Saiz (2010), Gorback and Keys (2020)), but importantly, also show at least as much variation at the tract level. We merge their public-use data package onto our micro dataset according to parental location (measured when the children were teenagers).

6.2. Reduced Form Estimates

Tables 1 through 2 summarize results from models which regress our housing mobility measures (α, β, δ on interactions between birth year cohort trends with baseline measures of housing supply elasticity in a child's county of birth. These models can be interpreted as the effect of housing supply elasticity on the cross-cohort trend in housing wealth mobility. Intuitively, these models seek to uncover how housing wealth mobility is shaped by different children having greater exposure to more elastic housing markets.

Relative to the cohort trend in housing wealth mobility experienced by children growing up in more elastic markets, we find a trend toward more immobility—i.e. lower absolute mobility and a higher dependence of child wealth on parent wealth—among children growing up in more inelastic markets. We are in the process of scaling these effects into responses of mobility to asset price, as well as understanding how these responses are generated by movements in direct versus indirect effects. These reduced-form estimates suggest that housing supply expansions can help stabilize housing price growth and promote intergenerational wealth mobility in the population at large, without much affecting Black-White mobility gaps

7. Implications

8. Conclusion

References

Baum-Snow, N. and L. Han (2024): "The Microgeography of Housing Supply," *Journal of Political Economy*, 132, 1897–1946.

Black, S. E. and P. J. Devereux (2011): "Recent Developments in Intergenerational Mobility," Elsevier, vol. 4 of *Handbook of Labor Economics*, 1487–1541.

Bond, B., J. D. Brown, A. Luque, and A. O'Hara (2014): "The Nature of the Bias When Studying Only Linkable Person Records: Evidence from the American Community Survey," *Center for Administrative Records Research and Applications Working Paper*.

Chetty, R. and N. Hendren (2018): "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects," *Quarterly Journal of Economics*, 113.

Chetty, R., N. Hendren, M. R. Jones, and S. R. Porter (2019): "Race and Economic Opportunity in the United States: an Intergenerational Perspective*," *The Quarterly Journal of Economics*, 135, 711–783.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014): "Where is the land of Opportunity? The Geography of Intergenerational Mobility in the United States *," *The Quarterly Journal of Economics*, 129, 1553–1623.

Collins, W. J. and M. H. Wanamaker (2022): "African American Intergenerational Economic Mobility since 1880," *American Economic Journal: Applied Economics*, 14, 84–117.

Gorback, C. S. and B. J. Keys (2020): "Global Capital and Local Assets: House Prices, Quantities, and Elasticities," Working Paper 27370, National Bureau of Economic Research.

Layne, M., D. Wagner, and C. Rothhaas (2014): "Estimating record linkage false

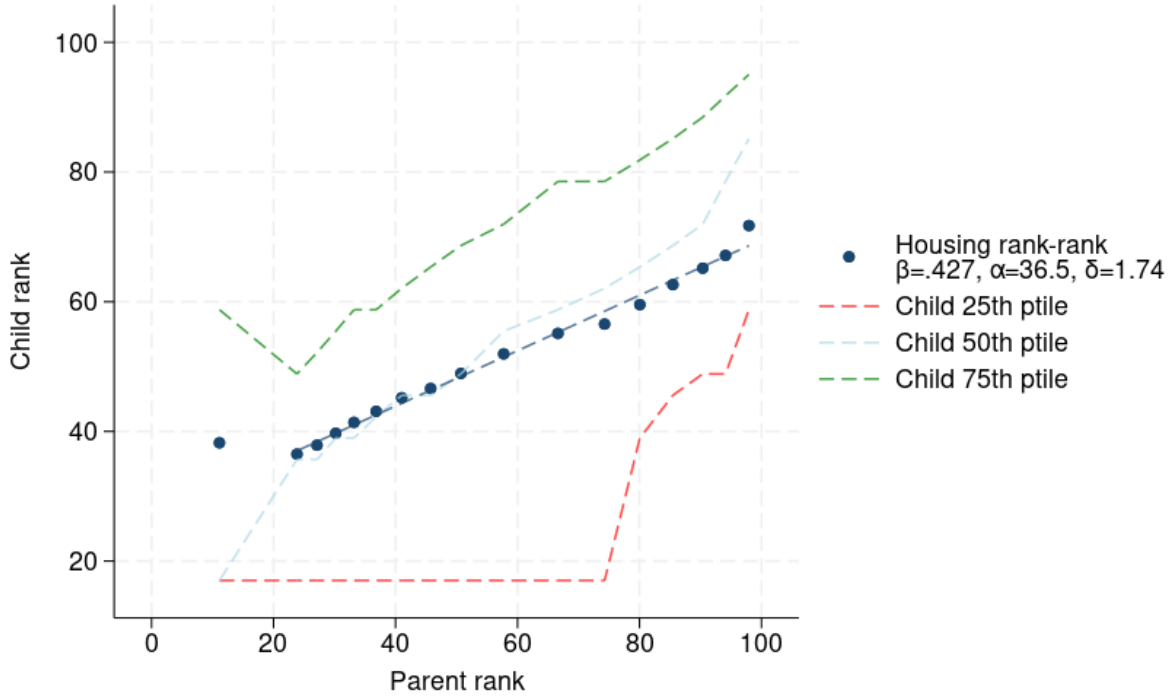
match rate for the Person Identification Validation System," *Center for Administrative Records Research and Applications Working Paper*.

Mogstad, M. and G. Torsvik (2023): "Family background, neighborhoods, and intergenerational mobility," in *Handbook of the Economics of the Family, Volume 1*, ed. by S. Lundberg and A. Voena, North-Holland, vol. 1 of *Handbook of the Economics of the Family*, 327–387.

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

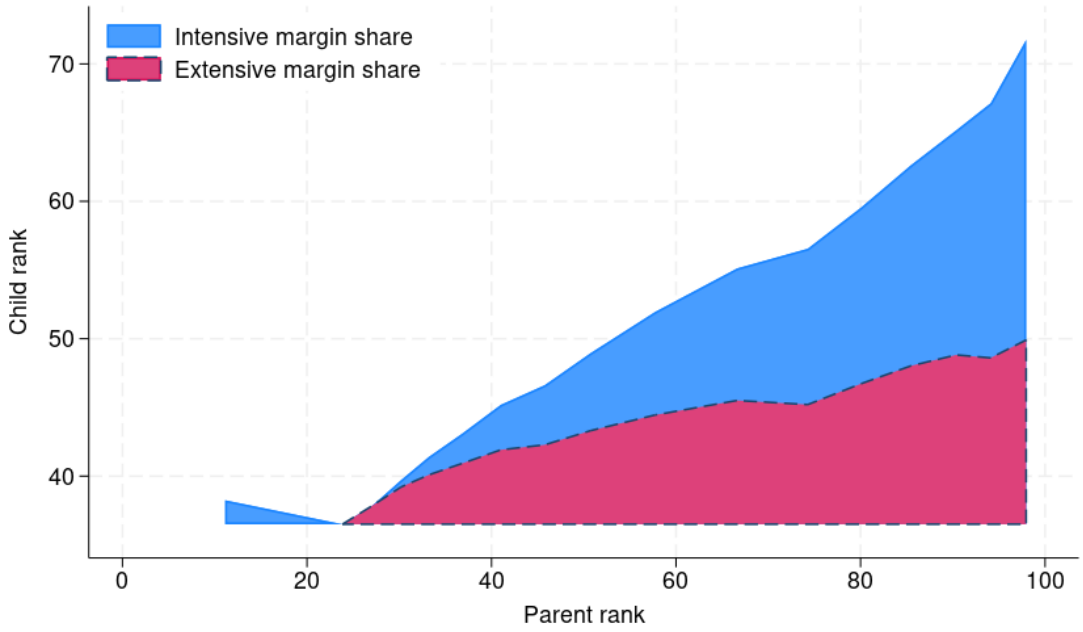
Tables and Figures

Figure 1: Intergenerational Mobility of Housing Assets - Rank-Rank relationship



Note: This figures shows the rank-rank relationship of housing assets. Blue dots represent the average child ranks conditional on parent rank; the blue dashed line is the linear rank-rank slope (β). Percentiles of the child housing rank distribution conditional on parent rank are also shown. α is the average child rank for children of parents in the bottom of the home value distribution and δ is the difference in the average child rank between children of parent renters and parents in the bottom of the housing value distribution.

Figure 2: Decomposition - home ownership rate v. home value



Note: This figure decomposes the rank-rank slope in housing assets into the portions attributable to increasing child home ownership rates across the parent distribution ("extensive margin") and the portion attributable to increasing home values conditional on ownership across the parent distribution ("intesnsive margin").

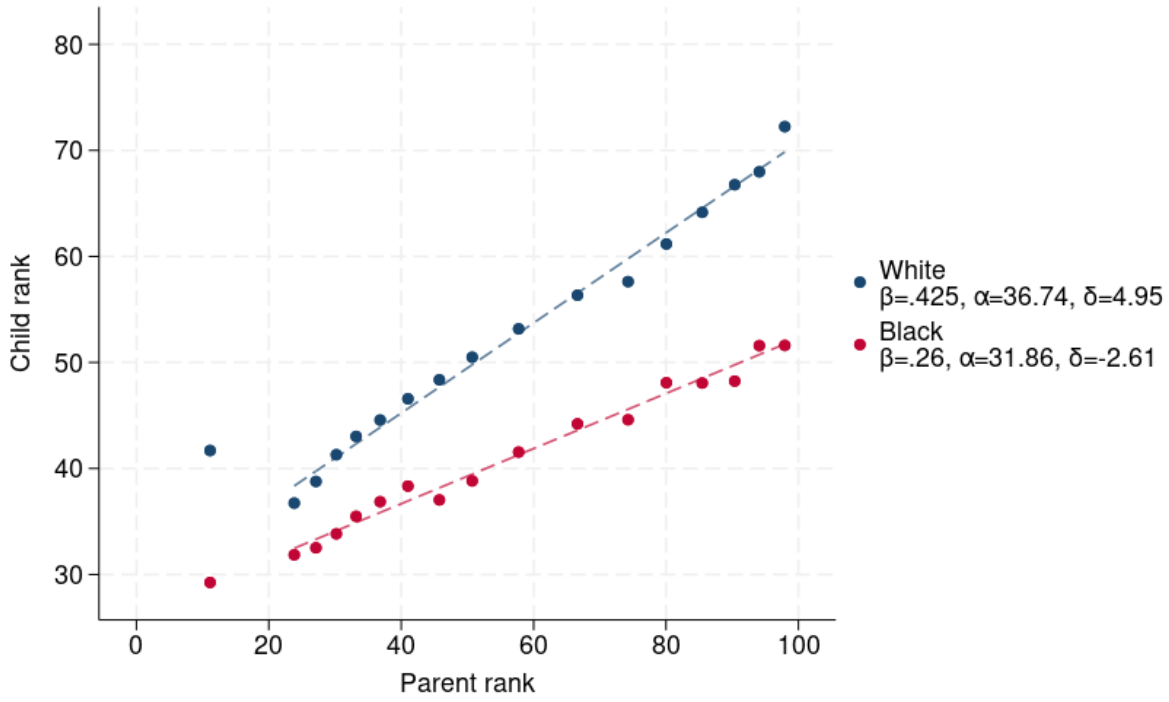
Table 1: Regressions of cross-cohort mobility trend on housing supply: full sample

	Housing wealth rank			*Pr(ownership)*		
	(1)	(2)	(3)	(1)	(2)	(3)
β 1978-86 trend*1 sd change in baseline hse	-0.0318	-0.0384	-0.0298	-0.0101	-0.0298	-0.0258
	(0.0046)	(0.0054)	(0.0058)	(0.0065)	(0.0078)	(0.0079)
α 1978-86 trend*1 sd change in baseline hse	0.0245	0.0245	0.0185	0.0156	0.0204	0.0171
	(0.0029)	(0.0034)	(0.0035)	(0.0041)	(0.0053)	(0.0050)
δ 1978-86 trend*1 sd change in baseline hse	-0.0192	-0.0106	-0.0060	-0.0175	-0.0060	0.0000
	(0.0029)	(0.0030)	(0.0029)	(0.0045)	(0.0054)	(0.0052)
N	2.33 mil	2.33 mil	2.33 mil	2.33 mil	2.33 mil	2.33 mil
R^2	0.160	0.161	0.161	0.101	0.102	0.102
county FE	X	X	X	X	X	X
(α,β,δ)*hse_diff*birth cohort FE	X	X	X	X	X	X
(α,β,δ)*division FE*birth cohort FE	X	X	X	X	X	X
(α,β,δ)*initial housing chars*birth cohort FE		X	X		X	X
(α,β,δ)*big county FE*birth cohort FE			X			X

Source: 2000 Census Long Form, Census Databank, 2019-2021 Black Knight property records. Housing supply elasticities come from Baum-Snow and Han (2024). See Section 2.2 for full sample description. Sample restricted to those who were observed as teenagers in the metro areas considered by Baum-Snow and Han (2024). This subsample comprises roughly 77% of the full sample (68% unweighted due to the Long Form oversampling non-metro areas).

Difference-in-difference regression estimates of equation (??) with robust standard errors clustered on county. Estimates show how the trend in mobility between the 1978 and 1986 cohorts changes if the housing supply elasticity (hse) in the teenage county increases by one standard deviation (i.e. 0.17).

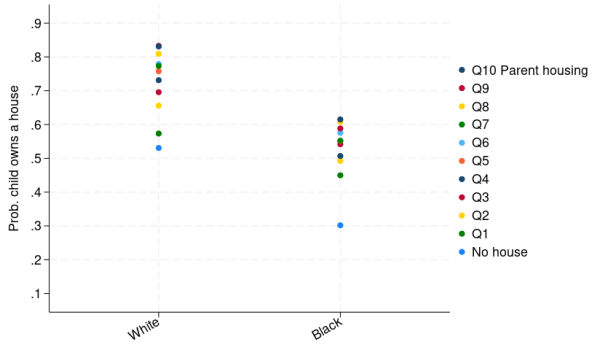
Figure 3: Intergenerational Mobility of Housing Assets by Race



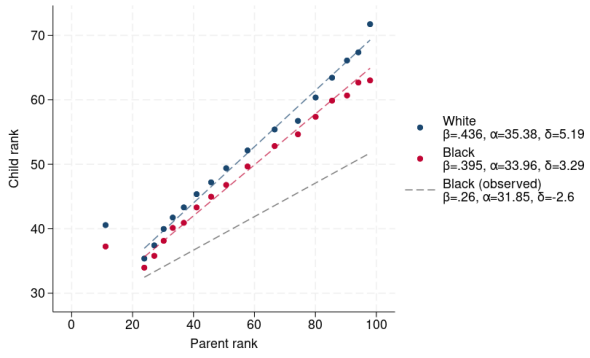
Note: This figures shows the rank-rank relationship of housing assets separately for Black and White Families. Dots represent the average child ranks conditional on parent rank and dashed lines show the linear rank-rank slopes (β); α is the average child rank for children of parents in the bottom of the home value distribution and δ is the difference in the average child rank between children of parent renters and parents in the bottom of the housing value distribution. Ranks are defined at the population level.

Figure 4: Home Ownership and Black-White IGM

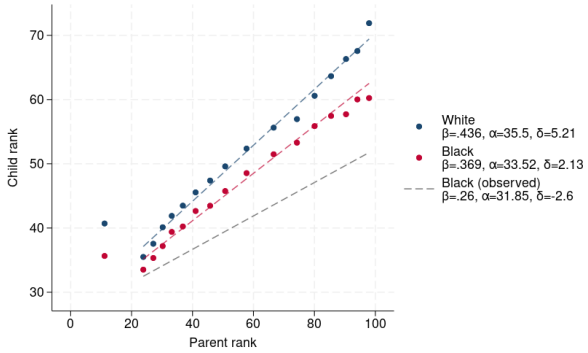
(a) Ownership probabilities by race



(b) Counterfactual ownership rates



(c) Counterfactual ownership - capitalized rent



Note: Panel A shows child home ownership rates across deciles of the parent housing asset distribution, separately for Black and White families. Panel B shows counterfactual rank-rank slopes for a simulation where the home ownership rates of Black children are equalized to those of white children within parent rank bins, and housing assets are equal to observed average housing assets for black home owners within each parental bin. Panel C shows another counterfactual exercise where Black and White child home ownership rates are equalized, but the new marginal Black home owners have are assigned estimated home values according to a capitalization approach of observed rent payments. For both, ranks are redefined for the population after the simulations and the corresponding rank-rank relationships for Black and White families are shown, as is the observed Black rank-rank relatinship (grey dashed line).

Figure 5: Geographic Distribution of Housing Wealth Mobility, α, δ and β Parameters

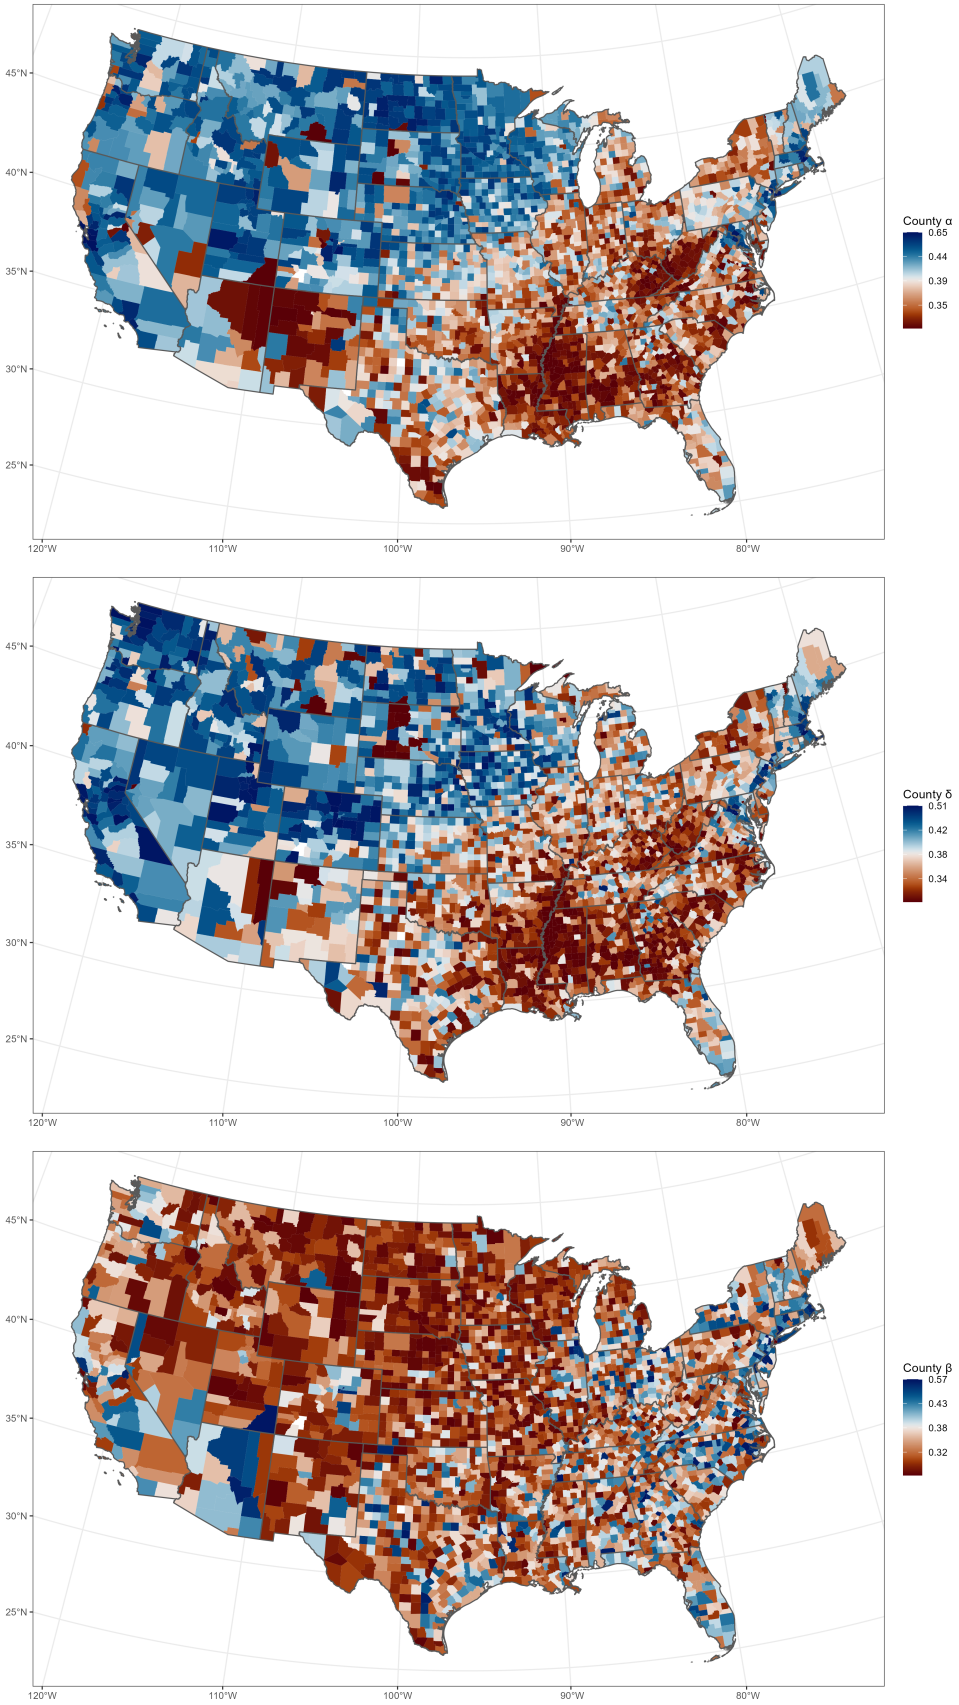


Figure 6: Geographic Distribution of Housing Wealth Mobility for White Individuals, α, δ and β Parameters

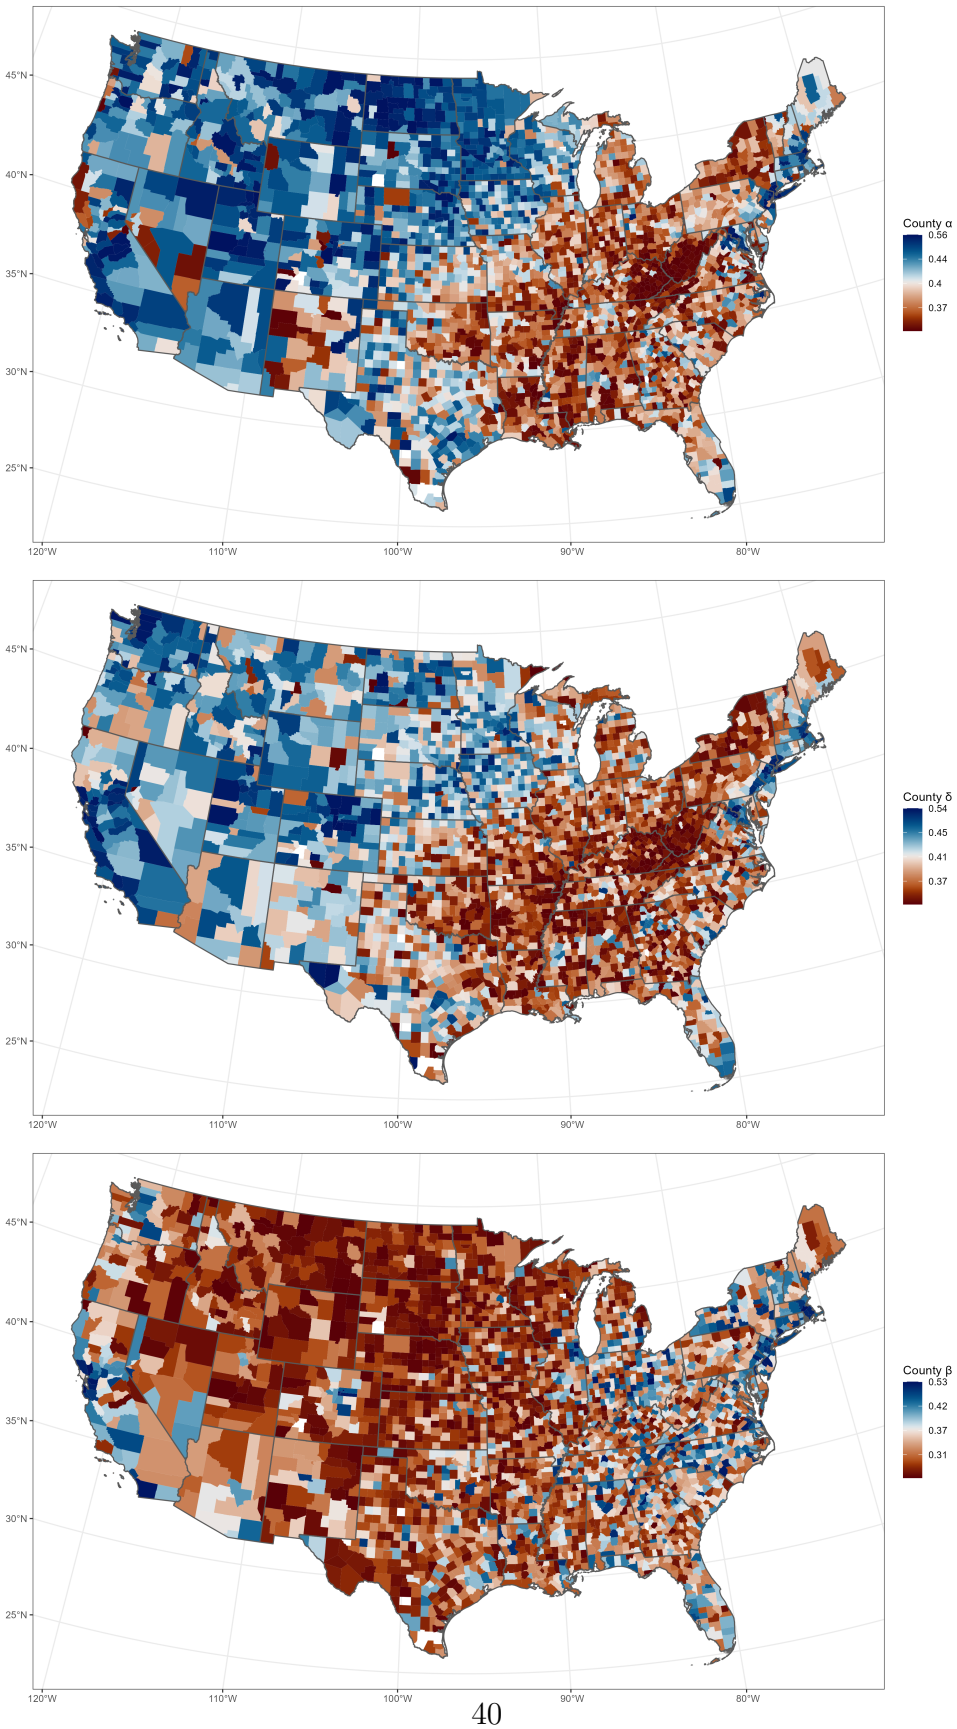


Figure 7: Geographic Distribution of Housing Wealth Mobility for Black Individuals, α, δ and β Parameters

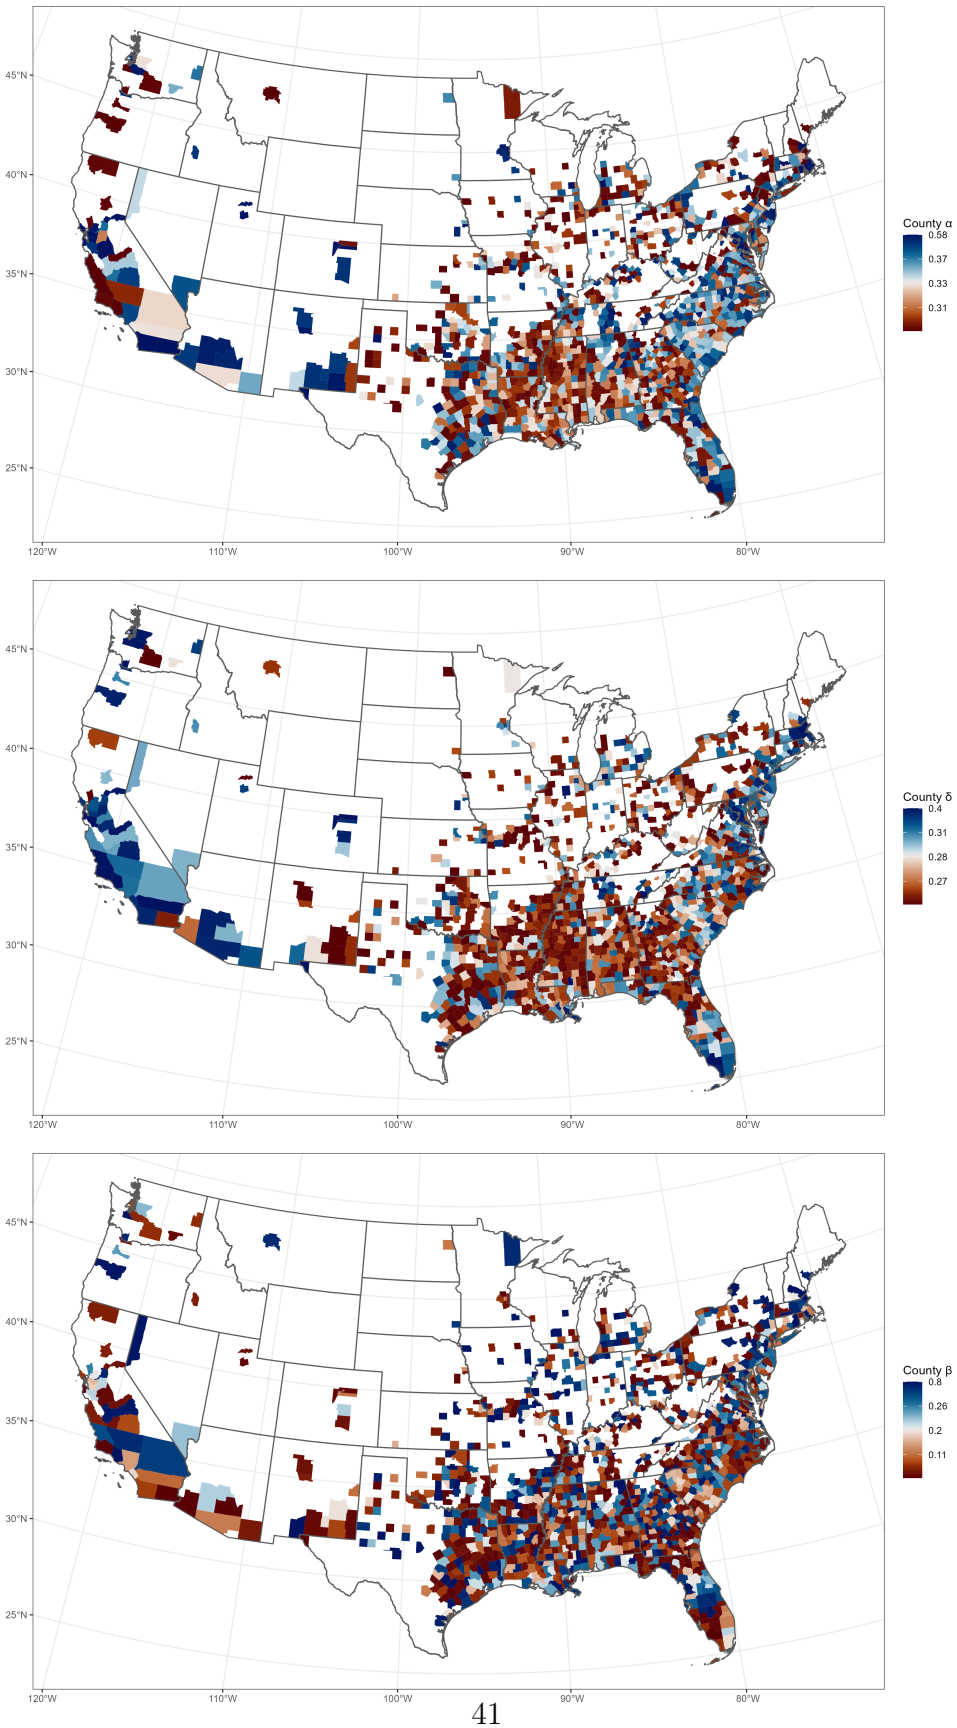
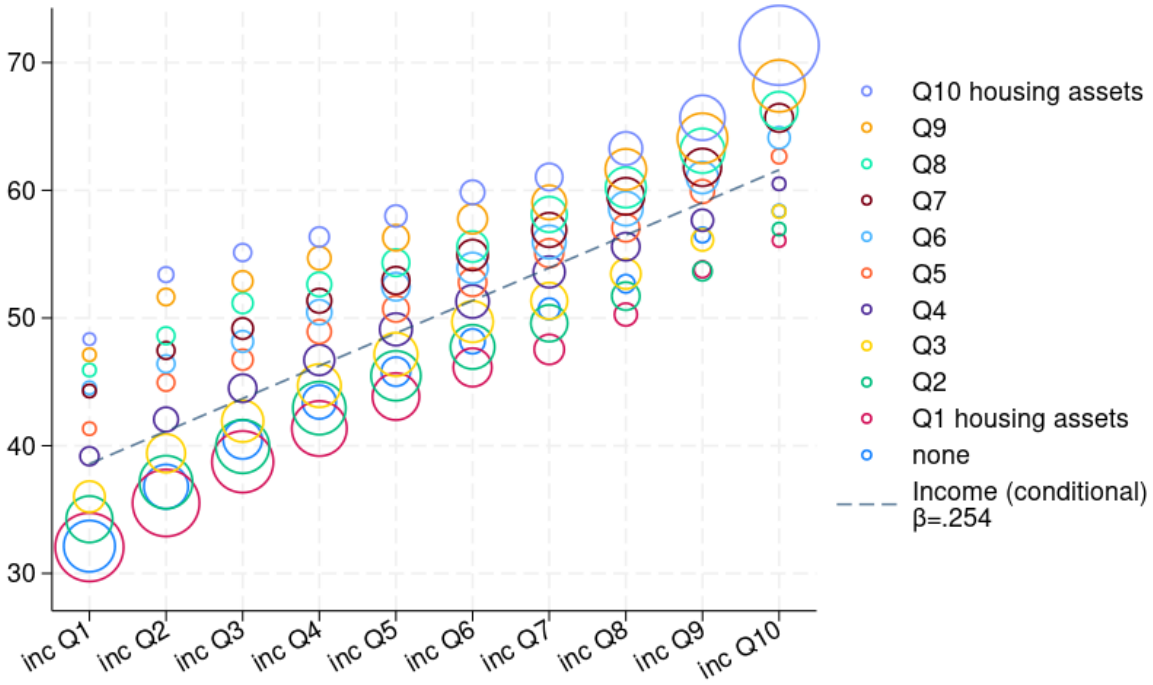


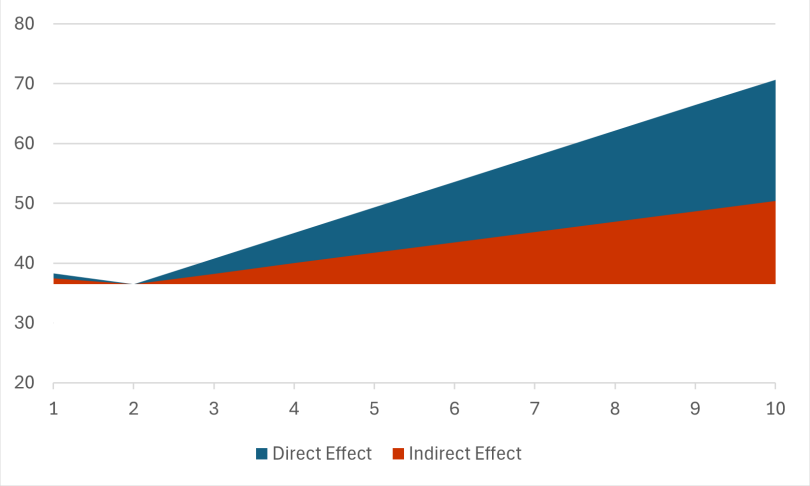
Figure 8: Child income and parent income and housing assets



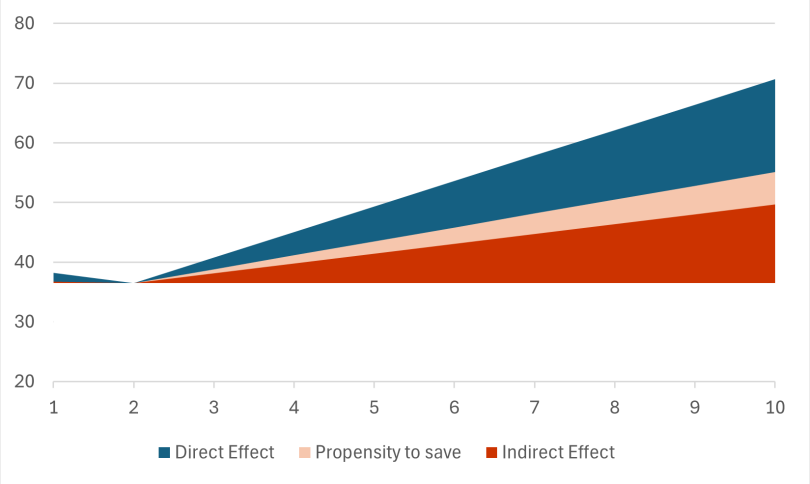
Note: This figure shows the non-parametric relationship between child income rank and parent income and housing asset ranks. Each bubble represents the average child income rank conditional on the decile of parent income and parent housing asset ranks. The size of each bubble represents the population of children in each cell.

Figure 9: Decomposition of Housing Asset Slope - Indirect and Direct Effects

(a) Direct v. Indirect Effects

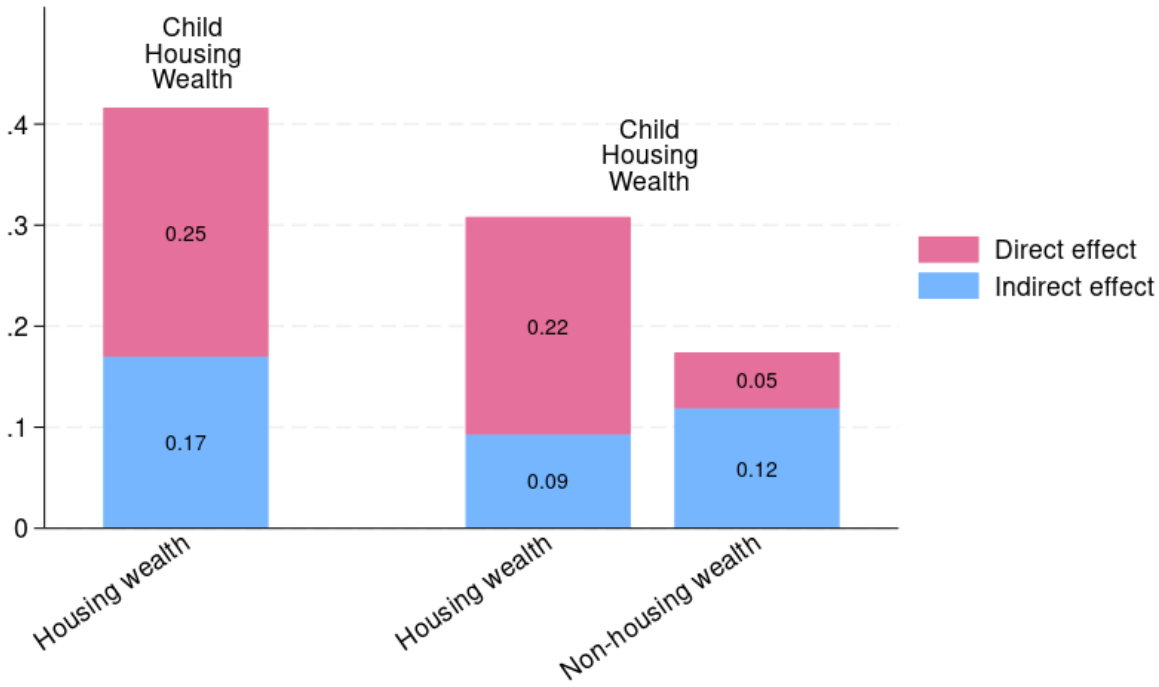


(b) Including the "propensity to save" channel



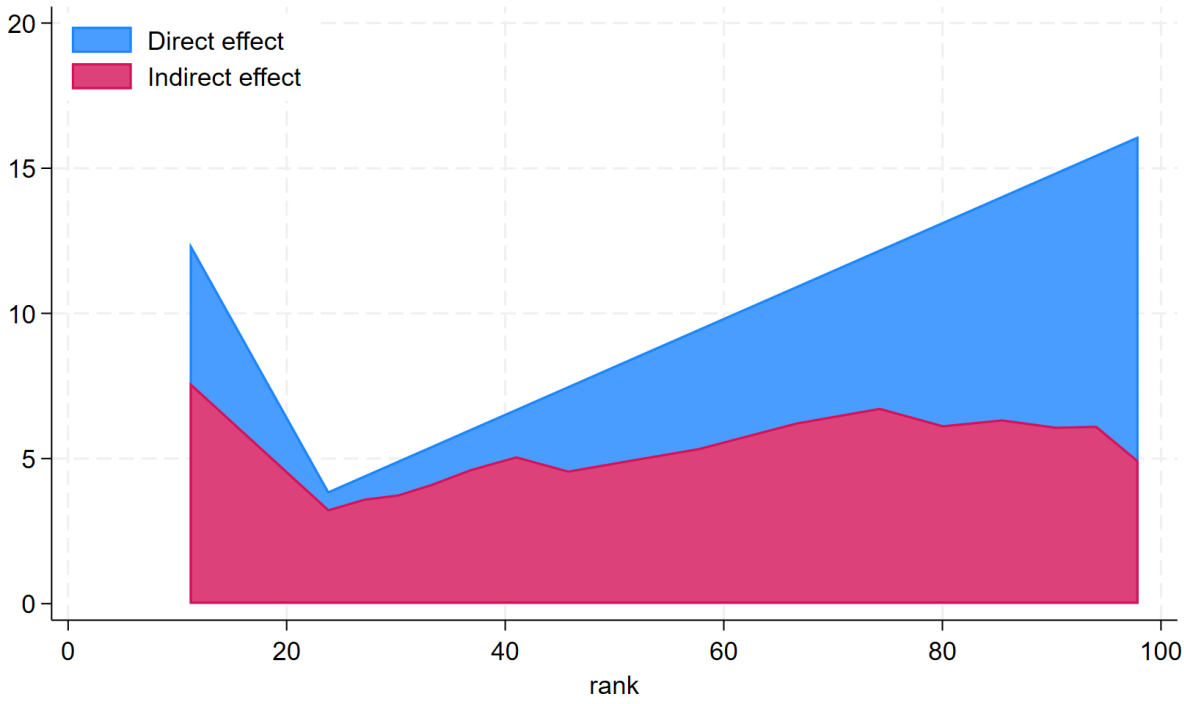
Note: Panel A shows a decomposition of the population rank-rank slope of housing assets into the Direct and Indirect Effects. The red area represents the share associated with indirect effects, i.e. the increase in child assets across the parent distribution due to the fact that children of wealthier parents earn more income. The blue is the direct effect, i.e. conditional on children earning the same amount, the increase in child assets due to having wealthier parents. Panel B shows the direct effect, indirect effect and the "propensity to save" channel — the interaction effect between child earnings and parent housing assets.

Figure 10: Direct and Indirect Effects by parent income and assets



Note: This figure shows a summary of the decomposition of the rank-rank slope in housing assets between a direct and indirect effect conditioning on parent housing alone (first bar). The second set of bars includes both parent housing and parent income ranks in the simultaneous equation model of parent resources on child earnings and parent resources on child housing conditional on child earnings. The indirect and direct effects associated with parent housing and parent income are shown.).

Figure 11: Direct and Indirect Effects - Decomposition of Black-White IG gaps



Note: This figure shows a decomposition between the gap in Black and White child average housing ranks across the parent distribution, decomposed into direct and indirect effects. The indirect effect shows the difference in Black and White ranks associated with Black and White children having different earnings conditional on parent rank. The direct effect shows the remaining gap in child ranks conditional on Black and White children having the same earnings within parent rank.

Table 2: Regressions of cross-cohort mobility trend on housing supply: White subsample

	Housing wealth rank			*Pr(ownership)*		
	(1)	(2)	(3)	(1)	(2)	(3)
β 1978-86 trend*1 sd change in baseline hse	-0.0288	-0.0277	-0.0218	0.0071	-0.0118	-0.0147
	(0.0033)	(0.0041)	(0.0044)	(0.0052)	(0.0057)	(0.0060)
α 1978-86 trend*1 sd change in baseline hse	0.0237	0.0190	0.0137	0.0034	0.0095	0.0101
	(0.0023)	(0.0029)	(0.0029)	(0.0037)	(0.0044)	(0.0044)
δ 1978-86 trend*1 sd change in baseline hse	-0.0260	-0.0080	-0.0023	-0.0155	-0.0001	0.0044
	(0.0030)	(0.0033)	(0.0031)	(0.0041)	(0.0049)	(0.0050)
N	1.64 mil	1.64 mil	1.64 mil	1.64 mil	1.64 mil	1.64 mil
R^2	0.143	0.143	0.143	0.069	0.070	0.070
county FE	X	X	X	X	X	X
(α,β,δ)*hse_diff*birth cohort FE	X	X	X	X	X	X
(α,β,δ)*division FE*birth cohort FE	X	X	X	X	X	X
(α,β,δ)*initial housing chars*birth cohort FE		X	X		X	X
(α,β,δ)*big county FE*birth cohort FE			X			X

Source: 2000 Census Long Form, Census Databank, 2019-2021 Black Knight property records. Housing supply elasticities come from Baum-Snow and Han (2024). See Section 2.2 for full sample description. Sample restricted to those who were observed as teenagers in the metro areas considered by Baum-Snow and Han (2024). This subsample comprises roughly 77% of the full sample (68% unweighted due to the Long Form oversampling non-metro areas).

Difference-in-difference regression estimates of equation (**??**) with robust standard errors clustered on county. Estimates show how the trend in mobility between the 1978 and 1986 cohorts changes if the housing supply elasticity (hse) in the teenage county increases by one standard deviation (i.e. 0.17).

Table 3: Regressions of cross-cohort mobility trend on housing supply: Black subsample

	Housing wealth rank			*Pr(ownership)*		
	(1)	(2)	(3)	(1)	(2)	(3)
β 1978-86 trend*1 sd change in baseline hse	-0.0136	-0.0218	-0.0228	0.0041	-0.0222	-0.0227
	(0.0099)	(0.0122)	(0.0135)	(0.0207)	(0.0245)	(0.0267)
α 1978-86 trend*1 sd change in baseline hse	0.0092	0.0136	0.0139	0.0031	0.0174	0.0175
	(0.0050)	(0.0063)	(0.0073)	(0.0117)	(0.0137)	(0.0155)
δ 1978-86 trend*1 sd change in baseline hse	-0.0061	-0.0095	-0.0082	-0.0087	-0.0140	-0.0053
	(0.0045)	(0.0056)	(0.0061)	(0.0099)	(0.0114)	(0.0125)
N	282 K	282 K	282 K	282 K	282 K	282 K
R^2	0.098	0.098	0.098	0.075	0.075	0.075
county FE	X	X	X	X	X	X
(α,β,δ)*hse_diff*birth cohort FE	X	X	X	X	X	X
(α,β,δ)*division FE*birth cohort FE	X	X	X	X	X	X
(α,β,δ)*initial housing chars*birth cohort FE		X	X		X	X
(α,β,δ)*big county FE*birth cohort FE			X			X

Source: 2000 Census Long Form, Census Databank, 2019-2021 Black Knight property records. Housing supply elasticities come from Baum-Snow and Han (2024). See Section 2.2 for full sample description. Sample restricted to those who were observed as teenagers in the metro areas considered by Baum-Snow and Han (2024). This subsample comprises roughly 77% of the full sample (68% unweighted due to the Long Form oversampling non-metro areas).

Difference-in-difference regression estimates of equation (**??**) with robust standard errors clustered on county. Estimates show how the trend in mobility between the 1978 and 1986 cohorts changes if the housing supply elasticity (hse) in the teenage county increases by one standard deviation (i.e. 0.17).

Table 4: Regressions of cross-cohort mobility trend on housing supply: heterogeneity by initial price level

	Housing wealth rank			*Pr(ownership)*		
	(1)	(2)	(3)	(1)	(2)	(3)
β 1978-86 trend*1 sd change in hse*initially cheap	-0.0284	-0.0388	-0.0299	-0.0037	-0.0287	-0.0245
	(0.0048)	(0.0057)	(0.0071)	(0.0067)	(0.0083)	(0.0084)
β 1978-86 trend*1 sd change in hse*initially pricey	-0.0291	-0.0332	-0.0242	-0.0178	-0.0208	-0.0167
	(0.0054)	(0.0067)	(0.0071)	(0.0086)	(0.0095)	(0.0095)
α 1978-86 trend*1 sd change in hse*initially cheap	0.0228	0.0245	0.0186	0.0128	0.0201	0.0167
	(0.0029)	(0.0034)	(0.0035)	(0.0045)	(0.0052)	(0.0050)
α 1978-86 trend*1 sd change in hse*initially pricey	0.0185	0.0209	0.0141	0.0135	0.0133	0.0097
	(0.0037)	(0.0045)	(0.0046)	(0.0063)	(0.0069)	(0.0067)
δ 1978-86 trend*1 sd change in hse*initially cheap	-0.0180	-0.0124	-0.0076	-0.0147	-0.0092	-0.0030
	(0.0027)	(0.0030)	(0.0029)	(0.0045)	(0.0050)	(0.0049)
δ 1978-86 trend*1 sd change in hse*initially pricey	-0.0001	-0.0033	0.0016	0.0110	0.0082	0.0147
	(0.0038)	(0.0039)	(0.0038)	(0.0065)	(0.0067)	(0.0064)
N	2.33 mil	2.33 mil	2.33 mil	2.33 mil	2.33 mil	2.33 mil
R^2	0.160	0.161	0.161	0.101	0.102	0.102
county FE	X	X	X	X	X	X
(α,β,δ)*hse_diff*birth cohort FE	X	X	X	X	X	X
(α,β,δ)*division FE*birth cohort FE	X	X	X	X	X	X
(α,β,δ)*initial housing chars*birth cohort FE		X	X		X	X
(α,β,δ)*big county FE*birth cohort FE			X			X

Source: 2000 Census Long Form, Census Databank, 2019-2021 Black Knight property records. Housing supply elasticities come from Baum-Snow and Han (2024). See Section 2.2 for full sample description. Sample restricted to those who were observed as teenagers in the metro areas considered by Baum-Snow and Han (2024). This subsample comprises roughly 77% of the full sample (68% unweighted due to the Long Form oversampling non-metro areas).

Difference-in-difference regression estimates of equation (**??**) with robust standard errors clustered on county. Estimates show how the trend in mobility between the 1978 and 1986 cohorts changes if the housing supply elasticity (hse) in the teenage county increases by one standard deviation (i.e. 0.17). Main estimates are allowed to differ according to whether the county-average parental home value is above ("pricey") / below ("cheap") the 85^{th} percentile of the national distribution of county averages. Slightly over half of the metro subsample lived in a pricey area as a teen.

A. Appendix: Robustness and Measurement

Figure A.1: Rank-Rank relationships, oldest v. youngest cohort

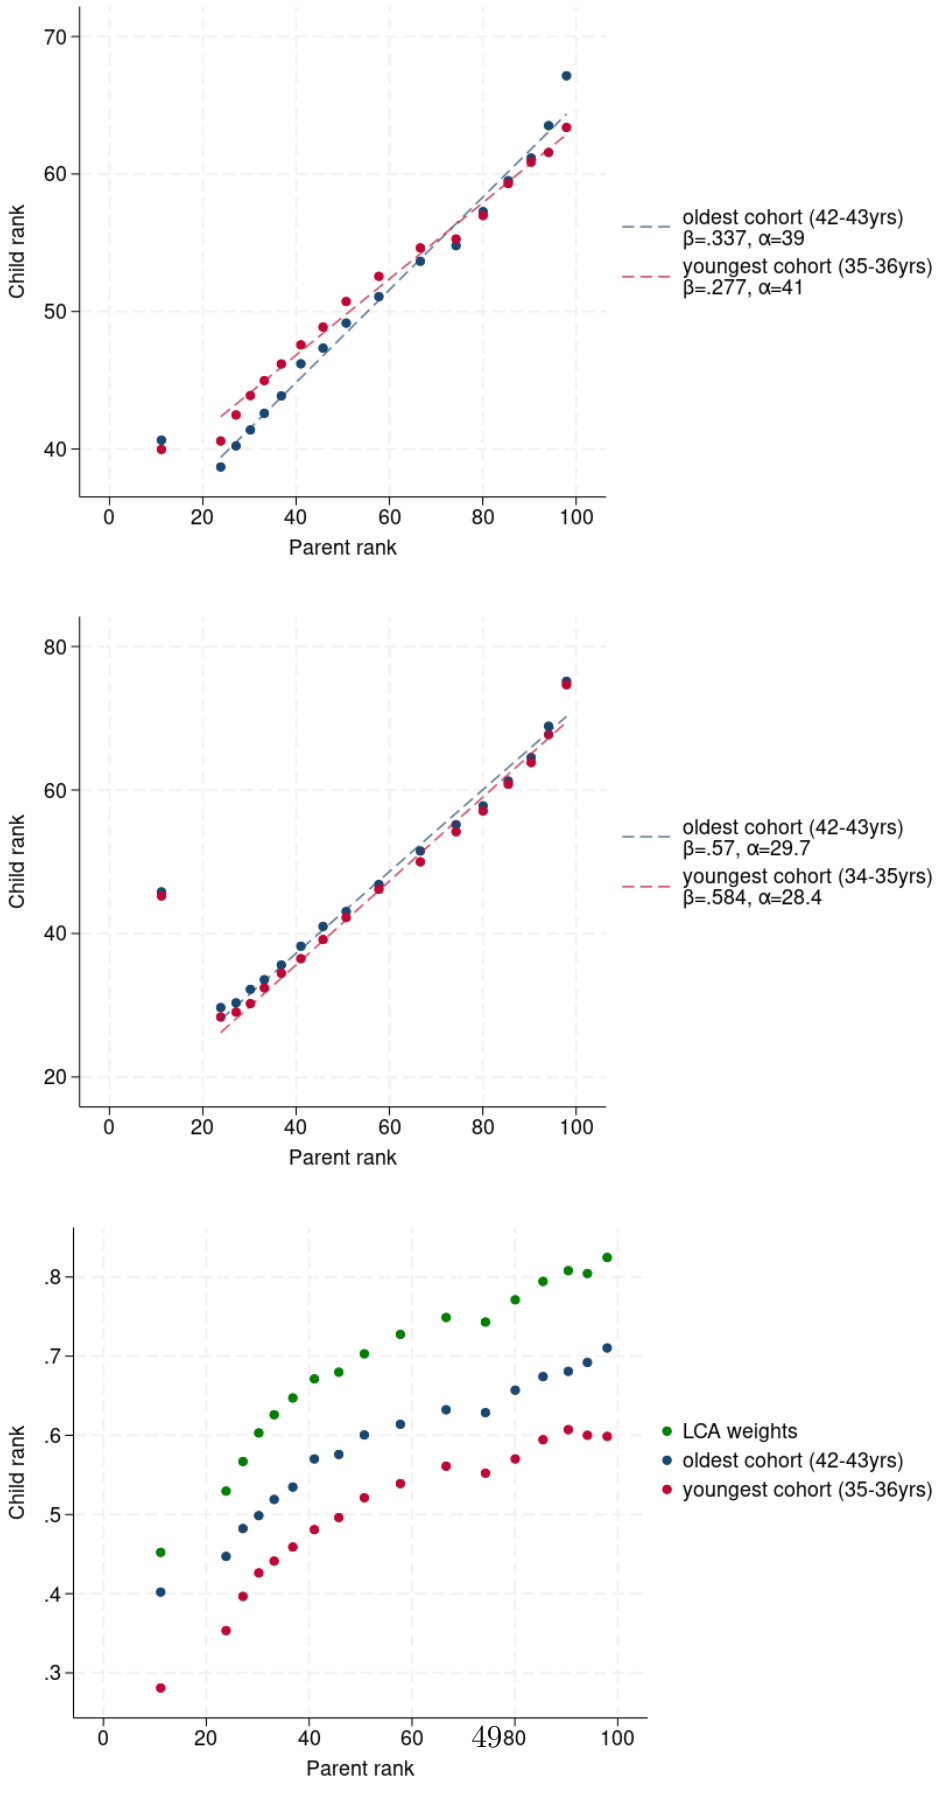
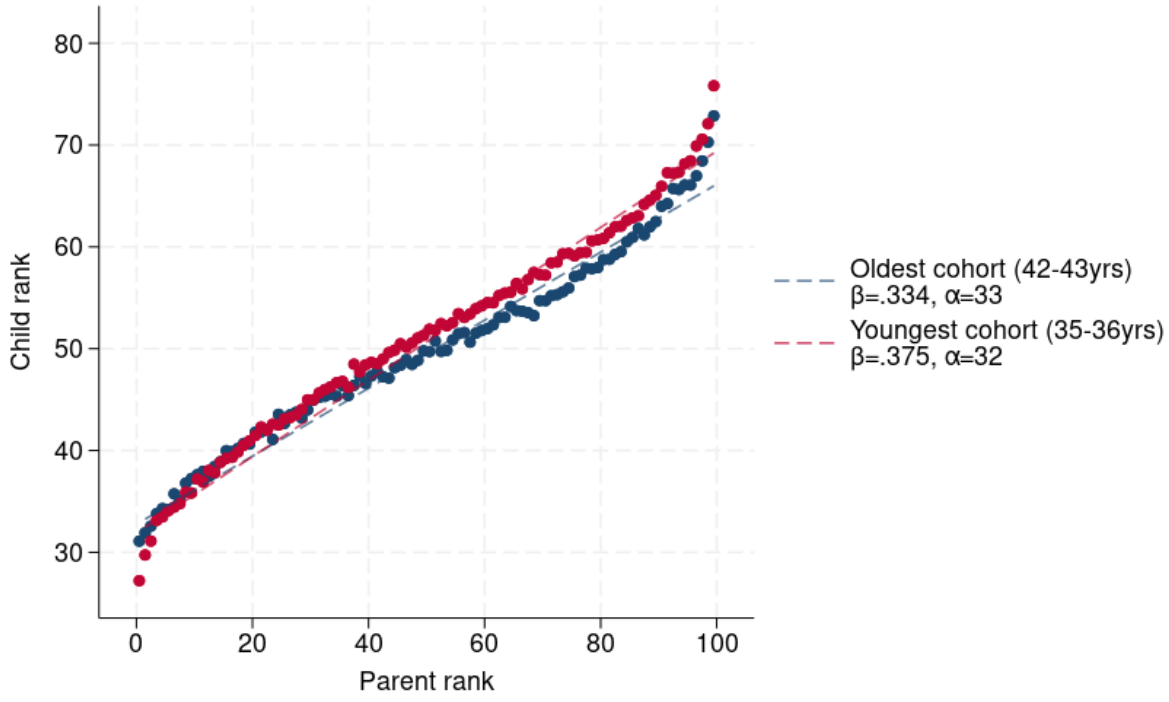


Figure A.2: Child income and parent income and housing assets



B. Additional Tables and Figures

Figure A.3: Rank-rank housing assets by race - no LCA weights

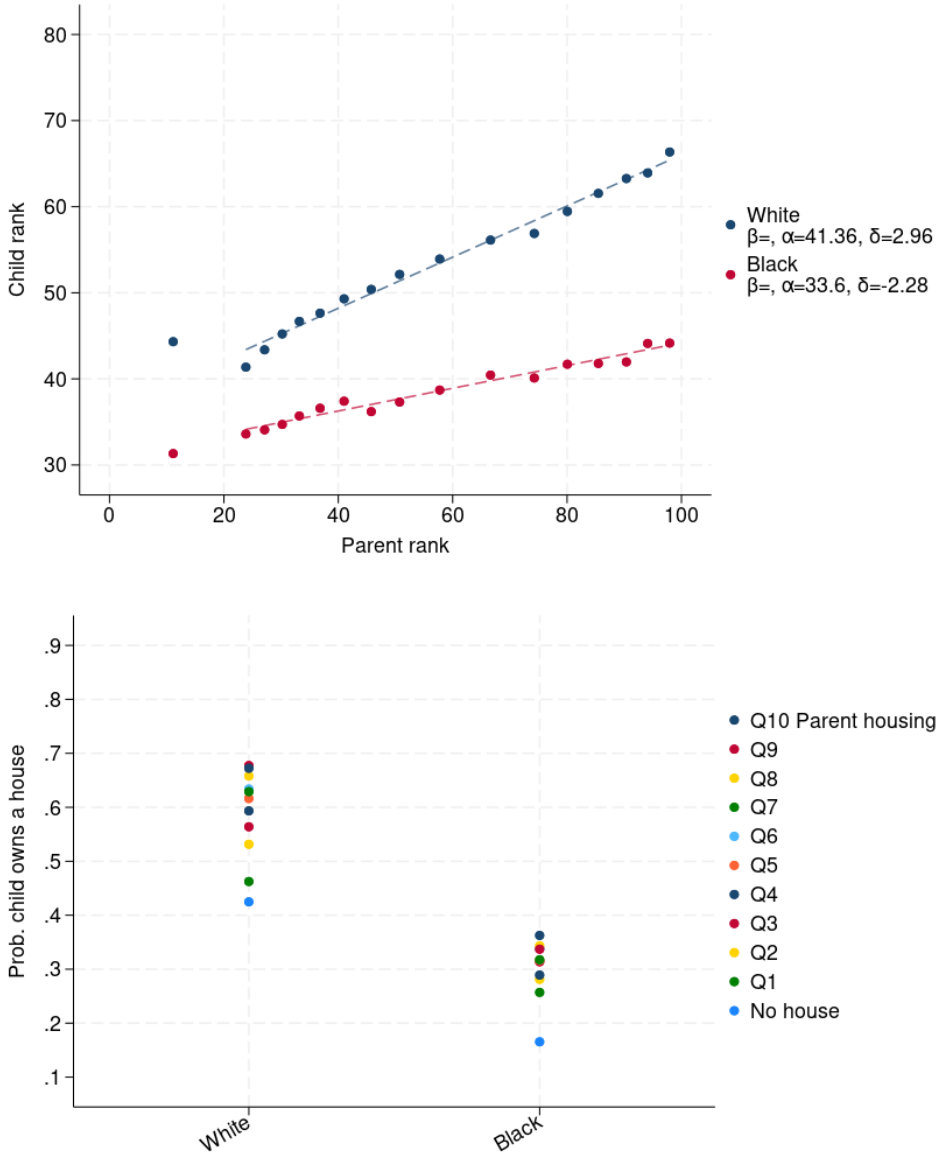
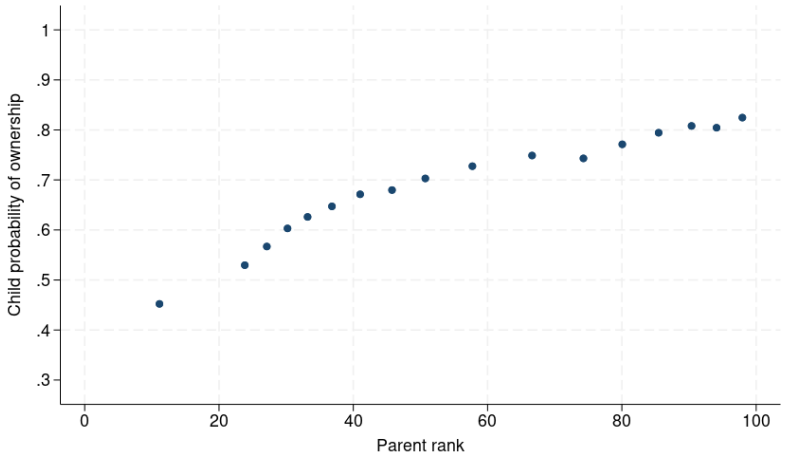
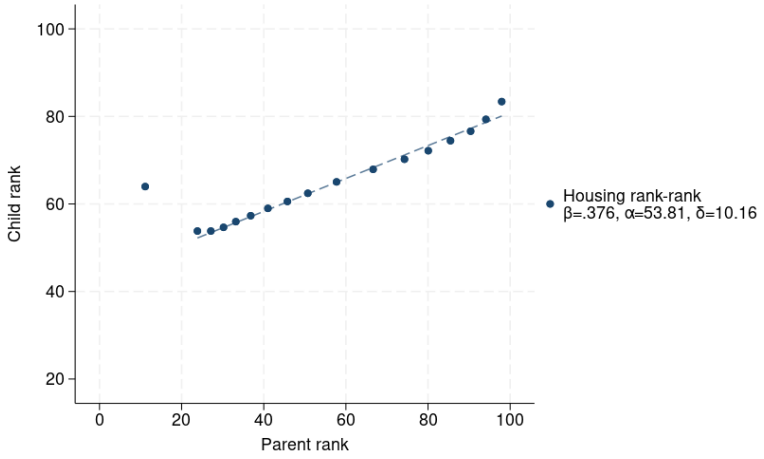


Figure B.4: IGM Housing Relationships - Extensive and Intensive Margins

(a) Probability of home ownership by parent rank



(b) Home value conditional on ownership by parent rank



Note:

B.1. Tables

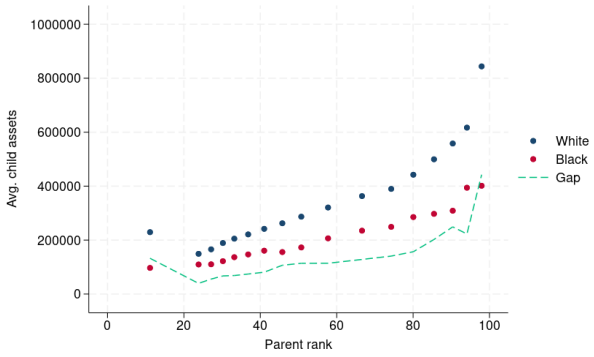
Table B.1: Ownership rates reported in ACS v. Black Knight data

2021		BK owner		*2020*		BK owner	
		No	Yes			No	Yes
ACS	No	0.353	0.058	ACS	No	0.330	0.069
owner	Yes	0.112	0.477	owner	Yes	0.113	0.488

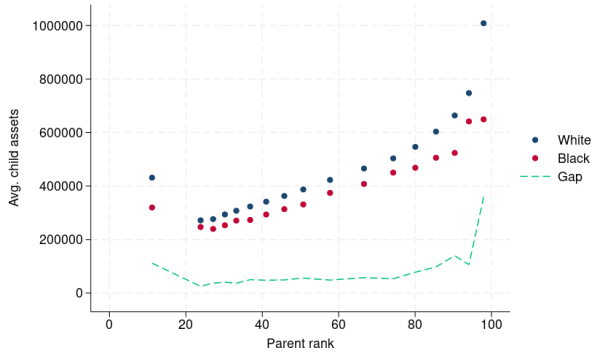
Source: American Community Survey linked to Black Knight property records.

Figure B.5: Average child assets by parent rank and race

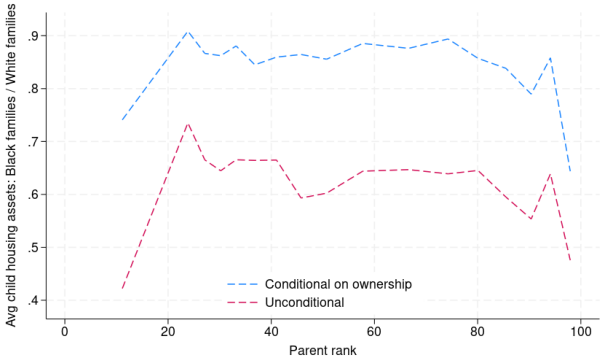
(a) Unconditional housing asset values



(b) Housing asset values conditional on ownership



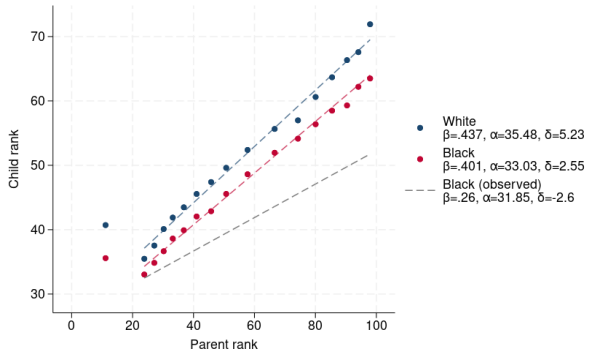
(c) Relative gaps in housing asset values



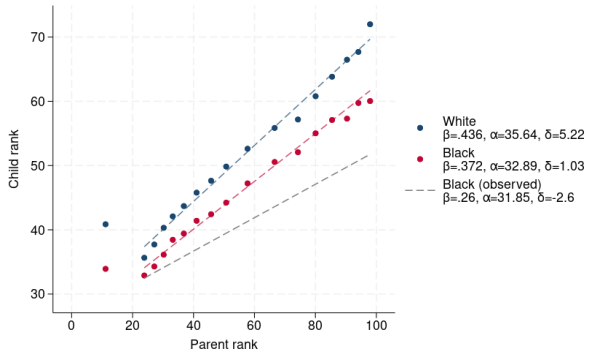
Note:

Figure B.6: Counterfactual Black ownership rates and home values

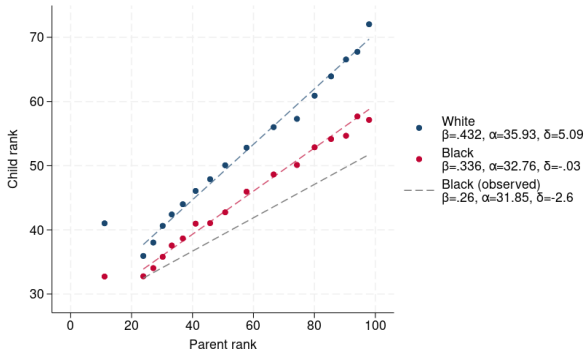
(a) Home values estimated using first-time owners



(b) Home values at 25^{th} ptile White value



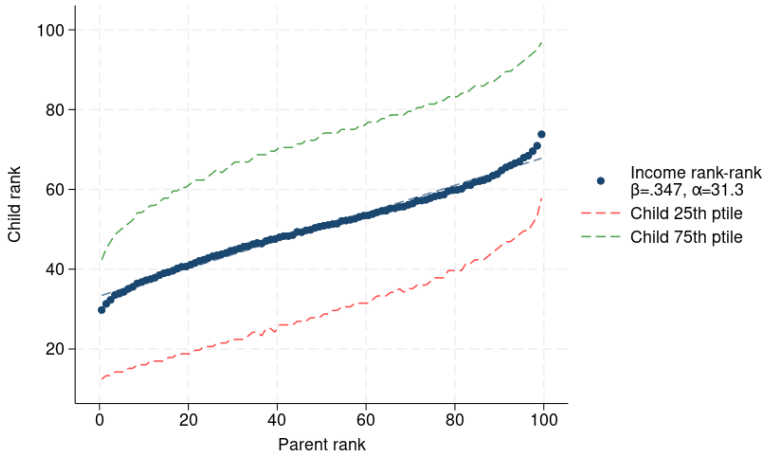
(c) Home values at 10^{th} ptile White value



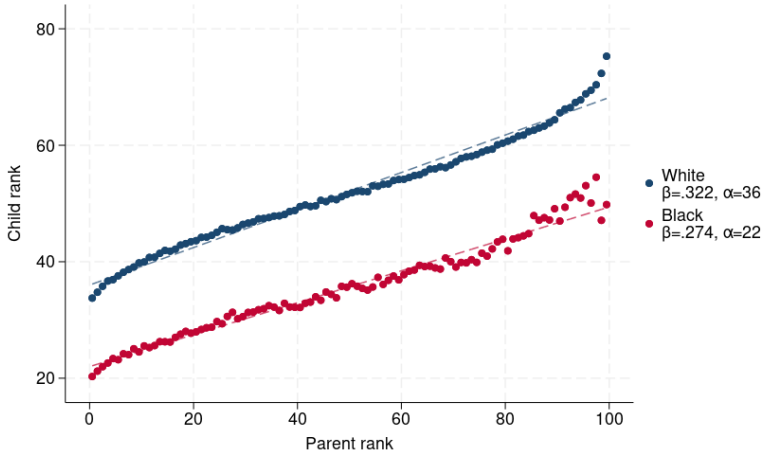
Note:

Figure B.7: IGM Income Rank-Rank Relationships

(a) Income rank-rank



(b) White v. Black families



Note:

Figure B.8: Rank-rank IGM - capitalized wealth

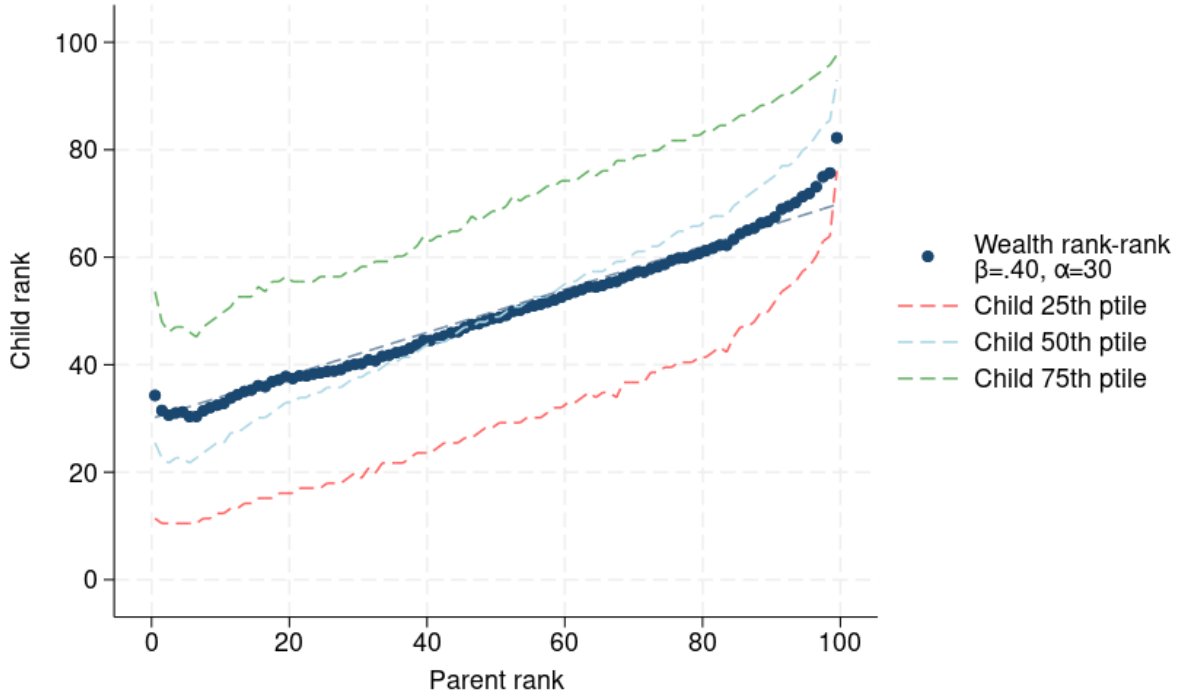


Figure B.9: Aggregate trends in Home Prices, CPI and Median Household Income

