Investigating the American Dream: The Role of Neighborhoods *

Alessandra Fogli † Veronica Guerrieri ‡ Marta Prato §

January 7, 2026

Abstract

In this chapter, we explore the relationship between residential segregation, income inequality, and intergenerational mobility. Using geo-coded NLSY data, we first document that US cities that display a higher level of residential segregation by income also exhibit a significantly lower degree of intergenerational mobility. We then develop a simple general equilibrium model with residential choice in the presence of local spillovers to explore this relationship. Children with higher innate productivity have higher return from the spillover's exposure, but only richer parents can afford the high-spillover neighborhood. We show that the decentralized equilibrium generates less intergenerational mobility than what a utilitarian planner would prescribe and that a simple transfer policy targeting low-income agents can improve welfare upon the equilibrium, but not restore efficiency. Finally, we show that when local spillovers evolve endogenously as a function of the distribution of families that sort in the two neighborhoods, the link between residential segregation and intergenerational mobility becomes even stronger.

Keywords: Intergenerational Mobility, Income inequality, Residential Segregation, Local Externalities
JEL Classification: E24, I2, O15, R2

*We thank the editors Steven Durlauf and Bhashkar Mazumder for helpful comments. We are also grateful to Bipul Verma for excellent research assistance. Word count: 12,800.

†Federal Reserve Bank of Minneapolis, afogli00@gmail.com
‡University of Chicago, Booth School of Business, vguerrie@chicagobooth.edu
§Bocconi University, IGIER, and CEPR, marta.prato@unibocconi.it

1 Introduction

Over the last five decades, while income inequality has risen sharply, US cities have also experienced a substantial increase in residential segregation by income. This phenomenon has important implications for the distribution of economic opportunities across neighborhoods. Using administrative data, Chetty and Hendren (2018b) show the substantial role that neighborhoods play in shaping the economic prospects of children. Children who grow up in wealthier neighborhoods benefit from better educational opportunities, stronger peer groups, more favorable social networks and social norms, which all contribute to greater upward mobility.

In this chapter, we explore how intergenerational mobility is influenced by the presence of neighborhood spillovers. The gap in opportunities across neighborhoods is priced in the cost of housing, which drives an increase in sorting by income within a city. This strengthens the connection between parental and children income, affecting intergenerational mobility and future income inequality at the city-wide level. In particular, children from high-income families, who are exposed to higher-spillover neighborhoods, are more likely to experience better economic outcomes, while children from low-income families, often confined to neighborhoods with fewer resources, face significant barriers to upward mobility.

We first document the simultaneous increase in income inequality and residential segregation by income experienced by the average US metro area between 1980 and 2010. Following Fogli et al. (forthcoming), we also show that there is positive correlation between income inequality and residential segregation across metro areas. Motivated by this evidence, we then use geo-coded NLSY data to explore the relationship between segregation and intergenerational mobility. In particular, we calculate transition probabilities across quartiles between the parents and the children's adult income for cities with different degrees of residential segregation. We find statistically significant lower intergenerational mobility for the 20% metros with highest segregation.

We then use a stylized general equilibrium model with two neighborhoods and local externalities to investigate this relationship. We start with a static version of the model where local spillovers are exogenous and parents differ only for their income and the innate productivity of their children. The tension in the residential choice of a parent comes from the fact that local spillovers are complementary to her child's productivity, but this is not related to parental income. This simple version of the model is already able to capture the two-way relationship between residential segregation and income inequality. We use a battery of comparative static exercises to shed light on it.

Next, we study a utilitarian planner who can decide how to allocate families across neighborhoods and can redistribute consumption across families. When innate productivity is independent from parental income, the planner allocation would place all the children with highest productivity in the neighborhood with higher spillover. This maximizes total future wages by exploiting the complementarity between innate productivity and spillovers. It also means that the planner would like to increase income inequality and then redistribute to decrease consumption inequality. At the same time, the planner allocation obtains perfect intergenerational mobility because the exposure to better opportunities does not depend on parental income. When we let children's innate productivity to be positively related to parental income, the planner allocation generates lower degrees of intergenerational mobility, although still higher than the decentralized equilibrium. We then show that a transfer policy that targets low-income families across the city is able to improve on the equilibrium allocation, without reaching the efficient outcome.

Finally, we consider a richer dynamic model where local spillovers evolve endogenously with the distribution of families across the neighborhoods. In particular, we assume that local spillovers depend on average parental income and average children's innate productivity in the neighborhood. We show that as we increase the weight on average parental income, residential segregation by income becomes stronger driving down intergenerational mobility. We also use a simple calibration to explore the response of the economy to an unexpected permanent skill premium shock, capturing the main source of the increase in inequality during the 1980s, in a similar spirit to Fogli et al. (forthcoming). As the skill premium increases, our model predicts that parental income becomes an even stronger driver of children's outcomes, reducing opportunities for upward mobility, and generating even higher future inequality.

2 Empirical Evidence

This section provides empirical evidence on the relationship between residential segregation by income, income inequality, and intergenerational mobility in the United States. We begin by using Census data to document trends in segregation and inequality over time at the aggregate level, showing that both have increased substantially between 1980 and 2010. We then examine the cross-sectional relationship between these two measures across metropolitan areas, demonstrating that cities with higher levels of income inequality also tend to exhibit greater residential segregation. Finally, we use geo-coded data from the National Longitudinal Survey of Youth (NLSY) to construct intergenerational mobility matrices for metropolitan areas with different levels of segregation, providing direct evidence that children growing up in more segregated cities face lower rates of upward mobility.

As measure of income inequality, we use the Gini coefficient. To quantify residential segregation by income, we use the dissimilarity index, which measures how evenly families from different income groups are distributed across neighborhoods within a metropolitan area. Specifically, we calculate the dissimilarity index for each metro area using the census tract as the geographical sub-unit of analysis. In each decade, we classify rich families as those above the metro's 80th percentile of the family-income distribution; all others are classified as poor. The dissimilarity index is then the share of families that would need to relocate across tracts for the within-metro tract composition of rich vs. poor families to be even.[1] We construct this measure for each decade from 1980 to 2010. Our analysis covers 380 metropolitan areas as defined by the 2003 Office of Management and Budget standards. To construct aggregate trend for both the Gini and the dissimilarity indices, we compute population-weighted averages across all metropolitan areas in our sample.

Figure 1: Inequality and Segregation over Time

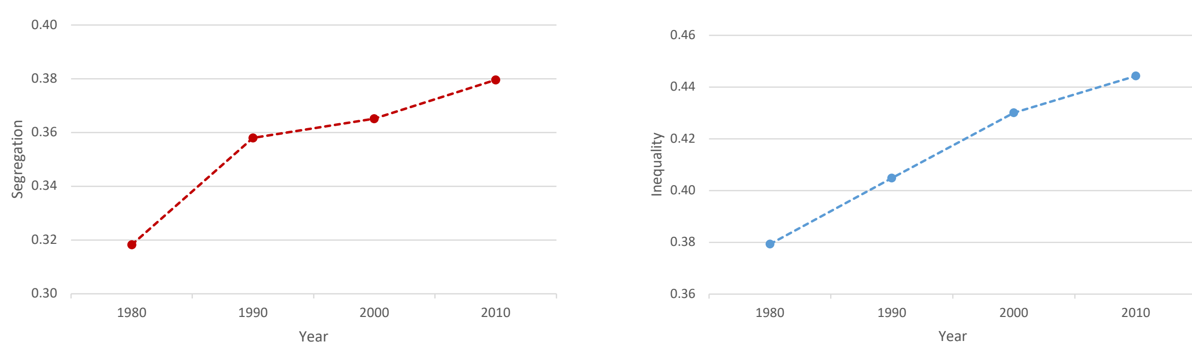

(a) Segregation over Time

(b) Inequality over Time

Figure 1 presents the evolution of residential segregation by income and income inequality at the aggregate level between 1980 and 2010. Panel (a) shows the dissimilarity index over time, which rises substantially from approximately 0.32 in 1980 to 0.38 in 2010. This 6-percentage-point increase indicates that residential segregation by income has intensified considerably over this three-decade period. Panel (b) displays the Gini coefficient of family income inequality, which exhibits a similar upward trajectory, increasing

[1]The dissimilarity index for metropolitan area j is computed using the following formula:

$$D(j) = \frac{1}{2} \sum_i \left| \frac{x_i(j)}{X(j)} - \frac{y_i(j)}{Y(j)} \right|, \tag{1}$$

where $X(j)$ and $Y(j)$ respectively denote the total number of poor and rich families in metro j, while $x_i(j)$ and $y_i(j)$ respectively denote the number of poor and rich families in census tract i within metro j.

from roughly 0.38 in 1980 to 0.44 in 2010. These parallel trends reveal that the spatial concentration of families by income across neighborhoods has grown in tandem with overall income inequality.

Figure 2: Inequality and Segregation by Family Type

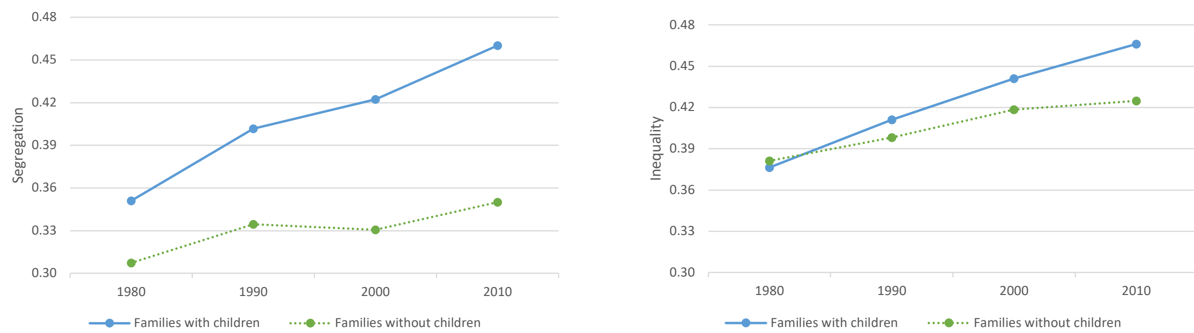

(a) Segregation by Family Type

(b) Inequality by Family Type

The relationship between segregation and inequality is even more pronounced when we focus on families with children. Figure 2 decomposes both segregation and inequality trends by family type, comparing families with children (blue solid line) and families without children (green dotted line). Panel (a) reveals striking differences in segregation levels and trends across these groups. In 1980, families with children exhibited a dissimilarity index of 0.35, already substantially higher than the 0.31 observed for families without children. By 2010, this gap had widened dramatically: the dissimilarity index for families with children reached 0.46, while families without children remained at only 0.35. Panel (b) shows a similar pattern for income inequality. While the Gini coefficient in 1980 was nearly identical for families with and without children (both approximately 0.38), by 2010 inequality for families with children had risen to 0.47, compared to only 0.42 for families without children. These patterns suggest that the presence of children is a key driver of both residential sorting decisions and income inequality dynamics. When choosing where to live, parents may be concerned about school quality, peer environments, social norms, and network opportunities, leading to sharper spatial stratification by income. The fact that both segregation and inequality have increased most sharply for families with children underscores the importance of understanding how neighborhood environments shape children's long-term economic outcomes.

The positive relationship between segregation and inequality is not limited to aggregate trends over time—it also appears in cross-sectional comparisons across metropolitan areas. Figure 3 displays scatter plots of the Gini coefficient against the dissimilarity index for all 380 metropolitan areas in our sample. Panel (a) shows the relationship in 1980, while panel (b) shows the relationship in 2010. Each bubble represents

Figure 3: Inequality and Segregation across US Metros

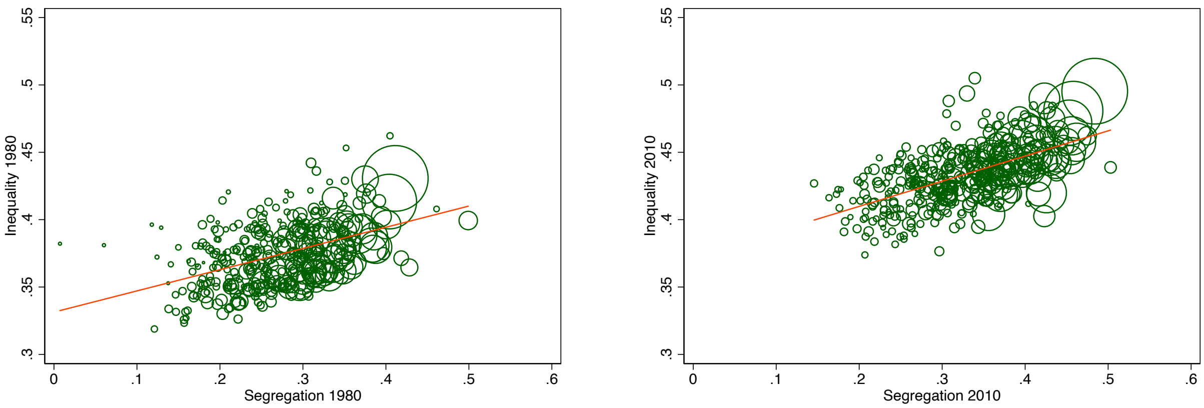

(a) 1980

(b) 2010

a metropolitan area, with size proportional to population, and the orange line represents the fitted linear relationship. Both panels show a clear positive correlation: cities with higher income inequality tend to exhibit greater residential segregation. The relationship appears in both 1980 and 2010, and if anything has become slightly steeper over time, consistent with a reinforcing dynamics between inequality and segregation. This relationship is statistically significant and robust to controls for demographic and industrial composition.

Together, these results indicate that inequality and segregation are closely linked both over time and across space. Metropolitan areas with higher levels of inequality consistently exhibit higher levels of segregation, and cities that experienced larger increases in inequality also experienced larger increases in segregation. This suggests the presence of reinforcing mechanisms between these two phenomena—a conjecture we will explore theoretically in the subsequent sections of this paper. But does this spatial sorting by income have consequences for economic mobility?

To address this question, we examine intergenerational income mobility patterns in metropolitan areas with varying levels of segregation. If neighborhood environments matter for children's long-term economic prospects—through channels such as school quality, peer effects, local role models, or exposure to opportunity—then we would expect children in more segregated cities to face greater persistence in their economic status across generations. We use geocoded data from the National Longitudinal Survey of Youth 1979 (NLSY79) to construct intergenerational mobility matrices that relate children's adult income to their parents' income. The NLSY79 is a nationally representative panel survey that began in 1979 with

12,686 respondents aged 14–22, who have been followed annually through 1994 and biennially thereafter. The restricted-access geocoded version of the NLSY79 allows us to link each respondent to their metropolitan area of residence during adolescence.

To measure intergenerational mobility, we proceed as follows. For each NLSY79 respondent, we calculate initial family income as the average of reported family income in the survey years 1979–1981 (when respondents were adolescents living with their parents). We calculate final income as the average of the respondent's own family income in the survey years 1998–2006 (when respondents were aged 33–47). All income values are deflated to 1978 dollars using the Consumer Price Index. We then assign each respondent to a quartile of the initial income distribution (based on their parents' income) and to a quartile of the final income distribution (based on their own adult income).

Using the geocoded data, we associate each respondent with the metropolitan area where they resided during the initial period (1979–1981). We merge this information with our measures of residential segregation for each metropolitan area in 1980. Given the high dispersion in segregation across metropolitan areas (with a standard deviation of 0.7), we then categorize them into two groups: high-segregation metros are those in the top 20th percentile of the dissimilarity index distribution, while low-segregation metros comprise the remaining 80%.

For each group of metropolitan areas (high-segregation and low-segregation), we construct a 4×4 intergenerational transition matrix. Each cell (i, j) of the matrix represents the probability that a child whose parents were in income quartile i ends up in income quartile j as an adult.

Figure 4: Intergenerational Mobility

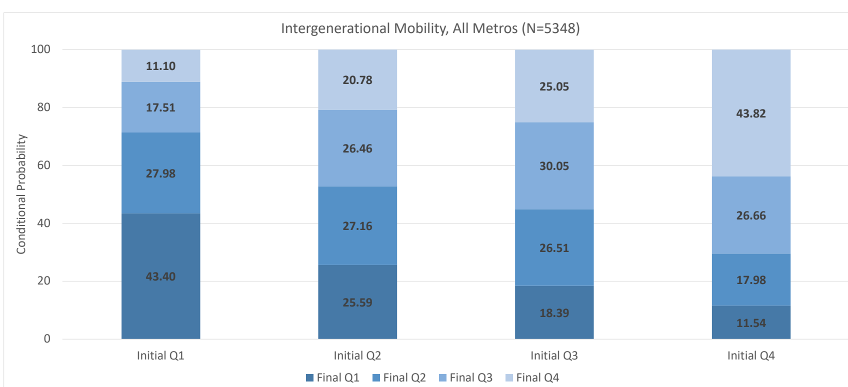

Figure 4 presents the intergenerational mobility matrix for all metropolitan areas in our sample, pooling across all levels of segregation. Each bar represents one of the four initial quartiles (Q1 through Q4), and the segments within each bar show the distribution of final adult income quartiles for children who started in that initial quartile. The figure reveals substantial persistence in income across generations: children born into the bottom quartile (Q1) have a 43.40% probability of remaining in the bottom quartile as adults, while children born into the top quartile (Q4) have a 43.82% probability of remaining in the top quartile. Conversely, the probability of moving from the bottom quartile to the top quartile is only 11.10%.

Figure 5: Intergenerational Mobility by Segregation

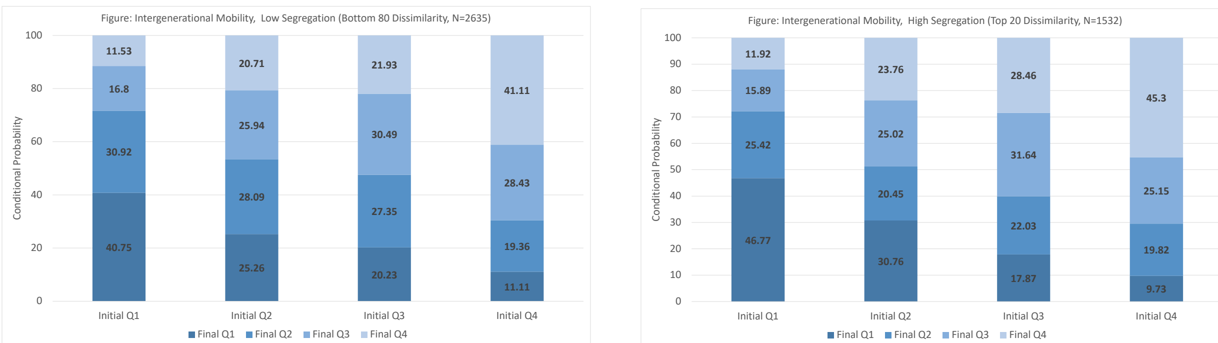

(a) Low Segregation

(b) High Segregation

Figure 5 decomposes these mobility patterns by level of residential segregation, comparing low-segregation metropolitan areas (panel a) with high-segregation metropolitan areas (panel b). The differences are striking and economically meaningful. In low-segregation cities, children born into the bottom quartile have a 40.75% probability of remaining there as adults. In high-segregation cities, this probability increases to 46.77%—nearly 6 percentage points higher. Similarly, among children born into the top quartile, the probability of remaining in the top quartile is 41.11% in low-segregation cities but rises to 45.30% in high-segregation cities. These patterns indicate that residential segregation is associated with substantially greater persistence of economic status across generations at both ends of the income distribution.

Table 1 formalizes these comparisons by presenting the differences in transition probabilities between high-segregation and low-segregation cities, along with standard errors and statistical significance tests. The table confirms that several of these differences are statistically significant. Most notably, the probability of remaining in the bottom quartile (Q1-Q1 persistence) is 6.02 percentage points higher in high-segregation cities, significant at the 10% level. Similarly, the probability of transitioning from Q1 to Q2 is 5.50 percentage points lower in high-segregation cities, also significant at the 10% level. At the top of the distribution, children born into Q4 in high-segregation cities are 6.53 percentage points more likely to fall to Q3 (significant at the 10% level).

Table 1: Difference in Transition Matrices: High Segregation - Low Segregation

	Initial Quartile			
Final Quartile	Q1	Q2	Q3	Q4
Q1	6.02*	5.50	-2.36	-1.38
	(3.14)	(3.43)	(3.08)	(2.26)
Q2	-5.50*	-7.64**	-5.32	0.46
	(3.01)	(3.16)	(3.38)	(3.02)
Q3	-0.91	-0.92	1.15	-3.28
	(2.46)	(3.40)	(3.78)	(3.39)
Q4	0.39	3.05	6.53*	4.19
	(2.28)	(3.27)	(3.59)	(3.82)

Notes: Standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Overall, these findings provide compelling evidence that residential segregation by income is associated with lower intergenerational mobility. Children growing up in highly segregated metropolitan areas face greater obstacles to upward mobility and are more likely to remain trapped in the economic circumstances of their birth. Conversely, children from affluent families in segregated cities enjoy greater advantages in maintaining their privileged status. These patterns are consistent with theories emphasizing the importance of neighborhood effects, peer influences, school quality, and exposure to economic opportunity during childhood. By concentrating poverty and affluence in separate neighborhoods, residential segregation appears to create divergent developmental environments that perpetuate inequality across generations.

This empirical evidence motivates the theoretical model developed in the next section, which explicitly incorporates local spillover effects and residential sorting to analyze a feedback mechanism between inequality and segregation and its consequences for intergenerational mobility.

3 A Simple Static Model

To explore the main mechanism behind the link between residential segregation, inequality, and intergenerational mobility, we use a stylized static model of a city with overlapping generations. This model is the simplified version of a wide class of models used in the literature and, in particular, builds on Fogli et al. (forthcoming).

There is a continuum of parents of measure 1, each with one child. A parent can be characterized by a pair (w,a), where $w \in [\underline{w},\overline{w}]$ denotes her own income, and $a \in [\underline{a},\overline{a}] =$ denotes her child's innate productivity. Let $F(w,a)$ denote the joint distribution of parental income and children's innate productivity. We start by assuming that the process for innate productivity a is *iid* and we will relax this assumption later on.

There are two neighborhoods $k \in \{A,B\}$ and each agent lives in a house of the same size and quality. We denote by R_k the rent in neighborhood k. We assume that in neighborhood A there is a fixed supply of houses H, while in neighborhood B there is a fully elastic supply of houses produced in a competitive market with marginal cost of construction normalized to zero.[2] This implies that the rent in neighborhood B is normalized to zero, that is, $R_B = 0$, while the rent in neighborhood A, R_A, is the only endogenous price. We assume that the houses in neighborhood A are owned by the parents in the top 20-th percentile of the income distribution, proportionally to their income.[3]

A key feature of the model is that there are local spillovers. In particular, the future income of a child with innate productivity a and parental wage w, who grows up in neighborhood $k \in \{A,B\}$ is equal to

$$w'(a,w,k) = b + a^{\alpha} S_k^{\beta},$$

where S_k denotes the local spillover in neighborhood k and captures the quality of public schools, peer effects, cultural and social norms, network effects, and so forth. For now we assume that S_A and S_B are exogenously given. Later, we will explore how they can evolve with the distribution of agents living in the neighborhood. It is important that our wage equation assumes that the local spillover is complementary to the child's innate productivity, that is, the local spillover has higher effect on children with higher a.[4]

As is common in this class of models, parents care about their own consumption and directly about their children's income (warm glow). In particular, they choose their consumption and the neighborhood where

[2]These assumptions can be easily generalized, as long as there is at least one neighborhood with no perfectly elastic supply of houses.

[3]That is, we assume that the total rents $R_A H$ are redistributed to parents with wage $w \geq \bar{w}$ with $F_w(\bar{w}) = 0.8$, with weights $s(w) = w / \int_{w \geq \bar{w}} w dF_w(w)$, where F_w is the marginal distribution with respect to w.

[4]This assumption is consistent with work by Sacerdote (2001), Imberman et al. (2012), and Lavy et al. (2012), who find that high productivity students are the ones who benefit the most from peer effects of other high productivity students. Another paper that speaks more specifically to the complementarity between innate productivity and spillover effects is Card and Giuliano (2016), which shows that high achievers from minority and disadvantaged groups show high returns when included in school tracking programs.

to raise their children to solve their optimal problem

$$
\begin{aligned}
U(w,a) = \max_{c,n} \log(c) + \log(w') \qquad \text{(P1)} \\
s.t.\ \ c + R_n \leq w + s(w) R_A H \mathbb{1}_{w > \bar{w}} \\
w' = b + a^{\alpha} S_n^{\beta}
\end{aligned}
$$

taking as given R_A.

The optimal residential policy that solves problem (P1) reduces to a cut-off policy, such that all parents with wage $w \geq \hat{w}(a)$ optimally choose to raise their children in neighborhood A and the others in neighborhood B, where the cut-off function $\hat{w}(a)$ is given by

$$
\hat{w}(a) = R_A \frac{b + a^{\alpha} S_A^{\beta}}{a^{\alpha}(S_A^{\beta} - S_B^{\beta})}. \qquad (2)
$$

Figure 6 plots such a function and shows that it is non-increasing in a. This is because for a given child's productivity, richer parents are more willing to pay for the higher spillover. At the same time, parents with the same wage but children with higher productivity have stronger incentive to pay for the higher spillover, given the complementarity between productivity and spillover.

Figure 6: Equilibrium Neighborhood Choice Threshold

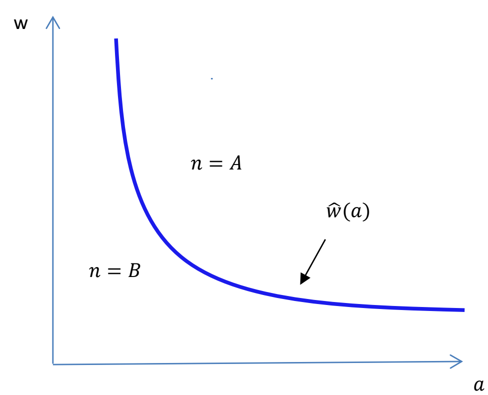

We can now define an equilibrium.

Definition. *For given initial distribution $F(w,a)$ and given $S_A > S_B$, an equilibrium residential cut-off policy $\hat{w}(a)$ and a rental rate R_A that satisfy the following conditions:*

1. *agents' optimization: the policy functions $n(w,a)$ solves problem (P1), for given R_A;*

2. *market clearing: R_A ensures housing market clearing in neighborhood A:*

$$H = \int \int_{w > \hat{w}(a)} F(w,a) dw da; \tag{3}$$

The fact that the optimal residential cutoff function is not-increasing implies that the equilibrium is going to feature residential segregation both by income and by innate productivity, as neighborhood A will be inhabited by richer families with children with higher productivity. An interesting implication of this model is that there could potentially be a feedback effect between income inequality and residential segregation by income.

As income inequality increases, there will be more rich parents willing to pay to live in neighborhood A. Under the lens of the model, this means that, for a given rental rate R_A, there will be a larger mass of families to the right of the cut-off function, pushing up the demand for housing. However, the supply of housing in neighborhood A is fixed, so, in general equilibrium, the rental rate in A, R_A, has to increase so that the housing demand does not change and the market clearing condition (3) is satisfied. Will residential segregation by income increase as a result of that? When the rental rate R_A increases there are two opposing effects. On the one hand, the cut-off function $\hat{w}(a)$ shifts up, increasing the degree of segregation by income. On the other hand, the curve becomes steeper, as $\partial^2 \hat{w}(a) / \partial a \partial R_A = -\alpha b / [a^{\alpha+1}(S_A^\beta - S_B^\beta)] < 0$. When the rental rate in A increases, parents with low-productivity children, everything else equal, prefer to live in B, given that they have a lower return from living in A and paying the higher rent, due to the complementarity between spillover and productivity. This dampens the increase in residential segregation by income driven by the first effect and tends to increase the degree of residential segregation by productivity. Below we show that with a simple calibration the first effect drastically dominates and a rise in inequality generates an increase in residential segregation by income.

It is interesting to explore the other direction of the relationship, that is, whether also an increase in residential segregation by income generates a rise in inequality. Remember that in our stylized model children income depends only on innate productivity and on the neighborhood spillover. As long as children innate productivity is *iid* and unrelated with parental income, this implies that children income inequality increases only when there is more residential segregation by productivity and high-productivity children are exposed to the higher spillover. This might or might not be related to an increase in residential segregation by income. However, in a more general version of the model where innate productivity is persistent across generations, then richer parents tend to have children with higher innate productivity and an increase in residential segregation by income naturally generates an increase in residential segregation by productivity as well. In Section 5 we consider a version of the model where spillovers are endogenous, in which case, an increase in residential segregation by income generates larger spillover gaps, and, in turn, a rise in future inequality, even if innate productivity is iid.

3.1 Comparative Statics

We now perform some comparative static exercises to further explore the relationship between inequality and residential segregation. In particular, we consider two exercises. First, we explore how the model responds to an increase in inequality, as captured by the volatility of parental income. Second, we investigate how the model responds to an increase in the spillover gap, which we interpret as an exogenous change in residential segregation. These exercises show the two-way relationship between inequality and residential segregation and its effects for intergenerational mobility.[5]

Our baseline measure of residential segregation by income is the dissimilarity index, where we define rich parents the ones in the top 20th percentile of the city-wide income distribution, and poor all the others. We use the same index to measure the degree of residential segregation by innate productivity, where we define high-productivity children the ones in the top 20-th percentile of the productivity distribution, and low-productivity children the others. Finally, from now on we use income volatility as our baseline measure of income inequality.[6]

Let us now perform the first exercise and explore the effects of changing the variance of the parental income distribution between 0.1 and 0.5. Panel (a) and (b) in Figure 7 show that an increase in parental income inequality generates an increase in residential segregation by income, while it generates a decrease in residential segregation by innate productivity. This is the case when the size of the high-spillover neighborhood is small enough relative to the share of rich families in the city. In this case, when income volatility increases and the right tail becomes fatter, the demand for living in neighborhood A increases, driving up the rental rate. Panel (c) in Figure 7 shows that the mass of parents who would choose to live

[5]For these exercises, we set the parameters of the wage function α, β and b and the housing supply H equal to the values of the same parameters we obtain from the calibration of the richer dynamic version of the model in Section 5. Moreover, we assume that children's innate productivity is drawn from a lognormal distribution with mean and standard deviation as obtained in the same calibration. We also set the exogenous spillovers in the two neighborhoods equal to the corresponding steady state values, which implies a spillover gap of 0.48. Finally, we assume that parental income is drawn from a lognormal distribution with mean 1 and standard deviation 0.5. We use a larger value of the standard deviation than the one obtained in the steady state to better illustrate the qualitative features of the model.

[6]Other measures of income inequality, such as, the Gini coefficient, the 90-10 ratio, or the Theil index, feature a similar qualitative behavior.

in A increases with income volatility, if we keep the rental rate in A fixed at the equilibrium level for the lowest income volatility value. This corresponds to the mass of agents (w,a) to the right of the cutoff function for the equilibrium with parental income volatility equal to 0.15. Panel (d) shows the resulting general equilibrium effect: as the demand pressure increases, the equilibrium rental rate in A also increases as a function of income volatility. As a result of that, more poor people are pushed out from neighborhood A. This effect increases residential segregation by income, but decreases sorting by innate productivity, because the poor parents with high-productivity children can no longer afford the higher rental rate in A.

Figure 7: Varying Parental Income's Volatility - Inequality, Segregation, and Housing Market

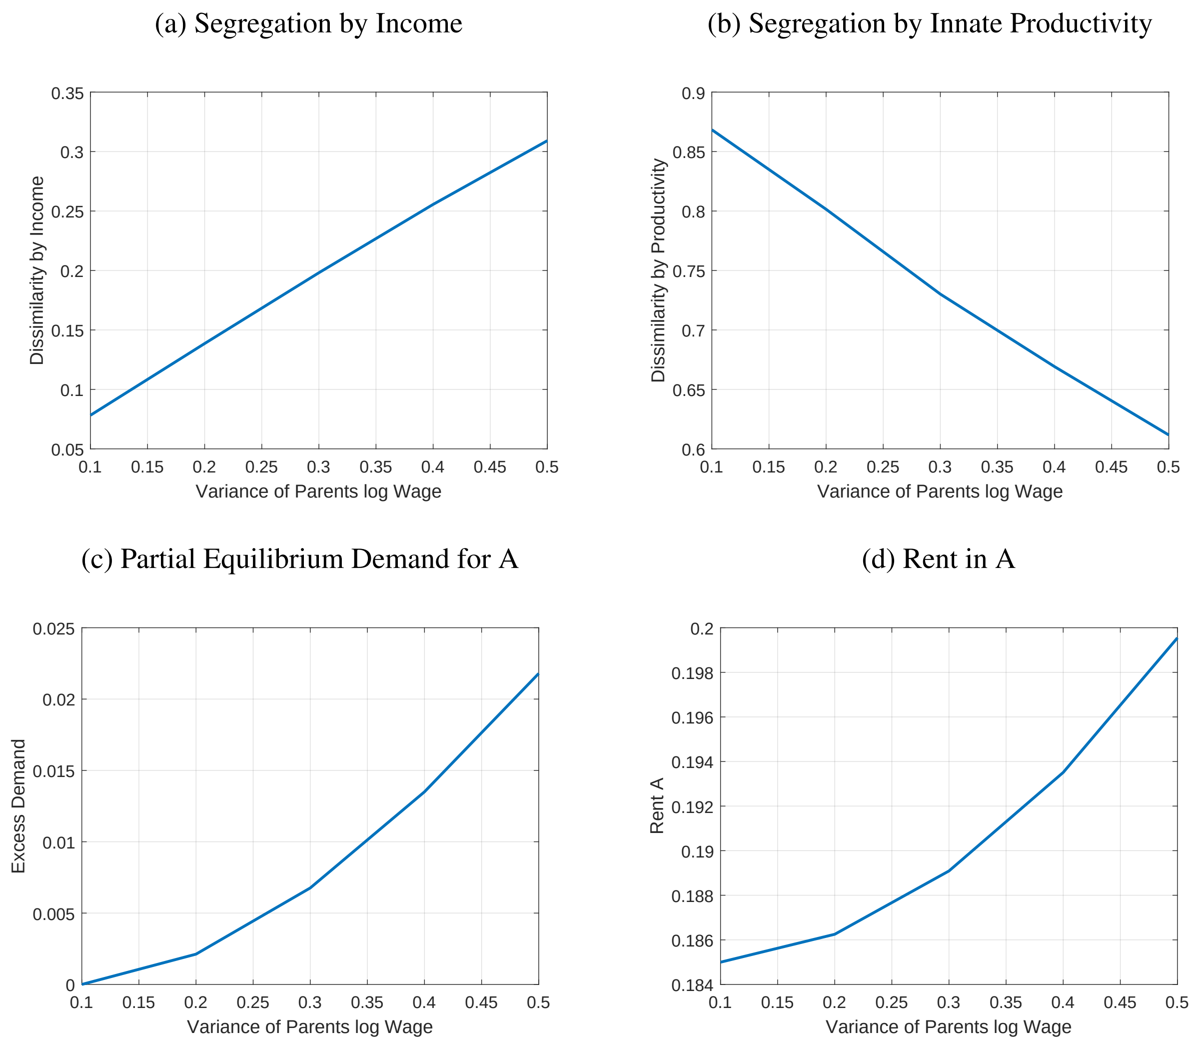

As more rich families live in neighborhood A, their children will be exposed to the highest spillover and, for the same innate productivity, will obtain a higher income than the children of poor families growing up in neighborhood B. This, in turn, will reduce intergenerational mobility, as shown in Figure 8. Panel (a) shows that the rank-rank correlation increases with parental income inequality, while panel (b) shows that the Q1-to-Q1 and Q4-to-Q4 transition probabilities increase with parental income inequality, while the Q1-Q4 decreases, overall reducing intergenerational mobility.[7]

Figure 8: Varying Parental Income's Volatility - Intergenerational Mobility

(a) Rank-Rank Correlation (b) Quartile Transitions

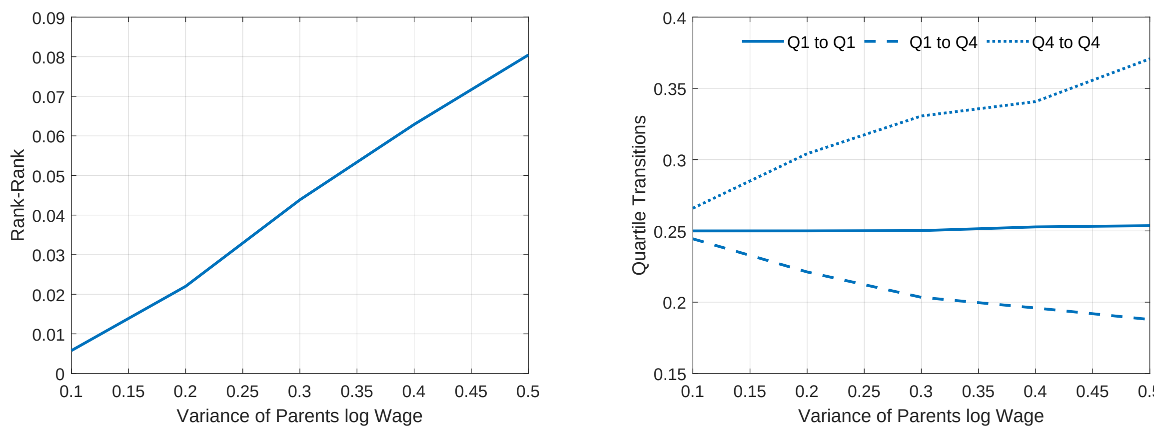

Overall, this first exercise confirms that in our model higher inequality can generate higher residential segregation by income. The next exercise explores the other direction of the relationship.

Our second exercise aims to explore the effect of an increase in residential segregation by income on inequality. In the context of our model, we choose the spillover gap between the two neighborhoods as an exogenous source of variation for residential segregation by income. Panel (a) in Figure 9 confirms that residential segregation by income increases with the spillover gap. Why is that? As the spillover gap increases, there is going to be higher demand to live in neighborhood A, as there is a higher return in raising children in A relative to B. Panel (c) confirms that, if we keep the rental rate fixed, the mass of parents who would like to live in A increases with the spillover gap. The figure plots the mass of families to the right of the cutoff function obtained in the equilibrium with the lowest spillover gap. However, in general equilibrium this generates a rise in the rental rate in A, as shown in panel (d), shifting the cut-off function to the right, and increasing residential segregation by income. Panel (a) also shows that, on the contrary, residential segregation by innate productivity decreases with the spillover gap. This is because,

[7]It is interesting to notice that children's income inequality also increases with parental income volatility. In some of the following exercises, we will explore this effet in more detail.

as the rental rate in A becomes higher, poor parents with high-productivity children won't be able to afford to live in neighborhood A and rich parents with low-productivity children will be.

Figure 9: Varying the Spillover Gap - Inequality, Segregation, and Housing Market

(a) Segregation (b) Inequality (c) Partial Equilibrium Demand for A (d) Rent in A

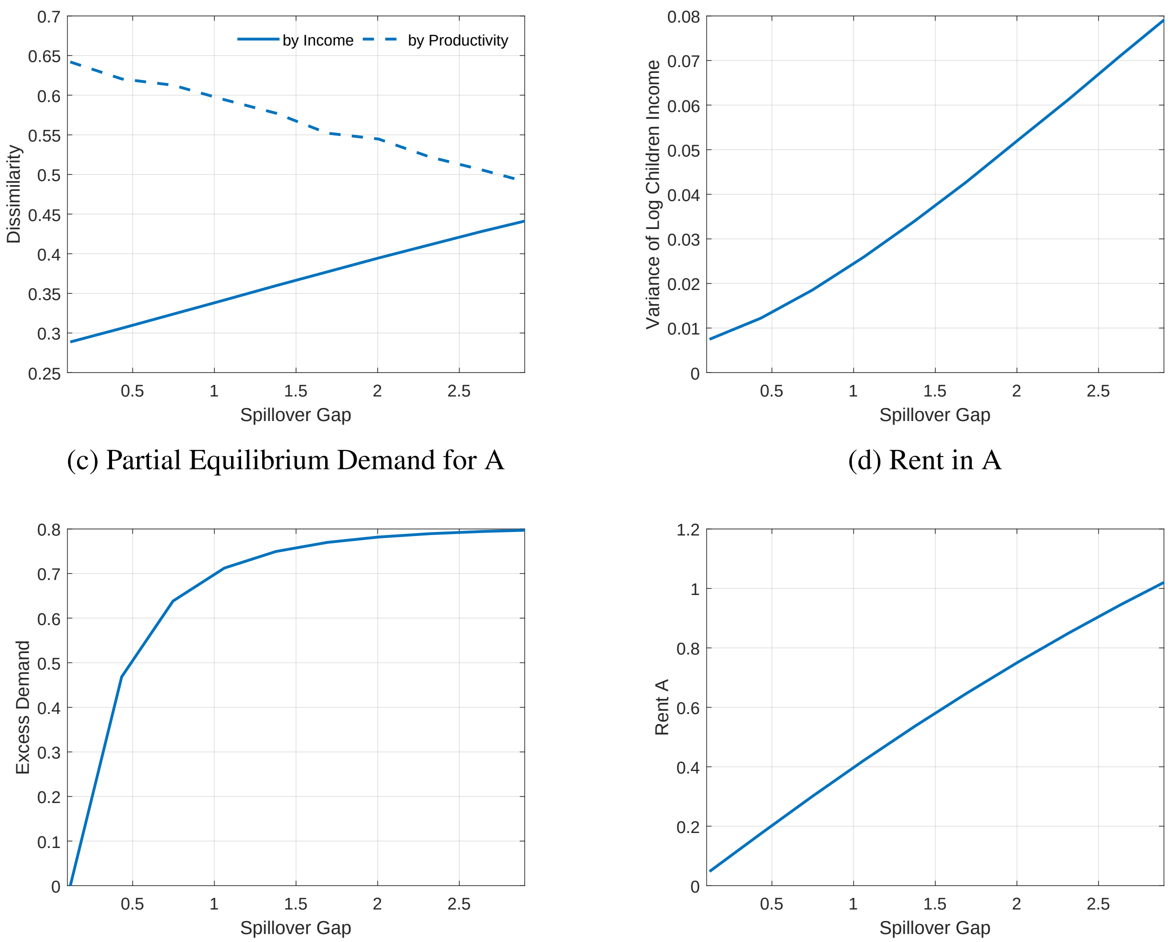

Panel (b) in Figure 9 shows the main result of this exercise, that is, that inequality in children's income increases as the spillover gap increases, which we interpret as residential segregation affecting future inequality. As the spillover gap increases, there is more inequality because the children growing up in A tend to be more productive and are exposed to an even higher spillover, and vice versa for the ones living in B.

Figure 10 shows that also intergenerational mobility decreases as the spillover gap increases. As before, we show different measure of intergenerational mobility: the rank-rank correlation and the transition prob-

Figure 10: Varying the Spillover Gap - Intergenerational Mobility

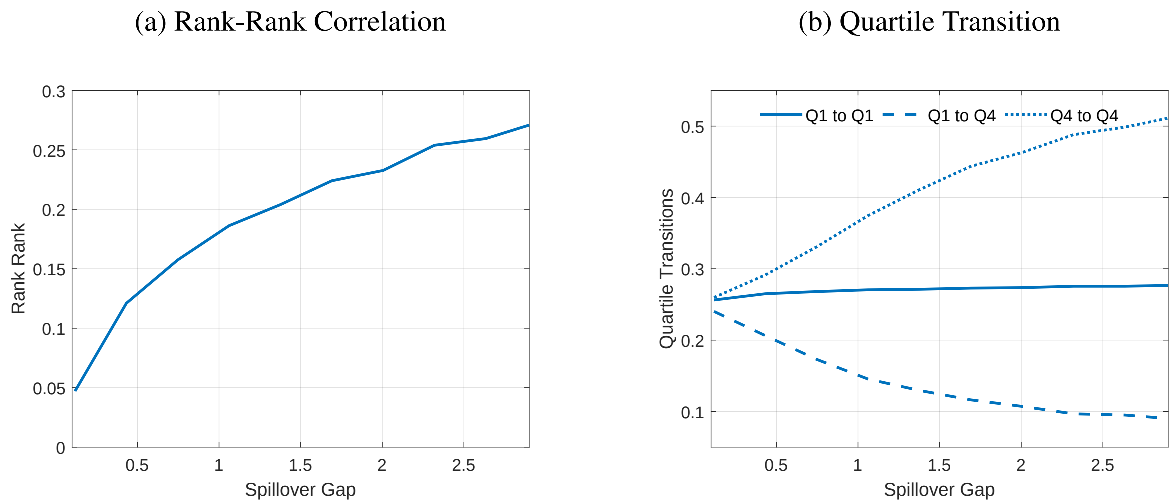

abilities for different quartile of the income distribution. Again, when the rental rate in A is higher and residential segregation by income is higher, more rich children will be exposed to the highest spillover and vice versa, reducing the opportunities of climbing up the social ladder.

Both these comparative exercises show that the model generates a positive correlation between income inequality, residential segregation by income, and the intergenerational correlation of incomes. These findings are consistent with the cross-sectional regularities documented in Section 2, e.g. the fact that metros with higher residential segregation by income display both higher income inequality and lower intergenerational mobility.

3.2 Utilitarian Planner

In this section, we analyze the welfare properties of the equilibrium. In particular, we study the optimal problem of a utilitarian planner that can decide both consumption and residential location of the agents taking into consideration the resource constraint and the housing supply constraint.

In particular, the planner chooses $c(w,a)$ and $n(w,a)$ to maximize the average utility of the parents

$$\int\int [u(c(w,a)) + g(b + a^{\alpha} S^{\beta}_{n(w,a)})] dF(w,a)$$

subject to the resource constraint

$$\int\int c(w,a) dF(w,a) \leq \int\int w dF(w,a)$$

and the housing supply constraint in neighborhood A, that is,

$$\int\int_{n(w,a)=A} dF(w,a) \leq H$$

Given that the two neighborhoods are different only because of the level of local spillover, which is complementary to the children's innate productivity, the planner chooses to allocate parents to different neighborhoods purely based on the level of the innate productivity of their children. In fact, Figure 11 shows that the planner residential choice can be characterized by a vertical cut-off function in the space ability-parental wage, so that all the parents with children with productivity above a cutoff $\hat{a}$ live in neighborhood A, and all the others live in neighborhood B. This is different from the equilibrium residential policy that depends not only on the child's innate productivity but also on the parental wage, as the latter determines the ability of the parent to pay for the higher rent in neighborhood A.

Figure 11: Planner vs. Equilibrium - Residential Choice

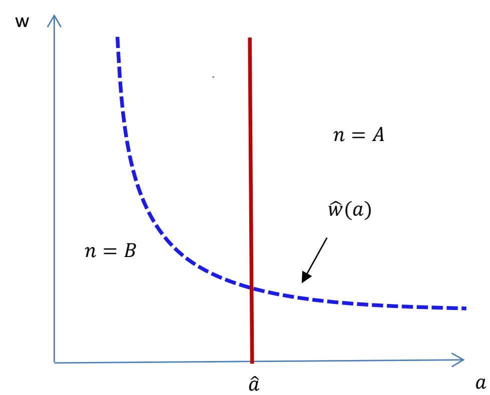

Figure 12 shows how the planner allocation compares with the equilibrium one as we increase the spillover gap. Panel (a) shows that the equilibrium is inefficient, as the planner is able to achieve higher welfare relative to the equilibrium allocation. Moreover, this welfare gap gets larger the larger is the spillover gap. This is because in equilibrium the spillover gap between the two neighborhoods generates a differential in rental rates and hence some degree of sorting by income, the more the larger is the spillover gap. However, in this simple model where innate productivity is unrelated to parental wage, the optimal degree of residential segregation by income is zero, as the planner wants to perfectly segregate parents only based on their children's productivity. Notice that when the spillover gap is zero, the planner can still achieve higher welfare by redistributing across families with children with different innate productivities,

Figure 12: Planner vs. Equilibrium - Welfare, Segregation, Income and Consumption Inequality

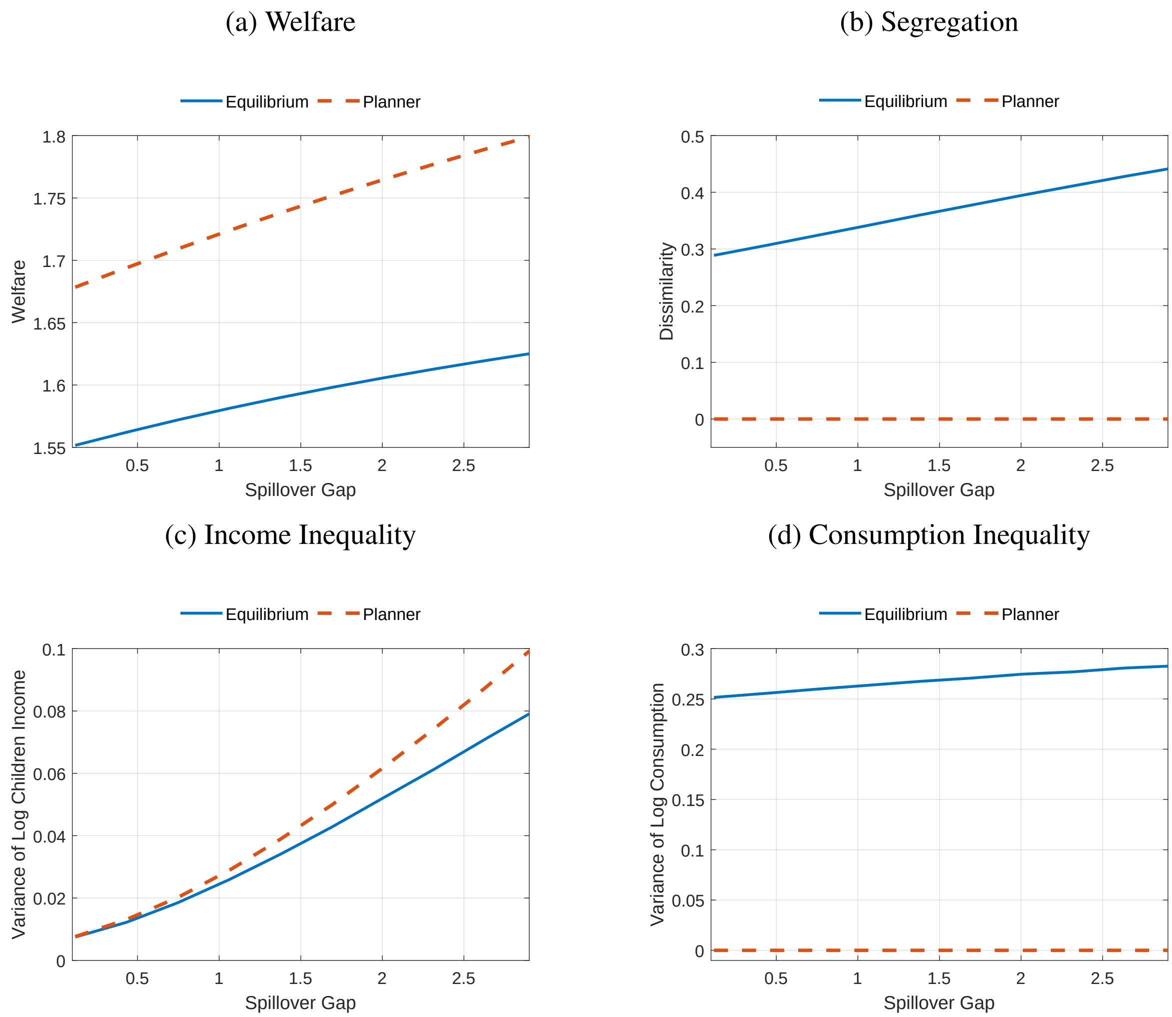

given that there is no possibility of insuring ex-ante. Panel (b) reproduces the result from Figure 9 that indeed residential segregation increases with the spillover gap, increasing the difference with the planner allocation.

Given that in the planner allocation residential sorting is purely driven by innate productivity, high-productivity children are exposed to the highest spillover and low-productivity children to the lowest. Given the complementarity between productivity and spillover, this means that the planner generates inequality in children's income. In particular, panel (c) shows that the equilibrium allocation generates lower children's income inequality than the planner, and the difference becomes larger when the spillover gap increases exactly because residential segregation by income increases, reducing the role of productivity in sorting. However, panel (d) shows that the planner uses redistribution to decrease consumption inequality relative to the equilibrium.

Figure 13 shows that the higher degree of children's income inequality in the planner allocation does not affect intergenerational mobility, as it is purely driven by the children innate productivity. In fact, the planner generates perfect intergenerational mobility, that is, zero rank-rank correlation and transition probabilities across quartiles of the income distribution that are all equal to 25%, irrespective of the spillover gap.

Figure 13: Planner vs. Equilibrium - Intergenerational Mobility

(a) Rank-Rank Correlation (b) Quartile Transitions

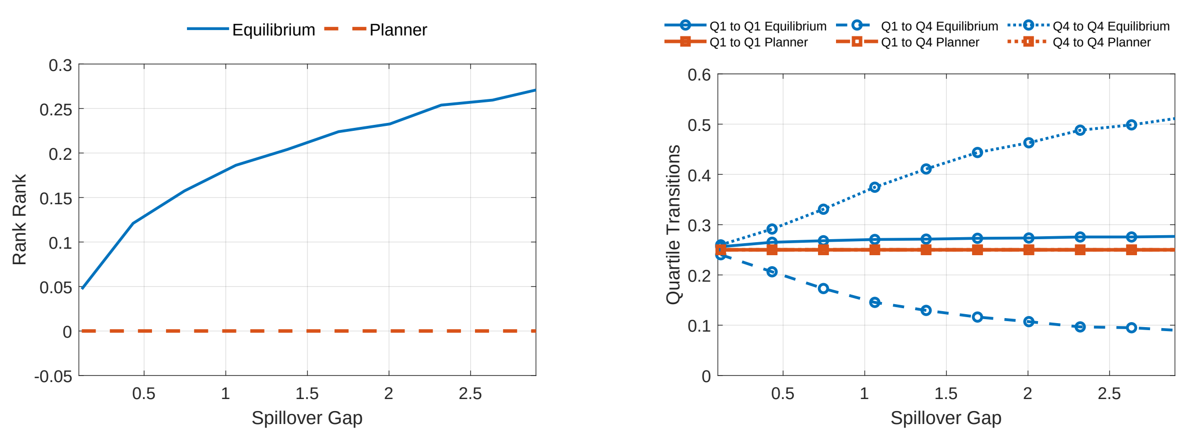

The planner allocation we just described is extreme in terms of consequences for intergenerational mobility and residential segregation by income because we assumed that innate productivity is *iid* and uncorrelated with parental wage. A less stark assumption would be to assume that innate ability has some degree of intergenerational transmission. In that case, even the planner allocation would generate some degree of residential segregation by income, as parental income is positively correlated with innate productivity. And for the same reason, it would generate imperfect intergenerational mobility. At the opposite extreme, when the correlation between innate productivity and parental wage is equal to 1, the equilibrium allocation is efficient.

We now assume that innate productivity is correlated with parental wage, with correlation parameter ρ. Figure 14 compares the planner and the equilibrium allocations as we increase ρ, where $\rho = 0$ corresponds to our baseline assumption of *iid* productivity.

Figure 14: Intergenerational Transmission of Ability - Inequality, Segregation, and Mobility

(a) Segregation (b) Inequality

(c) Rank Correlation (d) Quartile Transitions

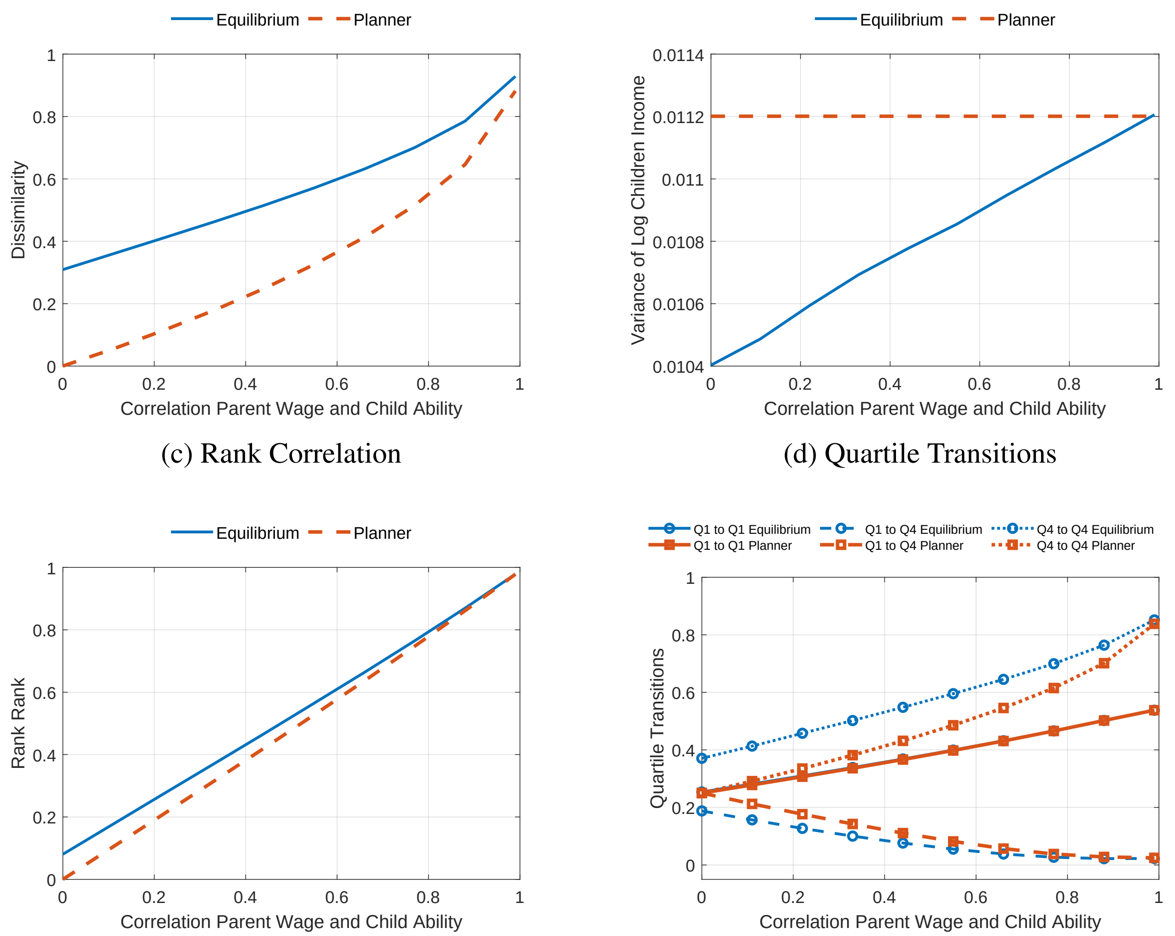

The figure shows that as ρ increases, the planner wants to increase the degree of residential segregation by income, given that parental income gets more correlated with innate productivity. However, as residential segregation by income increases, residential segregation by productivity decreases and income inequality decreases with it. Moreover, for the same reason, intergenerational mobility becomes lower as ρ increases and the residential location is more dependent on parental income.

4 Policy

In this section, we investigate whether a simple transfer policy can get the economy closer to the optimal allocation and improve intergenerational mobility. In particular, we consider a lump-sum transfer corresponding to 20% of the average wage given to all the parents in the lowest $25th$ percentile of the income distribution. The policy is financed with a proportional income tax on all remaining agents in the economy.

Figure 15 compares the welfare under this policy with the welfare obtained in the planner and equilibrium allocation, for different values of the spillover gap. The figure shows that the transfer policy is able to achieve a higher level of welfare relative to the equilibrium for any value of the spillover gap, although it never reaches the welfare level of the efficient allocation.

Figure 15: Welfare - Policy vs Planner vs Equilibrium

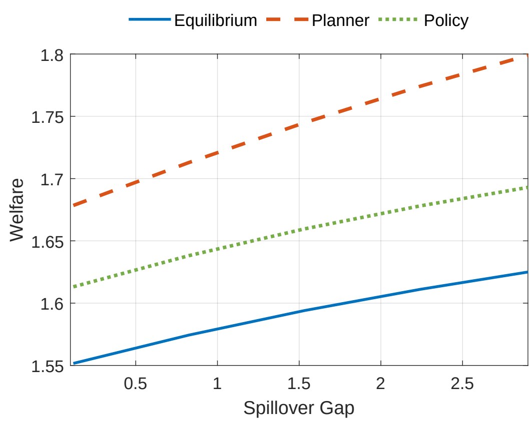

The policy is able to improve welfare, by bringing the economy closer to the optimal residential segregation of families. Panels (a) and (b) in Figure 16 show that the policy reduces residential segregation by income and increases residential segregation by innate productivity, which increases the efficiency of the economy. The policy is able to do so because in equilibrium some parents with high-productivity children might not have the resources to pay for the higher rental rate in neighborhood A, while the transfer allows them to do so. This increases the efficiency of the economy and income inequality, with it. Panel (c) shows that the policy achieves a higher level of income inequality relative to the equilibrium for any spillover gap value, although still lower than what the planner would prescribe. Finally, panel (d) shows that, on the contrary, the transfer reduces consumption inequality, contributing to the higher welfare level.

By allowing low-income parents with high-productivity children to afford the neighborhood with the highest spillover, the policy improves the future income prospects of children born in poorer families. This increases intergenerational mobility relative to the decentralized equilibrium, as shown in Figure 17. Panel (a) shows that the policy reduces the intergenerational rank correlation relative to the decentralized equilibrium. Consistently, Panel (b) shows that the policy reduces the probability that children from low-income

Figure 16: Transfer Policy - Segregation and Inequality

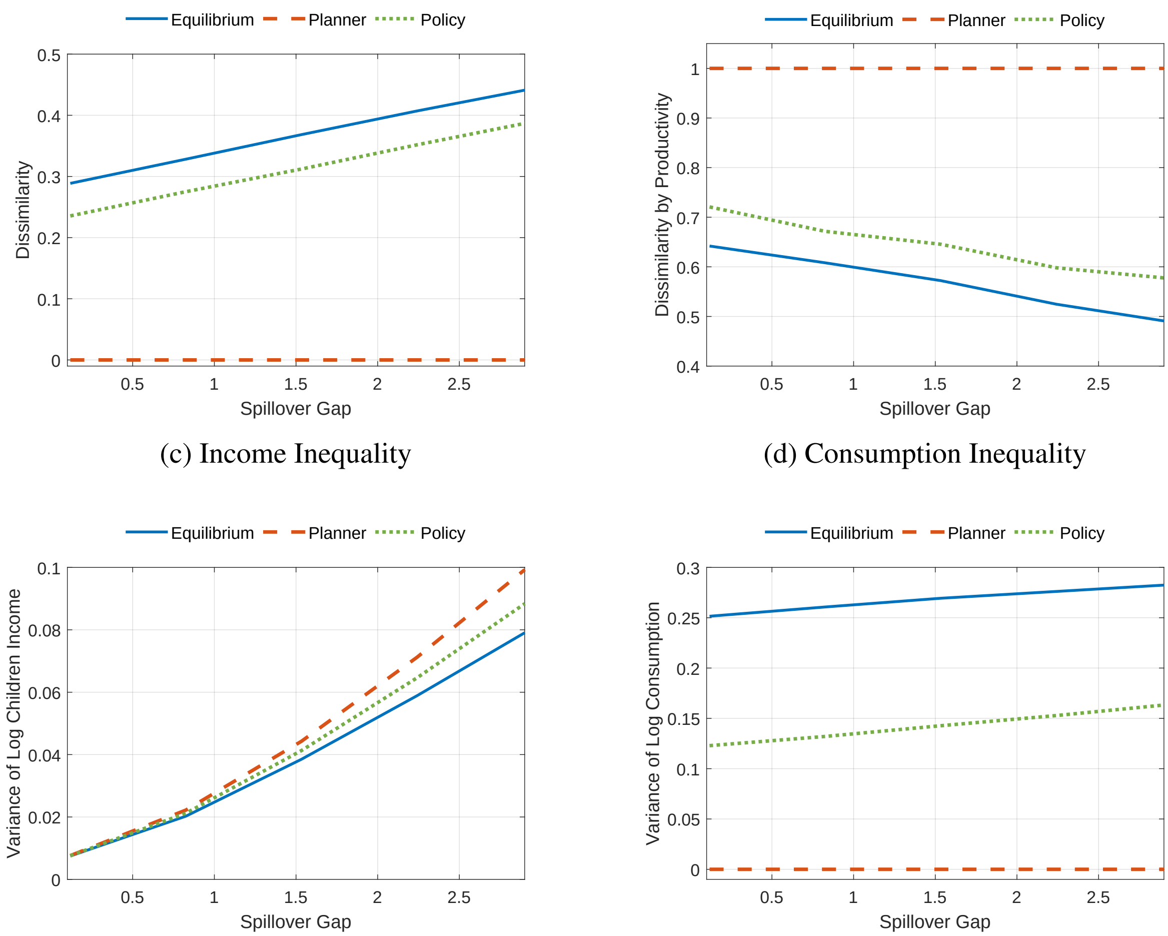

(a) Segregation by Income

(b) Segregation by Productivity

(c) Income Inequality

(d) Consumption Inequality

families remain in the lowest income quartile (Q1-Q1), while increases the likelihood that they move to higher income quartiles (Q1-Q4). However, the figure shows that all the measures of intergenerational mobility remain lower than the level that would be chosen by a planner.

In this section we have explored a policy that is not linked to the geography of the city. However, there is a relatively recent body of work that explores the effects of different types of neighborhood-specific policies in general equilibrium models with residential choice and local spillovers, such as (Agostinelli et al., 2024; Eckert and Kleineberg, 2024; Chyn and Daruich, forthcoming; Fogli et al., 2025b). These papers show that different neighborhood-specific policies might have different impact on the economy and it would be interesting to explore to which extent and in which context they can dominate a simple redistribution

Figure 17: Transfer Policy - Intergenerational mobility

(a) Rank Correlation (b) Quartile Transitions

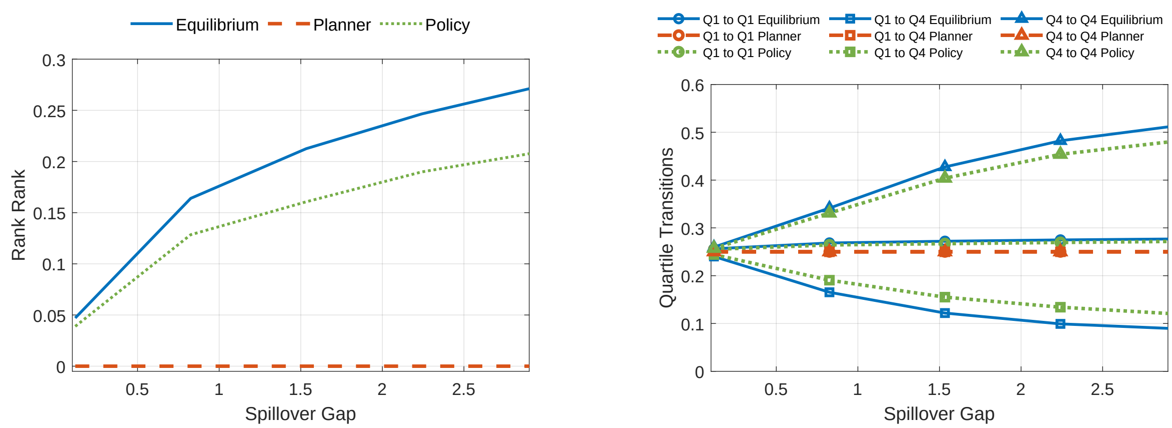

policy at the city-wide level.

5 Dynamic Model with Endogenous Spillovers

We now consider a dynamic version of the model where spillovers are endogenously determined. In particular, we assume that there are overlapping generations of agents, who live two periods: they are children in the first period and become parents in the second. A parent i at time t is characterized by a pair (w_{it}, a_{it}), where w_{it} represents her own wage and a_{it} represents her child's innate productivity. We assume that innate productivity is in part transmitted across generations and, in particular, that it follows an AR(1) process:

$$\log a_{it} = \rho \log a_{it-1} + \varepsilon_{it},$$

where a_{it-1} is the innate productivity of the parent and ε_{it} is a noise that is lognormally distributed with mean 1 and standard deviation σ_a. The joint distribution of parental wage and innate productivity evolves endogenously and is denoted by $F_t(w,a)$, with $F_0(w,a)$ taken as given. All parents with the same parental wage $w_{it} = w$ and the same child's innate productivity $a_{it} = a$ make the same optimal decision, and so from now on we will refer to a parent simply as a pair (w,a).

As in the static model, there are two neighborhoods, A and B, that differ purely because of different local spillovers at each period t, S_{At} and S_{Bt}. Moreover, neighborhood B has a perfectly elastic housing supply and we normalize the rental rate $R_{Bt} = 0$, while neighborhood A has a fixed housing supply H, so that the rental rate R_{At} is determined to clear the housing market in each period:

$$\int \int_{n_t(w,a)=A} dF_t(w,a) = H,$$

where $n_t(w,a)$ is the residential choice of parent (w,a).

Parents choose the neighborhood where to raise their child $n_t(w,a)$ and their own consumption $c_t(w,a)$, taking as given the local spillovers, S_{At} and S_{Bt}, and the rental rate in neighborhood A, R_{At}. The parent's utility is $log(c_t(w,a)) + log(W_{t+1}(w,a))$, where $W_{t+1}(w,a)$ is her child's wage, which is determined according to the following equation:

$$W_{t+1}(w,a) = (b + \eta a^{\alpha} S^{\beta}_{n_t(w,a)t}) w^{\gamma} \varepsilon_t. \tag{4}$$

Relative to the static model, we have now introduced a wage shock ε_t, which is normally distributed with standard deviation σ_ε and mean μ_ε, in order to generate enough income volatility to bring the model to the data.[8] We now also assume that the parental wage w directly affects the child's wage in order to generate empirically reasonable level of intergenerational correlation. Finally, we introduce the parameter η that we interpret in a broad way as skill premium, given that a can be thought of as the child's skill level.[9]

Finally, the novel ingredient of the model is the fact that local spillovers are endogenous. In general, they can be written as function of the distribution of agents living in the neighborhoods, that is,

$$S_{kt} = f(\Phi_{kt}),$$

where $\Phi_{kt}(w,a)$ denotes the distribution of parental wage and innate productivity (w,a) in neighborhood k at time t. In the model, a neighborhood's spillover affects future wage of children growing up in that neighborhood and captures different pecuniary and social externalities, that is, public schools' quality, peer effects, network effects, social norms, learning from neighbors' experience, and so forth. For simplicity, we are going to assume that the spillover takes a special form, that is, depends on a weighted average of parental income and children's productivity:

$$S_{kt} = \omega E_t[w_t|k] + (1-\omega)E_t[a_t|k].$$

In particular, we think that some form of local externalities are primarily driven by average parental income and others by average children's productivity. For example, on the one hand, the quality of public school

[8]Note that the parameter μ_ε can be normalized to an arbitrary value rescaling other parameters accordingly.

[9]See Fogli et al. (forthcoming) for a richer version of the model where parents also choose how much to invest in their children's education and the skill premium corresponds to the return to investment in education.

in the US is strongly affected by average parental income in the school district, given that public schools are primarily locally financed. On the other hand, peer effects are predominantly affected by average productivity of the children who live in the neighborhood. At the extremes, when $\omega = 1$ only parental income matters, and when $\omega = 0$ only children's productivity matters.

5.1 Simple Calibration

For the following numerical exercises, we use parameters obtained with a simple calibration of our model, following a simplified version of the calibration strategy in Fogli et al. (forthcoming). We calibrate the steady state of the model to the average US metro area in 1980, which is before the sharp increase in inequality experienced in the US.

In particular, we target 8 moments to pin down 8 parameters: the four parameters of the wage function, b, α, β, γ, the volatility of the wage shock ε_w, the two parameters of the AR(1) process for innate productivity, ρ and σ_a, and the weight ω for the spillover definition.

Table 2 shows the moments that we target.

Table 2: *Calibration Targets*

Description	Data	Model	Source
Income Volatility	0.16	0.15	Census 1980
Dissimilarity Index by Income	0.29	0.30	Census 1980
Rank-Rank Correlation	0.34	0.36	Chetty et al. (2014)
Q1-to-Q1 Transition Pr	0.46	0.44	NLSY
$(R_A - R_B)$ / Average Income	0.08	0.07	Census 1980
Share of Rich in A	0.43	0.45	Census 1980
Neighborhood Exposure 25th p	0.06	0.06	Chetty and Hendren (2018)
Neighborhood Exposure 75th p	0.05	0.05	Chetty and Hendren (2018)

The first two moments are average income inequality and average residential segregation by income. As our main measure of inequality, we use Census data to calculate the variance of log income. Given that our mechanism relies on the impact of local spillover on children's future income, we restrict the sample to families with children. We calculate income volatility at the metro level and then average them using population weights. As measure of residential segregation by income, we calculate the dissimilarity index for each metro area in 1980 and again average them using population weights. To calculate the dissimilarity index, we define rich the families in the top 20th of the income distribution and poor all the others, and we use two neighborhoods as geographic sub-unit of analysis. In particular, using census tract data, for each metro area, we rank the census tracts based on the share of rich that live there. Then, we define neighborhood A the group of census tracts with a share of rich above a threshold constructed such that roughly 20% of the population of the metro area lives in neighborhood A. Neighborhood B is the group of residual census tracts.

The next two moments are two measures of intergenerational mobility. First, we take the rank-rank income correlation between parents and children from Chetty et al. (2014). Second, we target the Q1-to-Q1 income quartile transition probability that we calculate using NLSY data.

Using the same definition of neighborhoods for each metro that we described above, we also target the share of rich families living in neighborhood A, and the average rental rate gap between the two neighborhoods, normalized by average income.

Finally, the two key moments that we target to discipline the strength of the local spillover effect in our model are the neighborhoods exposure estimates, for the 25th and for the 75th percentile of the income distribution, from Chetty and Hendren (2018b).[10] To map them into the model, we calculate the standard deviation of the expected future wage of the children of "movers" (i.e. parents who decide to live in a neighborhood different from the one where they grew up) and divide that by the average wage of the parents.[11]

Table 3 shows the calibrated parameter values. In this simple calibration, average children's productivity plays a more relevant role in determining the size of the local spillovers, as $\omega = .13$. In the next subsection we explore how this weight affects the equilibrium allocation.

5.2 Steady State Comparisons

We first explore how the steady state equilibrium changes when we vary the spillover specification. In particular, as we increase ω, the weight on average parental income increases at the expense of the average children's innate productivity.[12]

[10]We focus on their estimates for families moving across counties within the same commuting zone, given that we use the metro area as our geographic unit of analysis.

[11]The exposure effect is equal to $\sqrt{\frac{1}{2}\sum_{i\in\{A,B\},j\in\{A,B\}}(E(w'|i,j)-E(w'))^2}$, where $E(w'|i,j)$ is the expected income of children of movers from neighborhood i to neighborhood j and $E(w')$ is the average expected income of all movers.

[12]For this exercise, we use the parameters in Table3, but we set the value of $\beta = 1$ in order to solve for the steady state for all values of $\omega \in [0,1]$. Larger values of β lead do diverging dynamics that do not reach a steady state when

Table 3: *Parameters*

Parameter	Value	Description
α	0.69	Wage function parameter
β	1.20	Wage function parameter
γ	0.32	Wage function parameter
b	1.65	Wage function parameter
σ_ε	0.34	St. dev. of wage shock
σ_a	1.28	St. dev. of log ability
ρ	0.36	Autocorrelation of log ability
ω	0.13	Spillover function parameter

Panels (a) and (b) in Figure 18 show that as the weight on parental income ω increases, residential segregation by income increases, while residential segregation by innate productivity declines. This is because, as shown in panel (c), when the weight on parental income increases, the spillover gap between the two neighborhoods increases, and the rental rate to live in A with it. This creates a stronger force for sorting by income at the expenses of sorting by innate productivity. The key question is then: why does the spillover gap increases with ω? The reason is that the spillover gap today affects parental income tomorrow, and, in particular, as the spillover gap increases, high-productivity children exposed to neighborhood A will have even higher income when they become parents. This creates an amplification mechanism and further increases the future spillover gap.

As we have seen in Section 3.1, a larger spillover gap, by generating larger segregation by income, also generates higher income inequality, as shown in panel (d). Moreover, Figure 19 also shows that, in turn, this means that higher weight on parental income also means lower intergenerational mobility. Again, this is because as residential segregation by income increases, children with richer parents will have a higher probability of being exposed to the high spillover neighborhood and hence to obtain higher future income. Vice versa, high-productivity children of lower income families will be pushed out of the neighborhood that offers higher opportunities.

5.3 Skill Premium Shock

We now assume that the economy starts in steady state and is hit in 1980 by an unexpected permanent skill premium shock, that we represent as an increase in the parameter η in the wage equation (4). In particular, we increase η by 20% between 1980 and 1990.

combined with higher values of ω.

Figure 18: Varying ω - Segregation and Inequality

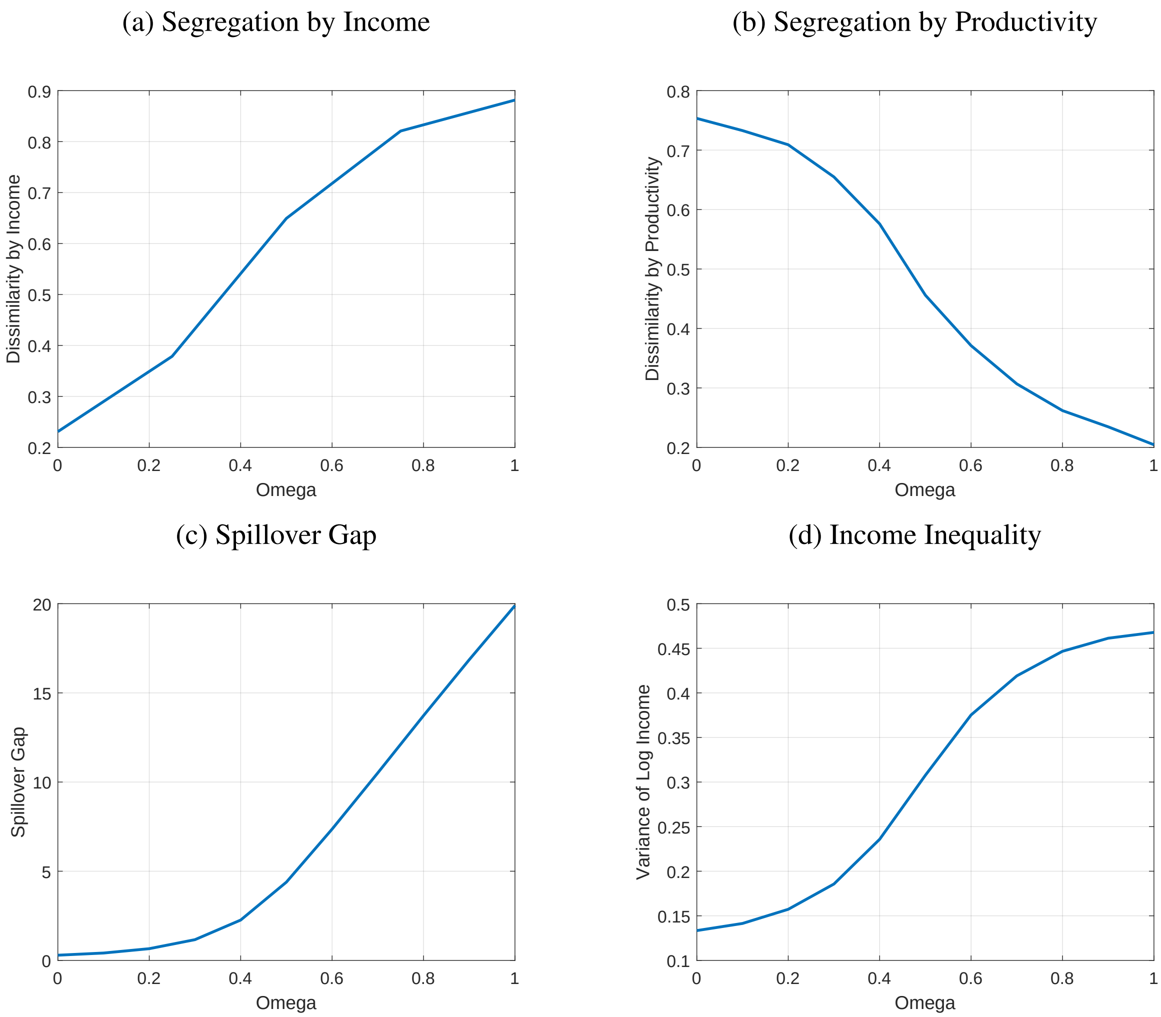

Let us first analyze how this shock affects the spillover gap between the two neighborhoods. Figure 20 shows that the spillover gap increases in response to the shock (solid blue line). As we discussed before, the spillover depends on the distribution of families who live in the neighborhood, both through average parental wage and through average children's productivity. To understand the role of the endogenous change in spillovers, we contrast the baseline model with a version of the model with fixed spillovers, where we keep the spillovers in the two neighborhoods equal to their steady state values and the spillover gap stays constant (red dashed line). In 1990, when the shock is introduced, the spillovers in the two neighborhoods change just because of the re-sorting of parents across them and, as a result, the spillover gap increases, even if only by very little. In particular, neighborhood A will have richer parents moving in and parents with high-productivity children moving out to B. This is because the skill premium shock

Figure 19: Varying ω - Intergenerational Mobility

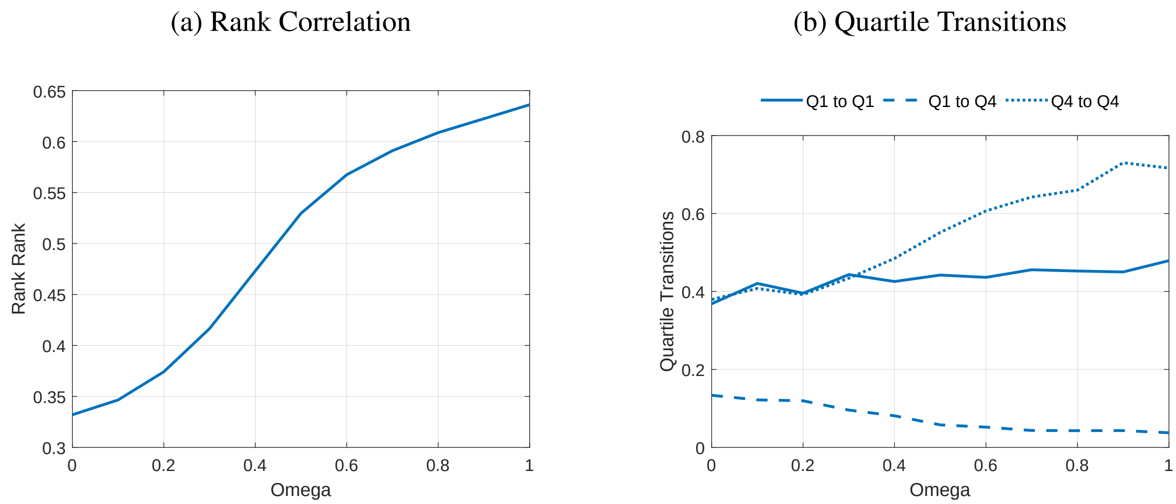

increases all children's future income, but more so for the ones who are exposed to the higher spillover and who have higher innate productivity. This implies that living in neighborhood A is now more attractive and this is reflected into a higher rental rate in A. As for the static model, a higher rental rate in A implies that now only richer parents can afford to live in A, and poorer parents, even if they have high-productivity children, have to move to neighborhood B. This means that the average parental income in neighborhood A increases, increasing the spillover in A and dampening the one in B. However, the average productivity decreases in A and increases in B, having opposite effects on the spillovers. On net, both spillovers increase, but the parental income effect dominates and the spillover gap increases as well. From 2000 on, there is an additional force that changes the spillover, that is, the fact that the parental income distribution is now different from the steady state one, because parents were exposed as children to a higher η. This implies that children who grew up in A have now even higher income relative to children who grew up in B and the spillover gap starts increasing more substantially, as shown in the figure.

Panel (a) in Figure 21 shows that residential segregation by income increases in response to the shock. The figure shows that part of this increase would happen even if the spillovers were fixed, while there is an amplification effect due to the endogeneity of the spillovers. Even if the spillover gap was fixed, as long as it is positive, the η shock would make neighborhood A more attractive, as η is complementary to the spillover. This would raise the rental rate in A and increase residential segregation by income. However, on top of that, the previous figure has shown that the spillover gap increases and this makes neighborhood A even more attractive and raises rental rates in A even more. This further increase residential segregation

Figure 20: Skill Premium Shock: Spillover Gap

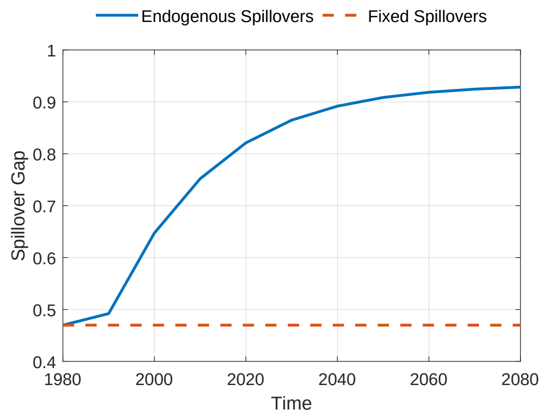

by income above the level that would arise if spillovers were fixed, the more the larger is the increase in the spillover gap. At the same time, panel (b) shows the other side of the medal, that is, that residential segregation by innate productivity declines in response to the shock. This happens exactly for the same reason: as income matters more for sorting, productivity matters less and, although neighborhood A still tend to have higher productivity children, this sorting effect is weaker, as some of the poorer parents have to move to B even if they have high productivity children. Panels (c) and (d) show the effects on income and consumption inequality. Both measures show an increase in inequality in response to the skill premium shock. The figure shows that there are both a direct effect, for given spillovers, and an indirect effect due to the endogenous response of the spillovers, which amplifies the inequality response. The direct effect is simply due to the fact that children exposed to the higher spillover will tend to be the richer parents in the future, but they will be even richer if η is larger. This mechanically increases inequality. This effect is amplified when the spillover gap increases, as shown in the figure.

Finally, Figure 22 shows that intergenerational mobility moves in the opposite direction of residential segregation by income and inequality. In fact, both the rank-rank coefficient and the quartile transition matrices show that intergenerational mobility declines in response to the skill premium shock, and more so because of the endogenous response of the spillover gap.

This exercises show how endogenous spillovers can amplify the effects of a skill premium shock on segregation, inequality, and intergenerational mobility. Fogli et al. (forthcoming) study the aggregate effects of a similar shock in a richer model with also educational investment. Using a battery of counterfactual exercises, they quantify the role of segregation in amplifying the rise in inequality due to a skill premium

Figure 21: Skill Premium Shock Dynamics - Inequality and Segregation

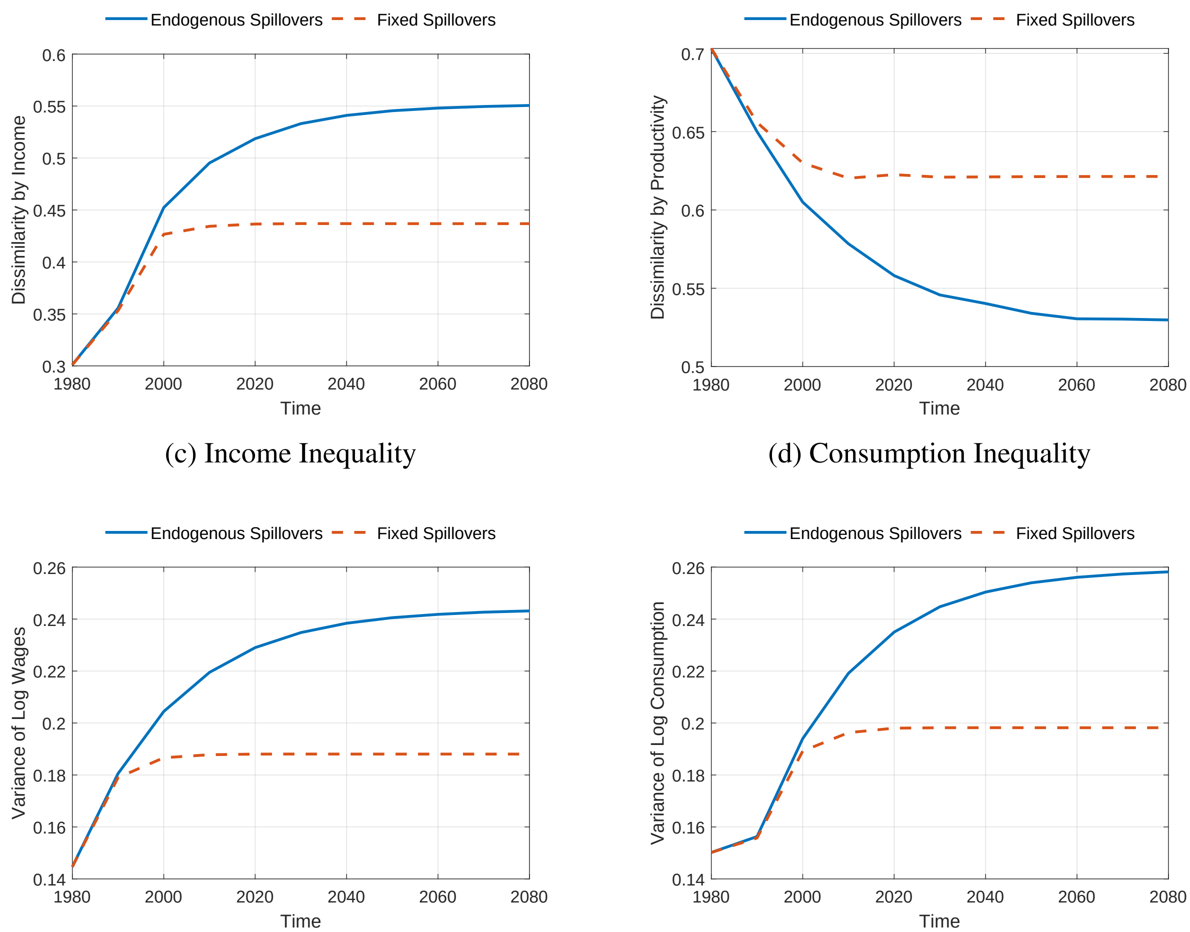

(a) Segregation by Income

(b) Segregation by Productivity

(c) Income Inequality

(d) Consumption Inequality

shock.

In the next section, we review different strands of literature exploring from different angles the effect of local spillovers on inequality, segregation, and intergenerational mobility.

6 Related Literature

The theoretical literature on neighborhood effects and intergenerational mobility traces its roots to foundational work on intergenerational income transmission through parental investment and credit constraints, dating back to Becker and Tomes (1979) and Loury (1981). While these early models emphasize intergenerational transmission within families, they largely abstract from the spatial dimension of inequality.

Figure 22: Skill Premium Shock Dynamics - Intergenerational Mobility

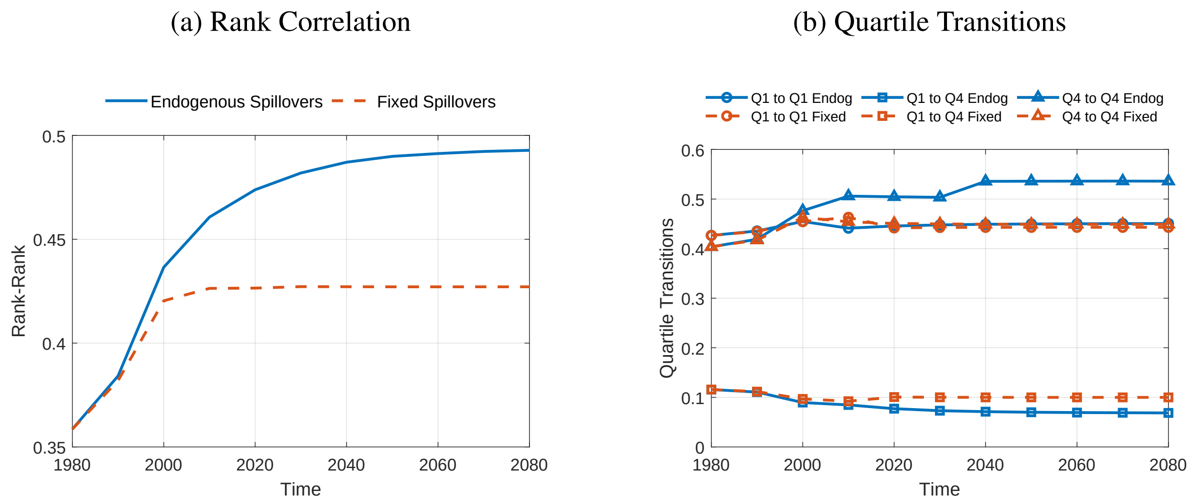

The next wave of theoretical work extends the analysis to neighborhoods, exploring how segregation and local spillovers emerge endogenously and shape long-run inequality. Durlauf (1996b) develops a dynamic model that examines multiple communities, where segregation results from locally financed public schools and social spillovers. His work emphasizes that segregation, driven by both fiscal and social mechanisms, perpetuates inequality over time through strong neighborhood feedback effects. In a related model, Durlauf (1996a) shows how residential stratification can lead to permanent relative income inequality, where inequality persists even in growing economies. In related work, Benabou (1996a) compares segregated and integrated communities, analyzing the impact on growth, inequality, and education. In a similar model, Benabou (1996b) further explores how social externalities and differences in school funding affect income distribution. Fernandez and Rogerson (1996, 1998) also focus on the role of public education financing in residential stratification. In particular, Fernandez and Rogerson (1998), using a calibrated dynamic model, analyze the static and long-term effects of educational finance reforms in the U.S., emphasizing how changes in funding could reduce disparities in both local mobility and broader income distribution.

With the advent of large-scale data, there has been a body of empirical work investigating intergenerational mobility and neighborhood exposure effects. Chetty and Hendren (2018a,b) use US administrative data to estimate how the neighborhood environment influences a child's income and overall well-being in adulthood. They find that moving to a higher-opportunity neighborhood during childhood can substantially improve future income. Chetty et al. (2016) complement these studies with insights from the Moving to Opportunity (MTO) program, which provides experimental evidence on the effects of relocating low-income families to better neighborhoods. They find that housing vouchers enable low-income families living in high-poverty neighborhoods to relocate to low-poverty areas, with significant improvements in children's long-term outcomes, further emphasizing the potential of neighborhood interventions. Moreover, Chetty et al. (2014) examine the geography of intergenerational mobility across the U.S., mapping regional variations and identifying areas where mobility is particularly high or low. Together, these papers contribute to a growing body of evidence showing that neighborhood factors—such as local economic conditions, school quality, and social capital—have a profound and lasting impact on intergenerational mobility, and they highlight the importance of targeted policies to improve opportunities in disadvantaged areas. Other studies, such as Rothstein (2019), have highlighted other dimensions of differences across neighborhoods, such as the labor market structure, job networks, or institutional factors.

This empirical work has spurred new quantitative research on neighborhood externalities. Durlauf and Seshadri (2017) examine how neighborhood-level externalities, which they represent as the average human capital in the community, affect long-term mobility. Their framework replicates the empirical "Great Gatsby Curve," linking higher inequality to lower mobility through the joint dynamics of human capital formation and neighborhood effects, because higher inequality strengthens segregation and amplifies neighborhood-based externalities. A similar relationship between inequality and intergenerational mobility also emerges in Fogli et al. (forthcoming). They develop a general equilibrium overlapping generation model with residential choice and educational investment in the presence of endogenous local spillovers. To discipline the strength of the local spillovers on children's future income they use the empirical estimates of neighborhood exposure from Chetty and Hendren (2018b). They use this calibrated model to quantify the role of residential segregation by income in amplifying the response of income inequality to a skill-premium shock matched to the increase in return to college experienced by the US in the 1980s.

More related work focuses on the aggregate effects of neighborhood-specific policies. Chyn and Daruich (forthcoming) analyze large-scale versions of housing voucher programs, such as the Moving to Opportunity (MTO) experiment, as well as place-based wage subsidies in a model of two neighborhoods, modeling a rich human capital accumulation process and labor market structure. They examine the general equilibrium trade-offs between these two types of policies, finding that vouchers generally produce larger welfare gains, though place-based wage subsidies may dominate in settings with tight housing supply. In a similar spirit, using a three-neighborhood model with endogenous spillovers, Fogli et al. (2025b) compare the aggregate implications of large-scale housing voucher programs with two types of neighborhood-specific policies: a place-based transfer and a place-based investment program, such as improvements in public school funding. Their results show that housing vouchers generates substantial income gains for children of recipient families, but scaling them up dampens these gains and induces large welfare losses for non-recipients. Place-based transfers deliver larger average welfare gains but are less effective in reducing inequality and segregation. Finally, place-based investment programs are able to raise average welfare while simultaneously reducing inequality, segregation, and improving intergenerational mobility over time. Davis et al. (2021) also study housing vouchers and they estimate a dynamic location-choice model. They emphasize the trade-off of restricting criteria to use the voucher to move to lower poverty neighborhoods: on the one hand, households who use the voucher end up in better locations, but on the other hand, less families end up taking up the voucher. In addition, they also show that general equilibrium effects, such as raising rents, dampen the net benefit effects of the policy.

Another group of recent papers use similar quantitative frameworks to focus on the effects of specific public school policies. Zheng and Graham (2022) develop a dynamic overlapping-generations model with neighborhood choice and endogenous school quality financed by local property taxes. They show that the tight link between housing markets and school funding amplifies spatial inequality: richer districts collect more revenues and provide better schools, limiting access for low-income families and reducing mobility. Their counterfactuals indicate that housing vouchers or redistribution of property tax revenues across districts can substantially increase long-run intergenerational mobility, although the full benefits take several generations to materialize. Eckert and Kleineberg (2024) embed labor markets with heterogenous workers in a similar spatial model. They show that equalizing school funding improves equality of educational outcomes, but it creates a mismatch between supply and demand of educated workers in different locations. Moreover, they show that subsidy and housing policies are less effective in improving equality in educational outcomes, but do not reduce output. Agostinelli et al. (2024) study residential sorting and school choice in a spatial equilibrium model of educational access. They find that both school-choice expansions and housing voucher programs generate limited aggregate welfare improvements once general equilibrium adjustments in housing and school composition are considered.

A related strand of the literature focuses on the interaction between income and racial residential segregation and their connection with intergenerational mobility. Mazumder (2021) and Davis and Mazumder (2017) document large and persistent racial gaps in intergenerational mobility in the United States, showing that black families experience substantially lower rates of upward mobility than white families even after controlling for parental income, and linking this pattern to rising inequality. Bayer et al. (2025) provide causal evidence that racial sorting remains a persistent force shaping household residential decisions in contemporary U.S. housing markets, as both black and white homeowners are more likely to move when receiving a new next-door neighbor of a different race. Bayer et al. (2021) further document that black and white households with similar incomes live in different neighborhoods, with large disparities in neighborhood resources. This is true at all income levels and across nearly all metropolitan areas. Aliprantis et al. (2024) provide complementary evidence, showing that even after accounting for income and wealth, black households live in substantially lower socioeconomic-status neighborhoods than white households. They attribute this pattern primarily to racial homophily, combined with the limited supply of high socioeconomic-status black neighborhoods. Davis et al. (2024) use a structural approach to estimate a dynamic neighborhood-choice model, and they also find that households exhibit preferences to live in neighborhoods where the vast majority of households are of the same race. Overall, these findings highlight that racial disparities in opportunity remain substantial.

Recent work has developed theoretical frameworks to examine drivers and implications of these empirical patterns. Aliprantis and Carroll (2018) develop a two-neighborhood model featuring residential choice and local human capital externalities. They calibrate their model using data from Chicago in 1960 to match the city's initial racial composition, and then use the model to analyze how removing legal forms of racial discrimination would affect segregation and intergenerational outcomes. Gregory et al. (2024) develop an overlapping-generations spatial equilibrium model with neighborhood spillovers similar to the one in this chapter, to analyze how racial segregation shapes the black–white gap in college attainment. In their framework, race affects outcomes through three channels: a wage gap, preferences over neighborhood racial composition, and barriers to mobility. They show that their model is able to rationalize the observed segregation and racial differences in educational attainment for St. Louis data. They also show that their model may generate multiple equilibria. Fogli et al. (2025a) propose a general equilibrium overlapping-generations model with local spillovers, where the racial bias arises endogenously. In particular, black and white agents are identical except for the initial income distribution that gives rise to differential access to opportunities. The key ingredient is a behavioral assumption: agents do not take into full consideration the fact that income depends on the neighborhood where agents grow up. Also their model generates multiple equilibria, opening the way for policy interventions that could shift the economy from a highly segregated equilibrium toward a more equitable one, and reducing racial economic gaps. Manysheva et al. (2025) provide related evidence from a different institutional context, developing a spatial model calibrated to post-Apartheid South Africa. They show that, despite the formal end of legal segregation, inherited spatial disadvantages, such as high commuting costs, unequal school quality, and limited access to credit, continue to reinforce inequality and slow convergence across racial groups, underscoring the persistent role of geography in shaping intergenerational outcomes.

Another related line of research examines how parenting style decisions influence children's outcomes and intergenerational mobility. Doepke and Zilibotti (2017) develop a model in which parenting styles adapt endogenously to socioeconomic conditions, generating feedback effects between inequality, parental behavior, and the transmission of opportunities. Building on this framework, Agostinelli et al. (forthcoming) analyze how parents shape their children's peer environments, showing that parenting style and peer choice interact to reinforce educational and income inequality. Their model closely parallels residential choice frameworks, as both decisions determine children's exposure to opportunity. In subsequent work, Agostinelli et al. (2022) study the unequal consequences of the shift to remote learning during the COVID-19 pandemic, finding that differential parental responses to school closures amplify preexisting educational and income gaps. Finally, Bellue (2025) extends this literature by linking parenting and neighborhood dynamics, showing how parental investments and local social interactions jointly influence children's long-term outcomes and the spatial transmission of inequality.

The modeling of residential choices also connects to a broader literature in urban and development economics that studies spatial sorting and neighborhood formation. Ferreira et al. (2017) develop a framework similar in spirit to neighborhood models of mobility to analyze the emergence and persistence of urban slums, calibrated to Brazilian data. Related work in urban economics emphasizes how spatial sorting arises endogenously from differences in local amenities. Guerrieri et al. (2013) highlight how neighborhood upgrading and gentrification can be driven by the endogenous evolution of amenities, while Couture et al. (2023) use a spatial equilibrium model with endogenous amenities and non-homothetic preferences to study within-city re-sorting after the 1990s. Moreover, Bilal and Rossi-Hansberg (2021) show that residential location choices can be viewed as a form of asset investment, linking housing markets, spatial inequality, and wealth accumulation.

7 Concluding Remarks

In this chapter, we have used a simple general equilibrium model of residential choice in the presence of local spillovers to show the tight link among residential segregation by income, inequality, and intergenerational mobility. The model predicts that cities with higher degrees of residential segregation to experience lower levels of intergenerational mobility, which is what we have documented for the US in the empirical section. Moreover, the model shows that this link is even stronger if we consider the fact that local spillovers are typically endogenous and depend on the sorting of families across neighborhoods.

We also use the static version of the model to show that a utilitarian planner would like less residential segregation by income in favor of more residential segregation by innate productivity and, hence, more intergenerational mobility. This would increase the efficiency of the economy, even if at the expenses of higher income inequality. However, the planner can then redistribute across families. We also show that a simple transfer policy that targets low-income families, without neighborhood-specific components, can improve upon the decentralized equilibrium and achieve higher welfare and higher intergenerational mobility. There is a growing body of literature focusing on neighborhood-specific policies and it would be interesting to explore their effects in this context.

Finally, we show how a dynamic version of the model responds to a skill premium shock and how the effects on segregation and inequality are amplified by endogenous evolution of the spillovers.

References

Agostinelli, Francesco, Matthias Doepke, Giuseppe Sorrenti, and Fabrizio Zilibotti, "When the great equalizer shuts down: Schools, peers, and parents in pandemic times," *Journal of Public Economics*, 2022, *206*, 104574.

_ , _ , _ , and _ , "It takes a village: the economics of parenting with neighborhood and peer effects," *Journal of Political Economy*, forthcoming.

_ , Paolo Martellini, and Margaux Luflade, "On the Spatial Determinants of Educational Access," *mimeo, NYU*, 2024.

Aliprantis, Dionissi and Daniel R Carroll, "Neighborhood dynamics and the distribution of opportunity," *Quantitative Economics*, 2018, *9* (1), 247–303.

_ , Daniel R. Carroll, and Eric R. Young, "What Explains Neighborhood Sorting by Income and Race?," *Journal of Urban Economics*, 2024, *141*, 103508.

Bayer, Patrick, Kerwin Kofi Charles, and Joon Yup Park, "Separate and Unequal: Race and the Geography of the American Housing Market," *mimeo, Duke University*, 2021.

_ , Marcus D. Casey, W. Ben McCartney, John Orellana-Li, and Calvin S. Zhang, "Distinguishing Causes of Neighborhood Racial Change: A Nearest-Neighbor Design," *American Economic Review*, 2025, *115* (11), 3999–4039.

Becker, Gary S and Nigel Tomes, "An equilibrium theory of the distribution of income and intergenerational mobility," *Journal of political Economy*, 1979, *87* (6), 1153–1189.

Bellue, Suzanne, "Residential and Social Mobility: A Quantitative Analysis of Parental Decisions with Social Learning," *mimeo, CREST-ENSAE-IP Paris*, 2025.

Benabou, Roland, “Equity and efficiency in human capital investment: the local connection,” *The Review of Economic Studies*, 1996, *63* (2), 237–264.

— , “Heterogeneity, stratification, and growth: macroeconomic implications of community structure and school finance,” *The American Economic Review*, 1996, pp. 584–609.

Bilal, Adrien and Esteban Rossi-Hansberg, “Location as an Asset,” *Econometrica*, 2021, *89* (5), 2459–2495.

Card, David and Laura Giuliano, “Can tracking raise the test scores of high-ability minority students?,” *American Economic Review*, 2016, *106* (10), 2783–2816.

Chetty, Raj and Nathaniel Hendren, “The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects,” *The Quarterly Journal of Economics*, 2018, *133* (3), 1107–1162.

— **and** — , “The impacts of neighborhoods on intergenerational mobility II: County-level estimates,” *The Quarterly Journal of Economics*, 2018, *133* (3), 1163–1228.

— **,** — **, and Lawrence F Katz**, “The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment,” *American Economic Review*, 2016, *106* (4), 855–902.

— **,** — **, Patrick Kline, and Emmanuel Saez**, “Where is the land of Opportunity? The Geography of Intergenerational Mobility in the United States *,” *The Quarterly Journal of Economics*, 09 2014, *129* (4), 1553–1623.

Chyn, Eric and Diego Daruich, “An Equilibrium Analysis of the Effects of Neighborhood-based Interventions on Children,” *American Economic Review*, forthcoming.

Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst, “Income Growth and the Distributional Effects of Urban Spatial Sorting,” *The Review of Economic Studies*, 04 2023, *91* (2), 858–898.

Davis, Morris A., Jesse Gregory, and Daniel A. Hartley, “Preferences over the Racial Composition of Neighborhoods: Estimates and Implications,” *mimeo, Rutgers University*, 2024.

Davis, Morris A, Jesse Gregory, Daniel A Hartley, and Kegon TK Tan, “Neighborhood effects and housing vouchers,” *Quantitative Economics*, 2021, *12* (4), 1307–1346.

Doepke, Matthias and Fabrizio Zilibotti, “Parenting with style: Altruism and paternalism in intergenerational preference transmission,” *Econometrica*, 2017, *85* (5), 1331–1371.

Durlauf, Steven N, “Neighborhood feedbacks, endogenous stratification, and income inequality,” *in Dynamic Disequilibrium Modelling: Proceedings of the Ninth International Symposium on Economic Theory and Econometrics, W. Barnett, G. Gandolfo, and C. Hillinger, eds.*, 1996.

— , “A theory of persistent income inequality,” *Journal of Economic growth*, 1996, *1* (1), 75–93.

— **and Ananth Seshadri**, “Understanding the Great Gatsby Curve,” in “NBER Macroeconomics Annual 2017, volume 32,” University of Chicago Press, 2017.

Eckert, Fabian and Tatjana Kleineberg, “The Geography of Opportunity: Education, Work, and Intergenerational Mobility Across US Counties,” *mimeo, UC San Diego*, 2024.

Fernandez, Raquel and Richard Rogerson, “Income distribution, communities, and the quality of public education,” *The Quarterly Journal of Economics*, 1996, *111* (1), 135–164.

— **and** — , "Public education and income distribution: A dynamic quantitative evaluation of education-finance reform," *American Economic Review*, 1998, pp. 813–833.

Ferreira, Pedro C, Alexander Monge-Naranjo, and Luciene Torres de Mello Pereira, "Of Cities and Slums," *Federal Reserve Bank of St Louis Working Paper 2016-022A*, 2017.

Fogli, Alessandra, Martin Garcia-Vazquez, Veronica Guerrieri, and Marta Prato, "Neighborhood segregation and endogenous racial bias," 2025.

— **, Veronica Guerrieri, Mark Ponder, and Marta Prato**, "The Macroeconomic Effect of Neighborhood Policies: a Dynamic Analysis," 2025.

— **,** — **,** — **, and** — , "The End of the American Dream? Inequality and Segregation in US Cities," *Journal of Political Economy*, forthcoming.

Gregory, Victoria, Julian Kozlowski, and Hannah Rubinton, "The Impact of Racial Segregation on College Attainment in Spatial Equilibrium," *FRB St. Louis Working Paper*, 2024, (2022-36).

Guerrieri, Veronica, Daniel Hartley, and Erik Hurst, "Endogenous gentrification and housing price dynamics," *Journal of Public Economics*, 2013, *100*, 45–60.

Imberman, Scott A, Adriana D Kugler, and Bruce I Sacerdote, "Katrina's children: Evidence on the structure of peer effects from hurricane evacuees," *American Economic Review*, 2012, *102* (5), 2048–82.

Lavy, Victor, M Daniele Paserman, and Analia Schlosser, "Inside the black box of ability peer effects: Evidence from variation in the proportion of low achievers in the classroom," *The Economic Journal*, 2012, *122* (559), 208–237.

Loury, Glenn C, "Intergenerational transfers and the distribution of earnings," *Econometrica: Journal of the Econometric Society*, 1981, pp. 843–867.

M. V. Davis Jonathan and Bhashkar Mazumder, "The Decline in Intergenerational Mobility After 1980," Working Paper WP-2017-05, Federal Reserve Bank of Chicago 2017.

Manysheva, Kristina, Marti Mestieri, and Johanna Schauer, "End of Apartheid, Not of Inequality: The Slow Transition in a Segregated Economy," *CEPR Discussion Paper No. 20472. CEPR Press, Paris & London*, 2025.

Mazumder, Bhashkar, "Race and Intergenerational Mobility in the United States," in Jean Kimmel, ed., *Intergenerational Mobility: How Gender, Race, and Family Structure Affect Adult Outcomes*, W. E. Upjohn Institute for Employment Research, 2021, pp. 171–191. Published April 29 2022; Kalamazoo, MI.

Rothstein, Jesse, "Inequality of educational opportunity? Schools as mediators of the intergenerational transmission of income," *Journal of Labor Economics*, 2019, *37* (S1), S85–S123.

Sacerdote, Bruce, "Peer effects with random assignment: Results for Dartmouth roommates," *The Quarterly journal of economics*, 2001, *116* (2), 681–704.

Zheng, Angela and James Graham, "Public Education Inequality and Intergenerational Mobility," *American Economic Journal: Macroeconomics*, 2022, *14* (3), 250–82.