

The End of the American Dream? Inequality and Segregation in US cities *

Alessandra Fogli
Minneapolis Fed

Veronica Guerrieri
University of Chicago

July 2018

Abstract

Since the '80s the US has experienced not only a steady increase in income inequality, but also a contemporaneous increase in residential segregation by income. Using US Census data, we document a positive correlation between income inequality and residential segregation between 1980 and 2010, both across time and across space, at the MSA level. We then develop a general equilibrium overlapping generations model where parents choose the neighborhood where to raise their children and invest in their children's human capital. In the model, segregation and inequality amplify each other because of a local spillover that affects the returns to education. We calibrate the model to 1980 using Census data and the micro estimates of the local spillover effect derived by Chetty and Hendren (2018b). We then hit the economy with a skill premium shock and show that 20% of the increase in inequality in the short run, and 29% in the long run can be attributed to the feedback effect of the local spillover.

*Email addresses: afogli00@gmail.com; vguerrie@chicagobooth.edu. For helpful comments, we are grateful to Roland Benabou, Jarda Borovicka, Steven Durlauf, Cecile Gaubert, Mike Golosov, Luigi Guiso, Erik Hurst, Francesco Lippi, Guido Lorenzoni, Guido Menzio, Alexander Monge-Naranjo, Fabrizio Perri, Richard Rogerson, numerous seminar participants, and, in particular, to Kjetil Storesletten for the useful discussion. For outstanding research assistance, we thank Gustavo Gonzalez, Hyunju Lee, Qi Li, Emily Moschini, Luis Simon, and, in particular, Yu-Ting Chiang, Mark Ponder, and Francisca Sara-Zaror.

1 Introduction

It is a well documented fact that over the last 40 years, the US has experienced a steady increase in income inequality. At the same time there has been a substantial increase in residential segregation by income and education. What is the link between inequality and residential segregation? In particular, has residential segregation contributed to amplify the increase in income inequality caused by underlying shocks, such as skill-biased technical change? In this paper, we build a model of human capital accumulation with local spillovers and residential choice that can be used to address these questions.

There has been a large theoretical literature in the '90s focusing on the relation between inequality and local externalities, starting from the seminal work by Benabou (1996a,b), Durlauf (1996a,b), Fernandez and Rogerson (1996, 1997, 1998). More recently, administrative data have been used to propose direct estimates of neighborhood spillover effects. In particular Chetty, Hendren and Katz (2016), and Chetty and Hendren (2018a,b) have shown that there are substantial income effects coming from the childhood exposure to better neighborhood. We bridge these two strands of literature, by proposing a general equilibrium model calibrated using the micro estimates from Chetty and Hendren (2018b) to understand the contribution of these types of externalities to the recent rise in inequality.

We first document a strong correlation between income inequality and residential segregation by income at the MSA level, both across time and across space. We use US Census tract data on family income between 1980 and 2010 to construct measures of inequality and residential segregation at the MSA level. To measure inequality, we use the Gini coefficient. To measure segregation, we use the dissimilarity index, which is a measure of how uneven is the distribution of two groups across geographical areas. In particular, we divide the population in two income groups, rich and poor, using the 80th income percentile, and compute the dissimilarity index across census tracts belonging to the same MSA. We also check the robustness of our main findings with alternative measures of income inequality, such as the 90/10 ratio, and of income segregation, such as the dissimilarity index calculated with different percentiles and the H^R index.¹ Using these data, we show that 1) average inequality and residential segregation have increased steadily since 1980;

¹The H^R index has been proposed by Reardon and Firebaugh (2002) and Reardon and Bischoff (2011), and also used by Chetty et al. (2014).

2) inequality and residential segregation in 1980 are correlated across MSA; 3) the change in inequality and residential segregation between 1980 and 2000 are correlated across MSA.

We also look at intergenerational mobility, by using a restricted-access geocoded version of the National Longitudinal Survey of Youth (NLSY79). We show that intergenerational mobility is lower in metro areas that display higher levels of residential segregation. In particular, we split the metro areas in two groups: above and below the median dissimilarity index in 1980. We then show the share of people who stay in the lowest quartile of the wage distribution conditional on their parents being in the same quartile goes from 40.7% in the less segregated metros to 47.3% in the more segregated ones, testifying a lower degree of intergenerational mobility in the more segregated metro areas. Similarly, the share of people who stay in the highest quartile of the wage distribution conditional on their parents being in the same quartile goes from 41.4% to 46.8%.

We then build a general equilibrium overlapping generation model with human capital accumulation and residential choice. The model generates a feedback effect between income inequality and residential segregation, that amplifies the response of inequality to underlying shocks. Agents live for two periods: first they are young and go to school and then they are old and become parents. There are two neighborhoods and parents choose both the neighborhood where they raise their children and the investment in their children's human capital. The key ingredient of the model is a local spillover: investment in human capital yields higher returns in neighborhoods with higher average level of human capital. In turns, this implies that the residential choice itself affects human capital accumulation. The local spillover generates sorting in equilibrium: richer parents with more talented children will pay higher rents to live in the neighborhood with higher average human capital. The spillover effect in the model is meant to capture a variety of mechanisms: differences in the quality of public schools, peer effect, social norms, learning from neighbors' experience and so forth. For our purposes here, it does not matter which of these effects are causing the spillover.

Next, we calibrate the steady state of the model to the US economy in 1980. In particular, we target the average level of income inequality and residential segregation by income, using the Census data just described. We also target the skill premium, the share of collage graduates, the ratio of collage graduates in the two neighborhoods, and the rank-rank correlation between children's and parents' income. A key target of our calibration is the effect of the local spillover. This is where we use the micro estimates obtained with the quasi-experiment in Chetty and Hendren

(2018b).

We then perform our main exercise. Assuming that the original shock to inequality comes purely from skill-biased technical change, we want to understand the role of the local spillover in amplifying the increase in inequality over time. In particular, we study the effects of an unexpected, one-time shock to the skill premium on inequality, segregation, and intergenerational mobility over time. If the skill premium increases, the relative spillover in the rich neighborhood increases, generating even more inequality. Despite the parsimony of the model, the exercise generates patterns for inequality and segregation that resemble the data. We then run some counterfactual exercises to understand the different forces at play. The main exercise is a counterfactual where we look at the response of the economy to the same shock, keeping the level of the spillover in the two neighborhoods at the steady state values. The exercise shows that the spillover feedback effect can contribute to 20% of the increase in inequality in the short run and to 29% in the long run. Moreover, it contributes to 18% and 15% of the increase in segregation in the short and long run respectively, and to 12% and 18% of the decrease in intergenerational mobility in the short and long run. The increase in the relative spillover of the rich neighborhood is in part due to the general equilibrium increase in the rental price of housing in the rich neighborhood. We show that the general equilibrium effect accounts for roughly 30% of the increase in the relative spillover.

Related Literature.

Our model builds on a large class of models with multiple communities, local spillovers, and endogenous residential choice, studying the effects of stratification (residential segregation in our language) on income distribution, going back to the fundamental work by Becker and Tomes (1979) and Loury (1981). Among the seminal papers in this literature, Benabou (1993) explores a steady state model where local complementarities in human capital investment, or peer effects, generate occupational segregation and studies its efficiency properties.² Durlauf (1996b) proposes a related dynamic model with multiple communities, where segregation is driven by both locally financed public schools and local social spillovers. The paper shows that economic stratification together with strong neighborhood feedback effects generate persistent inequality.³

²De Bartolome (1990) also studies efficiency properties of a similar type of model where communities stratification is driven by peer effects in education. In similar papers, the local social externalities take the form of role models (Streufert (2000)), or referrals by neighborhoods (see Montgomery (1991a,b)).

³Durlauf (1996a) uses a related model to study how it can generate permanent relative income inequality (opposed

Benabou (1996a) embeds growth with complementary skills in production in a similar model, where local spillovers are due both to social externalities (as peer effects) and locally financed public school. The paper analyzes the trade-off coming from the fact that stratification helps growth in the short run due to the complementarities in skills, while integration helps growth in the longer run, as generates less inequality, and hence heterogeneity in skills, over time. It also studies how alternative systems of education financing affect the economy. Fernandez and Rogerson (1996) also study the impact of a number of reforms on public education financing using a related model, with no growth, where residential stratification is purely driven by locally financed public education.⁴ Fernandez and Rogerson (1998) calibrate to US data a dynamic version of a similar model to analyze the static and dynamic effects of public school financing reforms. This paper is, to the best of our knowledge, the first to calibrate a model in this literature. Benabou (1996b) also studies the effects of public-school financing reforms in a similar model, but he allows for non-fiscal channels of local spillovers, like peers, role models, norms, networks, and so forth and show that disentangling between financial and social local spillover is important for assessing these types of policies.

Our model builds on the same idea of this class of papers that stratification, due to a local spillover, generates more inequality over time. We focus on a model that can be calibrated and brought to the data, while, most of the papers discussed, with the notable exception of Fernandez and Rogerson (1998), focus on the qualitative implications of the models. In that spirit, most of them analyze the two extreme scenarios of full stratification and full integration. Given our quantitative direction, we need to make the model more special under some dimensions in order to handle a continuous measure of segregation. In order to discipline the model with data on education, we also introduce an endogenous educational choice, that is absent in the previous papers. Moreover, differently from the literature, we model the local spillover as a black box, that can be interpreted as driven either by a financial or a social channel. While for normative questions that have been explored in the literature the specification of the spillover is clearly important, for positive questions like the ones analyzed in our paper, it is less so. This is why we prefer to leave the framework more flexible to possibly incorporate different types of local spillover effects.

The most related paper to our work is the contemporaneous work of Durlauf and Seshadri (2017).

to absolute low-income or poverty traps) in an economy where everybody's income is growing.

⁴In a similar framework, Fernandez and Rogerson (1997) study the effect of community zoning regulation on allocations and welfare.

They also build on this class of models to explore the idea that larger income inequality is associated to lower intergenerational mobility, the "Gatsby curve". The model in the paper is close to our model in many dimensions, although the calibration strategy and the main exercise are different and complement well each other.

More recently, Zheng (2017) also studies a model similar to ours where a local spillover generates residential sorting and calibrates it to the US data. The main objective of her paper is to study the effects of different public school allocation mechanisms, while we focus on quantifying the role of the local spillover in amplifying the changes in inequality over time. To this end, she models the local spillover explicitly in terms of public school quality and peer effect.

Another recent related paper is Ferreira, Monge-Naranjo and Torres de Mello Pereira (2017), who use a model close to ours to think about the emergence and persistence of urban slums and calibrate it to Brazilian data. They propose a model with overlapping generation of individuals with different skills, where local spillovers take the form of human capital externalities. They embed growth in the model to think about structural transformation together with urban evolution. They use the model to ask what are the effects of slums on human capital accumulation, structural transformation, urban development and mobility.

Besides the vast literature on city segregation, there are also papers that investigate the consequences of high levels of segregation in a cross section of countries. Alesina and Zhuravskaya (2011), using a measure of segregation similar to ours, show that countries where different linguistic and ethnic groups are more segregated across regions are characterized by significantly lower government quality.

Our work is also related to the literature investigating the evolution of race-based segregation in US cities and its consequences on individual outcomes. The seminal paper of Cutler and Glaeser (1997) shows that blacks living in more segregated metros have significantly worse outcomes than blacks living in less segregated cities. Given the correlation between income and race, these findings are relevant for our analysis. Interestingly, however, Cutler, Glaeser and Vigdor (1999) show that the American ghetto, rapidly expanding between 1890 and 1970 as blacks migrated to the cities, eventually started declining. Income-based segregation has progressively replaced race-based segregation in US cities.

The paper is organized as follows. In Section 2, we document the positive correlation between

inequality and segregation across space and time and between intergenerational mobility and segregation. Section 3 describes the model. In Section 4 we describe our calibration strategy and we present our main quantitative results. Section 5 proposes an extended version of the model. Section 6 concludes.

2 Empirical Evidence

Over the last forty years US cities have experienced a profound transformation in their socio-economic structure: poor and rich families have become increasingly spatially separated over time. As noted by Massey, Rothwell and Domina (2009), this is a new phenomenon in US cities, historically predominantly segregated on the basis of race.⁵ During the last third of the twentieth century, the United States moved toward a new regime of residential segregation characterized by decreasing racial-ethnic segregation and rising income segregation. Such a shift took place at the same time of a dramatic increase in income inequality. In this section we document the magnitude of the phenomenon. These measures will be used for our calibration exercise in the next section.

2.1 Segregation and Inequality over Time

The term segregation refers to the spatial distribution of different groups of the population in a geographic unit across geographic subunits. The groups can be defined according to different categories, such as race, education and income, and segregation can be measured at different geographic levels, such as state, county or metro. We are interested in measuring the residential segregation by income within US cities. To this end, we divide the populations in two groups, rich and poor, and we use the metro area as our geographic unit and the census tract (according to the definition of the Census 2000) as our subunit. We use census tract tabulations of family income and define rich all families above the p -th percentile of the metro income distribution, using the 80th percentile as our benchmark definition.

⁵Massey, Rothwell and Domina (2009) documents that from 1900 to 1970s what changed over time was the level at which racial segregation occurred, with the locus of racial separation shifting from the macro level (states and counties) to the micro level (municipalities and neighborhoods).

In general, segregation is a multidimensional concept, capturing different aspects of the spatial distribution of the population.⁶ In this paper we follow Massey, Rothwell and Domina (2009) and focus on the dimension known as evenness, that is the degree to which two or more groups are distributed evenly over a set of geographic units. Evenness is most commonly measured by the index of dissimilarity, which varies from 0 to 1, with the former value indicating perfect evenness and the latter maximum separation. A group is evenly distributed when each geographic subunit has the same percentage of group members as the population in the geographic unit. In our case, the dissimilarity index measures the percentage of rich and poor that would have to change residence for each census tract to have the same percentage of that group as the whole metro area.⁷

The dissimilarity index for metro j is then calculated as follows:

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

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


We calculate the dissimilarity index for a sample of approximately 380 metros over the period 1980-2010. We then aggregate the metro level dissimilarity indexes into a national one using metro level population weights. We plot the resulting measure of segregation at the national level in Figure 1. The graph shows that the distribution of income has become progressively more uneven across census tracts over time. If in 1980, roughly 30% of the population had to change residence to achieve an even distribution across census tracts in the average US city, in 2010 the population that needed to change residence increased to roughly 36%. The increase was especially large between 1980 and 1990 and again between 2000 and 2010.

Using the same data on family income at the tract level that we used to calculate the dissimilarity index, we also compute the Gini coefficient at the metro level and similarly aggregate the metro level statistics at the national level using metro population weights. Income data at the

⁶Massey and Denton (1988) grouped the measures into 5 key dimensions: evenness, exposure, concentration, centralization, and clustering.

⁷US Census Bureau, Appendix B: Measures of Residential Segregation". This is also the measure followed by Domina (2006). The two measures are equivalent when the two groups span the whole population, which is our case.

Figure 1: Inequality and Segregation over Time



census tract level are reported in bins and are top coded. Top-coded income data are a significant concern when calculating inequality measures. Some papers dealing with individual level income data, such as Armour, Burkhauser and Larrimore (2016), have addressed this issue by using to estimate a Pareto distribution for the top income bracket. However, this methodology is not feasible when dealing with binned, rather than continuous, income data. The methodology mostly used for binned data has been the one proposed by Nielsen and Alderson (1997), who use the Pareto coefficient from the last full income bracket to estimate the conditional mean of the top-coded bracket.⁸ However, such procedure does not exploit the fact that the Census reports the precise empirical average income by census tract. This information can be useful to improve the estimation of the top-coded distribution. We therefore follow a recent methodology proposed by von Hippel, Hunter and Drown (2017) who estimate the CDF of the income distribution non-parametrically and then use the empirical mean to fit the top-coded distribution.⁹ We plot the resulting estimate of the Gini coefficient in Figure 1 together with the dissimilarity index. Both measures show a significant increase over time, especially between 1980 and 1990, with the Gini coefficient rising from roughly .36 to roughly .42 over the entire period. The increase in spatial

⁸See for instance, Reardon and Bischoff (2011).

⁹For details see Appendix B.

segregation by income across neighborhoods happened at the same time of a steady increase in income inequality.

We now check the robustness of these patterns, using alternative measures of income segregation and income inequality.

Figure 2: Dissimilarity Index: different cutoffs




Figure 2 plots the dissimilarity index calculated using different percentiles to define the income groups. The red dashed line shows our benchmark dissimilarity index, while the solid blue line and the dotted green line show the the dissimilarity index constructed using the 10th and the 50th percentiles respectively. The figure shows that the dissimilarity index shifts up as the cut-off percentile decreases, suggesting that groups progressively more homogenous according to income are also characterized by higher levels of segregation. However, regardless of the level, all measures show an increasing trend over time. From now on, when we refer to the dissimilarity index, we refer to the average dissimilarity index across metro areas, population-weighted, that uses the 80th percentile as cut-off to define the rich and the poor.

Figure 3 shows our benchmark dissimilarity index against the H^R index, that is another common measure of income segregation proposed by Reardon and Firebaugh (2002) and Reardon and Bischoff (2011). To construct the H^R index in a given metro, first we define the information

theory index (or Theil index) $H(p)$ that measures the segregation of rich and poor individual across census tracts, where as before the rich are the individuals above the p -th percentile of the family income in the metro area and the poor are the others:

$$H(p) = 1 - \frac{1}{E(p)} \sum_{j=1}^J \frac{t_j}{T} E_j(p),$$

where $E(p) = -[p \log(p) + (1 - p) \log(1 - p)]$ is the entropy, which is another measure of evenness, at the metro level using p as the cutoff percentile, $E_j(p)$ is the same entropy index calculated for the census tract j , and t_j/T is the share of the metro population in census tract j . The corresponding income segregation index is

$$H^R = \frac{1}{\int_0^1 E(p) dp} \int_0^1 E(p) H(p) dp = 2 \int_0^1 E(p) H(p) dp.$$

Relative to the dissimilarity, this index has the advantage that does not rely on a single cut-off to define rich and poor.¹⁰ Figure 3 plots the average of the H^R indices at the metro level, weighted by population, and shows that it has also increased over the last four decades.¹¹

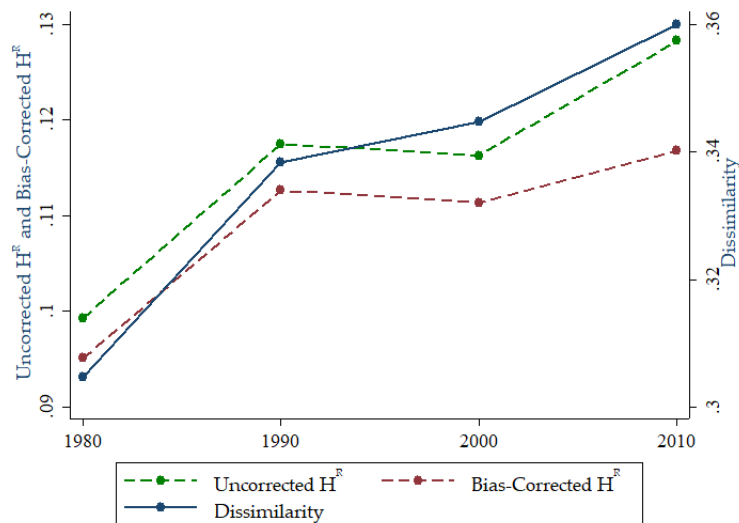
Figure 3 also shows another variant of the H^R index, the "bias-corrected H^R ", which has been calculated following the recent work by Reardon, Bischoff, Owens and Townsend (2018). They develop a methodology to correct the potential sample bias in the H^R index coming from the small number of observations that tend to overestimate the extent of segregation.¹² The "bias-corrected H^R " is systematically lower than its uncorrected counterpart and the difference is larger, as expected, in the last part of the sample when data from ACS are used. Nevertheless, segregation appears to have increased also after correcting for the potential bias.

¹⁰In practice, since income is available only in coarsened form (income data are reported into 16 income categories in U.S. census and ACS data), we follow Reardon and Bischoff (2011) and estimate $H(p)$ by first computing H at the set of finite values of that correspond to the percentiles of the thresholds used to coarsen the data; we then fit a polynomial function through the resulting points; and then using the fitted polynomial as an estimate of $H(p)$.

¹¹The increase in residential income segregation over time is a robust finding. Several sociologists have documented this fact using different measures of segregation. In particular, Jargowsky (1996) documents an increase in economic segregation for US metros between 1970 and 1990 using the Neighborhood Sorting Index, Watson (2009) finds an increase in residential segregation by income between 1970 and 2000 using the Centile Gap Index and, most recently, Reardon and Bischoff (2011) and Reardon, Bischoff, Owens and Townsend (2018) document this fact using the information theory index. Similar conclusions have also been reached by Owens (2018) who analyses segregation within US cities using school districts instead of census tracts as subunits of analysis.

¹²This issue is particularly salient for multigroup indexes cutting the distribution in many groups (reducing the number of observations for each) and when using ACS data which is characterized by a lower sampling rate than the decennial Census.

Figure 3: Segregation: different measures



Finally, figure 4 plots three other measures of income inequality that have been widely used in the literature:¹³ the 90/10 ratio that measures the ratio of the family income in the top 90th percentile of the population relative to the income in the bottom 10th percentile, and, similarly, the 50/10 ratio, and the 90/50 ratio.¹⁴ The figure shows that the 90/10 ratio that is a measure of overall inequality has increased steadily since 1980. Moreover, the 90/50 ratio has also been increasing, while the 50/10 ratio is flat or even slightly decreasing after 1990. This confirms that the rise in income inequality has been driven by the top of the distribution, as already shown by Autor, Katz and Kearney (2008) for individual wage inequality.

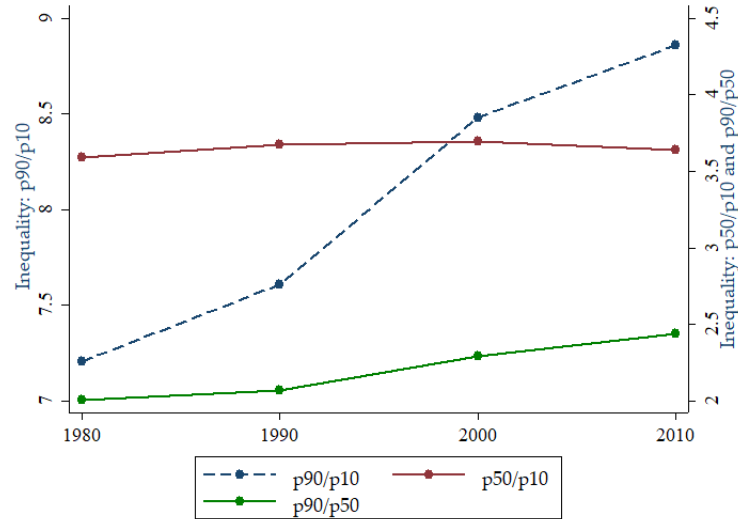
2.2 Segregation and Inequality Across US Metros

Next, we document that residential segregation and inequality are also correlated across space. Figure 5 shows the relationship between the Gini coefficient and the dissimilarity index across metro areas in 1980. The graph shows population-weighted data-points, with bubbles propor-

¹³The recent increase in US income inequality is also a well established fact. See, for example, Katz and Murphy (1992); Autor, Katz and Krueger (1998); Goldin and Katz (2001); Card and Lemieux (2001); Acemoglu (2002); Autor, Katz and Kearney (2008)

¹⁴The procedure implemented to calculate these ratios from binned data is described in Appendix B.

Figure 4: Inequality: different measures



tional to population weights, using 1980's metro-level population weights.¹⁵ The figure shows that there is a positive correlation between segregation and inequality. We estimate a regression coefficient of 1.367, with standard error of 0.095.¹⁶


Next, we show that there is a positive correlation also between changes in segregation and changes in inequality across space. Figure 6 plots the change at the metro level in the Gini coefficient between 1980 and 2000 against the change at the metro level in the dissimilarity index over the same time period. Again, the size of the bubble is proportional to the population of the metro area. The figure shows that the metro areas that experienced a larger increase in inequality between 1980 and 2000 are also those that experienced a larger increase in residential segregation over the same time period. The regression coefficient is 0.576 with a standard error of 0.087.¹⁷

¹⁵The weights are based on population counts within each metro in our sample from the U.S. census for that decade.

¹⁶The results of a regression of segregation on inequality across US metros in 1980 are reported in Table 10, Appendix B.

¹⁷The results of a regression of changes in segregation on changes in inequality across US metros between 1980 and 2000 are also reported in Table 10, Appendix B.

Figure 5: Inequality and Segregation across US Metros



2.3 Intergenerational Mobility and Segregation Across US Metros

Finally, we show that residential segregation across metros is also correlated with intergenerational mobility.¹⁸ To this end, we use a restricted-access geocoded version of the National Longitudinal Survey of Youth (NLSY79). The NLSY is a nationally representative sample of 12,686 young men and women who were 14-22 years old when they were first surveyed in 1979. These individuals were interviewed annually through 1994 and are currently interviewed on a biennial basis. We use this survey to construct intergenerational transition matrices for family income, calculated separately for observations in high-segregation locations and low-segregation areas.¹⁹ Using the geocoded version of the NLSY we are able to associate individual-level observations with their geographical location.²⁰

¹⁸This is consistent with the findings in Chetty, Hendren, Kline and Saez (2014).

¹⁹In constructing the matrices we follow closely Mazumder (2005) that shows significant differences in the mobility matrices for white and black individuals.

²⁰Summary statistics of the final dataset used to produce the transition matrices are included in Table 11, Appendix B.

Figure 6: Change in Inequality and Segregation 1980-2000




Combining the cross-sectional and supplemental samples, we create two variables: initial income and final income. Initial income is the average of the family income observed in the first 3 survey years. Final income is the average in the last 5 survey years. Each observation is assigned to a quartile in the initial income distribution and in the final income distribution separately.

Using the restricted-access geocoded NLSY data, we associate each observation of the NLSY79 subsample we constructed with a metro area where each individual grew up when young (1979-1981). Next we rank the metros from the main dataset by level of segregation in 1980, using the dissimilarity index calculated at the 80th percentile, and we create two groups: high-segregation metros are the one above the 50th percentile and low-segregation metros are the others. We can then construct two transitional matrices, as described above, one for each group.

Panels (a) and (b) in Figure 7 show the transitional matrices constructed for the low-segregation metros and for the high-segregation metros, respectively. The figures show that intergenerational mobility is in general stronger in areas that experience lower levels of residential segregation. In particular, we can focus on the two quadrants on the bottom left and top right, which tell us some-

Figure 7: Intergenerational Mobility Matrices



thing about upward and downward mobility, respectively. The bottom left quadrant of the matrix captures the persistence of poverty: the probability of ending up in the lowest quartile when one’s parents were in the lowest quartile. This probability is roughly 40% in the low-segregation metros while it is roughly 47% in the high-segregation metros.²¹ Similar differences characterize the persistence of wealth: the top right quadrant of the matrix represents the probability of ending up in the top quartile of the income distribution when one’s parents were in the same quartile. This probability is roughly 41% in the low-segregation areas, while it is roughly 47% in the high-segregation areas.

3 Model

We now propose a model of a metro area where families choose the neighborhood where to reside taking into consideration that there are local spillovers affecting their children’s future income.

²¹This difference is statistically significant at the 1% level. We present results of our hypothesis testing in Table 12, Appendix B. These results are robust to different combinations of groupings and different cutoffs for the dissimilarity index.

3.1 Set up

The economy is populated by overlapping generations of agents who live for two periods. In the first period, the agent is a child and accumulates human capital. In the second period, the agent is a parent. A parent at time t earns a wage $w_t \in [\underline{w}, \bar{w}]$ and has one child with ability $a_t \in [\underline{a}, \bar{a}]$. The ability of a child is correlated with the ability of the parent. In particular, $\log(a_t)$ follows the AR1 process

$$\log(a_t) = \rho \log(a_{t-1}) + v_t,$$

where v_t is normally distributed with mean zero and variance σ_v , and $\rho \in [0, 1]$ is the auto-correlation coefficient. The joint distribution of parents' wages and children's abilities evolves endogenously and is denoted by $F_t(w_t, a_t)$, with $F_0(w_0, a_0)$ given.

There are two neighborhoods, denoted by $n \in \{A, B\}$. Houses have all the same dimension and quality and the rent in neighborhood n at time t is denoted by R_t^n . For simplicity we make the extreme assumption that the housing supply is fixed and equal to H in neighborhood A and fully elastic in neighborhood B .²² In particular, we assume that the marginal cost of construction in neighborhood B is equal to 0 so that $R_t^B = 0$ for all t . The rental price in neighborhood A , R_t^A , is an endogenous equilibrium object.

In the baseline model we assume that there are two educational levels, that is, $e \in \{e^L, e^H\}$.²³ Also, let τ be the cost of investing in high education.

We assume that parents care both about their own consumption and about their children's future wage.²⁴ In particular, their preferences are given by $u(c_t) + g(w_{t+1})$, where u is a concave and continuously differentiable utility function, and g is increasing and continuously differentiable. An old agent with wage w_t and with a child of ability a_t chooses 1) how much to consume, $c_t(w_t, a_t) \in \mathbb{R}_+$; 2) where to live, $n_t(w_t, a_t) \in \{A, B\}$; and 3) how much to invest in the child's education, $e_t(w_t, a_t) \in \{e^L, e^H\}$. These choices affect the child's future wage, as explained below.

²²We make this assumption for simplicity, but one could introduce an intermediate level of housing elasticity in both neighborhoods.

²³In Section 5 we explore a version of the model with a continuous choice of education.

²⁴This assumption is common in this class of models. The assumption that agents cannot save (if not by investing in housing or kids' education) is for simplicity. The assumption that agents cannot borrow is for realism, given that typically people cannot borrow against children's future income. An alternative specification could have parents getting utility directly from their children's consumption, but with the introduction of a borrowing constraint.

A key ingredient of the model is the presence of a local spillover that affects the children's human capital accumulation, and hence future income. Children's wages are affected by their ability shock, by their education, by the neighborhood where they grow up because of the local spillover effect, and also directly by their parents' wage.²⁵ Formally, the child of an agent (w_t, a_t) who chooses to live in neighborhood n and get education level e is going to earn a wage

$$w_{t+1} = \Omega(w_t, a_t, e, S_t^n, \varepsilon_t), \quad (1)$$

where ε_t is an iid normally distributed noise with cdf Ψ , S_t^n is the local spillover effect in neighborhood n at time t , and Ω is non-decreasing in all its arguments. Children with higher ability, higher education, who live in neighborhoods with higher spillover, and have richer parents will accumulate more human capital, and hence earn higher wages. Because the residential and the educational choice are functions of the parents' wage and child's ability (w_t, a_t) , with a slight abuse of notation, we can write $w_{t+1} = w_{t+1}(w_t, a_t, \varepsilon_t)$. We will show that in equilibrium parents with higher wage, for given ability, will choose higher education and the neighborhood with higher spillover. This implies that wages will end up being increasing in parents' wages, also accounting for the indirect effect.

Let us now turn to the spillover. We assume that the spillover in neighborhood n at time t is equal to the average human capital of children growing up in that neighborhood, that is, equal to the children's average wage:

$$S_t^n = \frac{\int \int \int_{n_t(w_t, a_t)=n} w_{t+1}(w_t, a_t, \varepsilon_t) F_t(w_t, a_t) \Psi_t(\varepsilon_t) dw_t da_t d\varepsilon_t}{\int \int_{n_t(w_t, a_t)=n} F_t(w_t, a_t) dw_t da_t}. \quad (2)$$

Given that wages are increasing in ability and in parents' wage, neighborhoods with higher average human capital tend to be neighborhoods with higher ability children and richer parents. The idea is that children growing up in these neighborhoods will accumulate more human capital, for the same level of education and ability, because of pecuniary and social local externalities.²⁶ This formalization of the spillover can capture different sources of local externalities: neighborhoods with richer families have better public schools that are locally financed, children who grow

²⁵Parents' wage affect children's wage also indirectly through the educational and residential choices.

²⁶Alternative specifications could have the spillover equal to the average wage of the parents or to the average level of education of the children in the neighborhood. However, the first would miss the role of the innate ability and the second the role of parents' income. Also, in the baseline model, the second specification would not be particularly appealing because of the binary nature of the education level.

up there have better peers, parents who live there may invest more in education because they learn more successful stories, social norms are more conducive to educational investment, and so forth. The presence of the spillover effect implies that the rental rate in neighborhood A , R_t^A , also depends on the strength of the externality S_t^A , which is endogenous.

To simplify the analysis, we make two assumptions. First, we assume that the ability and the spillover affects a child's future wage only if he gets the high level of education.

Assumption 1 *The function $\Omega(w, a, e, S, \varepsilon)$ is constant in S and a if $e = e^L$, and is increasing in S and a if $e = e^H$.*

In the quantitative exercise, we will interpret the children with high education as college graduate and the ones with low education as less than college graduates. The assumption that the wage of children with low education level does not depend on the ability stands for the fact that abilities that are relevant in high-skill jobs may be different and more heterogenous than abilities that are relevant for low-skill jobs. The assumption that the spillover does not affect the wage of children with low education is extreme, but can be interpreted with the idea that the quality of schooling k-12th turns out to be important in determining future wages of college graduates more than no-college graduates. This second assumption simplifies the analysis because all parents living in the rich neighborhood also pay for their children's college, given that there would be no other reason to pay a higher rent in the first place. We will relax it in the general model.

Second, we assume that there are complementarities between the spillover and the children's ability, between education and ability, between parents' wage and spillover, and between parents' wage and education. In particular, we make the following assumption.

Assumption 2 *The composite function $g(\Omega(w, a, e, S, \varepsilon))$ has increasing differences in a and S , in a and e , in w and S , and in w and e .*

To sum up, a parent with wage w_t who has a child with ability a_t at time t solves the following problem

$$\begin{aligned}
 U(w_t, a_t) &= \max_{c_t, e_t, n_t} u(c_t) + E[g(w_{t+1})] & \text{(P1)} \\
 \text{s.t. } & c_t + R_t^n + \tau e_t \leq w_t \\
 & w_{t+1} = \Omega(w_t, a_t, e_t, S_t^n, \varepsilon_t),
 \end{aligned}$$

taking as given spillovers and rental rates in the two neighborhoods, S_t^n and R_t^n for $n = A, B$.

3.2 Equilibrium

We are now ready to define an equilibrium.

Definition 1 For a given initial wage distribution $F_0(w_0, a_0)$, an equilibrium is characterized by a sequence of educational and residential choices, $\{e_t(w_t, a_t)\}_t$ and $\{n_t(w_t, a_t)\}_t$, a sequence of rents in neighborhood A, $\{R_t^A\}_t$, a sequence of spillovers in neighborhoods A and B, $\{S_t^A\}_t$ and $\{S_t^B\}_t$, and a sequence of distributions $\{F_t(w_t, a_t)\}_t$ that satisfy:

1. agents optimization: for each t , the policy functions e_t and n_t solve problem (P1), for given R_t^n and S_t^n for $n = A, B$;
2. spillovers' consistency: for each t , equation (2) is satisfied for both $n = A, B$;
3. market clearing: for each t , the housing market clears in neighborhood A

$$H = \int \int_{n_t(w_t, a_t) = A} F_t(w_t, a_t) dw_t da_t; \quad (3)$$

4. wage dynamics: for each t condition (1) is satisfied.

From now on, we focus on equilibria where the housing market in neighborhood A clears with positive rents, that is, $R_t^A > 0$ for all t , which requires also $S_t^A > S_t^B$ for all t .²⁷

Assumptions 1 and 2 allow us to characterize the equilibrium in a fairly simple way, as shown in the following proposition.

Proposition 1 Under assumptions 1 and 2, for each time t there are two non-increasing cut-off functions $\hat{w}_t(a_t)$ and $\hat{\hat{w}}_t(a_t)$, with $\hat{w}_t(a_t) \leq \hat{\hat{w}}_t(a_t)$ such that

$$e_t(w_t, a_t) = \begin{cases} e^L & \text{if } w_t < \hat{w}_t(a_t) \\ e^H & \text{if } w_t \geq \hat{w}_t(a_t) \end{cases}, \quad (4)$$

and

$$n_t(w_t, a_t) = \begin{cases} B & \text{if } w_t < \hat{\hat{w}}_t(a_t) \\ A & \text{if } w_t \geq \hat{\hat{w}}_t(a_t) \end{cases}. \quad (5)$$

²⁷If $S_t^A \leq S_t^B$, nobody would like to live in A and the rental rate in A would be zero.

This proposition shows that in equilibrium the residential and the educational choices can be simply characterized by two monotonic cut-off functions.²⁸

Figure 8: Equilibrium Characterization




Figure 8 shows a graphical characterization of the equilibrium, for given spillovers and rental rates, with $R_t^A > 0$. On the x-axis there is the children's ability level a_t and on the y-axis the parents' wage w_t . For any given level of children's ability a_t , there are two thresholds for the parents' wage $\hat{w}_t(a_t)$ and $\hat{\hat{w}}_t(a_t)$, with $\hat{w}_t(a_t) \leq \hat{\hat{w}}_t(a_t)$, such that parents with wage $w_t < \hat{w}_t(a_t)$ choose to live in B and not to pay for a high level of education for their children, parents with wage $\hat{w}_t(a_t) \leq w_t < \hat{\hat{w}}_t(a_t)$ choose to live in B and pay for their children's high level education, and parents with wage $w_t \geq \hat{\hat{w}}_t(a_t)$ choose to live in A and pay for their children's high level education. The figure shows that children with richer parents and higher ability tend to be more educated and to live in neighborhood A. On the one hand, for given children's ability, richer parents are more willing to pay the cost of high-level education (e.g. college tuition) and the cost of a higher local externality (higher rental rate). On the other hand, for given wage, the higher the ability of a child, the more willing the parent is to pay for high-level education and for higher local externality because of the complementarities between ability and education and between ability and local spillover, respectively. For a given ability, a random child who grows up in B rather than A has lower probability of getting high-level education, both because parents living

²⁸Assumptions 1 and 2 are needed to obtain the monotonicity result.

in B are poorer on average and because the local spillover is weaker, hence making the incentive to pay for education even smaller.

The classic papers in this literature, building on Benabou (1996b) and Durlauf (1996b), typically focus on two extreme cases of segregation by income: either the two neighborhoods are equal to each other and have a representative distribution of income, or they are perfectly segregated, with all the richest agents residing in one and all the poorest in the other. Our model is richer in this dimension, as it allows us to obtain an intensive measure of segregation which we can match to the data. This is due to the presence of heterogeneity in ability: if all agents had the same ability level, the cut-off function $\hat{w}_t(a_t)$ in figure 8 would be horizontal and the two neighborhoods would feature full segregation by income. However, thanks to the heterogeneity in ability, the two cut-off functions are monotonically non-increasing in ability and some poorer parents with high ability children choose to live in A to exploit the complementarity with the higher spillover.

Our model also allows us to think about segregation by education. In our baseline model, given the binary choice of education, neighborhood A will always be fully segregated, in the sense that all children will get high-level education. However, neighborhood B will generically feature a mix of high- and low-level educated children. In particular, the degree of segregation by education is driven by the distance between the two cut-off functions $\hat{w}_t(a_t)$ and $\hat{w}_t(a_t)$. For some parameter configurations, these two functions can coincide, in which case there is perfect segregation by education, as all children living in A will get high-level education and all children in B will not.

3.3 Functional Forms

To study the model numerically, we now make some functional form assumptions. In particular, assume that $e^L = 0$, $e^H = 1$, $u(c) = g(c) = \log(c)$, and that the law of motion for the wages takes the form

$$\Omega(w, a, e, S^n, \varepsilon) = (b + ae(\beta_0 + \beta_1 S^n))w^\alpha \varepsilon. \quad (6)$$

On the one hand, this implies that the wage of children with low-level education ($e_t = 0$) is simply equal to $bw^\alpha \varepsilon_t$, and does not depend on either the children's ability or the neighborhood spillover, satisfying assumption 1. On the other hand, the wage of children with high-level education

($e_t = 1$) is a function of their ability as well as of the spillover. In particular, β_1 is the key parameter that determines the strength of the spillover. The specific functional form in (6) also satisfies assumption 2. In particular, ability is complementary both to education and to the local spillover.

With these functional forms, the household's problem reduces to

$$U(w_t, a_t) = \max_{e,n} \log(w_t - R_t^n - \tau e) + \log((b + a_t e(\beta_0 + \beta_1 S_t^n)) w_t^\alpha \varepsilon_t), \quad (\text{P2})$$

where $R_t^B = 0$, R_t^A is determined implicitly by $H = \int [1 - F_t(\hat{w}_t)] dG(a)$, and the spillover effects are given by

$$S_t^A = \frac{\int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \int_{\underline{a}}^{\bar{a}} \int_{\hat{w}_t(a)}^{\bar{w}} w_{t+1} dF(w_t, a_t) d\Psi(\varepsilon_t)}{\int_{\underline{a}}^{\bar{a}} \int_{\hat{w}_t(a)}^{\bar{w}} dF(w_t, a_t)}, \text{ and } S_t^B = \frac{\int_{\underline{\varepsilon}}^{\bar{\varepsilon}} \int_{\underline{a}}^{\bar{a}} \int_{\hat{w}_t(a)}^{\bar{w}} w_{t+1} dF(w_t, a_t) d\Psi(\varepsilon_t)}{\int_{\underline{a}}^{\bar{a}} \int_{\hat{w}_t(a)}^{\bar{w}} dF(w_t, a_t)}. \quad (7)$$

With the simple functional forms assumed, the cut-off functions that characterize the optimal education and residential choices can be characterized in closed form as we describe below.

For every ability level a , two cases are possible. In the first case, for a given ability, there is a positive measure of children who get high education in neighborhood B, and the two cut-offs are:

$$\hat{w}_t(a_t) = \tau \left[1 + \frac{b}{a_t(\beta_0 + \beta_1 S_t^B)} \right], \quad (8)$$

and

$$\hat{w}_t(a_t) = \tau + R_t^A \left[1 - \frac{b + a_t(\beta_0 + \beta_1 S_t^B)}{b + a_t(\beta_0 + \beta_1 S_t^A)} \right]^{-1}. \quad (9)$$

This case arises when the RHS of equation (8) is smaller than the RHS of equation (9). Equation (8) shows that the education cut-off \hat{w}_t is decreasing in the spillover in neighborhood B, that is, the higher is the spillover in B, the higher is the human capital accumulated by children getting high-level education, the higher is the willingness of parents living in B to pay for their children's education. Moreover, it shows that, for a given ability, the willingness of parents leaving in B to pay for education is higher when the parameters affecting the strength of the return to education and to spillover, β_0 and β_1 , are higher, and when the cost of education τ and/or the constant component of the income of low-educated children b are lower. Equation (9) shows that the location decision depends on the trade-off between the spillover advantage relative to the cost of leaving in neighborhood A.

In the second case, for a given ability, there is perfect segregation by education, that is, all children in B get the low education level. In this case, the cutoff functions coincide and are equal to

$$\hat{w}_t(a_t) = \hat{w}_t(a_t) = (\tau + R_t^A) \left[1 + \frac{b}{a_t(\beta_0 + \beta_1 S_t^A)} \right].$$

The higher is the cost of education τ or the lower is the return from education, which, for a given ability, depends on parameters β_0 and β_1 , and on the local spillover in A, the higher is the cutoff and the more the parents who decide to live in B. Moreover, more parents decide to live in B, the higher is the rental rate to live in A, R_t^A , and the higher is the constant component of the income of low-educated children.

4 Numerical Analysis

In this section we perform the main exercise of the paper. First, we calibrate the steady state of the model to the US economy in 1980. Then, we look at the effects of an unexpected, one-time, permanent shock to the skill premium on inequality, segregation, and intergenerational mobility over time and we quantify the contribution of the local spillover feedback to these dynamics.

4.1 Calibration

We now discuss how we choose the parameters so that the steady state equilibrium of the model matches salient features of the US economy in 1980. For the calibration we use the specific functional forms for the utility and the wage dynamics function that we have described in subsection 3.3, that is, $e^L = 0$, $e^H = 1$, $u(c) = g(c) = \log(c)$, and Ω satisfying equation (6). Table 1 shows the targets of our baseline calibration, which we are now going to discuss.

A crucial target for our calibration is what we call the “return to the spillover”, that is, the effect of the local spillover in the neighborhood where a child grows up on his future income. This effect is difficult to measure in the data. Fortunately, there has been a wave of recent research that uses micro data to estimate it. In particular, we use the results of Chetty and Hendren (2018b). Using tax returns data for all children born between 1980 and 1986, they estimate the causal effect of local spillovers on children’s future income, by looking at movers across US counties. Their baseline estimation implies that growing up in a 1 standard deviation better county from

Table 1: *Calibration Targets*

Description	Value	Source
Return to spillover	0.100	Chetty and Hendren (2018b)
Return to college	0.391	Goldin and Katz (2009), Census 1980
Gini coefficient	0.365	Census 1980, family income
Dissimilarity index	0.306	Census 1980, family income
Rank-rank correlation	0.340	Chetty, Hendren, Kline and Saez (2014)
Share of college grads	0.178	Census 1980
B/A ratio of college shares	0.400	Census 1980

birth would increase a child’s household income in adulthood by approximately 10%. This is the number that we target in our calibration. Let us explain how we do that. Given that nobody literally moves in our model, we map the “movers” in Chetty and Hendren (2018b) to the parents who decide to live in a neighborhood different from the one where they grew up, that is, the one chosen by their own parents. Then, we calculate the difference between the expected future income of the children of “movers” if they grew up in the rich neighborhood (A) and the expected future income of the same children if they grew up in the poor neighborhood (B). Finally, we divide this by the standard deviation of the spillover S^n across the two neighborhoods.²⁹

Another important target is the US skill premium in 1980s that is calculated using Census data, following Goldin and Katz (2009). In the model, we map the skill premium to the difference between the average log wage of educated agents and the average log wage of non educated agents in steady state.

As baseline measures of inequality and income segregation at the metro level, we use the gini coefficient and the dissimilarity index, where we define rich the households in the top 20th percentile of the metro income distribution, and poor the others. In particular, we target the average Gini coefficient and the average dissimilarity index for all metro areas in 1980. As described in Section 2, we use Census data to calculate both the Gini coefficient and the dissimilarity index at the metro level and then we aggregate them, weighting by population. As we discussed in Section 2, there are many alternative measures of income segregation that are used in the literature. Another measure that has been widely used in the more recent literature is the H^R index proposed by Reardon and Firebaugh (2002) and Reardon and Bischoff (2011), and also used by

²⁹This is simply equal to $\sqrt{H(1-H)}(S^A - S^B)$, where H is the housing supply in the rich neighborhood and S^A and S^B the steady state level of the spillovers in the two neighborhoods.

Chetty et al. (2014). We are going to check how well our model is able to match this measure as a validation of our calibration in the next section.

Another feature of the US data we want to target is the level of intergenerational mobility. To this end, we target the rank-rank correlation between log wages of parents and children estimated using administrative records by Chetty, Hendren, Kline and Saez (2014).³⁰ We chose this statistic instead than the log-log correlation because Chetty, Hendren, Kline and Saez (2014) argue that it provides a more robust summary of intergenerational mobility.³¹

Finally, given that the educational choice is binary, we interpret high-education level as college completion. We then target the share of college graduates in 1980 and the ratio of the share of college graduates in the poor relative to the rich neighborhoods that we obtain using Census data. In particular, we look at the number of people above 25 year old who completed college at the census tract level. To calculate the second statistic, we need to divide the census tracts in each metro area in two groups that correspond to neighborhood A and B in the model. In order to do so, first we rank the census tracts by median income. Then, we look at their population and define neighborhood A the richest census tracts with population above the 10th percentile (given that this is the percentile closest to the the calibrated value of H).

Table 2: Calibration Results

Target	Data	Model
Return to spillover	0.100	0.096
Return to college	0.391	0.393
Gini coefficient	0.365	0.365
Dissimilarity index	0.306	0.309
Rank-rank correlation	0.340	0.239
Share of college grads	0.178	0.179
B/A ratio of college shares	0.400	0.115

³⁰The rank-rank correlation is the relationship between the rank based on income of children relative to others in their birth cohort and the rank of parents based on income relative to others in the same birth cohort.

³¹We also use the NLSY data to calculate the rank-rank correlation and we obtain a value of .41, which is not too far from the number in Chetty, Hendren, Kline and Saez (2014). We prefer to use their number as main target, because of the coverage of their data. In alternative calibrations, we have also looked at the probability that a child ends up in the lowest quartile of the distribution when his parents are in the lowest quartile of the distribution and at the probability that a child ends up in the highest quartile of the distribution when his parents are in the same quartile. However, these statistics rely more heavily on the shape of the wage distribution which we do not target, given the stylized nature of the model.

Table 2 shows the results of our calibration, in terms of how well we match our targets for the steady state equilibrium of the model. The table shows that we are able to match well the return to spillover, the return to college, the Gini coefficient, the dissimilarity index, and the share of college graduates. We do less well on matching the rank-rank correlation and the relative share of college graduates in the two neighborhoods. Clearly the fact that we cannot match well all the moments comes from the stylized nature of our model. In particular, we overestimate intergenerational mobility (underestimate the rank-rank correlation) and we underestimate the relative share of college educated agents in the poor neighborhood. We believe that the more problematic assumption is the binary choice of education. The model imposes that the rich parents can at best pay for their children’s college, but they cannot invest more than that in education.³² This means that in order to match better intergenerational wage persistence, the model would require the poor to underinvest in college education, explaining that it is hard for our baseline model to match both these moments. In Section 5 we are going to explore an extended version of the model where we allow for a continuous educational choice.

Table 3: *Parameters*

Parameter	Value	Description
H	0.07	Size of neighborhood A
α	0.19	Wage function parameter
β_0	1.10	Wage function parameter
β_1	0.29	Wage function parameter
τ	0.31	Cost of education
b	1.12	Wage fixed component for no-college
ρ	0.70	Autocorrelation of ability
σ	0.50	Standard dev. of log innate ability
μ_a	-2.80	Average of log innate ability
μ_ε	0.42	Average of log wage noise shock
σ_ε	0.65	Standard dev. of log wage noise shock

Table 3 shows the parameters that we are using to calibrate the model, their calibrated value, and their description. Notice that the number of parameters is higher than the number of targets because the model is highly non-linear.

³²This also implies that all kids in the rich neighborhood go to college.

4.2 Skill Premium Shock

As the data show, the US experienced a steady increase in labor income inequality starting in 1980. While, there are different potential driving forces behind this increase, here we focus on skill-biased technical change, which is widely recognized to be a crucial source of inequality (see, for example, Katz and Murphy, 1992; Autor, Katz and Krueger, 1998; Goldin and Katz, 2001; Card and Lemieux, 2001; Acemoglu, 2002; Autor, Katz and Kearney, 2008).

In this spirit, our main exercise is to explore the response of the economy to an unexpected, one-time, permanent shock to the skill premium. In particular, we change β_0 and β_1 proportionally so as to match the increase in the skill premium in the data. According to Census data, the skill premium in the US increased from .39 in 1980 to .54 in 1990 and up to .57 in 2000. In the model, individuals live for two periods: in the first period, they are young and go to school, and in the second period, they are old and work. As noted by Fernandez and Rogerson (1998), in this class of models, individuals spend the same time in period 1 and 2, so we could target the length of a period to the working period or to the schooling period. Given our focus on human capital accumulation, we choose to interpret one period as 10 years.³³ We interpret period $t = 0$ as 1980, when the economy is in steady state. Then, we assume that at that time an unexpected, permanent shock hits proportionally β_0 and β_1 so that the skill premium goes from 0.39 in 1980 ($t = 0$) to 0.54 in 1990 ($t = 1$).

Table 4: *Response to a Skill Premium Shock*

	t = 0	t = 1	t = 2	t = 3
Return to college	0.393	0.538	0.557	0.562
Gini coefficient	0.365	0.393	0.399	0.400
Dissimilarity index	0.309	0.343	0.355	0.357
Rank-rank correlation	0.239	0.321	0.354	0.360
Rent in neighborhood A	0.100	0.376	0.628	0.715
A/B spillovers ratio	1.584	1.998	2.184	2.225

Table 4 shows the response of the economy to such a shock, one, two, and three periods ahead. In particular, we show the behavior of the return to college, the Gini coefficient, the dissimilarity

³³The schooling age could be interpreted as 10 or 15 years depending on which level of education one targets. We ended up choosing 10 years also considering that Census data are available every 10 years.

index, the rank-to-rank correlation, the rent in the rich neighborhood, and the ratio of the spillover in neighborhood A over the one in neighborhood B.

In response to the skill premium shock, both inequality and segregation increase. Although the model is stylized in many dimensions, these responses are not too far from what happened in the data. In particular, the Gini coefficient increased from 0.365 in 1980 to 0.393 in 1990 and to .4 in 2010.³⁴ This increase in inequality is somehow smaller than in the data, as the Gini coefficient reached 0.397 in 1990 and .426 in 2010. The smaller response is not surprising, as we only focused on one shock behind the increase in inequality. The dissimilarity index increased from 0.309 in 1980, to 0.343 and 0.360 in 1990 and 2010 respectively. In the data the same index achieved 0.338 and 0.360 in the same years, which is roughly what happened in our simulation. To visualize these results, panels (a) and (b) in Figure 9 show the behavior of the Gini coefficient and of the dissimilarity index, respectively, in response to the skill premium shock we analyze, together with their patterns in the data.

As mentioned before, another measure that has been widely used to measure income segregation is the H^R index, introduced by Reardon and Firebaugh (2002) and Reardon and Bischoff (2011). Our steady state calibration implies an H^R index of 0.094 in 1980, which is pretty close to the 0.099 in the data, and to the 0.095 of the “bias-corrected” version, calculated in Reardon, Bischoff, Owens and Townsend (2018).³⁵ Figure 10 shows the response of the H^R index to the skill premium shock in the model, compared to its pattern in the data.

In response to the skill premium shock, intergenerational mobility decreases, as highlighted by the rank-rank correlation going from 0.239 in 1980, to roughly 0.321 in 1990, up to 0.36 in 2010. Unfortunately, given the limited availability of data, it is hard to calculate a reliable time-series for the rank-rank correlation. However, Aaronson and Mazumder (2008) show some indirect evidence of a positive relationship between the skill premium and the IGE (intergenerational elasticity) that is consistent with our findings. Moreover, it is interesting to notice that although our calibration is not able to achieve a rank-rank correlation of 0.34 in 1980, it actually reaches a similar number in 2010. Chetty, Hendren, Kline and Saez (2014) focus on US citizens in the

³⁴We calculate the Gini coefficient of time t using the income distribution of the people who are born at time t . Looking at the parents’ income distribution at the same time would not be interesting as the shock happens after their human capital has been determined.

³⁵To calculate the H^R index in the model, we average the two-group Theil index defined for the 100 percentiles of the distribution, weighting for the entropy defined for each decile.

Figure 9: Responses to a skill premium shock




1980-1982 birth cohorts and they measure their income as mean total family income in 2011 and 2012, and their parents' income as mean family income between 1996 and 2000. This implies that a rank-to-rank correlation of 0.34 would actually probably be a better target for 2010, in line with the results of our exercise.

4.3 Main Counterfactual Exercise

We can now use the model to perform a number of counterfactual exercises. The first one addresses the main question that motivates the paper: how important are local spillovers in amplifying the effects of a shock to income inequality? Table 5 shows the response of our model economy to the skill premium shock, keeping the spillovers S^A and S^B at their initial steady state

Figure 10: H^R Index



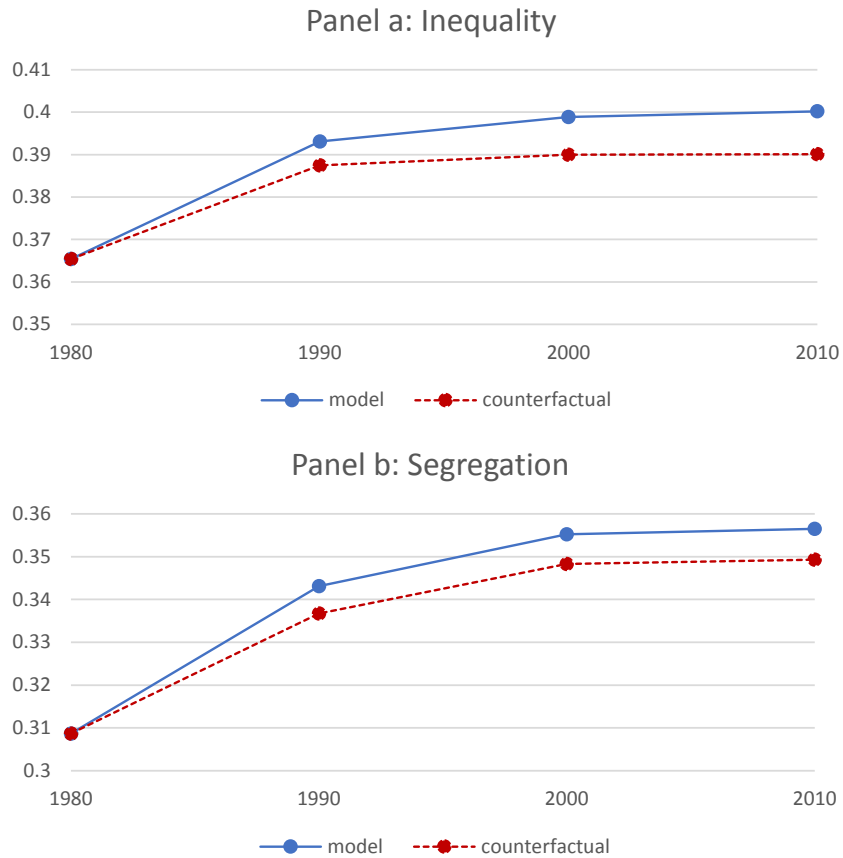
levels. We then compare this response to the full equilibrium response of the economy in Table 4, and interpret the differential response as the amplification effect due to the local externality. For comparison, the two responses of inequality and segregation to the shock are plotted in Figure 11.

Table 5: *Counterfactual with no Spillover Feedback*

	t = 0	t = 1	t = 2	t = 3
Return to college	0.393	0.512	0.523	0.523
Gini coefficient	0.365	0.388	0.390	0.390
Dissimilarity index	0.309	0.337	0.348	0.349
Rank-rank correlation	0.239	0.311	0.336	0.338
Rent in neighborhood A	0.100	0.202	0.270	0.288
A/B spillovers ratio	1.584	1.584	1.584	1.584

Comparing Tables 4 and 5, one can notice that in the first period the Gini coefficient increases by 0.028 in the equilibrium simulation and only by 0.022 in the counterfactual simulation with fixed spillovers. This means that the local externality accounts for 20% of the increase in inequality in the short run. If we look at long run responses, three periods after the shock, the increases are 0.035 in the equilibrium simulation and 0.025 in the counterfactual, implying a contribution of the spillover of 29%. We can do a similar decomposition for the responses of the dissimilarity

Figure 11: Counterfactual with no spillover feedback




index and obtain that the local externality accounts for 18% of the increase in segregation in the short run and for 15% in the long run. For intergenerational mobility—measured by the rank-rank correlation between parents and kids—the local externality contributes to 12% of the decrease in the short run and 18% in the long run.

The key mechanism behind the different responses in Tables 4 and 5 is due to the equilibrium responses of the local spillover in the two neighborhoods. The skill premium shock increases the spillover in both neighborhoods, but it increases it more in the rich neighborhood as shown in the last line of Table 4, which reports the ratio S_t^A/S_t^B . The same line in Table 5 shows that this ratio is constant in the counterfactual, as assumed by construction.

Two effects explain the response of S_t^A/S_t^B to the skill-premium shock in the equilibrium sim-

Figure 12: Decomposing the spillover feedback



ulation. First, there is a direct effect: children in the rich neighborhood benefit more from the increase in the skill premium because they all receive the high-level education and are exposed to the higher spillover level in neighborhood *A*. This mechanically increases the level of their human capital, and hence of the spillover in neighborhood *A* more than in neighborhood *B* in response to a skill-premium shock. Second, as the rental rate in the rich neighborhood increases, the degree of sorting by income increases. Although the more talented children will benefit more from the increase in skill premium, only richer families will be able to pay the higher cost of living in the rich neighborhood, irrespective of their children’s ability. This further raises the gap between the spillovers in the two neighborhoods.

The relative contribution of the two effects just described is illustrated in Figure 12. The black solid line shows the simulated path of the spillover ratio S_t^A / S_t^B in equilibrium, while the blue dashed line shows the direct effect of the skill-premium shock on the ratio. The direct effect is calculated by looking at the response of the spillover ratio, keeping fixed the educational and the residential choice of each household at their steady state values. The figure shows that the endogenous reallocation of households across neighborhoods plays a crucial role in producing a large increase in the spillover ratio.

4.4 Understanding the mechanism

Next, to better understand the different forces at play in our model, Table 6 shows two more counterfactual exercises that, together with the previous one, help us understanding the different effects of the skill premium shock on the economy.


Table 6: *More Counterfactual Exercises*

Counterfactual I: direct effect	t = 0	t = 1	t = 2	t = 3
Return to college	0.393	0.658	0.674	0.676
Gini coefficient	0.365	0.389	0.393	0.395
Dissimilarity index	0.309	0.309	0.363	0.376
Rank-rank correlation	0.239	0.293	0.319	0.326
Rent in neighborhood A	0.100	0.100	0.100	0.100
A/B spillovers ratio	1.584	1.584	1.584	1.584
Counterfactual II: partial equilibrium				
Return to college	0.393	0.535	0.553	0.555
Gini coefficient	0.365	0.390	0.394	0.394
Dissimilarity index	0.309	0.501	0.732	0.801
Rank-rank correlation	0.239	0.321	0.352	0.356
Rent in neighborhood A	0.100	0.100	0.100	0.100
A/B spillovers ratio	1.584	1.584	1.584	1.584

First, there is a standard direct effect of the increase in the skill premium. Keeping the spillovers, the house rental price, and the educational and residential choices as given, inequality mechanically increases because of two reasons. First, the income gap between college and non-college educated workers mechanically increases. Second, the return to the local spillover, which is complementary to education, is also higher, implying that children living in the rich neighborhood, all college educated, have an even higher income because of the higher spillover. The first counterfactual exercise in Table 6 shows the results of this direct effect, by displaying what would happen in response to the same skill premium shock we considered in subsection 4.2, if the spillovers, the rental prices and the equilibrium cut-off functions are kept unchanged. The table shows that inequality would increase substantially just because of this direct effect. However, segregation would not move much on impact if the residential choice was given and there was no general equilibrium effect.

The second effect comes from the change in the educational and residential choices. Panel (a)

Figure 13: Cut-off Response to Skill Premium Shock



in Figure 13 shows qualitatively the response of the educational and residential cut-off functions, respectively $\hat{w}_t(a)$ and $\hat{w}_t(a)$, to the skill premium shock. The figure shows that both cut-off functions shift to the left, so that more children of any ability would get higher education and would live in neighborhood A, if spillovers and rental rates were unchanged. The change in the educational choice is intuitive: the higher the skill premium, the higher the demand for college, conditional on any level of ability. Moreover, given that the local spillover is complementary to education, the higher the skill premium, the higher is the return to the spillover, and hence the higher is the demand to live in neighborhood A, conditional on any level of ability. These changes could potentially increase or decrease inequality depending on the original distribution. The second counterfactual exercise in Table 6 shows the partial equilibrium response to the same skill premium shock of subsection 4.2, that embeds the direct effect and the change in the cut-off functions, but keep the spillovers and the rental price in A as given. It turns out that, in our numerical exercise, the change in the cut-off functions increase inequality further.

The third effect is the general equilibrium effect, coming from the response of the rental rate in neighborhood A to clear the housing market, taking the spillovers as given. Panel (b) in Figure 13 shows that the residential cut-off function shifts somehow back to the right. As we explained above, taking as given the rental rate and the spillover effects, the demand to live in neighborhood A will increase because of the differential spillover and the complementarity between spillover and education, shifting the residential cut-off to the left. Given that the housing

supply in neighborhood A is fixed, this pushes up rental rates in that neighborhood, shifting the housing demand back to the right. In particular, the figure shows that the shift back is more pronounced for the poorer parents, who won't be able to afford the higher cost of living in the rich neighborhood, irrespective of their children's ability.³⁶ On net, this will generate the tilting that we see in panel (b) in Figure 13 that corresponds to more income segregation. The counterfactual exercise in Table 5 shows the general equilibrium effect of the skill premium shock, keeping the spillovers as given. The comparison with the partial equilibrium exercise shows that the rental rate in neighborhood A increases substantially. This is because of the increase in attractiveness of the high spillover that benefits children growing up in the rich neighborhood. The comparison also shows that in our numerical exercise the general equilibrium effect dampens the increase in inequality relative to the partial equilibrium effect.

4.5 House Prices


As we emphasized above, a key feature of our exercise is that the spillover in the rich neighborhood increases more relative to the poor neighborhood in response to the skill premium shock. This has the general equilibrium effect of increasing the rental rate in the rich neighborhood relative to the poor one. As a validating exercise, using Census data, we have looked at the behavior of house prices in rich and poor neighborhoods between 1980 and 2010. To define the rich and poor neighborhoods, we proceed as follows. We rank the census tracts by median income, we take the the richest tracts that contain 10% of the MSA population and we set the price in the rich neighborhood as the average price in these tracts.³⁷

Figure 14 shows the difference between house prices in the rich and the poor neighborhoods, normalized by median income. The figure shows that houses have become relatively more expensive over time in the rich neighborhoods, consistent with our model.

³⁶Given the complementarity between education and ability, the partial equilibrium shift to the left would be more pronounced for more able children. However, given our functional forms, the other effect dominates.

³⁷This is the same strategy that we used to calculate the share of college educated in the rich and poor neighborhoods that we used as one of the calibration targets. We choose 10% because is the closest percentile to the calibrated value for the size of the rich neighborhood, H .

Figure 14: House Price Difference between Rich and Poor Neighborhoods



5 Extended Model

In this section, we extend the model in two main dimensions: first, we introduce a residential preference shock and, second, we make the educational choice continuous. The introduction of the preference shock is important to obtain a more realistic setting where not all parents who live in the more expensive neighborhood choose high levels of education for their children. In our baseline model, the only reason to pay a higher rent to live in neighborhood A is to exploit the higher externality that affects the returns to education. In reality, residential choices are not purely driven by educational considerations. Families may prefer more expensive neighborhoods for a number of different reasons, such as better amenities, or higher status. By missing this feature of reality, the baseline model might generate a distribution of children growing up in neighborhood A biased towards too high innate ability. We then introduce a preference shock $\theta \in \{0, \bar{\theta}\}$ where $\bar{\theta} \geq 0$ and $\pi = \text{Prob}(\theta = \bar{\theta})$, so that families with $\theta = \bar{\theta}$ will enjoy their consumption more if they live in neighborhood A, that is, utility from current consumption becomes $\log((1 + \theta I_{n=A})c)$.

The introduction of a continuous educational choice seems potentially important for our quantitative exercise. The binary educational choice has the disadvantage that constraints rich parents on how much they can invest in their children's education, given that the best they can do is to pay for their college. As we mentioned in the previous section, this means that the baseline model generates too much intergenerational mobility, given that the only way parents have to invest in their children's education is to pay for college, but the data discipline the skill premium and hence bound how much rich parents can pay to increase their children expected income. This means that the binary choice also naturally bounds the possible increase in the spillover in response to an increase in inequality. Finally, this is consistent with the fact that there has been an increasing polarization between educational investment in rich and poor families.³⁸ We then assume that the educational choice is continuous, with $e \in [0, 1]$ and that the cost of education is linear τe .

With these two modifications, the problem of a household (w, a) becomes

$$U(w, a) = \max_{e, n} \log((1 + \theta I_{n=A})(w - R_t^n - \tau e)) + \log(b + ae(\beta_0 + \beta_1 S_t^n))\varepsilon. \quad (\text{P3})$$

In order to understand better the role of the educational choice, let us, for a moment, shut down the preference shock, that is, set $\hat{\theta} = 0$. In this case, the first order condition for the educational choice gives

$$e(w, a|n) = \frac{w - R_t^n}{2\tau} - \frac{b}{2a(\beta_0 + \beta_1 S_t^n)},$$

where $e(w, a|n)$ is the educational choice of a parent with wage w and a child with innate ability a conditional on living in neighborhood n and on $e(w, a|n)$ being positive. The expression shows that, as expected, education is increasing in income w , innate ability a , and in the spillover of the neighborhood S_t^n . It is also increasing in β_0 and β_1 , which determine the return to education. Moreover, it is decreasing in the cost of education τ and in b , which is the average wage of non-college educated workers.

The equilibrium definition is a natural extension of Definition 1 in Section 3.2, except that the agents' policy functions now also depend on the preference shock and the educational choice is now a continuous choice.

We now calibrate the model according to the same strategy described in Section 4. The only difference is that, given that now education level is a continuous variable, we define a cut-off \hat{e}

³⁸See Duncan and Murnane (2016)

such that individuals with an education level above \hat{e} are college educated, and the ones with education below are not. We choose \hat{e} so that, in 1980, 17.8% of the population is college educated, as in the data. When we look at the evolution of the ratio of college graduates in the two neighborhoods, we keep \hat{e} fixed to that level to define college and no-college graduates in the model.

Table 7: Calibration Results

Target	Data	Model
Return to spillover	0.100	0.097
Return to college	0.391	0.388
Gini coefficient	0.365	0.363
Dissimilarity index	0.306	0.304
Rank-rank correlation	0.340	0.338
Poor/rich college share	0.4	0.343

Table 7 shows how we match the targets with our calibration. Relative to the baseline model, we can now match the rank-rank correlation and, at the same time, getting closer to the ratio of college share in the two neighborhoods.

We then perform the same exercise that we have conducted in section 4 and hit the steady state economy with an unexpected and permanent shock that proportionally increases β_0 and β_1 in order to increase the skill premium from 0.39 in 1980 to 0.54 in 1990.

Table 8 shows the responses of the economy to such a shock. The results are qualitatively and quantitatively similar to the baseline model: in response to an increase in the skill premium, inequality and residential segregation increase, while intergenerational mobility decreases. However, we we perform the main counterfactual exercise, keeping the spillover levels in the two neighborhoods at their 1980 values, we obtain significantly larger effects of the spillover mechanism. In particular, the table shows that 42% of the increase in inequality in the short run and 58% in the long run can be attributed to the spillover effect. Moreover, the spillover effect accounts for 68% and 65% of the increase in residential segregation in the short and long run respectively. Figure 15 compares the response of inequality and segregation to the skill premium shock in the model with the analogous responses in the counterfactual exercise.

Table 8: *Response to a Skill Premium Shock*

	t = 0	t = 1	t = 2	t = 3
Return to college	0.388	0.537	0.574	0.581
Gini coefficient	0.363	0.384	0.394	0.397
Dissimilarity index	0.304	0.337	0.348	0.350
Rank-rank correlation	0.338	0.375	0.399	0.407
Rent in neighborhood A	0.844	1.140	1.542	1.736
A/B spillovers ratio	1.573	2.097	2.438	2.579
Counterfactual with no spillover feedback				
Return to college	0.388	0.528	0.535	0.535
Gini coefficient	0.363	0.375	0.377	0.378
Dissimilarity index	0.304	0.315	0.320	0.320
Rank-rank correlation	0.338	0.363	0.377	0.378
Rent in neighborhood A	0.844	0.893	1.002	1.034
A/B spillovers ratio	1.573	1.573	1.573	1.573


6 Concluding Remarks

We proposed a model where segregation and inequality amplify each other because of a local spillover that affects returns to education. We calibrated the model using US data in 1980, and using the micro estimates of neighborhood externalities that Chetty and Hendren (2018b) proposed using administrative data. We then hit the economy with an unexpected permanent shock to the skill premium and looked at the responses over time of inequality, residential segregation, and intergenerational mobility. We used the model to show that local spillovers contributed to roughly 20% of the increase in inequality in response to a skill-biased technical change shock in the short run, and to roughly 29% in the long run. We also show that in a more general version of the model, this contribution increases to 42% and 58% in the short and long run, respectively.

There are many directions to extend the model in future work. It would be interesting to have a larger number of neighborhoods to use the full richness of the data to discipline the model. Another interesting extension would be to endogenize the cost of education, that has also increased over time, and think about possible feedback effects. Also, it would be interesting to explore the normative implications of our results.

Another interesting direction that we are planning to explore in the future is to use the model to think about the correlation of inequality, segregation, and intergenerational mobility across metro

Figure 15: Responses to a skill premium shock



areas, also to understand their differential responses to a common skill premium shock.

References

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

Acemoglu, Daron, “Technical change, inequality, and the labor market,” *Journal of economic literature*, 2002, 40 (1), 7–72.

Alesina, Alberto and Ekaterina Zhuravskaya, “Segregation and the Quality of Government in a Cross Section of Countries,” *American Economic Review*, 2011, 101 (5), 1872–1911.

- Armour, Philip, Richard V. Burkhauser, and Jeff Larrimore**, “Using the Pareto Distribution to Improve Estimates of Topcoded Earnings,” *Economic Inquiry*, 2016, 54 (2), 1263–1273.
- Autor, David H, Lawrence F Katz, and Alan B Krueger**, “Computing inequality: have computers changed the labor market?,” *The Quarterly Journal of Economics*, 1998, 113 (4), 1169–1213.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney**, “Trends in US Wage Inequality: Revising the Revisionists,” *The Review of Economics and Statistics*, 2008, 90 (2), 300–323.
- Bartolome, Charles AM De**, “Equilibrium and inefficiency in a community model with peer group effects,” *Journal of Political Economy*, 1990, 98 (1), 110–133.
- 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.
- Benabou, Roland**, “Workings of a city: location, education, and production,” *The Quarterly Journal of Economics*, 1993, 108 (3), 619–652.
- , “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.
- Card, David and Thomas Lemieux**, “Can falling supply explain the rising return to college for younger men? A cohort-based analysis,” *The Quarterly Journal of Economics*, 2001, 116 (2), 705–746.
- 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*, 2014, *129* (4), 1553–1623.
- Cutler, David M. and Edward L. Glaeser**, “Are Ghettos Good or Bad?,” *Quarterly Journal of Economics*, 1997, *112* (3), 827–872.
- , —, and **Jacob L. Vigdor**, “The Rise and Decline of the American Ghetto,” *Journal of Political Economy*, 1999, *107* (3), 455–506.
- Domina, Thurston**, “Brain Drain and Brain Gain: Rising Educational Segregation in the United States, 1940-2000,” *City and Community*, 2006, *5* (4), 387–407.
- Duncan, Greg J and Richard J Murnane**, “Rising inequality in family incomes and children’s educational outcomes,” *RSF*, 2016.
- 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.
- 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 —, “Keeping People Out: Income Distribution, Zoning, and the Quality of Public Education,” *International Economic Review*, 1997, *38* (1), 23–42.
- 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,” *St Louis Fed, mimeo*, 2017.
- Goldin, Claudia and Lawrence F Katz**, “Decreasing (and then increasing) inequality in America: a tale of two half-centuries,” *The causes and consequences of increasing inequality*, 2001, pp. 37–82.
- Goldin, Claudia Dale and Lawrence F Katz**, *The race between education and technology*, Harvard University Press, 2009.
- Jargowsky, Paul A.**, “Take the Money and Run: Economic Segregation in U.S. Metropolitan Areas,” *American Sociological Review*, 1996, *61* (6), 984–998.
- Katz, Lawrence F and Kevin M Murphy**, “Changes in relative wages, 1963–1987: supply and demand factors,” *The quarterly journal of economics*, 1992, *107* (1), 35–78.
- Loury, Glenn C**, “Intergenerational transfers and the distribution of earnings,” *Econometrica: Journal of the Econometric Society*, 1981, pp. 843–867.
- Massey, Douglas S and Nancy A Denton**, “The dimensions of residential segregation,” *Social forces*, 1988, *67* (2), 281–315.
- , **Jonathan Rothwell, and Thurston Domina**, “The changing bases of segregation in the United States,” *The Annals of the American Academy of Political and Social Science*, 2009, *626* (1), 74–90.
- Mazumder, Bhashkar**, “Fortunate Sons: New Estimates of Intergenerational Mobility in the United States Using Social Security Earnings Data,” *Review of Economics and Statistics*, 2005, *87* (2), 235–255.
- Montgomery, James D**, “Social networks and labor-market outcomes: Toward an economic analysis,” *The American economic review*, 1991, *81* (5), 1408–1418.
- , *Social networks and persistent inequality in the labor market*, Center for Urban Affairs and Policy Research Evanston, 1991.

- Nielsen, Francois and Arthur S. Alderson**, “The Kuznets Curve and the Great U-Turn: Income Inequality in U.S. Counties, 1970 to 1990,” *American Sociological Review*, 1997, 62 (1), 12–33.
- Owens, Ann**, “Income Segregation between School Districts and Inequality in Students’ Achievement,” *Sociology of Education*, 2018, 91 (1), 1–27.
- Reardon, Sean F and Glenn Firebaugh**, “Measures of multigroup segregation,” *Sociological methodology*, 2002, 32 (1), 33–67.
- **and Kendra Bischoff**, “Income inequality and income segregation,” *American Journal of Sociology*, 2011, 116 (4), 1092–1153.
- Reardon, Sean F., Kendra Bischoff, Ann Owens, and Joseph B. Townsend**, “Has Income Segregation Really Increased? Bias and Bias Correction in Sample-Based Segregation Estimates,” *Demography*, 2018, *forthcoming*.
- Streufert, Peter**, “The effect of underclass social isolation on schooling choice,” *Journal of Public Economic Theory*, 2000, 2 (4), 461–482.
- von Hippel, Paul T., David J. Hunter, and McKalie Drown**, “Better Estimates from Binned Income Data: interpolated CDFs and Mean-Matching,” *Working Paper*, 2017.
- Watson, Tara**, “Inequality and the Measurement of Residential SEgregation by Income in American Neighborhoods,” *NBER Working Paper No. 14908*, 2009.
- Zheng, Angela**, “Public Education Inequality and Intergenerational Mobility in an Overlapping Generations Model,” *mimeo*, 2017.

Appendix

A Proof of Proposition 1

Given that we focus on equilibria with $R^A > R^B = 0$, we require $S^A > S^B$. Also, this together with assumption 1 implies that agents who chooses low education strictly prefer neighborhood B to neighborhood A, so nobody chooses $e = e^L$ and $n = A$. Hence, agents choose among three options: 1) high education and neighborhood A, for short HA ; 2) high education and neighborhood B, HB ; 3) low education and neighborhood B, LB .

Let us consider a given time t and drop the time subscript to simplify notation. Also, to simplify the notation, let us drop ε , given that it is iid, so does not play any role for the cut-off functions. Consider an agent with wealth w and ability a who chooses HA . It must be that he prefers that to choosing HB or LB , that is,

$$u(w - R^A - \tau) + g(\Omega(w, a, e^H, S^A)) \geq u(w - \tau) + g(\Omega(w, a, e^H, S^B)) \quad (10)$$

and

$$u(w - R^A - \tau) + g(\Omega(w, a, e^H, S^A)) \geq u(w) + g(\Omega(w, a, e^L, S^B)). \quad (11)$$

Take any $w' > w$. By concavity of u and $R^A > 0$, we have

$$u(w' - R^A - \tau) - u(w' - \tau) \geq u(w - R^A - \tau) - u(w - \tau)$$

and

$$u(w' - R^A - \tau) - u(w') \geq u(w - R^A - \tau) - u(w).$$

Combining these conditions with the assumption that the composite function $g(\Omega)$ has increasing differences in w and S and in w and e (from assumption 2), we obtain

$$u(w' - R^A - \tau) + g(\Omega(w', a, e^H, S^A)) \geq u(w' - R^B - \tau) + g(\Omega(w', a, e^H, S^B))$$

and

$$u(w' - R^A - \tau) + g(\Omega(w', a, e^H, S^A)) \geq u(w' - R^B) + g(\Omega(w', a, e^L, S^B))$$

for all $w' > w$ and given a . Let us call $w_1(a)$ and $w_2(a)$ the values of w that make respectively conditions (10) and (11) hold with equality for given a . We can then define the cutoff function

$$\hat{w}(a) = \max\{w_1(a), w_2(a)\}.$$

This proves that all agents with $w \geq \hat{w}(a)$ choose the option HA for given a . Using assumption 1 and 2 and the implicit function theorem, it is straightforward to show that both $w_1(a)$ and $w_2(a)$ are decreasing functions, and hence that $\hat{w}(a)$ is a decreasing function as well.

Next, consider an agent with wealth w and ability a who chooses LB . It must be that he prefers that to choosing HA or HB , that is,

$$u(w - R^B) + g(\Omega(w, a, e^L, S^B)) \geq u(w - R^A - \tau) + g(\Omega(w, a, e^H, S^A)) \quad (12)$$

and

$$u(w - R^B) + g(\Omega(w, a, e^L, S^B)) \geq u(w - R^B - \tau) + g(\Omega(w, a, e^H, S^B)). \quad (13)$$

Following analogous steps to before, we can show that, for given a , all agents with $w' < w$ will prefer LB to both HA and HB . Notice that the value w that makes equation (12) hold with equality is the cut-off value $w_2(a)$ defined above. Moreover, let us call $w_3(a)$ the value of w that makes condition (13) hold with equality for given a . We can then define the cutoff function

$$\hat{\hat{w}}(a) = \min\{w_2(a), w_3(a)\}.$$

This proves that all agents with $w \leq \hat{\hat{w}}(a)$ choose the option LB for given a . Using assumption 2 and the implicit function theorem, it is straightforward to show that $w_3(a)$ is also a decreasing function, and hence that $\hat{\hat{w}}(a)$ is a decreasing function as well. Given that both $\hat{w}(a)$ and $\hat{\hat{w}}(a)$ are decreasing functions, it must be that $\hat{w}(a) \geq \hat{\hat{w}}(a)$ for all a . If there was an a' such that $\hat{w}(a) < \hat{\hat{w}}(a)$, then all agents with $w \in (\hat{w}(a), \hat{\hat{w}}(a))$ would find strictly optimal both HA and LB , which is a contradiction. This proves that an equilibrium is characterized by two decreasing function $\hat{w}(a)$ and $\hat{\hat{w}}(a)$ with $\hat{w}(a) \geq \hat{\hat{w}}(a)$ for all a , such that all agents with (w, a) such that $w > \hat{w}(a)$ will choose $e = e^H$ and $n = A$ and all agents with (w, a) such that $w < \hat{\hat{w}}(a)$ will choose $e = e^L$ and $n = B$.

B Data Methodology

B.1 Segregation and Inequality over Time

Data sources and sample selection We use tract level income data from Decennial Censuses (1980 to 2010) and from the American Community Surveys (ACS) for the 5 year period spanning 2008-2012. Our sample includes metropolitan areas using the 2003 OMB definition. Table 9 reports our summary statistics at the metro and census tract level. The number of census tracts has increased over time reflecting the increase in the population.

Table 9: *Summary Statistics*

year	no_metros	no_CTs	ave_no_CTs
1980	379	42406	111.9
1990	380	48412	127.4
2000	380	53033	139.6
2010	380	59842	157.5

Segregation and Inequality Across US Metros In Table 10 we report the results of a regression of segregation on inequality at the metro level, first in levels and then in first differences. The results show a strong correlation between the two variables: metros with higher level of inequality in 1980 are also those that display higher level of segregation in the same year. This finding is robust to different definitions of the two variables and it holds for all decades in our sample. The relationship between the change in inequality and the change in segregation between 1980 and 2000 is not as strong but still statistically significant, meaning that the metros where inequality has grown the most are also those where segregation has increased the most.

B.2 Computing the Gini

The Gini coefficients in this paper were calculated following the method of von Hippel et. al. (2017). First, a non-parametric estimation of the income CDF was calculated for each metropolitan area. The non-parameteric CDF was calculated using the function `binsmooth`, provided by von Hippel et. al. in R. This function linearly interpolates between the upper bounds of each

Table 10: *Regression Analysis*

	(1)	(2)
	Dissimilarity 1980	Change Dissimilarity: 1980 to 2000
Gini 1980	1.367*** (0.095)	
Change Gini: 1980 to 2000		0.576*** (0.087)
_cons	-0.194*** (0.035)	0.010* (0.005)
N	378	378
r2	.3543	.1047
F	206.3	43.97

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

income bracket to calculate the CDF, preserving the empirical cumulative distribution for each bin. It then uses the empirical mean income to calculate the implied upper-bound for the support of the PDF, choosing the upper-bound and scale parameter so that the mean of the estimated CDF matches the empirical mean. Three methods are proposed to characterize the distribution of the top bracket: linear, Pareto, and exponential. The default method is linear and is what is used here. The binsmooth function returns a non-parameteric CDF function which can be used to calculate the Gini coefficient (and the conditional mean income of the top-coded bracket). Define:

$$\mu = \int xf(x)dx$$

Then the Gini coefficient is calculated as:

$$G = 1 - \frac{1}{\mu} \int_0^E (1 - F(x))^2 dx$$

These integrals must be calculated numerically, however because the CDF is piecewise linear, there is little approximation error. Importantly, the μ from the non-parametric CDF matches the empirical mean. After a Gini coefficient is calculated for each MSA, the weighted average of these coefficients is taken, using the count of family units in the MSA as weights.

B.3 Computing the Income Ratios

To calculate the various income ratios, the income brackets at the tract level were collapsed to the MSA level. The income brackets were then sorted by income level (by year and MSA) and the cumulative distribution of persons within each income bracket was calculated. Next, we find the income bracket associated with the bottom 10% of the population and the top 10% of the population and calculate the ratio of these two incomes. Because the income distribution is discrete, the exact cut-offs cannot be calculated. To deal with this, the cut-offs were defined to be the income level which had the minimum distance to the relevant threshold.

B.4 Intergenerational Mobility and Segregation Across US Metros

Computing initial and final income Data are from the restricted geocoded version of the NLSY. We keep only observations with observed family income for any of the 3 years 1978, 1979, 1980 and also for any of the 5 years 1997, 1999, 2001, 2003, 2005. Income recorded in year y refers to income in year $y - 1$, so these are for survey years 1979-1981 and 1998-2006 respectively. We deflate all income observations to 1978 dollars using the CPI series for all consumers (U.S. city average) downloaded from the BLS.

Summary Statistics NLSY sample In Table 11 we report the number of observations as well as the number of metros in each segregation quartile.

Table 11: *Geographic Units and Population per Segregation Quartile*

quartile	countmetro_quartile1980	countpop_quartile1980
1	40	768
2	40	966
3	40	1256
4	39	2347

Hypothesis Testing The transition matrices in Figure 7 compare intergenerational mobility across metros grouped by the dissimilarity index where the cutoffs is the top 20 percent of the

income distribution. Mobility is consistently lower in highly segregated metros compared to low segregated ones.

Table 12: *Hypothesis Testing*

Cutoff	Low	High	F statistic	Reject H_0 at 5% Level
p80	Below 50th p	Above 50th p	0.03	Yes
p80	Below 50th p	Above 75th p	0.033	Yes
p50	Below 50th p	Above 50th p	0.014	Yes
p50	Below 50th p	Above 75th p	0.016	Yes

Table 12 reports the significance levels for a Wald test, adjusted for survey weights.³⁹ Each row of the table reports the significance level of the test for the null hypothesis that the low segregation group has the same persistence of the first quartile of family income as the high segregation group. We define the low segregation group of metros as those in the bottom 50 percent according to the dissimilarity index. The first row shows results for the case in which we compare this group with the high segregation group defined as the top 50 percent and the second row shows results for the case in which we compare the same low segregation group with the high segregation group defined as the top 25 percent. Since the measured levels of persistence for the first quartile are always higher in the high segregation metro group, rejecting the null hypothesis means accepting the alternative hypothesis that the high segregation metros have a higher persistence of the first quartile of family income. We run the test for these two different definitions of "high" segregation metros and for two different cutoffs of the dissimilarity index. In all these instances we can reject the hypothesis that mobility is the same in the two group of metros at the 5 percent level.

Other Mobility Statistics We calculate three statistics by regressing child outcomes on parent outcomes, depending on the measure used to capture outcomes. For the levels regression, they are both in levels of earnings (corrected for inflation). For the IGE regression, they are both log levels of earnings, corrected for inflation. And for the rank-rank regression, parents and children

³⁹The null hypothesis H_0 is that the persistence of the first quartile of income is the same for high and low segregation metro groups.

Table 13: *Mobility Statistics*

Group	Levels	IGE	Rank-rank
All	3.315	0.493	0.413
High Seg (top 50 p)	3.38	0.52	0.44
Low Seg (bottom 50 p)	3.18	0.43	0.35

are both sorted by earnings within their cohort, assigned a ranking, and then the ranking of the children is regressed on the ranking of the parent.